

Do Context-Aware Translation Models Pay the Right Attention?

Kayo Yin¹ Patrick Fernandes^{1,2,3} Danish Pruthi¹ Aditi Chaudhary¹
André F. T. Martins^{2,3,4} Graham Neubig¹

¹Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA

²Instituto Superior Técnico & LUMIS (Lisbon ELLIS Unit), Lisbon, Portugal

³Instituto de Telecomunicações, Lisbon, Portugal

⁴Unbabel, Lisbon, Portugal

{kayoy, pfernand, ddanish, aschaudh, gneubig}@cs.cmu.edu
andre.t.martins@tecnico.ulisboa.pt

Abstract

Context-aware machine translation models are designed to leverage contextual information, but often fail to do so. As a result, they inaccurately disambiguate pronouns and polysemous words that require context for resolution. In this paper, we ask several questions: What contexts do human translators use to resolve ambiguous words? Are models paying large amounts of attention to the same context? What if we explicitly train them to do so? To answer these questions, we introduce **SCAT** (Supporting Context for Ambiguous Translations), a new English-French dataset comprising supporting context words for 14K translations that professional translators found useful for pronoun disambiguation. Using **SCAT**, we perform an in-depth analysis of the context used to disambiguate, examining positional and lexical characteristics of the supporting words. Furthermore, we measure the degree of alignment between the model’s attention scores and the supporting context from **SCAT**, and apply a guided attention strategy to encourage agreement between the two.¹

1 Introduction

There is a growing consensus in machine translation research that it is necessary to move beyond sentence-level translation and incorporate document-level context (Guillou et al., 2018; Läubli et al., 2018; Toral et al., 2018). While various methods to incorporate context in neural machine translation (NMT) have been proposed (Tiedemann and Scherrer (2017); Miculicich et al. (2018); Maruf and Haffari (2018), *inter alia*), it is unclear whether models rely on the “right” context that is actually sufficient to disambiguate difficult translations. Even when additional context

¹Our **SCAT** data and code for experiments are available at <https://github.com/neulab/contextual-mt>.

	Human
En	Look after her a lot. Okay. Any questions? Have we got her report? Yes, it 's in the infirmary already
Fr	Dorlotez-la. D'accord. Vous avez des questions ? On dispose de son rapport . Oui, il est à l'infirmerie.
Context-aware baseline	
En	Look after her a lot. Okay. Any questions? Have we got her report? Yes, it 's in the infirmary already.
Fr	Dorlotez-la. D'accord. Vous avez des questions ? On dispose de son rapport ? Oui, elle est déjà à l'infirmerie.
Model w/ attention regularization	
En	Look after her a lot. Okay . Any questions ? Have we got her report ? Yes it 's in the infirmary already.
Fr	Dorlotez-la. D'accord. Vous avez des questions ? On dispose de son rapport ? Oui, il est déjà à l'hôpital

Table 1: Translation of the ambiguous pronoun “it”. In French, if the referent of “it” is masculine (e.g., report) then “il” is used, otherwise “elle”. The model with regularized attention translates the pronoun correctly, with the largest attention on the referent “report”. Top 3 words with the highest attention are **highlighted**.

is provided, models often perform poorly on evaluation of relatively simple discourse phenomena (Müller et al., 2018; Bawden et al., 2018; Voita et al., 2019b,a; Lopes et al., 2020) and rely on spurious word co-occurrences during translation of polysemous words (Emelin et al., 2020). Some evidence suggests that models attend to uninformative tokens (Voita et al., 2018) and do not use contextual information adequately (Kim et al., 2019).

To understand plausibly *why* current NMT models are unable to fully leverage the disambiguating context they are provided, and *how* we can develop models that use context more effectively, we pose the following research questions: (i) In context aware translation, what context is intrinsically useful to disambiguate hard translation phenomena such as ambiguous pronouns or word senses?; (ii) Are context-aware MT models paying attention to the relevant context or not?; and (iii) If not, can we

encourage them to do so?

To answer the first question, we *collect annotations of context that human translators found useful* in choosing between ambiguous translation options (§3). Specifically, we ask 20 professional translators to choose the correct French translation between two contrastive translations of an ambiguous word, given an English source sentence and the previous source- and target-side sentences. The translators additionally highlight the words they found the most useful to make their decision, giving an idea of the context useful in making these decisions. We collect 14K such annotations and release **SCAT** (“Supporting Context for Ambiguous Translations”), the first dataset of human rationales for resolving ambiguity in document-level translation. Analysis reveals that inter-sentential target context is important for pronoun translation, whereas intra-sentential source context is often sufficient for word sense disambiguation.

To answer the second question, we *quantify the similarity of the attention distribution of context-aware models and the human annotations* in **SCAT** (§4). We measure alignment between the baseline context-aware model’s attention and human rationales across various model attention heads and layers. We observe a relatively high alignment between self attention scores from the top encoder layers and the source-side supporting context marked by translators, however, the model’s attention is poorly aligned with target-side supporting context.

For the third question, we explore a method to *regularize attention towards human-annotated disambiguating context* (§5). We find that attention regularization is an effective technique to encourage models to pay more attention to words humans find useful to resolve ambiguity in translations. Our models with regularized attention outperform previous context-aware baselines, improving translation quality by 0.54 BLEU, and yielding a relative improvement of 14.7% in contrastive evaluation. An example of translations from a baseline and our model, along with the supporting rationale by a professional translator is illustrated in Table 1.

2 Document-Level Translation

Neural Machine Translation. Current NMT models employ encoder-decoder architectures (Bahdanau et al., 2015; Vaswani et al., 2017). First, the encoder maps a source sequence $\mathbf{x} = (x_1, x_2, \dots, x_S)$ to a continuous representation $\mathbf{z} =$

(z_1, z_2, \dots, z_S) . Then, given \mathbf{z} , the decoder generates the corresponding target sequence $\mathbf{y} = (y_1, y_2, \dots, y_T)$, one token at a time. Sentence-level NMT models take one source sentence and generate one target sentence at a time. These models perform reasonably well, but given that they only have **intra-sentential context**, they fail to handle some phenomena that require **inter-sentential context** to accurately translate. Well-known examples of these phenomena include gender-marked anaphoric pronouns (Guillou et al., 2018) and maintenance of lexical coherence (Läubli et al., 2018).

Document-Level Translation. Document-level translation models learn to maximize the probability of a target document \mathbf{Y} given the source document \mathbf{X} : $P_\theta(\mathbf{Y}|\mathbf{X}) = \prod_{j=1}^J P_\theta(\mathbf{y}^j|\mathbf{x}^j, \mathbf{C}^j)$, where \mathbf{y}^j and \mathbf{x}^j are the j -th target and source sentences, and \mathbf{C}^j is the collection of contextual sentences for the j -th sentence pair. There are many methods for incorporating context (§6), but even simple **concatenation** (Tiedemann and Scherrer, 2017), which prepends the previous source or target sentences to the current sentence separated by a $\langle \text{BRK} \rangle$ tag, achieves comparable performance to more sophisticated approaches, especially in high-resource scenarios (Lopes et al., 2020).

Evaluation. BLEU (Papineni et al., 2002) is most widely used to evaluate MT, but it can be poorly correlated with human evaluation (Callison-Burch et al., 2006; Reiter, 2018). Recently, a number of neural evaluation methods, such as COMET (Rei et al., 2020), have shown better correlation with human judgement. Nevertheless, common automatic metrics have limited ability to evaluate discourse in MT (Hardmeier, 2012). As a remedy to this, researchers often use **contrastive test sets** for a targeted discourse phenomenon (Müller et al., 2018), such as pronoun anaphora resolution and word sense disambiguation, to verify if the model ranks the correct translation of an ambiguous sentence higher than the incorrect translation.

3 What Context Do Human Translators Pay Attention to?

We first conduct a user study to collect supporting context that translators use in disambiguation, and analyze characteristics of the supporting words.

<p>Source context: Yeah, there were nine or ten witnesses.</p> <p>Source sentence: No, they didn't talk.</p> <p>Source context you highlighted: Reset Highlights</p> <p>Source sentence you highlighted: Reset Highlights</p>	<p>Target context: Oui, il y avait neuf ou dix témoins.</p> <p>Target sentence: Non, _____ ne veulent rien dire. <input checked="" type="radio"/> ils <input type="radio"/> elles</p> <p>How confident are you? <input type="button" value="Not at all"/> <input type="button" value="Somewhat"/> <input type="button" value="Very"/></p> <p>Target context you highlighted: • témoins Reset Highlights</p> <p>Target sentence you highlighted: Reset Highlights</p>
---	---

Figure 1: The annotation page shown to translators.

3.1 Recruitment and Annotation Setup

We recruited 20 freelance English-French translators on Upwork.² The translators are native speakers of at least one of the two languages and have a job success rate of over 90%. Each translator is given 400 examples with an English source sentence and two possible French translations, and one out of 5 possible context levels: no context (**0+0**), only the previous source sentence as context (**1+0**), only the previous target sentence (**0+1**), the previous source sentence and target sentence (**1+1**), and the 5 previous source and target sentences (**5+5**). We vary the context level in each example to measure how human translation quality changes.

Translators provide annotations using the interface shown in Figure 1. They are first asked to select the correct translation out of the two contrastive translations, and then highlight word(s) they found useful to arrive at their answer. In cases where multiple words are sufficient to disambiguate, translators were asked to mark only the most salient words rather than all of them. Further, translators also reported their confidence in their answers, choosing from “not at all”, “somewhat”, and “very”.

3.2 Tasks and Data Quality

We perform this study for two tasks: pronoun anaphora resolution (PAR), where the translators are tasked with choosing the correct French gendered pronoun associated to a neutral English pronoun, and word sense disambiguation (WSD), where the translators pick the correct translation of a polysemous word. PAR, and WSD to a lesser extent, have been commonly studied to evaluate context-aware NMT models (Voita et al., 2018; Lopes et al., 2020; Müller et al., 2018; Huo et al., 2020; Nagata and Morishita, 2020).

²<https://www.upwork.com>

Pronoun Anaphora Resolution. We annotate examples from the contrastive test set by Lopes et al. (2020). This set includes 14K examples from the OpenSubtitles2018 dataset (Lison et al., 2018) with occurrences of the English pronouns “it” and “they” that correspond to the French translations “il” or “elle” and “ils” or “elles”, with 3.5K examples for each French pronoun type. Through our annotation effort, we obtain 14K examples of supporting context for pronoun anaphora resolution in ambiguous translations selected by professional human translators. Statistics on this dataset, **SCAT: Supporting Context for Ambiguous Translations**, are provided in Appendix A.

Word Sense Disambiguation. There are no existing contrastive datasets for WSD with a context window larger than 1 sentence, therefore, we automatically generate contrastive examples with context window of 5 sentences from OpenSubtitles2018 by identifying polysemous English words and possible French translations. We describe our methodology in Appendix B.

Quality. For quality control, we asked 8 internal speakers of English and French, with native or bilingual proficiency in both languages, to carefully annotate the same 100 examples given to all professional translators. We compared both the answer accuracies and the selected words for each hired translator against this control set and discarded submissions that either had several incorrect answers while the internal bilinguals were able to choose the correct answer on the same example, or that highlighted contextual words that the internal annotators did not select and that had little relevance to the ambiguous word. Furthermore, among the 400 examples given to each annotator, the first hundred are identical, allowing us to measure the inter-annotator agreement for both answer and supporting context selection.

First, for answer selection on PAR, we find 91.0% overall agreement, with Fleiss’ free-marginal Kappa $\kappa = 0.82$. For WSD, we find 85.9% overall agreement with $\kappa = 0.72$. This indicates a substantial inter-annotator agreement for the selected answer. In addition, we measure the inter-annotator agreement for the selected words by calculating the F1 between the word selections for each pair of annotators given identical context settings. For PAR, we obtain an average F1 of 0.52 across all possible pairs, and a standard deviation of

Context	PAR		WSD	
	Correct	Not confident	Correct	Not confident
0 + 0	78.4	27.0	88.7	7.0
1 + 0	90.6	13.2	88.7	6.5
0 + 1	93.0	9.2	87.5	6.7
1 + 1	93.6	6.7	87.1	6.5
5 + 5	95.9	2.8	88.7	5.9
No ante	75.4	33.8	–	–
Has ante	96.0	3.3	–	–

Table 2: Percentage of correct and zero-confidence answers by varying context level. $n+m$: n previous source and m previous target sentences given as context. Values in bold are significantly different from values above it ($p < 0.05$).

0.12. For WSD, we find an average F1 of 0.46 and a standard deviation of 0.12. There is a high agreement between annotators for the selected words as well.

3.3 Answer Accuracy and Confidence

Table 2 shows the accuracy of answers and the percentage of answers being reported as *not at all* confident for each of the 5 different context levels. For PAR, there is a large increase in accuracy and confidence when just one previous sentence in either language is provided as context compared to no context at all. Target-side context also seems more useful than source: only target-side context gives higher answer accuracy than only source-side context, while the accuracy does not increase significantly by having both previous sentences.

For WSD, we do not observe significant differences in answer accuracy and confidence between the different context levels (Figure 2). The high answer accuracy with 0+0 context and the low rate of zero-confidence answers across all settings suggest that the necessary disambiguating information is often present in the intra-sentential context. Alternatively, this may be partially due to characteristics of the automatically generated dataset itself: we found that some examples are misaligned so the previous sentences given as context do not actually correspond to the context of the current sentences, and therefore do not add useful information. We also observe that translators tend to report a high confidence and high agreement in incorrect answers as well. This can be explained by the tendency to select the masculine pronoun in PAR (Figure 3) or the prevailing word sense in WSD.

To properly translate an anaphoric pronoun, the translator must identify its antecedent and deter-




Figure 2: Distribution of confidence in answers per context level for PAR (left) and WSD (right).




Figure 3: Distribution of gender of selected pronouns per confidence level for PAR.

mine its gender, so we hypothesize that the antecedent is of high importance for disambiguation. In our study, 72.4% of the examples shown to annotators contain the antecedent in the context or current sentences. We calculate how answer accuracy and confidence vary between examples that do or do not contain the pronoun antecedent. We find that the presence of the antecedent in the context leads to larger variations in answer accuracy than the level of context given, demonstrating the importance of antecedents for resolution.

3.4 Analysis of the Highlighted Words

Next, we examine the words that were selected as rationales from several angles.

Distance. Figure 4 shows, for each context level, the number of highlighted words at a given distance (in sentences) from the ambiguous word. For PAR, when no previous sentences are provided, there are as many selected words from the source as the target context. With inter-sentential context, experts selected more supporting context from the target side. One possible reason is that the source and target sentences on their own are equally descriptive to perform PAR, but one may look for the coreference chain of the anaphoric pronoun in the target context to determine its gender, whereas the same coreference chain in the source context would not necessarily contain gender information. Moreover, the antecedent in the target side is more reliable




Figure 4: Sentence distance of the highlighted words for each context level for PAR and WSD.

than the source antecedent, since the antecedent can have multiple possible translations with different genders. For WSD, we find that inter-sentential context is seldom highlighted, which reinforces our previous claim that most supporting context for WSD can be found in the current sentences.

Part-of-Speech and Dependency. We use spaCy (Honnibal and Montani, 2017) to predict part-of-speech (POS) tags of selected words and syntactic dependencies between selected words and the ambiguous word. In Table 3a, we find that nominals are the most useful for PAR, which suggests that human translators look for other referents of the ambiguous pronoun to determine its gender. This is reinforced by Table 3b, where the antecedent of the pronoun is selected the most often.

For WSD, proper nouns and pronouns are not as important as nouns, probably because they do not carry as much semantic load that indicates the sense of the ambiguous word. Determiners, verbs and adpositions are relatively important since they offer clues on the syntactic dependencies of the ambiguous word on other words as well as its role in the sentence, and modifiers provide additional

POS	PAR			WSD		
	Source	Target	Total	Source	Target	Total
noun	1550	2952	4502	3340	937	4277
proper noun	136	4056	4192	192	304	496
pronoun	2676	389	3065	119	204	323
verb	247	859	1106	406	367	773
determiner	499	498	997	800	1091	1891
auxiliary	310	136	446	78	85	163
adjective	105	319	424	291	226	517
adposition	65	172	237	283	481	764
conjunction	71	63	134	83	92	175
numeral	37	39	76	22	440	462
particle	37	8	45	61	0	61

(a) Part-of-speech

DEP	PAR			WSD		
	Source	Target	Total	Source	Target	Total
antecedent	1297	3853	5150	–	–	–
determiner	304	1577	1881	497	407	904
modifier	168	272	440	258	501	759
conjunct	22	58	80	84	45	129
case marking	7	48	55	4	93	97

(b) Dependency relation

Table 3: Most frequent part-of-speech and dependency relation of highlighted words.

PAR	<i>Listen, these big celebrities, they do it different than anybody else? Jesus, you know if they knew you had hidden cameras in that bedroom...</i>
	<i>Dis-moi, ^{DET} ces ^{ante NOUN} vedettes, ^{PRON} elles le font différemment des autres? Bon Dieu, tu te rends compte que si elles/ils savaient que cette chambre cache des caméras...</i>
WSD	<i>Right this way. Your charm is only exceeded by your frankness. Suivez-moi. Ton charme/portebonheur n'a d'égal que ta franchise.</i>

Table 4: Examples of supporting context.

information about the ambiguous word.

The main difference between PAR and WSD is that for PAR, the key supporting information is *gender*. The source side does not contain explicit information about the gender of the ambiguous pronoun whereas the target side may contain other gendered pronouns and determiners referring to the ambiguous pronoun. For WSD however, the key supporting information is *word sense*. While the source and target sides contain around the same amount of semantic information, humans may prefer to attend to source sentences that express how the ambiguous word is used in the sentence.

4 Do Models Pay the Right Attention?

Next, we study NMT models and quantify the degree to which the model’s attention is aligned with the supporting context from professional translators.

Metric	baseline			attnreg-rand			attnreg-pre		
	Enc self	Dec cross	Dec self	Enc self	Dec cross	Dec self	Enc self	Dec cross	Dec self
Dot Product (\uparrow , uniform=0.04)	0.61	0.13	0.10	0.50	0.28	0.20	0.53	0.63	0.31
KL Divergence (\downarrow , uniform=3.6)	3.2	3.6	3.6	3.2	3.2	3.5	3.2	3.2	3.5
Probes Needed (\downarrow , uniform=12.6)	8.5	11.3	13.5	7.6	9.0	5.8	6.9	6.0	10.0

Table 5: Alignment between model attention and **SCAT**.

4.1 Model

We incorporate the 5 previous source and target sentences as context to the base Transformer (Vaswani et al., 2017) by prepending the previous sentences to the current sentence, separated by a $\langle \text{BRK} \rangle$ tag, as proposed by Tiedemann and Scherrer (2017).

4.2 Similarity Metrics

To calculate similarity between model attention and highlighted context, we first construct a human attention vector α_{human} , where 1 corresponds to tokens marked by the human annotators, and 0 otherwise. We compare this vector against the model’s attention for the ambiguous pronoun for a given layer and head, α_{model} , across three metrics:

Dot Product. The dot product $\alpha_{\text{human}} \cdot \alpha_{\text{model}}$ measures the total attention mass the model assigns to words highlighted by humans.

KL Divergence. We compute the KL divergence between the model attention and the normalized human attention vector $\text{KL}(\alpha_{\text{human-norm}} || \alpha_{\text{model}}(\theta))$, where the normalized distribution $\alpha_{\text{human-norm}}$ is uniform over all tokens selected by humans and a very small constant ϵ elsewhere such that the sum of values in $\alpha_{\text{human-norm}}$ is equal to 1.

Probes Needed. We adapt the “probes needed” metric by Zhong et al. (2019) to measure the number of tokens we need to probe, based on the model attention, to find a token highlighted by humans. This corresponds to the ranking of the first highlighted token after sorting all tokens by descending model attention. The intuition is that the more attention the model assigns to supporting context, the fewer probes are needed to find a supporting token.

4.3 Results

We compute the similarity between the model attention distribution for the ambiguous pronoun and the supporting context from 1,000 **SCAT** samples. In Table 5, for each attention type we report the best score across layers and attention heads. We also report the alignment score between a uniform

distribution and supporting context for comparison. We find that although there is a reasonably high alignment between encoder self attention and **SCAT**, decoder attentions have very low alignment with **SCAT**.

5 Making Models Pay the Right Attention

5.1 Attention Regularization

We hypothesize that by encouraging models to increase attention on words that humans use to resolve ambiguity, translation quality may improve. We apply **attention regularization** to guide model attention to increase alignment with the supporting context from **SCAT**. To do so, we append the translation loss with an *attention regularization loss* between the normalized human attention vector $\alpha_{\text{human-norm}}$ and the model attention vector for the corresponding ambiguous pronoun α_{model} :

$$\mathcal{R}(\theta) = -\lambda \text{KL}(\alpha_{\text{human-norm}} || \alpha_{\text{model}}(\theta))$$

where λ is a scalar weight parameter for the loss.

During training, we randomly sample batches from **SCAT** with $p = 0.2$. We train with the standard MT objective on the full dataset, and on examples from **SCAT**, we additionally compute the attention regularization loss.

5.2 Data

For document translation, we use the English and French data from OpenSubtitles2018 (Lison et al., 2018), which we clean then split into 16M training, 10,036 development, and 9,740 testing samples. For attention regularization, we retain examples from **SCAT** where 5+5 context was given to the annotator. We use 11,471 examples for training and 1,000 for testing.

5.3 Models

We first train a **baseline** model, where the 5 previous source and target sentences serve as context and are incorporated via concatenation. This baseline model is trained without attention regularization.

We explore two models with attention regularization: (1) *attnreg-rand*, where we jointly train on the MT objective and regularize attention on a randomly initialized model; (2) *attnreg-pre*, where we first pre-train the model solely on the MT objective, then we jointly train on the MT objective and regularize attention. We describe the full setup in Appendix C.

5.4 Evaluation

As described in Section 2, we evaluate translation outputs with BLEU and COMET. In addition, to evaluate the direct translation quality of specific phenomena, we translate the 4,015 examples from Lopes et al. (2020) containing ambiguous pronouns that were not used for attention regularization, and we compute the mean word f-measure of translations of the ambiguous pronouns and other words, with respect to reference texts.

We also perform contrastive evaluation on the same subset of Lopes et al. (2020) with a context window of 5 sentences (*Big-PAR*) and the contrastive test sets by Bawden et al. (2018), which include 200 examples on anaphoric pronoun translation and 200 examples on lexical consistency/word sense disambiguation. The latter test sets were crafted manually, have a context window of 1 sentence, and either the previous source or target sentence is necessary to disambiguate.

Context-aware models often suffer from error propagation when using previously decoded output tokens as the target context (Li et al., 2020a). Therefore, during inference, we experiment with both using the gold target context (**Gold**) as well as using previous output tokens (**Non-Gold**).

5.5 Overall Performance

Before delving into the main results, we note that we explored regularizing different attention vectors in the model (Appendix C.3) and obtain the best BLEU and COMET scores for *attnreg-rand* when regularizing the self-attention of the top encoder layer, cross-attention of the top decoder layer and self-attention of the bottom decoder layer. For *attnreg-pre*, regularizing self-attention in the top decoder layer gives the best scores. Thus, we use these as the default regularization methods below.

Moving on to the main results in Table 6, we observe that *attnreg-rand* improves on all metrics, which demonstrates that attention regularization is an effective method to improve translation quality. Although *attnreg-pre* does not improve gen-

eral translation scores significantly, it yields considerable gains in word f-measure on ambiguous pronouns and achieves some improvement over the baseline on contrastive evaluation on Big-PAR and PAR. Attention regularization with supporting context for PAR seems to especially improve models on similar tasks. The disparity between BLEU/COMET scores and targeted evaluations such as word f-measure and contrastive evaluation further suggests that general MT metrics are somewhat insensitive to improvements on specific discourse phenomena. For both models with attention regularization, there are no significant gains in WSD. As discussed in §3.4, WSD and PAR require different types of supporting context, so it is natural that regularizing attention using supporting context extracted from only one task does not always lead to improvement on the other.

5.6 Analysis

We now investigate how models trained with attention regularization handle context differently compared to the baseline model.

How does attention regularization influence alignment with human rationales? We revisit the similarity metrics from §4.2 to measure alignment with **SCAT**. In Table 5, the dot product alignment over attention in the decoder increases with attention regularization, suggesting that attention regularization guides different parts of the model to pay attention to useful context. Interestingly, although only the encoder self-attention was explicitly regularized for *attnreg-pre*, the model seems to also have learned better alignment for attention in the decoder. Moreover, *attnreg-pre* generally has better alignment than *attnreg-rand*, suggesting that models respond more to attention regularization once it has been trained to perform translation.

Which attention is the most useful? For each of *attnreg-rand* and *attnreg-pre*, we perform attention regularization on either the encoder self-attention, decoder cross-attention or decoder self-attention only. In Table 7, encoder self-attention seems to contribute the most to both translation performance and contrastive evaluation. Although *attnreg-rand* models achieve higher BLEU and COMET scores, *attnreg-pre* obtain higher scores on metrics targeted to pronoun translation. Attention regularization seems to have limited effect on WSD performance, the scores vary little between attention types.

Model	Gold		Non-Gold		F-measure		Contrastive Evaluation		
	BLEU	COMET	BLEU	COMET	Pronouns	Other	Big-PAR	PAR	WSD
baseline	33.5	41.7	29.0	37.1	0.36	0.44	90.0	60.0	55.0
attnreg-rand	33.9	42.3	30.2	39.3	0.42	0.44	91.1	69.0	56.0
attnreg-pre	32.3	36.9	29.3	34.6	0.48	0.44	90.6	62.5	52.5

Table 6: Overall results. Scores significantly better than baseline with $p < 0.05$ are bolded. COMET scores are multiplied by 100.

	Attention	B	C	F-meas		Contrastive Evaluation		
				Pron.	Other	Big-PAR	PAR	WSD
AR	Enc self	33.8	42.5	0.47	0.44	91.3	59.5	57.0
AR	Dec cross	33.8	41.7	0.47	0.43	89.8	65.0	55.5
AR	Dec self	33.1	41.6	0.41	0.45	88.9	66.0	53.5
AP	Enc self	32.3	37.4	0.48	0.44	91.5	66.5	52.5
AP	Dec cross	32.3	37.4	0.47	0.43	92.1	63.0	55.0
AP	Dec self	32.3	36.9	0.48	0.44	90.6	62.5	52.5

Table 7: Performance of models with various regularized attention. AR: attnreg-rand, AP: attnreg-pre, B: BLEU, C:COMET

mask	baseline	attnreg-rand	attnreg-pre
no mask	82.5	84.4	86.7
supporting	75.1	69.4	<u>55.0</u>
random	76.0	77.0	80.4
source	65.5	67.3	73.0
target	70.6	75.3	67.6
all	<u>65.3</u>	<u>67.1</u>	68.7

Table 8: Contrastive performance with various masks on the context. The lowest score for each model is underlined.

How much do models rely on supporting context? We compare model performance on contrastive evaluation on **SCAT** when it is given full context, and when we mask either the supporting context, random context words with $p = 0.1$, the source context, the target context, or all of the context. In Table 8, we find that *baseline* varies little when the supporting context is masked, which again suggests that context-aware baselines do not use the relevant context, although they do observe a drop in contrastive performance when the source and all context are masked. Models with attention regularization, especially *attnreg-pre* observe a large drop in contrastive performance when supporting context is masked, which indicates that they learned to rely more on supporting context. Furthermore, for *attnreg-pre*, the score after masking supporting context is significantly lower than when masking all context, which may indicate that having irrelevant context can have an adverse ef-

fect. Another interesting finding is that both *baseline* and *attnreg-rand* seem to rely more on the source context than the target context, in contrast to human translators. This result corroborates prior results where models have better alignment with supporting context on attention that attends to the source (encoder self-attention and decoder cross-attention), and regularizing these attention vectors contributes more to translation quality than regularizing the decoder self-attention.

6 Related Work

6.1 Context-Aware Machine Translation

Most current context-aware NMT approaches enhance NMT by including source- and/or target-side surrounding sentences as context to the model. Tiedemann and Scherrer (2017) concatenate the previous sentences to the input; Jean et al. (2017); Bawden et al. (2018); Zhang et al. (2018) use an additional encoder to extract contextual features; Wang et al. (2017) use a hierarchical RNN to encode the global context from all previous sentences; Maruf and Haffari (2018); Tu et al. (2018) use cache-based memories to encode context; Miculicich et al. (2018); Maruf et al. (2019) use hierarchical attention networks; Chen et al. (2020) add document-level discourse structure information to the input. While Maruf et al. (2019); Voita et al. (2018) also find higher attention mass attributed to relevant tokens in selected examples, our work is the first to guide model attention in context-aware NMT using human supervision and analyze its attention distribution in a quantitative manner.

However, recent studies suggest that current context-aware NMT models often do not use context meaningfully. Kim et al. (2019) claim that improvements by context-aware models are mostly from regularization by reserving parameters for context inputs, and Li et al. (2020b) show that replacing the context in multi-encoder models with random signals leads to similar accuracy as using the actual context. Our work addresses the above

disparities by collecting human supporting context to regularize model attention heads during training.

6.2 Attention Mechanisms

Though attention is usually learned in an unsupervised manner, recent work supervises attention with word alignments (Mi et al., 2016; Liu et al., 2016), event arguments and trigger words (Liu et al., 2017; Zhao et al., 2018), syntactic dependencies (Strubell et al., 2018) or word lexicons (Zou et al., 2018). Our work is closely related to a large body of work that supervises attention using human rationales for text classification (Barrett et al., 2018; Bao et al., 2018; Zhong et al., 2019; Choi et al., 2020; Pruthi et al., 2020). Our work, however, is the first to collect human evidence for document translation and use it to regularize the attention of NMT models.

7 Implications and Future Work

In this work, we collected a corpus of supporting context for translating ambiguous words. We examined how baseline context-aware translation models use context, and demonstrated how context annotations can improve context-aware translation accuracy. While we obtain promising results for context-aware translation by testing one method for attention regularization, our publicly available **SCAT** dataset could enable future research on alternative attention regularizers. Moreover, our analyses demonstrate that humans rely on different types of context for PAR and WSD in English-French translation, similar user studies can be conducted to better understand the usage of context in other ambiguous discourse phenomena, such as ellipsis, or other language pairs. We also find that regularizing attention using **SCAT** for PAR especially improves anaphoric pronoun translation, suggesting that supervising attention using supporting context from different tasks may help models resolve other types of ambiguities.

One caveat regarding our method for collecting supporting context from humans is the difference between *translation*, translating text from the input, and *disambiguation*, choosing between translation candidates. During translation, humans might pay more attention to the source sentences to understand the source material, but during disambiguation, we have shown that human translators rely more often on the target sentences. One reason why the model benefits more from increased attention on source may be because the model is trained and

evaluated to perform translation, not disambiguation. A future step would be to explore alternative methods for extracting supporting context, such as eye-tracking during translation (O’Brien, 2009).

8 Acknowledgements

We would like to thank Emma Landry, Guillaume Didier, Wilson Jallet, Baptiste Moreau-Pernet, Pierre Gianferrara, and Duy-Anh Alexandre for helping with a preliminary English-French translation study. We would also like to thank Nikolai Vogler for the original interface for data annotation, and the anonymous reviewers for their helpful feedback. This work was supported by the European Research Council (ERC StG DeepSPIN 758969), by the P2020 programs MAIA and Unbabel4EU (LISBOA-01-0247-FEDER-045909 and LISBOA-01-0247-FEDER-042671), and by the Fundação para a Ciência e Tecnologia through contract UIDB/50008/2020.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Yujia Bao, Shiyu Chang, Mo Yu, and Regina Barzilay. 2018. [Deriving machine attention from human rationales](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1903–1913, Brussels, Belgium. Association for Computational Linguistics.
- Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, and Anders Søgaard. 2018. [Sequence classification with human attention](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 302–312, Brussels, Belgium. Association for Computational Linguistics.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. [Re-evaluating the role of Bleu in machine translation research](#). In *11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy. Association for Computational Linguistics.

- Junxuan Chen, Xiang Li, Jiarui Zhang, Chulun Zhou, Jianwei Cui, Bin Wang, and Jinsong Su. 2020. [Modeling discourse structure for document-level neural machine translation](#). In *Proceedings of the First Workshop on Automatic Simultaneous Translation*, pages 30–36, Seattle, Washington. Association for Computational Linguistics.
- Seungtaek Choi, Haeju Park, Jinyoung Yeo, and Seungwon Hwang. 2020. [Less is more: Attention supervision with counterfactuals for text classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6695–6704, Online. Association for Computational Linguistics.
- Zi-Yi Dou and Graham Neubig. 2021. Word alignment by fine-tuning embeddings on parallel corpora. *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Denis Emelin, Ivan Titov, and Rico Sennrich. 2020. Detecting word sense disambiguation biases in machine translation for model-agnostic adversarial attacks. *arXiv preprint arXiv:2011.01846*.
- Liane Guillou, Christian Hardmeier, Ekaterina Lapshinova-Koltunski, and Sharid Loáiciga. 2018. [A pronoun test suite evaluation of the English–German MT systems at WMT 2018](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 570–577, Belgium, Brussels. Association for Computational Linguistics.
- Christian Hardmeier. 2012. Discourse in statistical machine translation. a survey and a case study. *Discours. Revue de linguistique, psycholinguistique et informatique. A journal of linguistics, psycholinguistics and computational linguistics*.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- 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, Online. 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*.
- 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, Hong Kong, China. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 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.
- Samuel Lüubli, 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, Brussels, Belgium. Association for Computational Linguistics.
- Bei Li, Hui Liu, Ziyang Wang, Yufan Jiang, Tong Xiao, Jingbo Zhu, Tongran Liu, and Changliang Li. 2020a. [Does multi-encoder help? a case study on context-aware neural machine translation](#). *arXiv preprint arXiv:2005.03393*.
- Bei Li, Hui Liu, Ziyang Wang, Yufan Jiang, Tong Xiao, Jingbo Zhu, Tongran Liu, and Changliang Li. 2020b. [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.
- 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)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Lemao Liu, Masao Utiyama, Andrew Finch, and Ei-ichiro Sumita. 2016. [Neural machine translation with supervised attention](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3093–3102, Osaka, Japan. The COLING 2016 Organizing Committee.
- Shulin Liu, Yubo Chen, Kang Liu, and Jun Zhao. 2017. [Exploiting argument information to improve event detection via supervised attention mechanisms](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1789–1798, Vancouver, Canada. Association for Computational Linguistics.

- António Lopes, M. Amin Farajian, Rachel Bawden, Michael Zhang, and André F. T. Martins. 2020. [Document-level neural MT: A systematic comparison](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 225–234, Lisboa, Portugal. European Association for Machine Translation.
- 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, Melbourne, Australia. Association for Computational Linguistics.
- Sameen Maruf, André F. T. 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Haitao Mi, Zhiguo Wang, and Abe Ittycheriah. 2016. [Supervised attentions for neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2283–2288, Austin, Texas. Association for Computational Linguistics.
- Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. [Document-level neural machine translation with hierarchical attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium. Association for Computational Linguistics.
- Mathias Müller, Annette Rios, 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, Brussels, Belgium. Association for Computational Linguistics.
- Masaaki Nagata and Makoto Morishita. 2020. [A test set for discourse translation from Japanese to English](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3704–3709, Marseille, France. European Language Resources Association.
- Sharon O’Brien. 2009. Eye tracking in translation process research: methodological challenges and solutions. *Methodology, technology and innovation in translation process research*, 38:251–266.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Danish Pruthi, Bhuwan Dhingra, Livio Baldini Soares, Michael Collins, Zachary C. Lipton, Graham Neubig, and William W. Cohen. 2020. [Evaluating explanations: How much do explanations from the teacher aid students?](#)
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 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.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. [Linguistically-informed self-attention for semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038, Brussels, Belgium. Association for Computational Linguistics.
- 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Antonio Toral, Sheila Castilho, Ke Hu, and Andy Way. 2018. [Attaining the unattainable? reassessing claims of human parity in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 113–123, Brussels, Belgium. Association for Computational Linguistics.
- Zhaopeng Tu, Y. Liu, Shuming Shi, and T. 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, pages 5998–6008. Curran Associates, Inc.

- Elena Voita, Rico Sennrich, and Ivan Titov. 2019a. [Context-aware monolingual repair for 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 877–886, Hong Kong, China. Association for Computational Linguistics.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019b. [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.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 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.
- Longyue Wang, Zhaopeng Tu, Andy Way, and Qun Liu. 2017. [Exploiting cross-sentence context for neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2826–2831, Copenhagen, Denmark. Association for Computational Linguistics.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Yue Zhao, Xiaolong Jin, Yuanzhuo Wang, and Xueqi Cheng. 2018. [Document embedding enhanced event detection with hierarchical and supervised attention](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 414–419, Melbourne, Australia. Association for Computational Linguistics.
- Ruiqi Zhong, Steven Shao, and Kathleen R. McKeown. 2019. [Fine-grained sentiment analysis with faithful attention](#). *CoRR*, abs/1908.06870.
- Yicheng Zou, Tao Gui, Qi Zhang, and Xuanjing Huang. 2018. [A lexicon-based supervised attention model for neural sentiment analysis](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 868–877, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

	0+0	1+0	0+1	1+1	5+5
# Examples	1,709	1,616	1,689	1,742	12,616
Answer accuracy (%)	78.4	90.7	93.31	93.7	97.2
Not at all confident (%)	27.6	13.9	9.4	6.9	2.8
Highlighted current source	953	636	574	533	1,169
Highlighted current target	1,019	913	941	965	6,198
Highlighted context source	–	500	–	226	711
Highlighted context target	–	–	619	709	6,011

Table 9: Statistics on **SCAT** for various context levels

Source	Target	Class
nail	clou, ongle <i>metal nail, fingernail</i>	non-synonymous
fork	diapason, fourche, fourchette <i>tuning fork, pitchfork, kitchen fork</i>	non-synonymous
mistakes	erreurs, fautes <i>errors, faults</i>	synonymous
heater	chauffage, radiateur <i>heating, radiator</i>	synonymous

Table 10: Examples of ambiguous word groups.

A SCAT: Supporting Context for Ambiguous Translations

The dataset contains 19,372 annotations in total on 14,000 unique examples of ambiguous anaphoric pronoun translations. In Table 9, for each context level, we report the total number of examples, the overall answer accuracy, the percentage of *not at all* confident answers, and the number of examples that contain a highlighted word the current/context source/target sentences.

B Generating Data for Word Sense Disambiguation

To automatically generate contrastive examples of WSD, we identify English words that have multiple French translations. To do so, we first extract word alignments from OpenSubtitles2018 using AWESOME-align (Dou and Neubig, 2021) and obtain:

$$A_m = \{\langle x_i, y_j \rangle : x_i \in \mathbf{x}_m, y_j \in \mathbf{y}_m\},$$

where for each word pair $\langle x_i, y_j \rangle$, x_i and y_j are semantically similar to each other in context.

For each pair $\langle x_i, y_j \rangle \in A_m$, we compute the number of times the lemmatized source word type ($v_x = \text{lemma}(x_i)$) along with its POS tag ($t_x = \text{tag}(x_i)$) is aligned to the lemmatized target word type ($v_y = \text{lemma}(y_j)$): $c(v_x, t_x, v_y)$. Then, we extract tuples of source types with its POS tags $\langle v_x, t_x \rangle$ that have at least two target words that have been aligned at least 50 times ($|\{v_y | c(v_x, t_x, v_y) \geq 50\}| \geq 2$). Finally, we filter out the source tuples which have an entropy

$H(v_x, t_x)$ less than a pre-selected threshold z . This entropy is computed using the conditional probability of a target translation given the source word type and its POS tag as follows:

$$p := p(v_y | v_x, t_x) = \frac{c(v_x, t_x, v_y)}{c(v_x, t_x)}$$

$$H(v_x, t_x) = \sum_{v_y \in \text{trans}(v_x, t_x)} -p \log_e p$$

where $\text{trans}(v_x, t_x)$ is the set of target translations for the source tuple $\langle v_x, t_x \rangle$ and $p(v_y | v_x, t_x)$ is the conditional probability of a given target translation v_y for the source word type v_x and its POS tag t_x .

Out of the 394 extracted word groups, we manually validate and retain 201 groups and then classify them into 64 synonymous and 137 non-synonymous word groups (Table 10). We create contrastive translations by extracting sentence pairs containing an ambiguous word pair, and replacing the translation of the polysemous English word by a different French word in the same group. For word groups with synonymous French words, we only retain examples where the French word appears within the previous 5 sentences to enforce lexical consistency, as otherwise the different French words may be interchangeable.

C Experimental Setup

C.1 Data preprocessing

We use the English and French data from the publicly available OpenSubtitles2018 dataset (Lison et al., 2018). We first clean the data by selecting sentence pairs with a relative time overlap between source and target language subtitle frames of at least 0.9 to reduce noise. Each data is then encoded with byte-pair encoding (Sennrich et al., 2016) using SentencePiece (Kudo and Richardson, 2018), with source and target vocabularies of 32k tokens.

C.2 Training configuration

We follow the Transformer base (Vaswani et al., 2017) configuration in all our experiments, with $N = 6$ encoder and decoder layers, $h = 8$ attention heads, hidden size $d_{\text{model}} = 512$ and feedforward size $d_{\text{ff}} = 2048$. We use the learning rate schedule and regularization described in Vaswani et al. (2017). We train using the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.98$.

Model	Attention and Layer	Gold		Non-Gold		F-measure		Contrastive Evaluation		
		BLEU	COMET	BLEU	COMET	Pronouns	Other	Big-PAR	PAR	WSD
AR	Enc self L1	33.8	42.5	29.9	39.0	0.47	0.44	91.3	59.5	57.0
AR	Dec cross L1	33.8	41.7	30.5	39.3	0.47	0.43	89.8	65.0	55.5
AR	Dec self L1	33.1	41.6	29.7	38.6	0.41	0.45	88.9	66.0	53.5
AR	Enc self L1 + Dec cross L1 + Dec self L6	33.9	42.3	30.2	39.3	0.42	0.44	91.1	69.0	56.0
AP	Enc self L1	32.3	37.4	28.8	34.4	0.48	0.44	91.5	66.5	52.5
AP	Dec cross L1	32.3	37.4	28.7	34.2	0.47	0.43	92.1	63.0	55.0
AP	Dec self L1	32.3	36.9	29.3	34.6	0.48	0.44	90.6	62.5	52.5
AP	Enc self L1 + Dec cross L1 + Dec self L6	32.3	36.9	29.3	34.6	48.0	44.1	90.6	62.5	52.5

Table 11: Results of all models with regularized attention. AR: attnreg-rand, AP: attnreg-pre

C.3 Attention regularization setups

For both *attnreg-rand* and *attnreg-pre*, we experiment performing regularization on different model attentions at different layers. For the attention regularization loss, In all experiments, we compute the attention regularization loss on the first attention head with $\lambda = 10$, and we divide the loss by the length of the input. We give the results for all setups in Table 11.

C.4 Evaluation

To calculate BLEU scores we use SacreBLEU BLEU+case.mixed+numrefs.1+smooth.exp+tok.13 a+version.1.4.14 (Post, 2018) and for COMET we use *wmt-large-da-estimator-1719*³. We test for statistical significance with $p < 0.05$ using bootstrap sampling on a single run (Koehn, 2004).

³<https://github.com/Unbabel/COMET>