

Improving Long Context Document-Level Machine Translation

Christian Herold Hermann Ney

Human Language Technology and Pattern Recognition Group
Computer Science Department
RWTH Aachen University
D-52056 Aachen, Germany
{herold|ney}@cs.rwth-aachen.de

Abstract

Document-level context for neural machine translation (NMT) is crucial to improve the translation consistency and cohesion, the translation of ambiguous inputs, as well as several other linguistic phenomena. Many works have been published on the topic of document-level NMT, but most restrict the system to only local context, typically including just the one or two preceding sentences as additional information. This might be enough to resolve some ambiguous inputs, but it is probably not sufficient to capture some document-level information like the topic or style of a conversation. When increasing the context size beyond just the local context, there are two challenges: (i) the memory usage increases exponentially (ii) the translation performance starts to degrade. We argue that the widely-used attention mechanism is responsible for both issues. Therefore, we propose a constrained attention variant that focuses the attention on the most relevant parts of the sequence, while simultaneously reducing the memory consumption. For evaluation, we utilize targeted test sets in combination with novel evaluation techniques to analyze the translations in regards to specific discourse-related phenomena. We find that our approach is a good compromise between sentence-level NMT vs attending to the full context, especially in low resource scenarios.

1 Introduction

Machine translation (MT) is the task of mapping some input text onto the corresponding translation in the target language. MT systems typically operate on the sentence-level and utilize neural networks trained on large amounts of bilingual data (Bahdanau et al., 2014; Vaswani et al., 2017). These neural machine translation (NMT) systems perform remarkably well on many domains and language pairs, sometimes even on par with professional human translators. However, when the automatic translations are evalu-

ated on the document-level (e.g. the translation of a whole paragraph or conversation is evaluated), they reveal shortcomings regarding consistency in style, entity-translation or correct inference of the gender, among other things (Läubli et al., 2018; Müller et al., 2018; Thai et al., 2022). The goal of document-level NMT is to resolve these shortcomings by including context information as additional input when translating a sentence.

In recent years, many works have been published on the topic of document-level NMT. However, most of these works focus only on including a few surrounding sentences as context. When the context size is increased beyond that, typically a degradation of overall translation performance is reported. Additionally, the transformer architecture as the quasi standard in NMT seems sub optimal to handle long sequences as input/output, since the memory complexity increases quadratically with the sequence length. This is due to the attention mechanism, where each token in a sequence needs to attend to all other tokens.

In this work, we propose a constrained attention variant for the task of document-level NMT. The idea is to reduce the memory consumption while at the same time focusing the attention of the system onto the most relevant parts of the sequence. Our contributions are two-fold:

1. We observe that the attention patterns become less focused on the current sentence when increasing the context-size of our document-level NMT systems. Therefore we propose a constrained attention variant that is also more memory efficient.
2. We utilize a targeted evaluation method to assess automatic translations in regards to consistency in style and coreference resolution. We find that our document-level NMT approach performs among the best across all language-pairs and test scenarios.

2 Related Work

Many works have been published on the topic of document-level NMT. The widely used baseline approach consists of simply concatenating a few adjacent sentences and feeding this as an input to the MT system, without modifying the system architecture in any way (Tiedemann and Scherrer, 2017; Bawden et al., 2018; Agrawal et al., 2018; Talman et al., 2019; Nguyen et al., 2021; Majumde et al., 2022). Also, several modifications to this baseline concatenation approach have been proposed. Ma et al. (2020) introduce segment embeddings and also partially constrain the attention to the tokens of the current sentence. Zhang et al. (2020) propose to calculate the self-attention both on the sentence- and on the document-level and then combine the two representations. Fernandes et al. (2021) and Lei et al. (2022) both mask out tokens in the current sentence to increase context utilization while Yang et al. (2023) remove tokens from the context if they are not attended. Typically, slight improvements in BLEU are reported as well as more significant improvements on targeted test sets e.g. for coreference resolution.

Apart from the simple concatenation method, there exist other approaches to document-level NMT. They include using a single document-embedding vector (Macé and Servan, 2019; Stojanovski and Fraser, 2019; Jiang et al., 2020; Huo et al., 2020), multiple encoders (Jean et al., 2017; Bawden et al., 2018; Wang et al., 2017; Voita et al., 2018; Zhang et al., 2018), hierarchical attention (Miculicich et al., 2018; Maruf et al., 2019; Tan et al., 2019; Wong et al., 2020), translation caches (Maruf and Haffari, 2018; Tu et al., 2018; Kuang et al., 2018) or dynamic evaluation (Mansimov et al., 2021). However, these approaches are less versatile and require significant changes to the model architecture, often introducing a significant amount of additional parameters. Furthermore, recent works have concluded that the baseline concatenation approach first proposed by Tiedemann and Scherrer (2017) performs as good - if not better - than these more complicated approaches (Lopes et al., 2020; Sun et al., 2022).

While the concatenation approach works well for short context sizes, when used with a larger number of context sentences, typically performance degradation is reported: Scherrer et al. (2019) saw a severe performance degradation when using input sequences with a length of 250 tokens. Liu et al.

(2020) could not get their system to converge when using context sizes of up to 512 tokens. They improve training stability by adding additional monolingual data via pre-training. Bao et al. (2021) also report that their systems with context length of more than 256 tokens fail to converge. They propose to partially constrain the attention to the current sentence, similar to Zhang et al. (2020). Sun et al. (2022) try to translate full documents with the concatenation approach but could not get their system to converge during training. Their solution is to mix document- and sentence-level data, which reportedly improves system convergence. Li et al. (2022) report severe performance degradation for context sizes longer than 512 tokens. They argue this is due to insufficient positional information and improve performance by repeatedly injecting this information during the encoding process. However, increasing the context size seems to not always result in performance degradation. In their works, Junczys-Dowmunt (2019) and Saleh et al. (2019) train systems with a context size of up to 1000 tokens without degradation in translation quality, which stands in contrast to the works mentioned above and which we will discuss again in the context of our own results. We want to point out that all of the approaches mentioned above still have the problem of quadratically increasing resource requirements, which poses a big challenge even on modern hardware.

Since our proposed approach consists of modifying the attention matrix in the model architecture, we give a brief overview of previous works related to this concept. The works of Ma et al. (2020), Zhang et al. (2020) and Bao et al. (2021) are most closely related and were already mentioned above. All three papers restrict the attention (partially) to the current sentence and combine sentence- and document-level attention context vectors for the final output. However, this means all approaches still suffer from the quadratic dependency on the number of input tokens. Luong et al. (2015) were among the first to propose using the attention concept for the task of MT. They also proposed using a sliding-window with target-to-source alignment for attention similar to us. However, they only work on sentence-level NMT and to the best of our knowledge, this approach was never before transferred to document-level NMT. Shu and Nakayama (2017) and Chen et al. (2018) both extend the approach of Luong et al. (2015) while still working

solely on sentence-level NMT. Our approach is also related to the utilization of relative positional encoding, which was introduced by Shaw et al. (2018) and later extended by Yang et al. (2018) to be applicable for cross-attention. The work by Indurthi et al. (2019) should also be mentioned, where they pre-select a subset of source tokens on which to perform attention on. Again, all of the above mentioned works only perform experiments on sentence-level NMT. The works of Child et al. (2019), Sukhbaatar et al. (2019) and Guo et al. (2020) are also related, since they use attention windows similar to us for tasks other than MT.

Finally, we briefly want to touch on the subject of automatic evaluation of document-level MT systems. Many works only report results on general MT metrics like BLEU (Papineni et al., 2002), sometimes matching n-grams across sentence-boundaries. However, it has been argued that these metrics do not capture well the very specific improvements that could be expected by including document-level context and that the reported improvements rather come from regularization effects and comparing to sub optimal baseline performance (Kim et al., 2019; Li et al., 2020; Nguyen et al., 2021). Several targeted test suites have been released to better assess the improvements gained by document-level NMT (Müller et al., 2018; Bawden et al., 2018; Voita et al., 2019; Jwalapuram et al., 2019). These test suites have some limitations, for example they are language-specific and they are based on just scoring predefined contrastive examples without scoring the actual translations. More recently, Jiang et al. (2022) and Currey et al. (2022) have released frameworks that allow to score the actual MT hypotheses in regards to their consistency regarding specific aspects of the translation.

3 Methodology

Here, we explain the baseline concatenation approach (Section 3.1), the more refined method that we are comparing ourselves against (Section 3.2) as well as our own approach (Section 3.3). We also discuss our different evaluation approaches in Section 3.5.

3.1 The Baseline Concatenation Approach

The baseline concatenation approach is very simple and follows Tiedemann and Scherrer (2017) using the vanilla transformer architecture (Vaswani

Model	Context	Attn. [%]	BLEU
sent.-level	0 sent.	100.0	32.8
concat adj.	1 sent.	76.0	33.1
	1000 tok.	46.6	29.5

Table 1: Percentage of attention on the n -th source sentence during decoding the n -th target sentence, as well as overall translation quality measured in BLEU, for the newstest2018 test set of the NEWS task.

et al., 2017). Assume we are given a document $\mathcal{D} = (F_n, E_n)_1^N$ consisting of N source-target sentence-pairs (F_n, E_n) . If we want our model to have a context length of k sentences, we simply concatenate the current input sentence with its $k - 1$ predecessor sentences and the input to the model would be

$$F_{n-k} \langle \text{sep} \rangle F_{n-k+1} \dots \langle \text{sep} \rangle F_n \langle \text{eos} \rangle$$

while on the target side we include the preceding sentences as a prefix

$$E_{n-k} \langle \text{sep} \rangle E_{n-k+1} \dots \langle \text{sep} \rangle E_{n-1} \langle \text{sep} \rangle.$$

We use a special token $\langle \text{sep} \rangle$ as a separator between adjacent sentences and $\langle \text{eos} \rangle$ denotes the end of the sequence. This is done to make it easier for the model to distinguish between the sentence that needs to be translated and the context. Furthermore, we use a special token $F_0 = E_0 = \langle \text{bod} \rangle$ to denote the start of a document. Since we use the vanilla transformer architecture with self-attention and cross-attention components, the memory usage is $\mathcal{O}(L^2)$ with L being the sequence length.

When we train full document-level systems, we simply concatenate all sentences in the document using again the special $\langle \text{sep} \rangle$ token. Due to hardware limitations, if the length of the target-side of the document exceeds 1000 tokens, we split the document into smaller parts of roughly equal length (i.e. a document of length 1500 tokens would be split into two parts with ca. 750 tokens each).

In a preliminary study, we train systems using no context (sentence-level), just a single sentence as context as well as the maximum context size of 1000 tokens. When looking at the percentage of attention that is paid to the n -th source sentence F_n when decoding the n -th target sentence E_n (extracted from cross-attention module, see Table 1) we find that this percentage becomes lower as the context size increases. This finding motivates us to explore approaches that bias the attention towards the current sentence.

3.2 LST-attention

This method was proposed by Zhang et al. (2020) and is called Long-Short Term (LST) attention. The authors find that their approach outperforms the baseline concatenation approach but they only use a maximum of 3 sentences as context. Nevertheless we deem this approach promising, since it also focuses the attention onto the current sentence. The input to the system is augmented in the same way as described in Section 3.1. Given some queries $Q \in \mathbb{R}^{I \times d}$, keys $K \in \mathbb{R}^{J \times d}$ and values $V \in \mathbb{R}^{J \times d}$, Zhang et al. (2020) formulate their restricted version of the attention as¹

$$A(Q, K, V) = \text{softmax} \left(\frac{Q \cdot K^\top}{\sqrt{d}} + M \right) V$$

with d being the hidden dimension of the model and $M \in \mathbb{R}^{I \times J}$ being the masking matrix. This masking matrix is defined as

$$M_{i,j} = \begin{cases} 0 & , s(i) = s(j) \\ -\text{inf} & , \text{otherwise} \end{cases}$$

where $s(\cdot) \in 1, \dots, N$ is a function that returns the sentence index that a certain position belongs to. This means we are restricting the attention to be calculated only within the current sentence. For self-attention in the encoder and the decoder, Zhang et al. (2020) calculate both the restricted and the non-restricted variant and then combine the output context-vectors via concatenation and a linear transformation. The cross-attention between encoder and decoder remains unchanged in this approach and the memory consumption remains $\mathcal{O}(L^2)$.

3.3 window-attention

This method is proposed by us. We can use the same formulation as above to describe this approach by simply changing the definition of the attention mask to

$$M_{i,j} = \begin{cases} 0 & , b_i - w \leq j \leq b_i + w \\ -\text{inf} & , \text{otherwise} \end{cases} \quad (1)$$

where w is the window size and $b_i \in 1, \dots, J$ is a target-source alignment. This means a certain query vector q_i is only allowed to attend to the key vectors k_j that surround the position b_i that

¹In practice we use multi-head attention in all our architectures, but we omit this in the formulas for sake of simplicity. Also, for all methods, causal masking is applied in the decoder self-attention just like in the vanilla transformer.

this query vector is aligned to. We replace all self-attention and cross-attention modules in our network with this window-attention variant. Please note that in practice we do not calculate this mask, but instead we first select the corresponding key-vectors for each query and then calculate attention only between these subsets which reduces the memory consumption from $\mathcal{O}(L^2)$ to $\mathcal{O}(L \cdot w)$. We also want to point out that with this approach, the context is not as restricted as it seems on first glance. For any individual attention module, the context is restricted to $2 \cdot w$ or w for self-attention in the encoder and decoder respectively. However, since in the transformer architecture we stack multiple layers, we get a final effective context size of $2 \cdot w \cdot \text{num_enc_layers} + w \cdot \text{num_dec_layers}$.

This approach requires us to define an alignment function $b_i : [1, I] \rightarrow [1, J]$. For self-attention, we assume a 1-1 alignment so the alignment function is the identity function $b_i = i$. For cross-attention, during training we use a linear alignment function

$$b_i = \text{round} \left(\frac{J}{I} \cdot i \right)$$

where J is the number of tokens in the source document and I is the number of tokens in the target document. This is not possible during decoding, as we do not know the target document length beforehand. Therefore, we propose three different ways to approximate the alignment during decoding:

1. 1-1 alignment: $b_i = i$
2. linear alignment: $b_i = \text{round}(\text{train_ratio} \cdot i)$ where we define *train_ratio* as the average source-target ratio over all documents in the training data.
3. sent-align: assume we have already produced N' full target sentences (i.e. we have produced $N' \langle \text{sep} \rangle$ tokens) up to this point, then

$$b_i = \begin{cases} \sum_{n=1}^{N'} J_n + 1 & , e_{i-1} == \langle \text{sep} \rangle \\ b_{i-1} + 1 & , \text{otherwise} \end{cases}$$

with J_n being the length of the n -th source sentence in the input document. In simple terms, when starting to decode a new sentence, we always force-align to the beginning of the corresponding source sentence.

We also test the *window-attention* approach with relative positional encoding in the self-attention

instead of absolute positional encoding, which in this framework only requires a small modification to Equation 1:

$$M_{i,j} = \begin{cases} r_{i-j} & , b_i - w \leq j \leq b_i + w \\ -\text{inf} & , \text{otherwise} \end{cases}$$

where $r_{i-j} \in \mathbb{R}^d$ are additional learnable parameters of the network.

3.4 Decoding

During decoding, given a document F_1^N , we want to find the best translation \hat{E}_1^N according to our model. We can not perform exact search due to computational limitations, therefore we have to use approximations. There exist multiple approaches for decoding with a document-level NMT system and since we could not determine a single best approach from the literature, we describe and compare two competing approaches.

Full Segment Decoding (FSD) (Liu et al., 2020; Bao et al., 2021; Sun et al., 2022): we split the document into non-overlapping parts $F_1^k, F_{k+1}^{2k}, \dots, F_{N-k}^N$ and translate each part separately using

$$\hat{E}_{i-k}^i = \operatorname{argmax}_{E_{i-k}^i} \{p(E_{i-k}^i | F_{i-k}^i)\},$$

which is approximated using standard beam search on the token level (we use beam size 12 for all experiments). For the full document-level systems, we simply use

$$\hat{E}_1^N = \operatorname{argmax}_{E_1^N} \{p(E_1^N | F_1^N)\}.$$

Sequential Decoding (SD) (Miculicich et al., 2018; Voita et al., 2019; Garcia et al., 2019; Fernandes et al., 2021): we generate the translation sentence by sentence, using the previously generated target sentences as context:

$$\hat{E}_i = \operatorname{argmax}_{E_i} \left\{ p(E_i | \hat{E}_{i-k}^{i-1}, F_{i-k}^i) \right\}.$$

3.5 Evaluation

For all tasks we report BLEU (Papineni et al., 2002) and TER (Snover et al., 2006) using the SacreBLEU (Post, 2018) toolkit. In addition, for the two En-De tasks (*NEWS* and *OS*) we analyze the translations in regards to ambiguous pronouns and style. For

pronouns, the goal is to measure how well a system can translate the English 3rd person pronoun ‘it’ (and its other forms) into the correctly gendered German form (which can be male, female or neuter depending on the context). For style, the goal is to measure, how well a system can translate the 2nd person pronoun ‘you’ (and its other forms) into the correct style in German. For example, ‘you’ (singular) can be translated into ‘sie’ or ‘du’ in German, depending if the setting is formal or informal. We employ several strategies to determine the systems ability to disambiguate these phenomena.

We utilize the ContraPro test suite (Müller et al., 2018) and report the contrastive scoring accuracy for pronoun resolution. The test suite contains 12,000 English sentences, each with the correct German reference as well as 2 contrastive German references where the pronoun has been changed to a wrong gender. We score all test cases with the NMT system and each time the system gives the best score to the true reference, it gets a point. In the end we report the scoring accuracy, i.e. the number of points the system has gathered divided by 12,000.

Additionally we also report F1 scores for pronoun and style translation, the method for which is inspired by Jiang et al. (2022). We use parts of speech (POS) taggers as well as language-specific regular expressions to identify ambiguous pronouns/formality in the test sets. We then compare the occurrences in the reference against the occurrences in the hypothesis, calculate precision and recall and then finally the F1 score for both pronoun as well as formality translation. The exact algorithm as well as detailed data statistics for the test sets are given in Appendix A.1. We report pronoun translation F1 score for both *NEWS* and *OS* tasks and the formality translation F1 score only for the *OS* task, since in the *NEWS* test set there are not enough examples of ambiguous formality cases. Our extension to the work of Jiang et al. (2022) can be found here: <https://github.com/christian3141/BlonDe>.

4 Experiments

We perform experiments on three document-level translation benchmarks. We call them **NEWS** (En→De) with newstest2018 as test set, **TED** (En→It) with tst2017 as test set and **OS** (En→De) where the test set is simply called test. *NEWS* is a collection of news articles, *TED* is a collection of

transcribed TED talks and their respective translations and *OS* consists of subtitles for movies and TV shows. Especially the latter holds many examples for discourse between different entities. For the details regarding data conditions, preparation and training, we refer to Appendix A.2.

4.1 GPU Memory efficiency

First, we compare the GPU memory consumption of the baseline *concat-adj.* approach against the *window-attention* approach for various input sequence lengths. The results are shown in Table 2. As expected, the memory usage increases at a much

# target tokens	concat-adj.	window-attn	
		$w = 10$	$w = 20$
736	2.3 GB	2.4 GB	3.5 GB
1472	5.8 GB	3.9 GB	5.9 GB
2208	10.9 GB	5.2 GB	8.5 GB

Table 2: GPU-memory consumption for the different approaches when training on a single document of specified number of target tokens.

higher rate for the *concat-adj.* approach, while the *window-attention* approach scales roughly linearly, the slope being a function of the window-size w .

4.2 Comparison of Decoding Strategies

After training all models on the *NEWS* task according to Appendix A.2, we test the different search strategies for each of the systems, the result of which can be found in Table 3. For the baseline *concat-adj.* approach as well as the *LST-attn* approach, *FSD* works best. However, we still see significant performance degradation for the systems

Model	Context	Search Strategy	BLEU
sent.-level	0 sent.	-	32.8
concat-adj.	2 sent.	<i>FSD</i>	33.4
		<i>SD</i>	33.0
	1000 tok.	<i>FSD</i>	29.5
		<i>SD</i>	23.1
LST-attn	1000 tok.	<i>FSD</i>	30.0
		<i>SD</i>	22.2
window-attn	1000 tok.	<i>FSD</i>	31.5
		<i>SD</i>	33.1

Table 3: Results for employing the different search strategies for translating the newstest2018 test set of the *NEWS* task.

using long context information. For *concat-adj.* and *LST-attn* with 1000 tokens context size, *SD* performs very poorly. This is because when beginning translating a document, the input sequences are very short and the systems can not appropriately handle that. However, *FSD* sometimes leads to sentence-misalignment while translating a document, resulting in a lower BLEU score as well. For the *window-attention* approach (rel. pos. enc., sent-align, window-size 20) we find that the *SD* decoding strategy works best. Since this approach seems to be able to better handle short input sequences, *SD* performs better than *FSD*, since it seems more robust to sentence-misalignment. Moving forward, all numbers reported will be generated with the best respective decoding approach, i.e. *SD* for *window-attention* and *FSD* for all other approaches.

4.3 Hyperparameter Tuning

Our *window-attention* approach has three hyperparameters that need to be tuned: (i) positional encoding variant (ii) alignment variant during search (iii) window size. Again, we use the *NEWS* task for tuning and the results for the different variants can be found in Table 4.

In terms of positional encoding, *relative* works significantly better than *absolute* for the *window-attention* system. We also test relative positional encoding (window-size 20) for the baseline *concat-adj.* method, but here the training did not converge. This is, because for long input sequences the system without explicit target-source alignment can no longer distinguish the token ordering on the source side (on the target side it is still possible due to the causal attention mask). The only way to resolve this would be to drastically increase the window-size for the relative positions, however, this would add a significant amount of additional parameters to the network so we decide against this. In terms of alignment, using the *sent-align* variant significantly outperforms the other approaches. For the *window-size*, 20 works best. An important finding is, that if we make the window too large, we start losing performance, probably due to the less focused attention problem discussed in Section 3.1.

4.4 Final Performance Comparison

In Table 5 we report the translation performance of the different document-level approaches on all three translation benchmarks measured in terms of BLEU and TER. None of the document-level systems can consistently outperform the sentence-

Model	pos. enc.	Alignment	window-size	BLEU	TER
concat-adj.	abs.	-	-	29.5	53.7
	rel.	-	-	N/A	N/A
window-attn	abs.	1-1	20	29.7	51.7
		train avg.	20	28.1	55.3
		sent-align	10	28.3	53.7
			20	30.3	50.9
			30	29.4	52.2
	rel.	1-1	20	31.9	49.8
		train avg.	20	30.5	53.2
		sent-align	10	30.6	51.8
			20	33.1	48.1
			30	32.8	48.4

Table 4: Results for the different hyperparameter settings of the *window-attention* system reported on the newstest2018 test set of the *NEWS* task. All systems have context size 1000 tokens.

Model	Context	NEWS		TED		OS	
		newstest2018		tst2017		test	
		BLEU	TER	BLEU	TER	BLEU	TER
sent.-level (external)	0 sent.	[†] 32.3	-	[‡] 33.4	-	*37.3	-
sent.-level (ours)		32.8	49.0	34.2	46.3	37.1	43.8
concat adj.	2 sent.	33.4	48.6	34.3	46.3	38.2	43.9
	1000 tok.	29.5	53.7	32.1	48.4	38.1	46.0
LST-attn	1000 tok.	30.0	53.1	29.8	54.5	38.5	45.1
window-attn	1000 tok.	33.1	48.1	34.6	45.8	38.3	44.4

Table 5: Results for the different document-level approaches in terms of BLEU and TER on the three translation benchmarks. Best results for each column are highlighted. External baselines are from [†] Kim et al. (2019), [‡] Yang et al. (2022) and *Huo et al. (2020).

level baseline on all tasks. On the *OS* test set, there is a disagreement between BLEU and TER which we think comes from the fact that the average sentence-length on this test set is quite short. The hypothesis of the sentence-level system is the shortest of all hypotheses and also shorter than the reference which gets punished more heavily by BLEU than TER. Out of all full-document approaches, *window-attention* performs best and is on par with the sentence-level baseline and the document-level system using only 2 sentences as context. For full-document translation, *LST-attn* performs better than the baseline concatenation approach but still falls behind the sentence-level system especially on the *NEWS* and *TED* tasks. One possible reason for why these approaches work better on *OS* is, that for this task we have much more training data available than for *NEWS* and *TED*. We argue that this could also be the reason for the conflicting results reported by Junczys-Dowmunt

(2019) and Saleh et al. (2019) compared to the other works who report performance degradation for longer context sizes (see Section 2). However, we leave a detailed analysis of this for future work.

Next, we analyze the ability of the systems to translate ambiguous pronouns and to translate in a consistent style using the methods explained in Section 3.5. The results for the two En→De tasks can be found in Table 6. For both *NEWS* and *OS*, all document-level systems can significantly improve over the sentence-level baseline in terms of pronoun translation. We also find that a context longer than two sentences does not seem to help for the pronoun task. This is actually to be expected since typically the distance between noun and pronoun is not that large and according to Müller et al. (2018), the overwhelming majority of ContraPro test cases do not require more than two sentences as context. For the correct translation of the style however, the larger context size is clearly beneficial,

Model	Context	NEWS			OS			
		ContraPro			ContraPro			test
		BLEU	Scoring Acc. Pronoun	Pronoun Trans. F1	BLEU	Scoring Acc. Pronoun	Pronoun Trans. F1	Formality Trans. F1
sent.-level	0 sent.	18.4	48.2	44.5	29.7	45.8	40.3	59.4
concat adj.	2 sent.	19.6	67.9	54.1	31.2	81.8	63.2	61.7
	1000 tok.	15.4	61.9	47.8	29.9	83.1	64.6	70.1
LST-attn	1000 tok.	16.8	61.4	51.3	29.1	83.3	64.8	70.9
window-attn	1000 tok.	19.6	63.0	51.9	31.4	83.9	66.5	67.9

Table 6: Results for the different document-level approaches in terms of pronoun and formality translation. Best results for each column are highlighted.

source	reference
What’s between you and Dr. Webber - is none of my business...	Was zwischen dir und Dr. Webber ist, geht mich nichts an...
- You don’t owe me an apology.	Du schuldest mir keine Entschuldigung.
You owe Dr. Bailey one.	Du schuldest Dr. Bailey eine.
We were taking a stand for Dr. Webber.	Wir haben uns für Dr. Webber eingesetzt.
I don’t understand why...	Ich verstehe nicht wieso...
Dr. Webber doesn’t need you to fight his battles.	Dr. Webber braucht dich nicht, um seine Schlachten zu kämpfen.
What you did stands to hurt this entire hospital.	Was du getan hast, hat dem ganzen Krankenhaus geschadet.
Your first priority needs to be this place and its patients.	Deine oberste Priorität muss diesem Haus und seinen Patienten gelten.
sentence-level-hypothesis	window-mask-hypothesis
Was zwischen Ihnen und Dr. Webber ist, geht mich nichts an...	Was zwischen dir und Dr. Webber ist, geht mich nichts an...
- Du schuldest mir keine Entschuldigung.	- Du schuldest mir keine Entschuldigung.
Sie schulden Dr. Bailey etwas.	- Du schuldest Dr. Bailey eine.
Wir haben für Dr. Webber Partei ergriffen.	Wir haben für Dr. Webber Stellung bezogen.
Ich verstehe nicht, warum...	Ich verstehe nicht, warum...
Dr. Webber braucht Sie nicht, um seine Schlachten zu schlagen.	Dr. Webber braucht dich nicht, um seine Schlachten zu kämpfen.
Was du getan hast, verletzt das gesamte Krankenhaus.	Was du getan hast, könnte das ganze Krankenhaus verletzen.
Ihre oberste Priorität muss dieser Ort und seine Patienten sein.	Deine oberste Priorität muss dieser Ort und seine Patienten sein.

Table 7: Example translation of a snippet from the OpenSubtitles test set. Formal 2nd person pronouns are marked in red and informal ones are marked in blue.

as the system with just 2 sentences as context can barely outperform the sentence-level baseline. To correctly infer the style of a conversation, ideally the whole dialog should be part of the context, especially the beginning of the conversation. In Table 7, we show a snippet of the test set of the OS task together with the translations of the sentence-level system and the window-attention system. This example highlights the need for long-context NMT systems especially for the task of dialogue translation, since there we need to stay consistent in terms of style, which the sentence-level system can not manage. Overall, the LST-attn approach performs best for the task of formality translation, but the other full-document systems are not far behind.

5 Conclusion

In this work, we focus on methods to increase the context-size for document-level NMT systems.

We point out the shortcomings of the baseline approaches to long-context document-level NMT and in turn propose to modify the attention component to be more focused and also to be more memory efficient. We compare our approach against approaches from literature on multiple translation tasks and using different targeted evaluation methods. We confirm the improved memory efficiency of the proposed method. We find that for some discourse phenomena like pronoun translation, the longer context information is not necessary. For other aspects, like consistent style translation, the longer context is very beneficial. It seems that the baseline concatenation approach needs large amounts of training data to perform well for larger context sizes. We conclude that our approach performs among the best across all tasks and evaluation methods, with the additional benefit of reduced memory consumption for long input sequences.

Acknowledgements

This work was partially supported by the project HYKIST funded by the German Federal Ministry of Health on the basis of a decision of the German Federal Parliament (Bundestag) under funding ID ZMVI1-2520DAT04A, and by NeuroSys which, as part of the initiative “Clusters4Future”, is funded by the Federal Ministry of Education and Research BMBF (03ZU1106DA).

Limitations

This work is about document-level NMT, we focus specifically on methods that improve the model performance for long input sequences. Due to constrained resources, this work has several limitations. To be able to train all methods including the inefficient baseline approach, we have to limit the context size to 1000 tokens. While we do a comparison to existing approaches, other approaches have been proposed to improve the performance of systems with long context information, which we do not compare against. We run experiments on three different tasks, but two of them are low resource and two of them translate into German, which was necessary because we only had access to German language experts for preparing the evaluation.

References

- Ruchit Rajeshkumar Agrawal, Marco Turchi, and Matteo Negri. 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.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Guangsheng Bao, Yue Zhang, Zhiyang Teng, Boxing Chen, and Weihua Luo. 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.
- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 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.
- Mauro Cettolo, Marcello Federico, Luisa Bentivogli, Jan Niehues, Sebastian Stüker, Katsutho Sudoh, Koichiro Yoshino, and Christian Federmann. 2017. Overview of the iwslt 2017 evaluation campaign. In *Proceedings of the 14th International Workshop on Spoken Language Translation*, pages 2–14.
- Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2018. [Syntax-directed attention for neural machine translation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers.
- Anna Currey, Maria Nädejde, Raghavendra Pappagari, Mia Mayer, Stanislas Lauly, Xing Niu, Benjamin Hsu, and Georgiana Dinu. 2022. Mt-geneval: A counterfactual and contextual dataset for evaluating gender accuracy in machine translation. *arXiv preprint arXiv:2211.01355*.
- Akhbardeh Farhad, Arkhangorodsky Arkady, Biesialska Magdalena, Bojar Ondřej, Chatterjee Rajen, Chaudhary Vishrav, Marta R Costa-jussa, España-Bonet Cristina, Fan Angela, Federmann Christian, et al. 2021. Findings of the 2021 conference on machine translation (wmt21). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88. Association for Computational Linguistics.
- Patrick Fernandes, Kayo Yin, Graham Neubig, and André FT Martins. 2021. Measuring and increasing context usage in context-aware 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 6467–6478.
- Eva Martínez Garcia, Carles Creus, and Cristina España-Bonet. 2019. Context-aware neural machine translation decoding. In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 13–23.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Xiangyang Xue, and Zheng Zhang. 2020. [Multi-scale self-attention for text classification](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7847–7854.
- 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 Fifth Conference on Machine Translation*, pages 604–616.
- Sathish Reddy Indurthi, Insoo Chung, and Sangha Kim. 2019. [Look harder: A neural machine translation model with hard attention](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3037–3043, Florence, Italy. Association for Computational Linguistics.

- Sebastien Jean, Stanislas Lauly, Orhan Firat, and Kyunghyun Cho. 2017. Does neural machine translation benefit from larger context? *arXiv preprint arXiv:1704.05135*.
- Shu Jiang, Hai Zhao, Zuchao Li, and Bao-Liang Lu. 2020. Document-level neural machine translation with document embeddings.
- Yuchen Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Sennrich, Ryan Cotterell, Mrinmaya Sachan, and Ming Zhou. 2022. **BlonDe: An automatic evaluation metric for document-level machine translation**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1550–1565, Seattle, United States. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt. 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.
- Prathyusha Jwalapuram, Shafiq Joty, Irina Temnikova, and Preslav Nakov. 2019. Evaluating pronominal anaphora in machine translation: An evaluation measure and a test suite. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2964–2975.
- Yunsu Kim, Duc Thanh Tran, and Hermann Ney. 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.
- Shaohui Kuang, Deyi Xiong, Weihua Luo, and Guodong Zhou. 2018. **Modeling coherence for neural machine translation with dynamic and topic caches**. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 596–606, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Taku Kudo. 2018. **Subword regularization: Improving neural network translation models with multiple subword candidates**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 66–75. Association for Computational Linguistics.
- Samuel Läubli, Rico Sennrich, and Martin Volk. 2018. Has machine translation achieved human parity? a case for document-level evaluation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4791–4796.
- Yikun Lei, Yuqi Ren, and Deyi Xiong. 2022. **CoDoNMT: Modeling cohesion devices for document-level neural machine translation**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5205–5216, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- 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. *arXiv preprint arXiv:2005.03393*.
- Yachao Li, Junhui Li, Jing Jiang, Shimin Tao, Hao Yang, and Min Zhang. 2022. P-transformer: Towards better document-to-document neural machine translation.
- Pierre Lison, Jörg Tiedemann, and Milen Kouylekov. 2018. Opensubtitles2018: Statistical rescoring of sentence alignments in large, noisy parallel corpora. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. **Multilingual Denoising Pre-training for Neural Machine Translation**. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- António V Lopes, M Amin Farajian, Rachel Bawden, Michael Zhang, and André FT Martins. 2020. Document-level neural mt: A systematic comparison. In *22nd Annual Conference of the European Association for Machine Translation*, pages 225–234.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. **Effective approaches to attention-based neural machine translation**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Shuming Ma, Dongdong Zhang, and Ming Zhou. 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.
- Valentin Macé and Christophe Servan. 2019. **Using whole document context in neural machine translation**. In *Proceedings of the 16th International Conference on Spoken Language Translation*, Hong Kong. Association for Computational Linguistics.
- Suvodeep Majumde, Stanislas Lauly, Maria Nadejde, Marcello Federico, and Georgiana Dinu. 2022. A baseline revisited: Pushing the limits of multi-segment models for context-aware translation. *arXiv preprint arXiv:2210.10906*.
- Elman Mansimov, Gábor Melis, and Lei Yu. 2021. **Capturing document context inside sentence-level neural**

- machine translation models with self-training. In *Proceedings of the 2nd Workshop on Computational Approaches to Discourse*, pages 143–153, Punta Cana, Dominican Republic and Online. Association for Computational Linguistics.
- Sameen Maruf and Gholamreza Haffari. 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.
- Sameen Maruf, André FT Martins, and Gholamreza Haffari. 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.
- Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. Document-level neural machine translation with hierarchical attention networks. *arXiv preprint arXiv:1809.01576*.
- Mathias Müller, Annette Rios Gonzales, Elena Voita, and Rico Sennrich. 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.
- Toan Q. Nguyen, Kenton Murray, and David Chiang. 2021. Data augmentation by concatenation for low-resource translation: A mystery and a solution. In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 287–293, Bangkok, Thailand (online). Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Matt Post. 2018. A call for clarity in reporting bleu scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191.
- Fahimeh Saleh, Alexandre Berard, Ioan Calapodescu, and Laurent Besacier. 2019. Naver labs Europe’s systems for the document-level generation and translation task at WNGT 2019. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 273–279, Hong Kong. Association for Computational Linguistics.
- Yves Scherrer, Jörg Tiedemann, and Sharid Loáigiga. 2019. Analysing concatenation approaches to document-level NMT in two different domains. In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 51–61, Hong Kong, China. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 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.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Raphael Shu and Hideki Nakayama. 2017. An empirical study of adequate vision span for attention-based neural machine translation. page 1.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231.
- Dario Stojanovski and Alexander Fraser. 2019. Combining local and document-level context: The LMU Munich neural machine translation system at WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 400–406, Florence, Italy. Association for Computational Linguistics.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. 2019. Adaptive attention span in transformers. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 331–335, Florence, Italy. Association for Computational Linguistics.
- Zewei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Lei Li. 2022. Rethinking document-level neural machine translation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3537–3548.
- Aarne Talman, Umut Sulubacak, Raúl Vázquez, Yves Scherrer, Sami Virpioja, Alessandro Raganato, Arvi Hurskainen, and Jörg Tiedemann. 2019. The University of Helsinki submissions to the WMT19 news translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 412–423, Florence, Italy. Association for Computational Linguistics.

- Xin Tan, Longyin Zhang, Deyi Xiong, and Guodong Zhou. 2019. [Hierarchical modeling of global context for document-level neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1576–1585, Hong Kong, China. Association for Computational Linguistics.
- Katherine Thai, Marzena Karpinska, Kalpesh Krishna, Bill Ray, Moira Inghilleri, John Wieting, and Mohit Iyyer. 2022. Exploring document-level literary machine translation with parallel paragraphs from world literature. *arXiv preprint arXiv:2210.14250*.
- Jörg Tiedemann and Yves Scherrer. 2017. Neural machine translation with extended context. In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 82–92.
- Zhaopeng Tu, Yang Liu, Shuming Shi, and Tong Zhang. 2018. Learning to remember translation history with a continuous cache. *Transactions of the Association for Computational Linguistics*, 6:407–420.
- 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 *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. [When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion](#).
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. Context-aware neural machine translation learns anaphora resolution. *arXiv preprint arXiv:1805.10163*.
- Longyue Wang, Zhaopeng Tu, Andy Way, and Qun Liu. 2017. Exploiting cross-sentence context for neural machine translation. *arXiv preprint arXiv:1704.04347*.
- KayYen Wong, Sameen Maruf, and Gholamreza Haffari. 2020. Contextual neural machine translation improves translation of cataphoric pronouns. *arXiv preprint arXiv:2004.09894*.
- Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. 2018. [Modeling localness for self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4449–4458, Brussels, Belgium. Association for Computational Linguistics.
- Jian Yang, Yuwei Yin, Shuming Ma, Liqun Yang, Hongcheng Guo, Haoyang Huang, Dongdong Zhang, Yutao Zeng, Zhoujun Li, and Furu Wei. 2023. Hanoi: Enhancing context-aware translation via selective context.
- Jian Yang, Yuwei Yin, Liqun Yang, Shuming Ma, Haoyang Huang, Dongdong Zhang, Furu Wei, and Zhoujun Li. 2022. Gtrans: Grouping and fusing transformer layers for neural machine translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 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.
- Pei Zhang, Boxing Chen, Niyu Ge, and Kai Fan. 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.

A Appendix

A.1 Pronoun and Formality Translation Evaluation

Here, we explain how we calculate the pronoun translation and formality translation F1 scores.

Pronouns

For each triplet (F_n, E_n, \hat{E}_n) (source, hypothesis, reference) of our test data we first check if it contains a valid ambiguous pronoun. That means, in the source sentence there must be an English 3rd person pronoun in the neutral form and it also must be labeled as a pronoun by the English POS-tagger. We also check if a 2nd or 3rd person plural pronoun is present in the source and if that is the case, we do not consider female pronouns on the target side, since we could not distinguish if e.g. ‘sie’ is the translation of ‘it’ or ‘they’. This would require a word alignment between source and hypothesis/reference which we do not have. If we found the example to be valid, we then check for occurrences of 3rd person pronouns in the male, female and neuter forms, in both reference and hypothesis using a German POS-tagger as well as language-specific regular expressions. After going through the complete test data (F_n, E_n, \hat{E}_n) sentence-by-sentence we calculate an F1 score for pronoun translation:

$$F1_{pro} = \frac{2 \cdot P_{pro} \cdot R_{pro}}{P_{pro} + R_{pro}}$$

with precision $P_{pro} =$

$$\frac{\sum_{n=1}^N \sum_x \min(\text{CP}(F_n, E_n, x), \text{CP}(F_n, \hat{E}_n, x))}{\sum_n \sum_x \text{CP}(F_n, E_n, x)}$$

and recall $R_{pro} =$

$$\frac{\sum_{n=1}^N \sum_x \min(\text{CP}(F_n, E_n, x), \text{CP}(F_n, \hat{E}_n, x))}{\sum_n \sum_x \text{CP}(F_n, \hat{E}_n, x)}$$

where $\text{CP}(\cdot, \cdot, \cdot)$ counts the number of valid pronoun occurrences and $x \in \{\text{male}, \text{female}, \text{neuter}\}$.

Formality

We follow almost exactly the same steps as for detecting the pronoun translations described above. The only differences are that we check for validity slightly differently and instead of pronouns we check for occurrences of formal/informal style. For sentence-pairs where 3rd person female/neuter or

3rd person plural pronouns are present, we do not count the formal occurrences, since we might not be able to distinguish the German translations in these cases. We calculate an F1 score for formality translation using

$$F1_{for} = \frac{2 \cdot P_{for} \cdot R_{for}}{P_{for} + R_{for}}$$

with precision $P_{for} =$

$$\frac{\sum_{n=1}^N \sum_x \min(\text{CP}(F_n, E_n, x), \text{CP}(F_n, \hat{E}_n, x))}{\sum_n \sum_x \text{CP}(F_n, E_n, x)}$$

and recall $R_{for} =$

$$\frac{\sum_{n=1}^N \sum_x \min(\text{CP}(F_n, E_n, x), \text{CP}(F_n, \hat{E}_n, x))}{\sum_n \sum_x \text{CP}(F_n, \hat{E}_n, x)}$$

where $\text{CP}(\cdot, \cdot, \cdot)$ counts the number of valid pronoun occurrences and $x \in \{\text{formal}, \text{informal}\}$.

The POS-taggers we use are `en_core_web_sm`² for English and `de_core_news_sm`³ for German. For both languages, spaCy claims an accuracy of 97% for POS-tagging and in our testing we did not find even a single error in pronoun-tagging. For calculating the Pronoun Translation F1 score we use the same ContraPro test set as described in Section 3.5 with the correct references. For calculating the Formality Translation F1 score, we use the test set from the OS En-De task. The statistics for both test sets are reported in Table 8. In the ContraPro test set, for each gender class we have exactly 4,000 examples. The fact that we identify more than 4,000 valid examples for the pronoun case means, that in some cases we identify multiple pronouns per sentence. All in all, we find the classes to be relatively balanced for these test sets.

A.2 Dataset Statistics and Experimental Setups

For the **NEWS En→De** task, the parallel training data comes from the NewsCommentaryV14 corpus⁴. As validation/test set we use the WMT newstest2015/newstest2018 test sets from the WMT news translation tasks (Farhad et al., 2021). For the **TED En→It** task, the parallel training data comes from the IWSLT17 Multilingual Task (Cettolo et al., 2017). As validation set we

²<https://spacy.io/models/en>

³<https://spacy.io/models/de>

⁴<https://data.statmt.org/news-commentary/v14/>

	Pronoun Trans. F1 score			Formality Trans. F1 score	
	neuter	male	female	formal	informal
# examples	4565	4688	4001	416	605

Table 8: Number of valid examples for specific ambiguous pronoun/style translation in the reference of our test sets.

use the concatenation of IWSLT17.TED.dev2010 and IWSLT17.TED.tst2010 and as test set we use IWSLT17.TED.tst2017.mltlng. For the **OS En→De** task, the parallel training data comes from the OpenSubtitlesV2018 corpus (Lison et al., 2018). We use the same train/validation/test splits as Huo et al. (2020) and additionally remove all segments that are used in the ContraPro test suite (Müller et al., 2018) from the training data. The data statistics for all tasks can be found in Table 9.

task	dataset	# sent.	# doc.
NEWS	train	330k	8.5k
	valid	2.2k	81
	test	3k	122
	ContraPro	12k	12k
TED	train	232k	1.9k
	valid	2.5k	19
	test	1.1k	10
OS	train	22.5M	29.9k
	valid	3.5k	5
	test	3.8k	5
	ContraPro	12k	12k

Table 9: Data statistics for the different document-level translation tasks.

Since in the original release of ContraPro only left side context is provided, we extract the right side context ourselves from the OpenSubtitlesV2018 corpus based on the meta-information of the segments. For translation of the ContraPro test set, as well as for scoring the contrastive references, we take both the left- and the right-side context into account. For the full-document systems, we cap the context size for the ContraPro test set to 4 sentences for computational reasons.

We tokenize the data using byte-pair-encoding (Sennrich et al., 2016; Kudo, 2018) with 15k joint merge operations (32k for OS En→De). The models are implemented using the fairseq toolkit (Ott et al., 2019) following the transformer base architecture (Vaswani et al., 2017) with dropout 0.3 and

label-smoothing 0.2 for **NEWS En→De** and **TED En→It** and dropout 0.1 and label-smoothing 0.1 for **OS En→De**. This resulted in models with ca. 51M parameters for *NEWS* and *TED* and ca. 60M parameters for *OS* for both the sentence-level and the document-level systems.

Let us assume that the training data \mathcal{C} consists of M documents \mathcal{D}_m and each document consists of source-target sentence pairs $(F_{n,m}, E_{n,m})$. The goal of training is to find the optimal model parameters $\hat{\theta}$ which minimize the loss function:

$$\hat{\theta} = \arg \min_{\theta} L(\theta)$$

When training the local context models, we define the loss function:

$$L(\theta) = -\frac{1}{|\mathcal{C}|} \sum_{m=1}^M \sum_{n=1}^{N_m} \log p_{\theta}(E_{n-k,m}^{n,m} | F_{n-k,m}^{n,m}).$$

When we take full documents as input to the model, the loss function simply becomes

$$L(\theta) = -\frac{1}{M} \sum_{m=1}^M \log p_{\theta}(E_{1,m}^{N_m,m} | F_{1,m}^{N_m,m}).$$

All systems are trained until the validation perplexity does no longer improve and the best checkpoint is selected using validation perplexity as well. Training took around 24h for *NEWS* and *TED* and around 96h for *OS* on a single NVIDIA GeForce RTX 2080 Ti graphics card. Due to computational limitations, we report results only for a single run. For the generation of segments (see Section 3.4), we use beam-search on the token level with beam-size 12 and length normalization.