

# Analyzing the Attention Heads for Pronoun Disambiguation in Context-aware Machine Translation Models

Paweł Mąka and Yusuf Can Semerci and Jan Scholtes and Gerasimos Spanakis

Department of Advanced Computing Sciences

Maastricht University

{pawel.maka, y.semerci, j.scholtes, jerry.spanakis}@maastrichtuniversity.nl

## Abstract

In this paper, we investigate the role of attention heads in Context-aware Machine Translation models for pronoun disambiguation in the English-to-German and English-to-French language directions. We analyze their influence by both observing and modifying the attention scores corresponding to the plausible relations that could impact a pronoun prediction. Our findings reveal that while some heads do attend the relations of interest, not all of them influence the models' ability to disambiguate pronouns. We show that certain heads are *underutilized* by the models, suggesting that model performance could be improved if only the heads would attend one of the relations more strongly. Furthermore, we fine-tune the most promising heads and observe the increase in pronoun disambiguation accuracy of up to 5 percentage points which demonstrates that the improvements in performance can be solidified into the models' parameters.

## 1 Introduction

In Context-Aware Machine Translation (MT), the context sentences are available to the system and can be used to maintain coherence of the translation and to resolve ambiguities (Agrawal et al., 2018; Bawden et al., 2018; Müller et al., 2018; Voita et al., 2019b). Both the source-side (sentences in the source language) and target-side context (the previously translated sentences) can be used as context. Although many novel architectures have been proposed to tackle the task of Context-aware MT (Tu et al., 2017; Bawden et al., 2018; Miculicich et al., 2018; Maruf et al., 2019; Huo et al., 2020; Zheng et al., 2021), we limit our investigation to the standard Transformer (Vaswani et al., 2017) architecture (often referred to as the *single-encoder* architecture), because of its simplicity and demonstrated high performance (Sun et al., 2022; Majumde et al., 2022; Gete et al., 2023; Post and

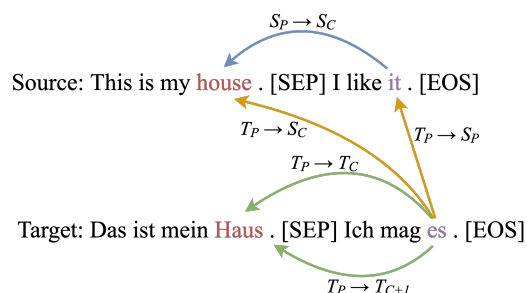


Figure 1: The types of relations we investigate.  $S_P$  and  $T_P$  mark a context-dependent word (e.g., pronoun) on the source- and target-side respectively, and  $S_C$  and  $T_C$  mark the source- and target-side context cue (e.g., antecedent). For the target side, the words represent tokens predicted by the model, and  $T_{C+1}$  is the token corresponding to the antecedent as the input. Colors of arrows designate self-attention (blue), cross-attention (yellow), and decoder-attention (decoder self-attention, green).

Junczys-Dowmunt, 2023; Mohammed and Niculae, 2024).

For the Transformer model to be successful in the task of Context-aware MT, it has to integrate the contextual information and utilize it to produce the translation autoregressively. During inference, the contextual information has to come through the mechanism of multi-head attention. Motivated by the findings that attention heads can learn to perform seemingly specific functions (Clark et al., 2019; Voita et al., 2019c; Jo and Myaeng, 2020; Olsson et al., 2022), we hypothesize that certain heads in a model can be crucial for the context utilization. We study the case of pronoun disambiguation, where the selection of the correct pronoun is dependent on the antecedent (contextual cue), which can be present in the previous (context) sentences. Therefore, we analyze the translation models through the lens of the attention given by the attention heads to the plausible relations that could influence the prediction of a pronoun:

pronoun-antecedent on the source and target sides, and pronoun-pronoun and pronoun-antecedent between target and source sides. Figure 1 illustrates the relations we consider.

Our analysis starts with measuring the attention scores given by each head to the relations we consider when evaluated on the contrastive datasets - ContraPro (Bawden et al., 2018) for the English-to-German direction and Large Contrastive Pronoun Testset (LCPT; Lopes et al. (2020)) for the English-to-French direction. In these datasets, the model is presented with the task of ranking several translations of the same source sentence with the same context. The provided translations differ only partially, and the provided context is required to choose the correct translation. Next, we correlate the scores with the model being correct when evaluated on the contrastive example. Furthermore, we artificially modify the attention scores of the heads and measure the difference in the accuracy on these datasets. Lastly, we fine-tune the selected heads to attend the relations more strongly.

## 2 Related Work

### 2.1 Context-aware Machine Translation

A direct approach to include previous sentences as context in MT is to concatenate them with the current sentence. This method is often called the *single-encoder* architecture, as the basic encoder-decoder architecture is used (Tiedemann and Scherrer, 2017; Ma et al., 2020; Zhang et al., 2020). The *multi-encoder* approach is to encode the context sentences by a separate encoder (Jean et al., 2017; Miculicich et al., 2018; Maruf et al., 2019; Huo et al., 2020; Zheng et al., 2021). Existing studies also investigated more exotic architectures (Miculicich et al., 2018; Bao et al., 2021; Chen et al., 2022; Maka et al., 2024), post-processing translation (Voita et al., 2019b,a), and employing a memory mechanism (Feng et al., 2022; Bulatov et al., 2022). Yin et al. (2021) applies attention regularization that concentrates on the context phrases marked as important by human translators which is similar to Head Tuning method in our work.

In recent years, the single-encoder architecture has seen increased prominence in the literature because of its simplicity and robust performance (Majumde et al., 2022; Gete et al., 2023; Post and Junczys-Downmunt, 2023; Mohammed and Niculae, 2024), even on long context sizes (of up to 2000 tokens) when data augmentation was used (Sun et al.,

2022). For this reason, this paper focuses on this architecture.

Studies have shown that the sentence-level metrics (such as BLEU (Papineni et al., 2002)) are not well suited to measure the translation quality with respect to context usage (Hardmeier, 2012; Wong and Kit, 2012). To address this issue, researchers introduced new metrics (Fernandes et al., 2021, 2023) and context-sensitive datasets (Wicks and Post, 2023; Fernandes et al., 2023), including contrastive datasets (Müller et al., 2018; Bawden et al., 2018; Voita et al., 2019b; Lopes et al., 2020). In this study we employ two of them: ContraPro (Müller et al., 2018) and Large Contrastive Pronoun Testset (LCPT; Lopes et al. (2020)).

### 2.2 Explaining Models

Explaining models' decisions or behaviors is an important issue from the safety and ethical considerations (Madsen et al., 2022). Additionally, we argue that they can also be valuable from the engineering perspective, where the shortcomings of the models can potentially be addressed when they are brought to light. Numerous methods explaining Natural Language Processing models have been proposed (Bau et al., 2018; Toneva and Wehbe, 2019; Ferrando et al., 2022; Langedijk et al., 2024; Meng et al., 2024).

Even though some works have argued against using raw attention scores as explanations themselves (Jain and Wallace, 2019) the attention mechanism is an important part of the Transformer and many researchers concentrated their efforts to understand its influence on the model's behavior (Wiegrefe and Pinter, 2019; Abnar and Zuidema, 2020; Kobayashi et al., 2020, 2021; Bogoychev, 2021; Gheini et al., 2021; Mohebbi et al., 2023). Our work is motivated by the findings suggesting that some heads have specific roles or functions (Clark et al., 2019; Voita et al., 2019c; Olsson et al., 2022) that can be linked to linguistic relations (Vig and Belinkov, 2019; Tenney et al., 2019; Jo and Myaeng, 2020).

Several works have investigated the models in MT (Goindani and Shrivastava, 2021; Voita et al., 2021; Sarti et al., 2023; Mohammed and Niculae, 2024). In contrast, our work focuses on identifying where in the model's architecture the contextual information is integrated. We also investigate the influence of the alterations to the model (i.e., modifying the attention scores corresponding to the plausible relations) on the context usage.

### 3 Methods

We denote the tokenized source and target sentences of example  $d$  in the dataset as  $S^d$  and  $T^d$  respectively. We define the sets containing the indices of contextually important tokens (context cues) on the source  $S_C$  and target side  $T_C$ , and the sets containing the indices of contextually dependent tokens on the source  $S_P$  and target side  $T_P$ . On the target side, the token indices denote the predicted tokens, rather than the input tokens, which are shifted right during training and contrastive evaluation. This way we analyze the model during the inference step when the context-dependent tokens  $T_P$  are being predicted. Consequently, the context cue tokens on the target side  $Y_C$  also mark the steps where those tokens are being generated by the model. Because the input tokens have been found to retain their identity throughout the layers of the Transformer (Brunner et al., 2020), we add the set  $T_{C+1}$  referring to the contextually important tokens on the input. In the case of ambiguous pronouns, the context cues tokens are the antecedents and the contextually dependent tokens are the pronouns, both on the source and target side.

We investigate the relations between the following token sets (see Figure 1 for an example):

- from source dependent to source important tokens ( $S_P \rightarrow S_C$ ) - in the encoder’s self-attention;
- from target dependent to source important tokens ( $T_P \rightarrow S_C$ ) - in the decoder’s cross-attention;
- from target dependent to source dependent tokens ( $T_P \rightarrow S_P$ ) - in the decoder’s cross-attention;
- from target dependent to target important tokens ( $T_P \rightarrow T_C$ ) - in the decoder’s self-attention;
- from target dependent to target important tokens on the input ( $T_P \rightarrow T_{C+1}$ ) - in the decoder’s self-attention.

We examine the models through three methods. *Attention Scores* (Section 3.1) and *Score-Accuracy Correlation* (Section 3.2) are based on observing the model’s behavior during the task of contrastive disambiguation. The method *Modifying Heads* (Section 3.3) introduce a form of disturbance into the functioning of an attention head.

#### 3.1 Attention Scores

To find the heads that learned to pay high attention to the relations of interest we measure and average the attention scores  $Z^{l,h}$  of each head  $h$  in every layer  $l$  for all relations of interest. For the contextually

important and dependent phrases that span over multiple tokens, we take the maximum score.

#### 3.2 Score-Accuracy Correlation

The fact that a particular attention head - on average - pays attention to the relation of interest does not necessarily mean that the head is crucial or helpful for disambiguation. Therefore, we want to measure how the head’s attention scores of a relation correspond to the model correctly disambiguating a particular example. We define a variable  $I$  as follows:

$$I^d = \begin{cases} 1, & \text{if correctly scored,} \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where  $d$  is an example from the contrastive dataset  $\mathcal{D}$ . We calculate the point-biserial correlation coefficient between attention scores given by each head  $Z^{l,h}$  and the variable  $I$ . The accuracy on the whole dataset can be calculated as  $\sum_{d \in \mathcal{D}} I^d / |\mathcal{D}|$ .

#### 3.3 Modifying Heads

The goal of Modifying Heads is to adjust the behavior of the head in a controlled manner. It modifies attention scores from a particular token ensuring that the total attention score given to a target subset of tokens equals a desired value  $C$  while preserving the pre-softmax attention scores  $H$  for all other target tokens. Modifying Heads allows us to experimentally test the behavior of the model if one (or more) of its heads were better or worse at the function of attending to the context-informative tokens. For simplicity, we jointly label the attending tokens (left-hand-side of the arrow  $\rightarrow$ ) as  $\mathcal{Y}$  and attended tokens (right-hand-side of the arrow  $\rightarrow$ ) as  $\mathcal{X}$ . Consequently,  $\mathcal{Y} \rightarrow \mathcal{X}$  can represent all investigated relations inside the corresponding attention module. Modifying Heads can be formulated as follows:

$$\tilde{H}_{i,j}^{l,h,d} = \log \left( \frac{C}{|\mathcal{X}^d|(1-C)} \sum_{k \in \mathcal{X}^d \setminus \mathcal{X}^d} \exp(H_{i,k}^{l,h,d}) \right) \quad (2)$$

$$\forall i \in \mathcal{Y}^d, j \in \mathcal{X}^d,$$

where  $\tilde{H}$  represent the updated pre-softmax attention scores,  $H$  are the original pre-softmax scores,  $k \in \mathcal{X}^d \setminus \mathcal{X}^d$  are the attended tokens not present in the subset of interest  $\mathcal{X}^d$ . The derivation of the formula can be found in Appendix A.

## 4 Experiments

All our experiments are implemented<sup>1</sup> in *Huggingface transformers* framework (Wolf et al., 2020).

### 4.1 Models

In this study, we use pre-trained single-encoder models for two reasons: to assess how models’ abilities learned on intra-sentential phenomena will translate into the inter-sentential regime, and to analyze the robustly trained and widely tested models instead of training models from random initialization. The first model we experiment with is *OPUS-MT en-de*<sup>2</sup> (Tiedemann and Thottingal, 2020; Tiedemann et al., 2023). It is a relatively small (6 layers, 8 heads) encoder-decoder Transformer model trained on English-to-German translation. The second model is *No Language Left Behind* (NLLB-200) (NLLB Team et al., 2022), which is a multilingual MT model. We use the small distilled version with approximately 600 million parameters<sup>3</sup>, which consists of 12 encoder and decoder layers and 16 heads.

### 4.2 Context-aware Fine-tuning

We fine-tuned both models for Context-aware MT by concatenating the previous sentences with the current sentence on both the source and target side. Each sentence is separated by the [SEP] token (before fine-tuning, we expanded the vocabulary of the OpusMT en-de model with this token). We trained two versions of the OpusMT en-de model with maximum context sizes - the number of context sentences - of one and three (referred to as context-aware-1 and context-aware-3 models), and one version of NLLB-600M with the maximum context size of one. We train the models with all the context sizes from zero to the maximum context size while keeping the context size the same on the source and target side for each example. During inference, we provide the models with the number of sentences equal to the maximum context size it was trained with.

For Opus-MT en-de, we used the train subset of the IWSLT 2017 (Cettolo et al., 2017) English-to-German dataset, and interleaved English-to-

German and English-to-French train subsets for NLLB-200. IWSLT 2017 is a document-level dataset (the details of the datasets are presented in Appendix C). Further details and the hyperparameters of the fine-tuning process are presented in Appendix D.

### 4.3 Contrastive Datasets

We used two contrastive datasets: ContraPro (Müller et al., 2018) for the English-to-German direction, and the Large Contrastive Pronoun Test-set (LCPT; Lopes et al. (2020)) for the English-to-French direction. Both are based on the OpenSubtitles 2018 dataset (Lison et al., 2018), and consist of the source sentence, the source- and target-side context with several translations differing only in a pronoun that requires context to be correctly translated. The details of the contrastive datasets are presented in Appendix C.

For each model, we first calculated the average attention scores assigned by each of the model’s heads to the relations of interest (as described in Section 3.1) for the examples from the contrastive datasets and calculated the point-biserial correlation between the measured attention scores given by each head to the relations of interest and the model correctly disambiguating a particular example (see Section 3.2). Next, we applied Modifying Heads method (see Section 3.3) to each head of the models. We used the following values of  $C$  (eq.2) [0.01, 0.25, 0.5, 0.75, 0.99] (uniformly probing the available range) for OpusMT en-de models, and [0.01, 0.99] for other models after observing that the relation between model’s ContraPro accuracy and modified value was mostly monotonic for all heads (see Appendix G for the detailed results). Additionally, we disabled each head by assigning the same probability to all tokens (see Appendix B for the detailed formulation) but we found that the results match the results for modifying the head to 0.01 without distinguishing the relation of interest.

## 5 Results

The base results in terms of accuracy and BLEU (Papineni et al., 2002) on contrastive datasets and BLEU on test subset of the IWSLT 2017 dataset are presented in Table 1. Since sentence-level models receive only the current sentence as input, we evaluate them by extracting examples from contrastive datasets where the antecedent is in the same sentence as the pronoun. This way we ensure that the

<sup>1</sup>The code for this paper can be found on Github <https://github.com/Pawel-M/context-mt-attention-analysis>.

<sup>2</sup><https://huggingface.co/Helsinki-NLP/opus-mt-en-de>

<sup>3</sup><https://huggingface.co/facebook/nllb-200-distilled-600M>

Model	Con- text	ContraPro Accuracy	ContraPro BLEU	IWSLT en-de	LCPT Accuracy	LCPT BLEU	IWSLT en-fr
Opus MT en-de	0	◇81.46%	◇30.24	32.42	-	-	-
	1	78.35%	31.00	34.40	-	-	-
	3	79.08%	31.35	34.59	-	-	-
NLLB-200	0	◇80.88%	◇29.46	32.31	◇95.02%	◇33.64	43.34
	1	77.95%	29.18	34.01	91.89%	31.49	43.86

Table 1: The results of the unmodified models. For the sentence-level models (with the context size of zero) we report accuracy and BLEU on contrastive datasets with the "antecedent distance" of zero (marked with ◇; meaning that the antecedent is located in the current sentence).

sentence-level models are tested exclusively on examples that they are capable to disambiguate. We do not filter the examples for context-aware models to allow the comparison between the tested models (as well as the results in the literature).

We analyzed the results in terms of average attention scores (Section 3.1), correlations between attention scores and correctly disambiguating an example (Section 3.2), and modifying heads (Section 3.3) to 0.01 and 0.99. The correlation is only partially corresponding to the changes in accuracy when modifying the attention scores, meaning that we find heads that demonstrate low correlation but have a high impact on the model’s performance when modified. We ignore the losses in accuracy when modified to 0.99 as possibly resulting from the decreased attention scores for other token-to-token relations, not investigated in this work. Similarly, we discard the increases in accuracy when modified to 0.01. We broadly categorized heads’ behavior for a particular relation into the following groups (note that the same head can be in different groups for different relations of interest):

- **irrelevant** - those heads do not attend the relation of interest and do not respond to the modifications;
- **attending and fully responsive** - those heads do attend the relation, the model’s accuracy decreases when modified to 0.01, and increases when modified to 0.99;
- **attending and negatively responsive** - the heads do attend the relation, and respond negatively to modifying to 0.01 but do not show improvement when modified to 0.99; we interpret those heads as already at the peak of their ability to help the model with pronoun disambiguation;
- **attending and positively responsive** - similar to the previous category but with only the positive response to modifying to 0.99; we hypothesize that the heads in this category do not attend the

relation to a sufficient degree for some examples and the correctly disambiguated examples are also supported by other heads in the model;

- **attending and non-responsive** - those heads do attend the relation but are not responsive to any modification; our interpretation is that those heads are important for other tasks than pronoun disambiguation;
- **non-attending and positively responsive** - the heads in this category are not attending the relation but positively influence the model’s accuracy when modified to 0.99.

In the following sections, we identify the heads by *a-l-h*, where *a* marks the attention module and can take values from {*e*, *c*, *d*} meaning the encoder-, cross-, and decoder-attention respectively, *l* and *h* mark the layer and head number. For example, *d-2-11* identifies the head number 11 in the self-attention module in the second decoder layer.

## 5.1 OpusMT en-de

We performed the analysis of the models based on OpusMT en-de (sentence-level, context-aware-1, and context-aware-3) on the ContraPro dataset. The sentence-level model was evaluated only on the examples where the antecedent was located in the current sentence. For the context-aware models, we employed the full dataset. We present the results for all three models in Figure 2, where for each head we show the metrics introduced in Section 3 (correlations, accuracy when modified to 0.01, and modified to 0.99) in relation to the averaged attention scores. The expanded results can be found in Appendix G. Additionally, to give a sense of the distribution of the results, we show histograms of the observed values of the metrics with annotated values corresponding to prominent heads in Appendix I.

Most heads do not pay - on average - a large at-

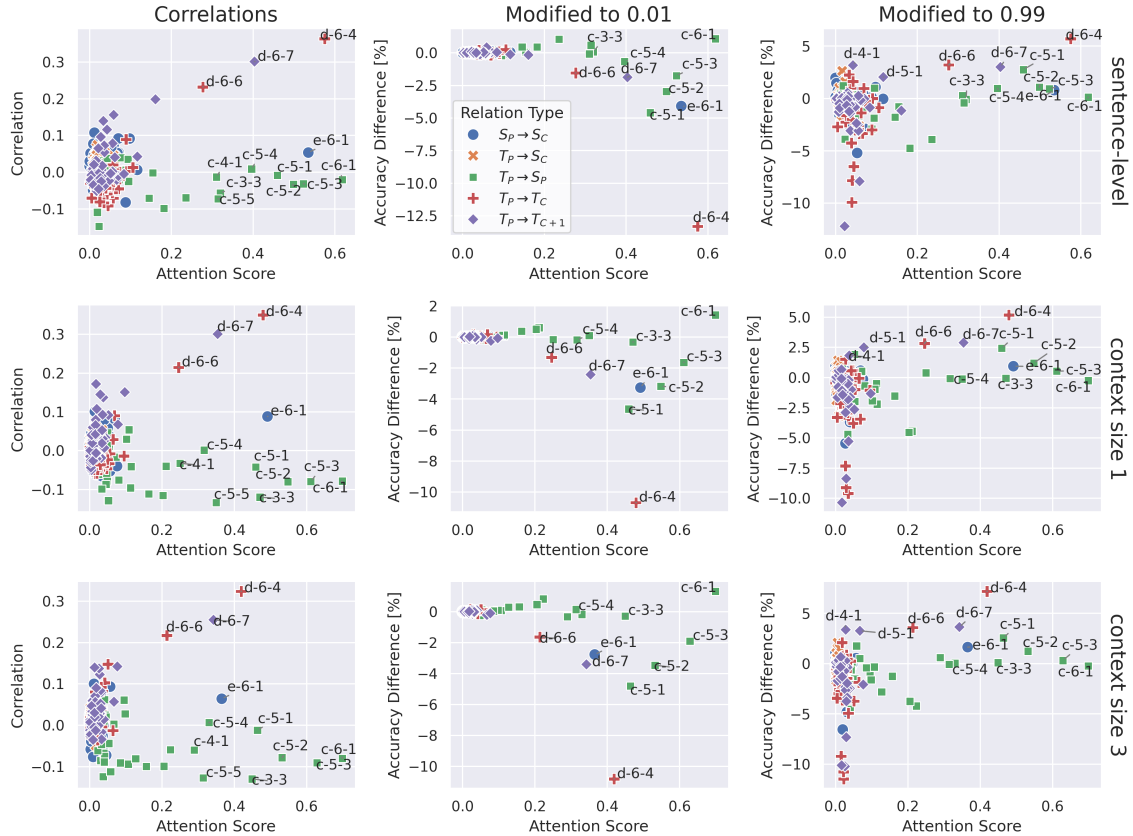


Figure 2: Results in terms of calculated metrics (correlations, difference in accuracy when modified to 0.01, and modified to 0.99; as columns) in relation to the averaged attention scores for the three models based on **OpusMT en-de** (sentence-level, context-aware-1, and context-aware-3; as rows).

attention to any of the relations of interest. Only a single encoder head (*e-6-1*) was found to assign high attention scores to the  $S_P \rightarrow S_C$  relation. It was negatively responsive for all models and positively responsive for the context-aware-3 model (increasing the accuracy by 1.6 percentage points). Head *e-2-8* was non-attending and positively responsive (improvement of up to 2 percentage points).

For cross-attention, we found that the  $T_P \rightarrow S_P$  relation is attended by several heads (most prominently heads *c-6-1*, *c-5-3*, *c-5-2*, *c-5-1*, *c-5-4*, and *c-3-3*). Out of them only heads *c-5-1* and *c-5-2* (to a lesser extent) are responsive. While no head put large attention to the  $T_P \rightarrow S_C$  relation, head *c-5-8* was positively responsive.

The most responsive heads in the whole model are located in the decoder’s attention. Two heads are attending and are responsive to the  $T_P \rightarrow T_C$  relation. The *d-6-4* head shows the highest correlation (above 0.3 for all models) and responsiveness (accuracy difference of below  $-10$  and above 5 percentage points for modifying to 0.01 and 0.99 respectively). The *d-6-6* head also shows rela-

tively high correlation and responsiveness but not to the same extent. For the  $T_P \rightarrow T_{C+1}$  relation, the *d-6-7* head is attending and responsive with high correlation. The negative responsiveness is higher for context-aware models compared to the sentence-level model. This indicates that attending the context-informative token on the input can be beneficial but is not crucial to the sentence-level model. The context-aware models learn to rely on this relation more heavily. Additionally, heads *d-4-1* and *d-5-1* are non-attentive and only positively responsive.

## 5.2 NLLB-200

The results of the analysis of the NLLB-200 context-aware models are presented in Figure 3. The expanded results, including the sentence-level model, can be found in Appendix H. The histograms of the results are presented in Appendix I.

### 5.2.1 English-to-German

For the English-to-German direction, only a single encoder head - *e-12-4* - was found to attend the




Figure 3: Results in terms of calculated metrics (correlations, difference in accuracy when modified to 0.01, and modified to 0.99; as columns) in relation to the averaged attention scores for the English-to-German and English-to-French directions (as rows) for the context-aware model based on **NLLB-200**.

$S_P \rightarrow S_C$  relation but was non-responsive. In the cross-attention modules, the  $T_P \rightarrow S_P$  relation is attended by several heads in both models (most notably by  $c-7-3$ ,  $c-8-3$ ,  $c-6-16$ ,  $c-10-13$ , and  $c-5-3$  for context-aware model only) but showed low correlation and responsiveness. For the context-aware model, two heads visibly changed their behavior - head  $c-9-1$  moderately attends the relation but does respond to modifying, and head  $c-8-15$  is non-attentive and only positively responsive. The  $T_P \rightarrow S_C$  relation was weakly attended by a single head ( $c-8-9$ ) which showed only the modest positive response.

Similar to OpusMT models, the most attentive and responsive heads are located in the decoder attention. Head  $d-10-3$  is moderately attending the  $T_P \rightarrow T_C$  relation, have relatively high correlation, and is highly responsive (improving the accuracy by almost 3 percentage points). Surprisingly, when this head is modified to 0.99 for the  $T_P \rightarrow T_{C+1}$  relation it improves the accuracy of the context-aware model. The  $d-9-9$  head attends  $T_P \rightarrow T_{C+1}$  relation and exhibits the highest correlation among all heads of the model and responds negatively. Several heads are non-attentive but positively responsive. Most notably, head  $d-10-6$  improves accuracy by 2 and 3 percentage points for sentence-level and context-aware models respectively, and head  $d-9-10$  by approximately 2 percentage points for both models. Interestingly, head  $d-9-12$  seems to pay attention to both relations and responds strongly

to modifying to 0.01. While it improves the accuracy for both relations when modified to 0.99, the improvement is larger (by 3 percentage points in both models) for the  $T_P \rightarrow T_C$  relation than for the  $T_P \rightarrow T_{C+1}$  relation with larger gains for context-aware model. Another head that pays attention to both relations is  $d-8-5$  but it shows moderate correlation and responsiveness.

## 5.2.2 English-to-French

For the English-to-French direction, the accuracy of the unmodified models on LCPT reaches 95% for the sentence-level model<sup>4</sup> and almost 92% for the context-aware model, which is considerably higher than on ContraPro. This could explain the lower ranges of responses to modifying heads and correlation coefficients observed in the results.

In contrast to the English-to-German direction, we observed a higher number encoder heads attending the  $S_P \rightarrow S_C$  relation (e.g.,  $e-12-4$ ,  $e-8-4$ ,  $e-6-13$ , and  $e-12-8$ ). Nevertheless, none of them responded to modifying. Similarly, several heads ( $c-6-16$ ,  $c-8-3$ ,  $c-5-3$ ,  $c-7-11$ ) were attending but non-responsive to the  $T_P \rightarrow S_P$  relation. For the  $T_P \rightarrow S_C$  relation, head  $c-8-6$  was non-attending and positively responsive. This behavior was exhibited by two decoder heads -  $d-10-3$  and  $d-10-6$  - for the  $T_P \rightarrow T_C$  relation. The  $T_P \rightarrow T_{C+1}$  relation was attended by three heads -  $d-8-5$  and  $d-9-12$

<sup>4</sup>For the sentence-level model we only consider examples where the antecedent is in the current sentence.




Figure 4: The difference in accuracy on ContraPro of the models based on OpusMT en-de (sentence-level, context-aware-1, and context-aware-3) for modifying and tuning selected heads.

heads also responding to modification, and  $d-9-9$  head being non-responsive. Additionally, head  $d-10-3$  was non-attending and positively responsive. Notice that this head is also responding to the  $T_P \rightarrow T_C$  relation, albeit to a lesser extent.

Several heads show similar behavior in both language directions. Those include:  $e-12-4$  in encoder-attention,  $c-8-3$ ,  $c-6-16$ ,  $c-5-3$ , and  $c-8-9$  in cross-attention, and  $d-9-12$ ,  $d-8-5$ , and  $d-10-3$  in decoder-attention.

### 5.3 Discussion

Taking into account the results for all models we make the following observations.

- **Some heads appear to have a function identifiable through analysis.** This is confirmed by previous research (Clark et al., 2019; Voita et al., 2019c; Tenney et al., 2019; Jo and Myaeng, 2020; Olsson et al., 2022), however we additionally demonstrate that adjustments to the functioning of heads (whether improved or diminished) leads to noticeable changes in models’ performance (in terms of the accuracy of pronoun disambiguation).

- **The decoder-attention (corresponding to the target-side context) has the highest impact on the pronoun disambiguation accuracy.** Note that we used the gold context (provided with the examples). In a real-world system, the context would come from the model’s predictions. It is interesting to note that the most relevant heads are located in higher layers. Our intuition is that in the decoder the output token is presumably decided in the layers closer to the output and only with this information can heads attempt to find the corresponding antecedent.

- **In the decoder-attention, the important context tokens are the tokens corresponding to both the antecedent being predicted (the  $T_P \rightarrow T_C$  relation) and being passed to the model as input (the  $T_P \rightarrow T_{C+1}$  relation) with most heads spe-**

cializing in attending one of the two and only some heads being able to utilize both.

- **Attending to the relation by a head does not necessarily imply it has an impact on the pronoun disambiguation accuracy.** The most responsive heads already attend the relations of interest, but there exist heads that are not utilized by the model that could improve the performance if were attending the relations.

Moreover, in the multi-lingual NLLB-200 model, we found several cross-lingual heads, exhibiting the same behavior for both tested language directions. Finally, we quantified the overlap of the improvement coming from modifying pairs of heads simultaneously to be lower than 30% for all models (see Appendix E).

## 6 Tuning Heads

After identifying the most responsive heads, we fine-tuned them to assess to what extent the augmented behavior of the heads can be solidified into the models’ parameters without compromising overall translation quality. To obtain the dataset containing the pronoun-antecedent pairs we applied the CTXPRO toolset (Wicks and Post, 2023) to the IWSLT 2017 en-de dataset (unrelated to the ContraPro dataset). The details can be found in Appendix F.

We tuned selected heads of the OpusMT en-de models. The results are shown in Figure 4. The models exhibit higher accuracy for each tuned head without reducing the translation quality (see Appendix F for the extended results). The improvement is the highest for the sentence-level model and reduces with increased context size. The low number of examples with antecedent distance of two or more in the training dataset could explain the reduction in improvement for the context-aware-3 model. Alternatively, the representations of the



tokens that are passed to the tuned head are insufficient for the head to attend the relation. We leave further investigation of this phenomenon and the exploration of the most efficient methods to tune the heads for future research.

## 7 Conclusions

In this paper, we researched the influence of the attention heads in the Context-aware MT models on the task of pronoun disambiguation in English-to-German and English-to-French language directions. We measured and modified the attention scores corresponding to the relations that could influence the prediction of a pronoun: pronoun-antecedent on the source and target sides, pronoun-pronoun, and pronoun-antecedent between the target and source sides.

We found that some heads do attend the relations of interest but not all of them influence the pronoun disambiguation capabilities of the models. We showed that some heads are *underutilized* by the models - the models' performance could improve if they attended the relations. We confirmed that the target-side context is more impactful than the source-side context. Additionally, we note that the target-side context cue tokens can be useful to the model both when passed to the model (as the input of the decoder) as well as when they are predicted (as the output of the decoder). Lastly, we showed that the heads can be fine-tuned to attend the relations improving the models.

## 8 Limitations

We only investigated the pre-trained models that we fine-tune for the Context-aware MT. The attention heads could behave differently if the models were trained from the random initialization. Additionally, we employed the contrastive datasets to analyze the models. The behavior of the models can differ from the generative setting. We also used the gold target context. In the real-world scenario, the models would base its predictions on the previously generated target context. Lastly, we only considered the pronoun disambiguation task and only two language directions: English-to-German and English-to-French.

## Acknowledgments

The research presented in this paper was conducted as part of VOXReality project<sup>5</sup>, which was funded

<sup>5</sup><https://voxreality.eu/>

by the European Union Horizon Europe program under grant agreement No 101070521.

## References

- Samira Abnar and Willem Zuidema. 2020. [Quantifying attention flow in transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197, Online. Association for Computational Linguistics.
- Ruchit Agrawal, Marco Turchi, and Matteo Negri. 2018. Contextual handling in neural machine translation: Look behind, ahead and on both sides. In *Proceedings of the 21st Annual Conference of the European Association for Machine Translation*, pages 31–40.
- 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, Online. Association for Computational Linguistics.
- Anthony Bau, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James R. Glass. 2018. [Identifying and controlling important neurons in neural machine translation](#). *ArXiv*, abs/1811.01157.
- 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.
- Nikolay Bogoychev. 2021. [Not all parameters are born equal: Attention is mostly what you need](#). In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 363–374, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Gino Brunner, Yang Liu, Damian Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. 2020. On identifiability in transformers. In *8th International Conference on Learning Representations (ICLR 2020)(virtual)*. International Conference on Learning Representations.
- Aydar Bulatov, Yury Kuratov, and Mikhail Burtsev. 2022. Recurrent memory transformer. *Advances in Neural Information Processing Systems*, 35:11079–11091.
- Mauro Cettolo, Marcello Federico, Luisa Bentivogli, Jan Niehues, Sebastian Stüker, Katsuhito Sudoh, Koichiro Yoshino, and Christian Federmann. 2017. [Overview of the IWSLT 2017 evaluation campaign](#). In *Proceedings of the 14th International Conference*

- on Spoken Language Translation, pages 2–14, Tokyo, Japan. International Workshop on Spoken Language Translation.
- Linqing Chen, Junhui Li, Zhengxian Gong, Min Zhang, and Guodong Zhou. 2022. [One type context is not enough: Global context-aware neural machine translation](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 21(6).
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Yukun Feng, Feng Li, Ziang Song, Boyuan Zheng, and Philipp Koehn. 2022. [Learn to remember: Transformer with recurrent memory for document-level machine translation](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1409–1420, Seattle, United States. Association for Computational Linguistics.
- Patrick Fernandes, Kayo Yin, Emmy Liu, André Martins, and Graham Neubig. 2023. [When does translation require context? a data-driven, multilingual exploration](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 606–626, Toronto, Canada. Association for Computational Linguistics.
- Patrick Fernandes, Kayo Yin, Graham Neubig, and André F. T. 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, Online. Association for Computational Linguistics.
- Javier Ferrando, Gerard I. Gállego, Belen Alastruey, Carlos Escolano, and Marta R. Costa-jussà. 2022. [Towards opening the black box of neural machine translation: Source and target interpretations of the transformer](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8756–8769, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Harritxu Gete, Thierry Etchegoyhen, and Gorka Labaka. 2023. [What works when in context-aware neural machine translation?](#) In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 147–156, Tampere, Finland. European Association for Machine Translation.
- Mozhdeh Gheini, Xiang Ren, and Jonathan May. 2021. [Cross-attention is all you need: Adapting pretrained Transformers for machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1754–1765, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Akshay Goindani and Manish Shrivastava. 2021. [A dynamic head importance computation mechanism for neural machine translation](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 454–462, Held Online. INCOMA Ltd.
- Christian Hardmeier. 2012. [Discourse in statistical machine translation: A survey and a case study](#). *Discours-Revue de linguistique, psycholinguistique et informatique*, 11.
- 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.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). 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 3543–3556, Minneapolis, Minnesota. 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*.
- Jae-young Jo and Sung-Hyon Myaeng. 2020. [Roles and utilization of attention heads in transformer-based neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3404–3417, Online. Association for Computational Linguistics.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online. Association for Computational Linguistics.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2021. [Incorporating Residual and Normalization Layers into Analysis of Masked Language Models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4547–4568, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Anna Langedijk, Hosein Mohebbi, Gabriele Sarti, Willem Zuidema, and Jaap Jumelet. 2024. [Decoder-Lens: Layerwise interpretation of encoder-decoder transformers](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4764–4780, Mexico City, Mexico. 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).
- 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.
- 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, Online. Association for Computational Linguistics.
- Andreas Madsen, Siva Reddy, and Sarath Chandar. 2022. [Post-hoc interpretability for neural nlp: A survey](#). *ACM Comput. Surv.*, 55(8).
- 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*.
- Paweł Maka, Yusuf Semerci, Jan Scholtes, and Gerassimos Spanakis. 2024. [Sequence shortening for context-aware machine translation](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1874–1894, St. Julian’s, Malta. 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.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2024. Locating and editing factual associations in gpt. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, Red Hook, NY, USA. Curran Associates Inc.
- 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.
- Wafaa Mohammed and Vlad Niculae. 2024. [On measuring context utilization in document-level MT systems](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1633–1643, St. Julian’s, Malta. Association for Computational Linguistics.
- Hosein Mohebbi, Willem Zuidema, Grzegorz Chrupała, and Afra Alishahi. 2023. [Quantifying context mixing in transformers](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3378–3400, Dubrovnik, Croatia. 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.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barraud, Gabriel Mejia-Gonzalez, Prangthip Hansanti, John Hoffman, Semaire Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova Dassarma, Tom Henighan, Benjamin Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, John Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom B. Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Christopher Olah. 2022. [In-context learning and induction heads](#). *ArXiv*, abs/2209.11895.
- 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 and Marcin Junczys-Dowmunt. 2023. Escaping the sentence-level paradigm in machine translation. *arXiv preprint arXiv:2304.12959*.
- Gabriele Sarti, Grzegorz Chrupała, Malvina Nissim, and Arianna Bisazza. 2023. Quantifying the plausibility of context reliance in neural machine translation. *arXiv preprint arXiv:2310.01188*.

- Noam M. Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). *ArXiv*, abs/1804.04235.
- Zwei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Lei Li. 2022. [Re-thinking document-level neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3537–3548, Dublin, Ireland. Association for Computational Linguistics.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Jörg Tiedemann, Mikko Aulamo, Daria Bakshandaeva, Michele Boggia, Stig-Arne Grönroos, Tommi Nieminen, Alessandro Raganato Yves Scherrer, Raul Vazquez, and Sami Virpioja. 2023. [Democratizing neural machine translation with OPUS-MT](#). *Language Resources and Evaluation*, (58):713–755.
- 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.
- Jörg Tiedemann and Santhosh Thottingal. 2020. OPUS-MT — Building open translation services for the World. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation (EAMT)*, Lisbon, Portugal.
- Mariya Toneva and Leila Wehbe. 2019. *Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain)*. Curran Associates Inc., Red Hook, NY, USA.
- Zhaopeng Tu, Yang Liu, Zhengdong Lu, Xiaohua Liu, and Hang Li. 2017. [Context gates for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 5:87–99.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Jesse Vig and Yonatan Belinkov. 2019. [Analyzing the structure of attention in a transformer language model](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 63–76, Florence, Italy. Association for Computational Linguistics.
- 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, Rico Sennrich, and Ivan Titov. 2021. [Analyzing the source and target contributions to predictions in neural 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 1126–1140, Online. Association for Computational Linguistics.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019c. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Rachel Wicks and Matt Post. 2023. [Identifying context-dependent translations for evaluation set production](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 452–467, Singapore. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). 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 11–20, Hong Kong, China. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Billy T. M. Wong and Chunyu Kit. 2012. [Extending machine translation evaluation metrics with lexical cohesion to document level](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1060–1068, Jeju Island, Korea. Association for Computational Linguistics.
- Kayo Yin, Patrick Fernandes, Danish Pruthi, Aditi Chaudhary, André F. T. Martins, and Graham Neubig. 2021. [Do context-aware translation models pay](#)

the right attention? 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 788–801, Online. Association for Computational Linguistics.

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.

Zaixiang Zheng, Xiang Yue, Shujian Huang, Jiajun Chen, and Alexandra Birch. 2021. Towards making the most of context in neural machine translation. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3983–3989.

## A Derivation of Modifying Heads Formula

In this section, we will provide a short derivation of the Modifying Heads equation (eq. 2) (see Section 3.3). Let us consider the  $i$ -th token in the left-hand side set  $\mathcal{Y}$  of the relation of interest  $\mathcal{Y} \rightarrow \mathcal{X}$  for the example  $d$  from the dataset. For a selected layer  $l$  and head  $h$ , the goal is to find the values of pre-softmax attention scores  $\tilde{H}_{i,j}^{l,h,d}$  for all tokens  $j$  from the right-hand side set  $\mathcal{X}$  of the relation of interest such that the sum of the resulting attention scores (post-softmax)  $\tilde{Z}_{i,j}^{l,h,d}$  would be equal to a desired value  $C$ . We also assume that the attention scores would be spread equally among the tokens from the right-hand side set  $\mathcal{X}$ . This can be formulated as:

$$C = \sum_{j' \in \mathcal{X}^d} \tilde{Z}_{i,j'}^{l,h,d} = |\mathcal{X}^d| \tilde{Z}_{i,j}^{l,h,d} \quad (3)$$

$$\forall i \in \mathcal{Y}^d, j \in \mathcal{X}^d,$$

The attention scores are calculated using the following equation:

$$Z_{i,k}^{l,h,d} = \frac{\exp(H_{i,k}^{l,h,d})}{\sum_{k' \in \mathcal{X}^d} \exp(H_{i,k'}^{l,h,d})} \quad (4)$$

$$\forall i \in \mathcal{Y}^d, k \in \mathcal{X}^d,$$

where  $X_d$  are the tokens from which the the right-hand side set  $\mathcal{X}$  are selected. We expand eq.3 using eq.4 and split the sum in the denominator into two

parts - inside  $\mathcal{X}^d$  and outside  $k \in X^d \setminus \mathcal{X}^d$ :

$$C = |\mathcal{X}^d| \tilde{Z}_{i,j}^{l,h,d}$$

$$= \frac{|\mathcal{X}^d| \exp(\tilde{H}_{i,j}^{l,h,d})}{|\mathcal{X}^d| \exp(\tilde{H}_{i,j}^{l,h,d}) + \sum_{k \in X^d \setminus \mathcal{X}^d} \exp(H_{i,k}^{l,h,d})}$$

$$\forall i \in \mathcal{Y}^d, j \in \mathcal{X}^d. \quad (5)$$

This equation can be reformulated as:

$$\exp(\tilde{H}_{i,j}^{l,h,d}) = \frac{C}{(1-C)|\mathcal{X}^d|} \sum_{k \in X^d \setminus \mathcal{X}^d} \exp(H_{i,k}^{l,h,d})$$

$$\forall i \in \mathcal{Y}^d, j \in \mathcal{X}^d, \quad (6)$$

Finally, we apply the logarithm to both sides of eq.6 and come to the equation for Modifying Heads from Section 3.3:

$$\tilde{H}_{i,j}^{l,h,d} = \log\left(\frac{C}{|\mathcal{X}^d|(1-C)} \sum_{k \in X^d \setminus \mathcal{X}^d} \exp(H_{i,k}^{l,h,d})\right)$$

$$\forall i \in \mathcal{Y}^d, j \in \mathcal{X}^d. \quad (7)$$

## B Disabling Heads Method

In addition to the methods described in Section 3 we disabled heads of the models. Disabling heads means that we assign equal attention score to all key  $K$  tokens, thus they cannot function as before. It can be done for all query  $Q$  tokens or only for the selected ones. To minimize the disturbance to the model, we only disable heads for the Q tokens in the  $\mathcal{Y}$  set. This allows the head to function normally for all other tokens but prevents it from attending to the contextually informative tokens from the contextually dependent tokens. Regardless, it still does not distinguish between relations of interest. For example, disabling a head in the cross-attention prevents it from attending both  $T_P \rightarrow S_P$  and  $T_P \rightarrow S_C$  relations reducing the granularity of the analysis.

For an example  $d$  from the dataset, disabling heads is defined as:

$$\tilde{Z}_{i,k}^{l,h,d} = \frac{1}{|X^d|}$$

$$\forall i \in \mathcal{Y}^d, k \in X^d, \quad (8)$$

where  $\tilde{Z}$  are the updated attention scores (after softmax is applied),  $X^d$  is the whole set of attended tokens, including the attended tokens-of-interest  $\mathcal{X}^d$ . Implementation-wise, we set pre-softmax attention scores  $H$  to zero and rely on the softmax

function to spread attention equally among the tokens.

### C Details of the Datasets

In this section, we present the details of the datasets used in our investigation. The number of documents, sentences, and tokens in the subsets of the IWSLT 2017 (Cettolo et al., 2017) dataset for English-to-German and English-to-French directions can be seen in Table 2.

Dataset	Docs	Sent/Doc	Tok/Sent
En-De Train	1698	121.4	21.9
En-De Valid	62	87.6	20.6
En-De Test	12	90.0	20.8
En-Fr Train	1914	121.6	22.0
En-Fr Valid	66	88.2	20.9
En-Fr Test	12	100.8	21.4

Table 2: The details of the IWSLT 2017 datasets.

The number of examples in the ContraPro and LCPT contrastive datasets for the different distances of the antecedent are presented in Table 3.

Antecedent Distance	ContraPro	LCPT
0	2400	5986
1	7075	4566
2	1510	1629
3	573	880
>3	442	939
All	12000	14000

Table 3: The number of examples in the ContraPro and LCPT datasets with different distances of the antecedent.

### D Details of Context-aware Fine-tuning

We trained the models with Adafactor optimizer (Shazeer and Stern, 2018) on a single GPU (NVIDIA GeForce RTX 3090 24GB) for 10 epochs and select the checkpoint with the highest BLEU (Papineni et al., 2002) on the ContraPro dataset (Müller et al., 2018) (translating the sentences with the provided target context). The models acquired the highest BLEU score after two epochs for OpusMT en-de with a context window of one, and only after one epoch for OpusMT en-de with a context window of three and NLLB-200 (with a context window of one).

The hyper-parameters are presented in Table 4. For OpusMT en-de with the context size of three, we decreased the batch size to 16 and increased the gradient accumulation steps to 16. For NLLB-200, we decreased the batch size further to 8 and increased the gradient accumulation steps to 32. This was done to reduce the memory requirements during training. We trained for 10 epochs and selected the checkpoint with the highest BLEU score on the ContraPro dataset. We tuned learning rate on OpusMT en-de model for context size of one and tried the following values:  $2e-6$ ,  $5e-6$ ,  $1e-5$ ,  $2e-5$ ,  $5e-5$ . We chose the value of  $1e-5$  for learning rate as leading to the highest BLEU on the ContraPro dataset (when generating translations). We used this value for the fine-tuning of all models.

Hyper-parameter	Value
Optimizer	Adafactor
Learning Rate	$1e-5$
LR Scheduler	Linear
LR Warmup Ratio	0
Weight Decay	0.01
Max Gradient Norm	1.0
Batch Size	32*
Gradient Accumulation Steps	8*
Num Epoch	10

Table 4: The hyper-parameters of the Context-aware fine-tuning of the tested models.

\* For Opus MT en-de with the context size of three and NLLB-200, we decreased the batch size to 16 and 8, and increased the gradient accumulation steps to 16 and 32 respectively.

### E Modifying Multiple Heads

For all investigated models, we found that several heads are improving the accuracy of the model when modified to attend the relations of interest. To assess to what extent the improvement of different heads is overlapping we selected several heads with highest improvement in accuracy when modified to 0.99 and modified each pair of heads together to 0.99 for the relations they responded. We experimented with two models: OpusMT en-de with context size of three (we selected 8 heads), and context-aware NLLB-200 model on English-to-German and English-to-French directions (we selected 7 and 5 heads respectively). We report the percentage overlap of the accuracy improvement of modifying both heads against the sum of improvements of modifying each head separately,

calculated as:

$$O_{\mathcal{M}} = \frac{\sum_{(l,h) \in \mathcal{M}} (A_{l,h} - A_B) - (A_{\mathcal{M}} - A_B)}{\sum_{(l,h) \in \mathcal{M}} (A_{l,h} - A_B)}, \quad (9)$$

where  $\mathcal{M}$  is the list of heads being simultaneously modified,  $A_B$  is the base accuracy of the model,  $A_{\mathcal{M}}$  is the accuracy observed when all heads are modified to 0.99, and  $A_{l,h}$  is the accuracy of the model with a single  $h$ -th head in  $l$ -th layer being modified to 0.99.

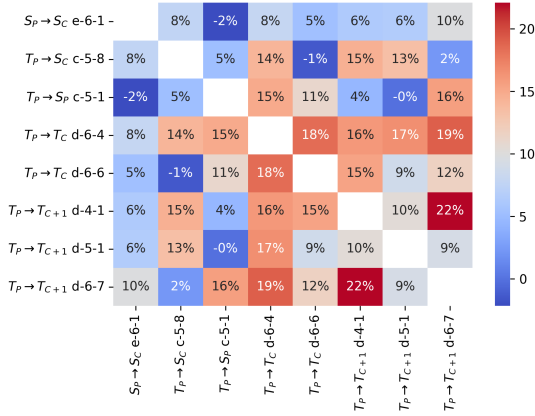


Figure 5: The percentage overlap of the improvement in accuracy on the ContraPro dataset for selected heads of the context-aware-1 model based on **OpusMT en-de**.




Figure 6: The percentage overlap of the improvement in accuracy on the **ContraPro** dataset (English-to-German) for selected heads of the context-aware model based on **NLLB-200**.

We present the results for the OpusMT en-de model in Figure 5. The corresponding results for NLLB-200 model can be found in in Figures 6 and 7 for English-to-German and English-to-French directions respectively. We do not report the result for the cases where the same head was selected based

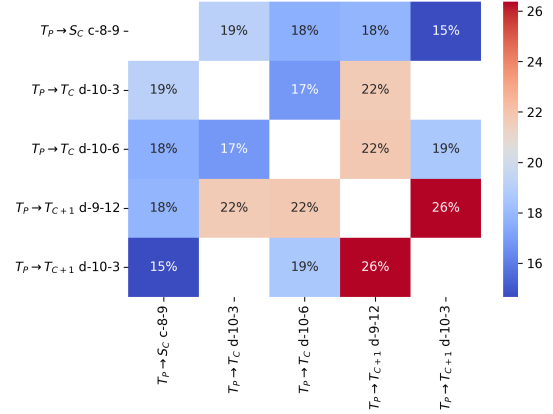


Figure 7: The percentage overlap of the improvement in accuracy on the **LCPT** dataset (English-to-French) for selected heads of the context-aware model based on **NLLB-200**.

on two different relations (e.g., head  $d-9-12$  for the  $T_P \rightarrow T_C$  and  $T_P \rightarrow T_{C+1}$  relations). It can be seen that some overlap is present for most of the heads but it is typically lower than 20% for English-to-German direction in both models and slightly higher (from 15% to 26%) for English-to-French direction in NLLB-200. This means that heads perform complementary (to an extent) functions in the models, and the redundancy of the function they perform is limited. Furthermore, improving heads in tandem could lead to the model with a better performance.

## F Details of Head Tuning

We trained the selected heads of a model to generate the attention scores matching the scores resulting from modifying to 0.99. We froze all parameters of the model apart from the parameters responsible for the Q and K transformation in a targeted layer and updated only the parameters corresponding to the targeted head. Although we experimented with the calculation of the loss based on post-softmax attention scores  $Z$ , we observed a slight decrease in BLEU on the IWSLT 2017 test-set. Therefore, we decided to use the pre-softmax attention scores  $H$  to calculate the loss. We formulate the target pre-softmax attention scores for a token  $i \in \mathcal{Y}$  as:

$$\hat{H}_{i,j}^{l,h,d} = \begin{cases} \tilde{H}_{i,j}^{l,h,d}, & \text{if } j \in \mathcal{X}, \\ \bar{H}_{i,j}^{l,h,d}, & \text{if } j \notin \mathcal{X}, \end{cases} \quad (10)$$

where  $\tilde{H}$  is calculated according to eq. 2 (see Section 3.3), and  $\bar{H}$  is the attention score obtained from

the unmodified model. In order to obtain the values of  $\bar{H}$ , at the start of training, we make a copy of the model parameters and perform a forward pass twice, first for the trained model and second for the unmodified model. We calculate the MSE loss between the model’s and target pre-softmax scores and backpropagate it to the model’s parameters.

We obtain training examples by applying the CTXPRO (Wicks and Post, 2023) toolset to the training subset of the IWSLT 2017 en-de dataset. The toolset applies predefined rules to find context-related phenomena in the dataset. We extracted the examples marked by the following rules:

- ACC.NEUT.SING,
- DAT.FEM.SING,
- NOM.FEM.SING,
- NOM.NEUT.SING,
- DAT.NEUT.SING,
- DAT.MASC.SING,
- NOM.MASC.SING,
- ACC.FEM.SING,
- ACC.MASC.SING.

The resulting number of examples with different antecedent distances are presented in Table 5. The hyper-parameters used in fine-tuning can be found in Table 7.

Antecedent Distance	CTXPRO
0	3816
1	2883
2	845
3	355
>3	448
All	8347

Table 5: The number of examples in the CTXPRO dataset.

## G Expanded Results of the OpusMT en-de Models

In this section, we present the expanded results for the models based on OpusMT en-de (sentence-level, context-aware-1, and context-aware-3). The accuracy of the context-aware model with the context size of one on the ContraPro contrastive dataset when modifying heads to different values (from the set [0.01, 0.25, 0.5, 0.75, 0.99]) for five relations of interest can be seen in Figure 8. The results show that the changes in accuracy are monotonic.

The raw results for all measured quantities (av-

erage attention scores, correlation between scores and being correct on the contrastive dataset examples, the difference in accuracy when disabling heads, and modifying head to 0.01 and 0.99) for all relations of interest are presented in Figures 9, 10, and 11 for the three models (sentence-level, context-aware-1, and context-aware-3) respectively.

## H Expanded Results of the NLLB-200 Models

Here, we show the expanded results for the two multi-lingual models (sentence-level and context-aware) based on NLLB-200. In Section 5.2 we presented the measured metrics introduced in Section 3 (correlations, accuracy when modified to 0.01, and modified to 0.99) in relation to the averaged attention scores only for the context-aware model. The results for both models and both language directions can be seen in Figure 12.

We show the raw results for all relations of interest in: Figure 13 for the sentence-level model on ContraPro (English-to-German) dataset, Figure 14 for the context-aware model on ContraPro dataset, Figure 15 for the sentence-level model on LCPT (English-to-French) dataset, and Figure 16 for the context-aware model on LCPT dataset.

## I Histograms of the Results

For each model, we generated histograms of the measured metrics and annotate with arrows the values for the most notable heads. Figures 17, 18, and 19 show the histograms for OpusMT en-de models (sentence-level, context-aware-1, and context-aware-3 respectively). The histograms for sentence-level NLLB-200 can be found in Figures 20 and 22 for the English-to-German and English-to-French directions respectively. The corresponding histograms of the context-aware NLLB-200 model are presented in Figures 21 and 23 for English-to-German and English-to-French respectively.



Model	Relation	Tuned Head	ContraPro Accuracy	Modified Accuracy	ContraPro BLEU	IWSLT BLEU
Sentence-level	-	-	81.46%	-	30.24	32.42
	$T_P \rightarrow T_C$	<i>d-6-4</i>	86.42%	87.17%	30.67	32.39
	$T_P \rightarrow T_C$	<i>d-6-6</i>	84.08%	84.67%	30.40	32.37
	$T_P \rightarrow T_{C+1}$	<i>d-6-7</i>	84.25%	84.46%	30.40	32.38
	$T_P \rightarrow S_P$	<i>c-5-1</i>	83.75%	84.21%	30.52	32.31
Context-aware-1	-	-	78.35%	-	31.00	34.40
	$T_P \rightarrow T_C$	<i>d-6-4</i>	82.08%	83.53%	31.17	34.50
	$T_P \rightarrow T_C$	<i>d-6-6</i>	80.01%	81.18%	31.06	34.55
	$T_P \rightarrow T_{C+1}$	<i>d-6-7</i>	80.69%	81.24%	31.02	34.47
	$T_P \rightarrow S_P$	<i>c-5-1</i>	80.68%	80.78%	31.15	34.47
Context-aware-3	-	-	79.08%	-	31.35	34.59
	$T_P \rightarrow T_C$	<i>d-6-4</i>	82.04%	86.28%	31.62	34.58
	$T_P \rightarrow T_C$	<i>d-6-6</i>	80.18%	82.68%	31.41	34.56
	$T_P \rightarrow T_{C+1}$	<i>d-6-7</i>	80.94%	82.72%	31.40	34.58
	$T_P \rightarrow S_P$	<i>c-5-1</i>	81.23%	81.64%	31.52	34.59

Table 6: The results of fine-tuning selected heads to attend the specified relation in terms of the accuracy and BLEU on ContraPro (the *Tuned Accuracy* and *ContraPro BLEU* columns respectively), BLEU on IWSLT 2017 en-de dataset (the *IWSLT BLEU* column) compared to the accuracy of the model with modifying the head to 0.99 (the *Modified Accuracy* column). The first row for each model shows the results of the unmodified and not fine-tuned model. The Accuracy and BLEU on the ContraPro dataset for the sentence-level models are calculated only on the examples where the antecedent is in the current sentence.

Hyper-parameter	Value
Optimizer	Adafactor
Learning Rate	1e-3
LR Scheduler	Linear
LR Warmup Ratio	0
Weight Decay	0.01
Max Gradient Norm	1.0
Batch Size	12
Gradient Accumulation Steps	8
Num Epoch	10

Table 7: The hyper-parameters of the fine-tuning of selected heads.

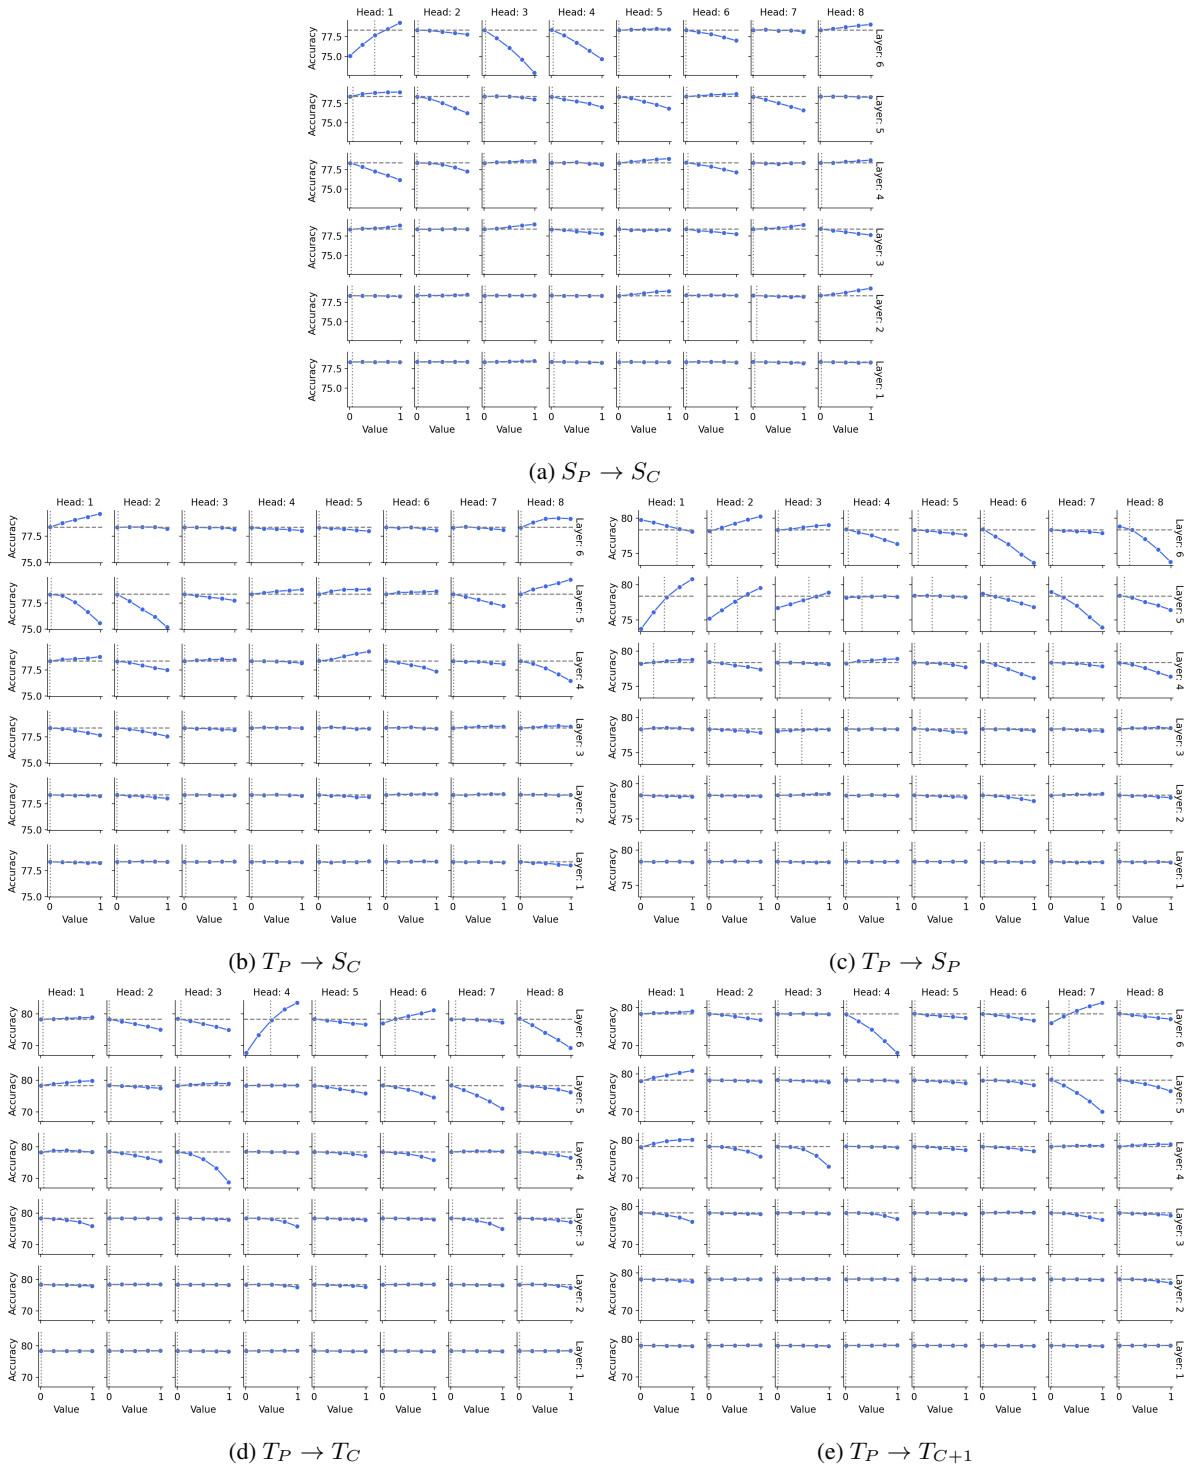


Figure 8: The accuracy of the **OpusMT en-de context-aware-1** model on the ContraPro dataset when modifying each of the attention heads to different values (out of  $[0.01, 0.25, 0.5, 0.75, 0.99]$ ) for the investigated relations of interest.

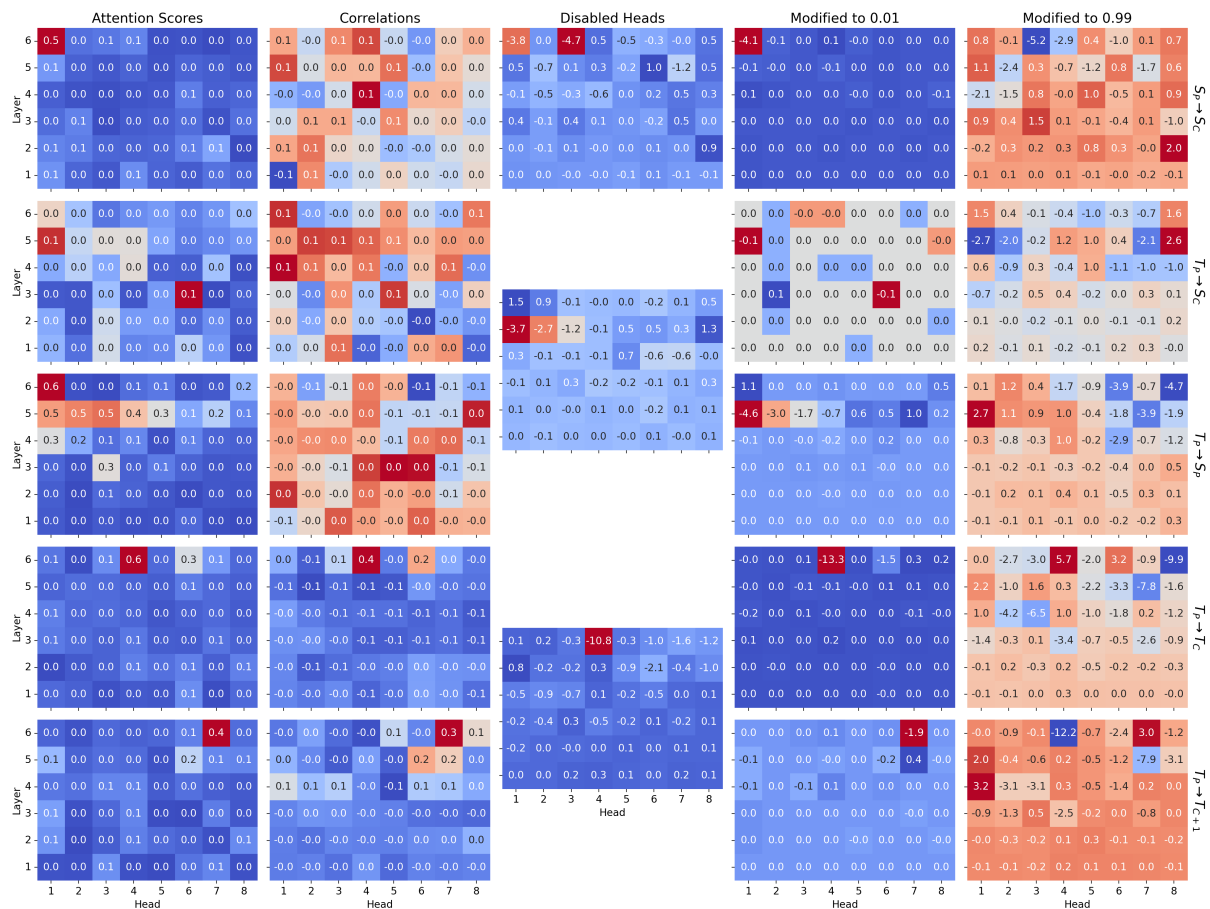


Figure 9: Measured average attention score, correlation between attention scores and accuracy on ContraPro dataset, and the difference in accuracy for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the **sentence-level OpusMT en-de** model for all relations of interest (as rows).

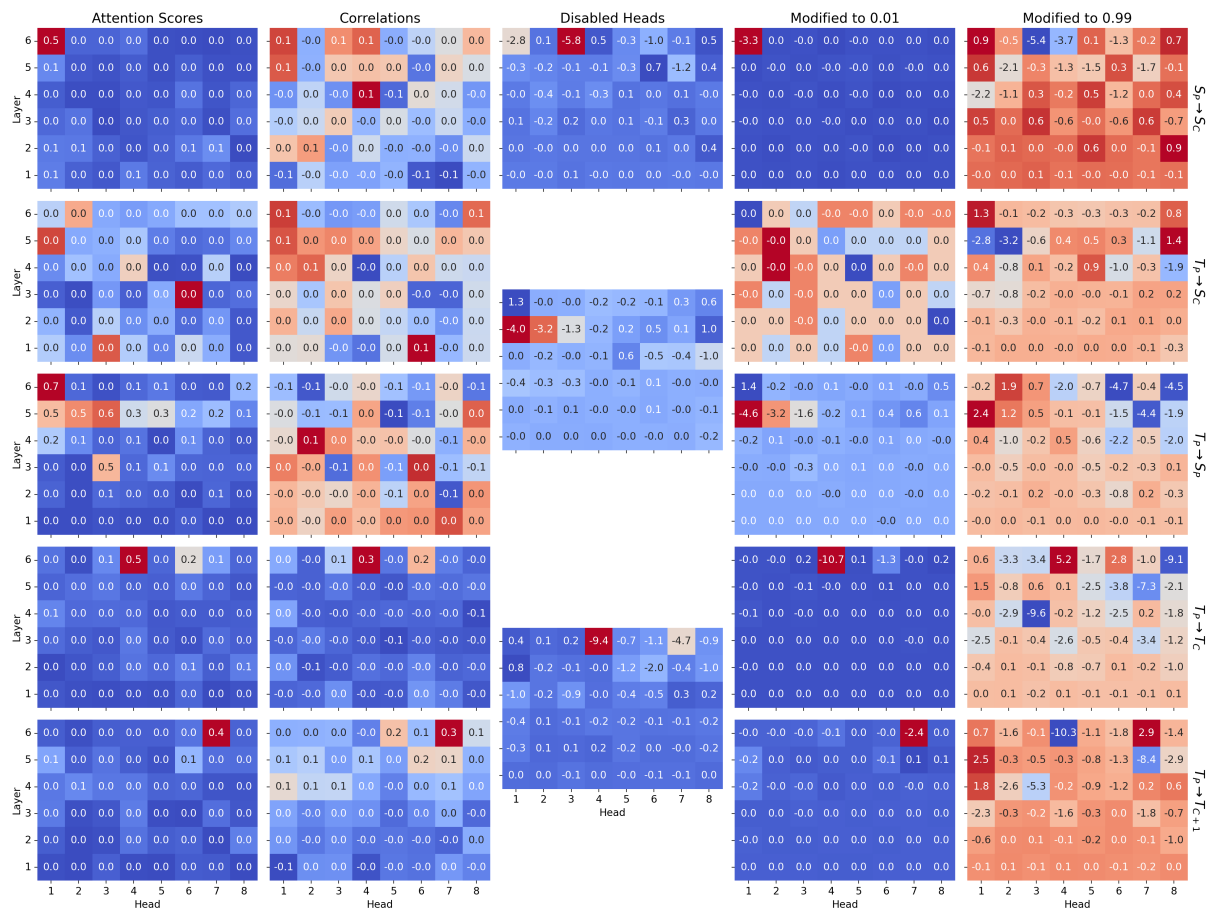


Figure 10: Measured average attention score, correlation between attention scores and accuracy on ContraPro dataset, and the difference in accuracy for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the **context-aware OpusMT en-de** model with the **context size of one** for all relations of interest (as rows).

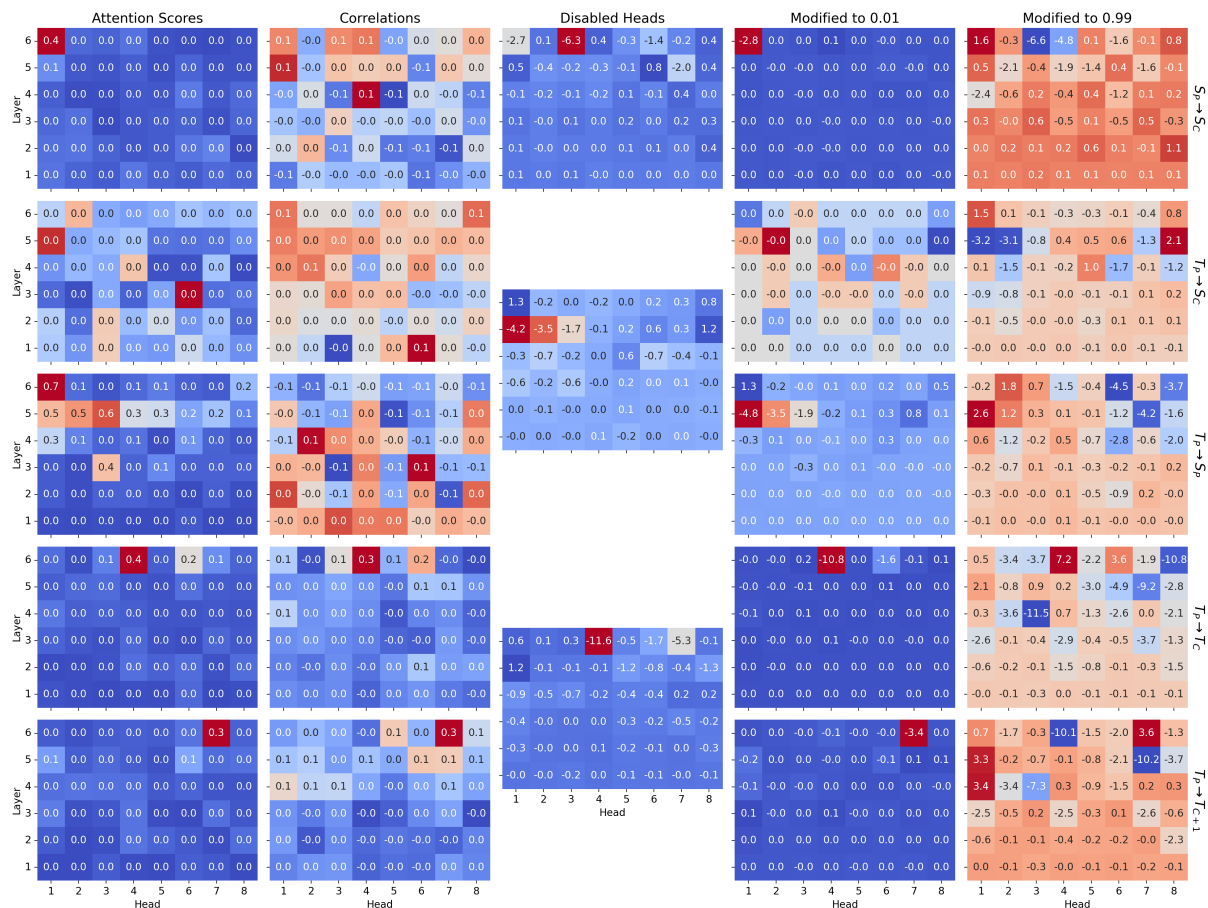


Figure 11: Measured average attention score, correlation between attention scores and accuracy on ContraPro dataset, and the difference in accuracy for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the context-aware OpusMT en-de model with the context size of three for all relations of interest (as rows).

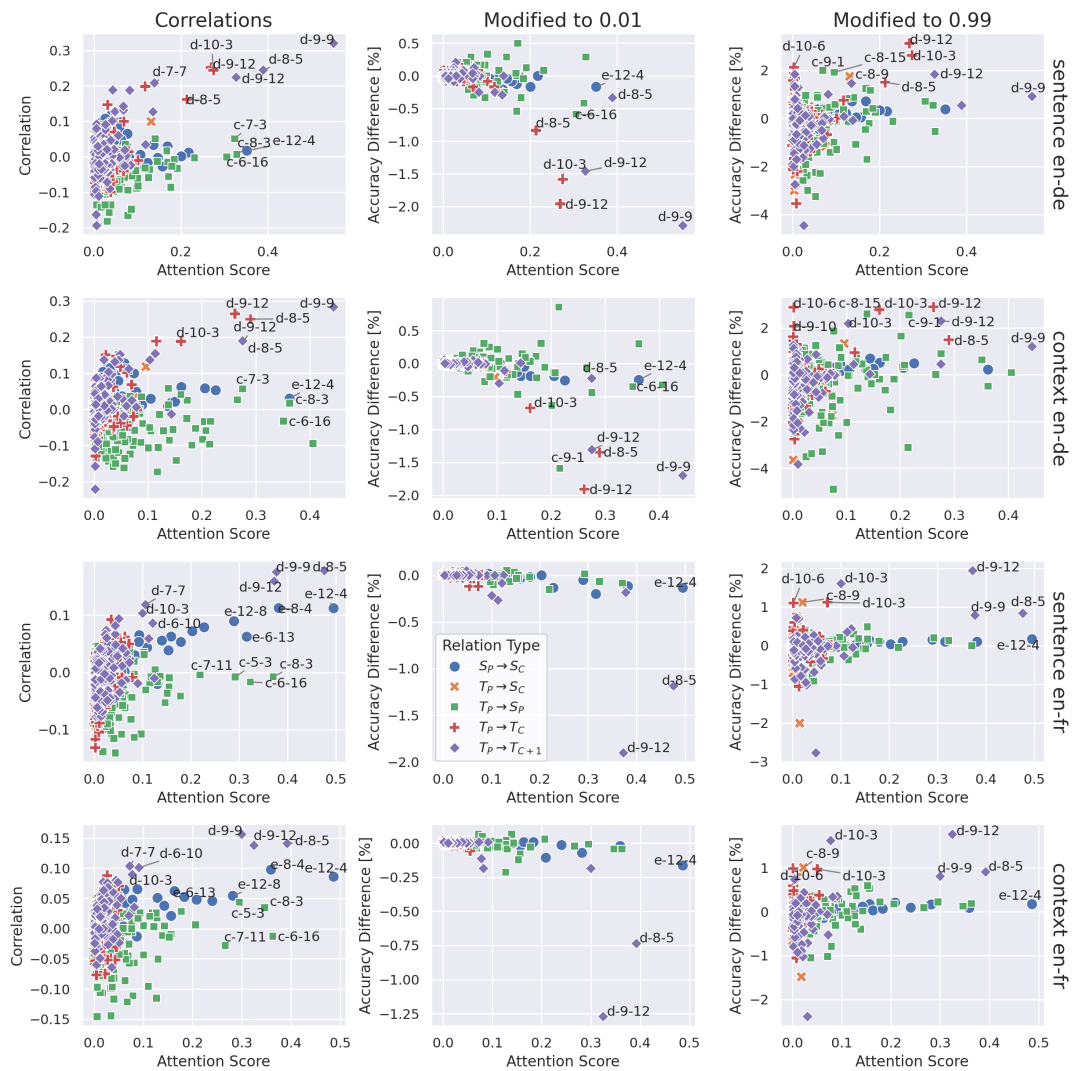


Figure 12: Results in terms of calculated metrics (correlations, difference in accuracy when modified to 0.01, and modified to 0.99; as columns) in relation to the averaged attention scores for the English-to-German and English-to-French directions for both sentence-level and context-aware models (as rows) based on **NLLB-200**.




Figure 13: Measured average attention score, correlation between attention scores and accuracy on **ContraPro** (English-to-German) dataset, and the difference in accuracy for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the **sentence-level NLLB-200** model for all relations of interest (as rows).




Figure 14: Measured average attention score, correlation between attention scores and accuracy on **ContraPro** (English-to-German) dataset, and the difference in for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the **context-aware NLLB-200** model for all relations of interest (as rows).






Figure 15: Measured average attention score, correlation between attention scores and accuracy on LCPT (English-to-French) dataset, and the difference in accuracy for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the sentence-level NLLB-200 model for all relations of interest (as rows).

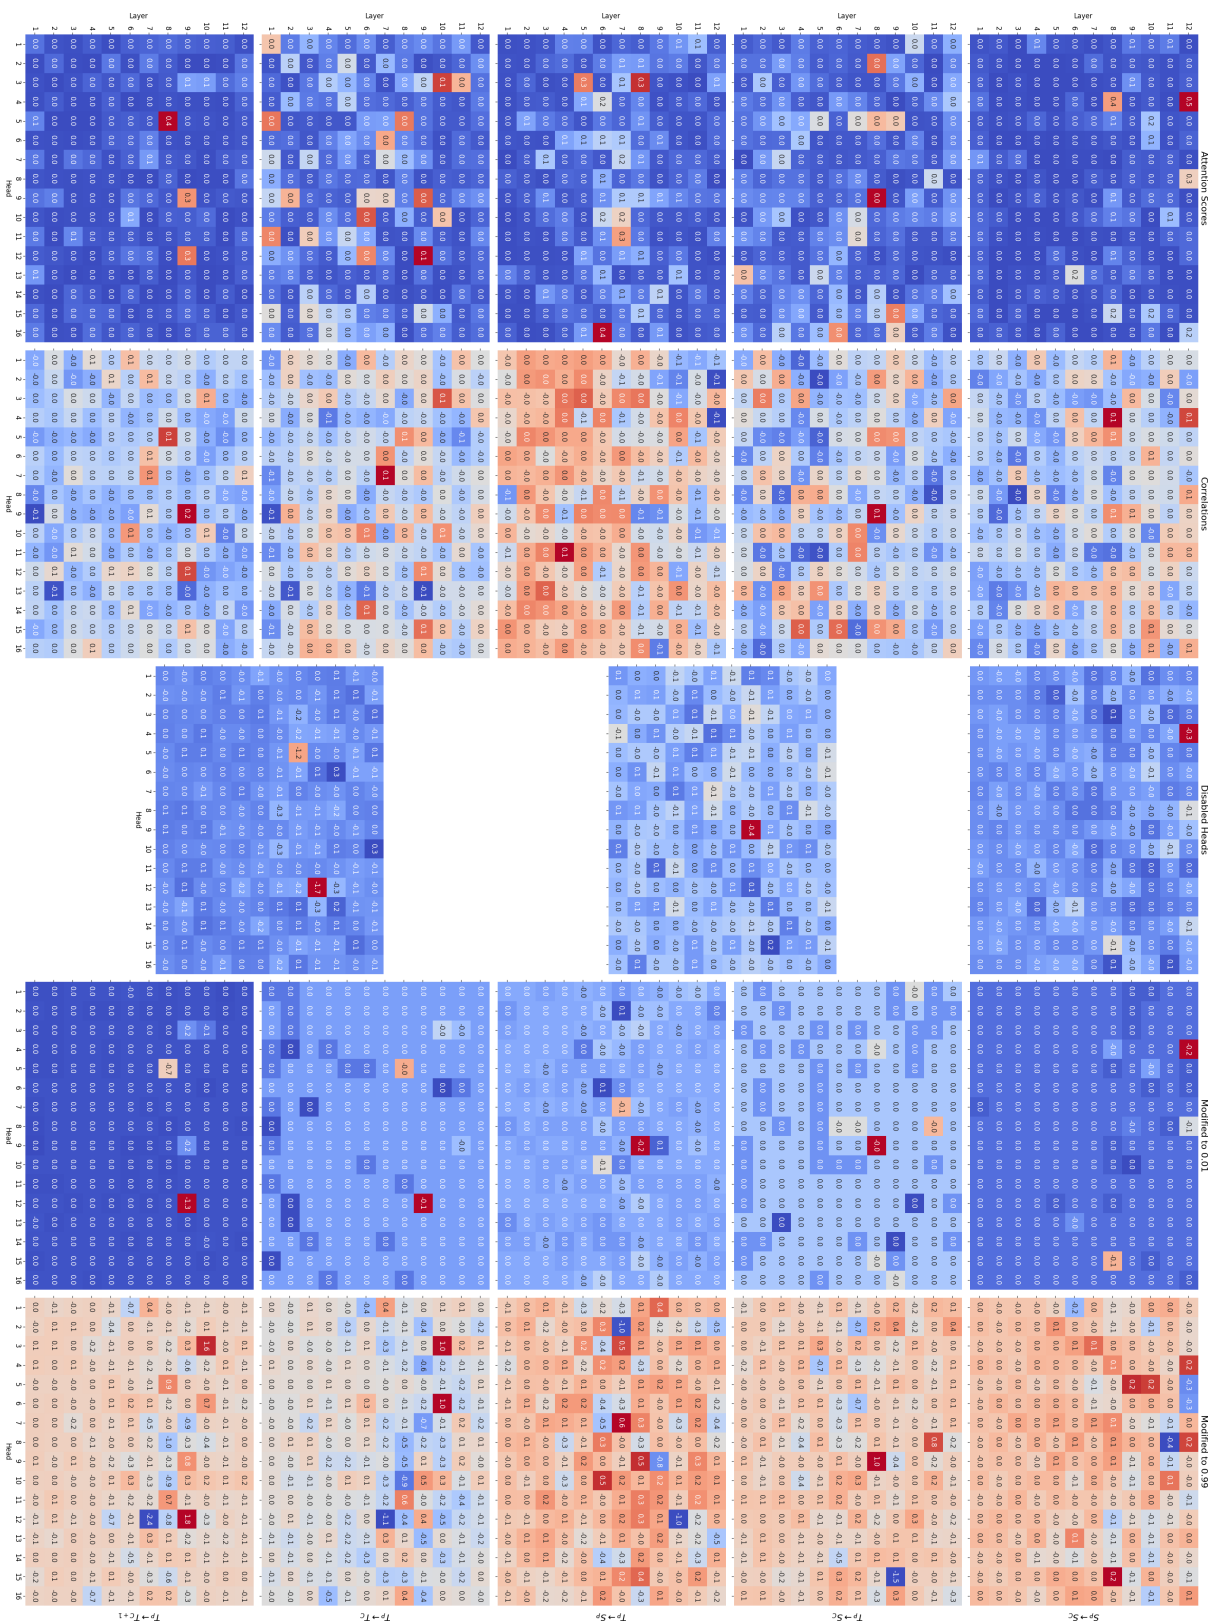


Figure 16: Measured average attention score, correlation between attention scores and accuracy on LCPT (English-to-French) dataset, and the difference in accuracy for: Disabling Heads, Modifying Heads to 0.01, and Modifying Heads to 0.99 (as columns) for each head of the **context-aware NLLB-200** model for all relations of interest (as rows).

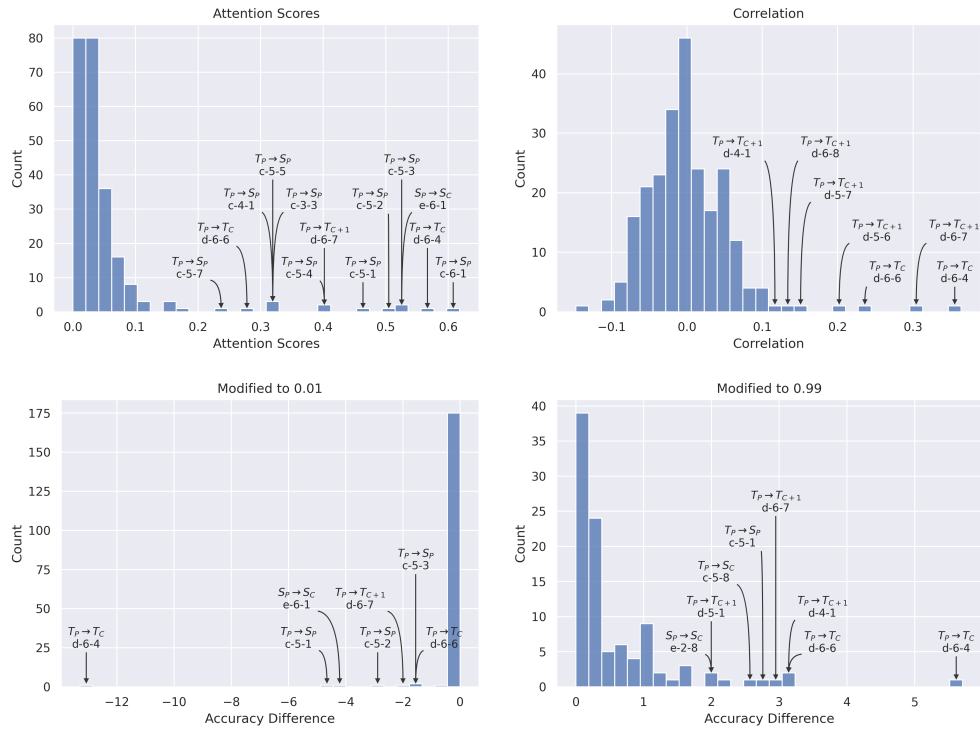


Figure 17: Histograms of average attention scores, correlations between attention scores and accuracy on ContraPro dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **sentence-level OpusMT en-de** model with the values of noticeable heads annotated with arrows.

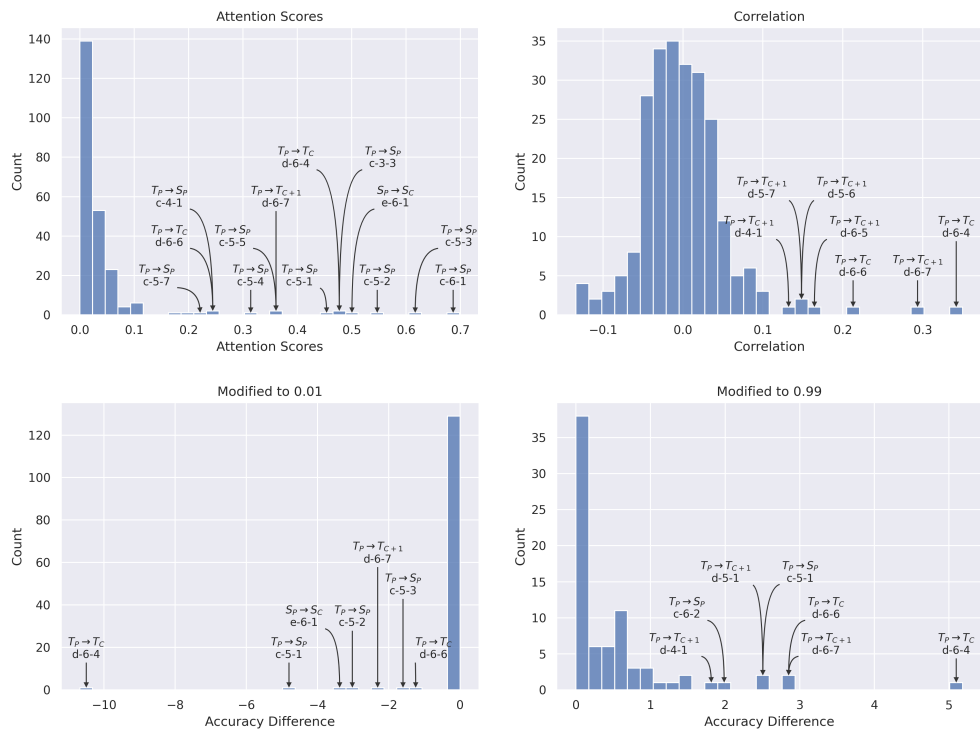


Figure 18: Histograms of average attention scores, correlations between attention scores and accuracy on ContraPro dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **context-aware OpusMT en-de** model with the **context size of one** with the values of noticeable heads annotated with arrows.

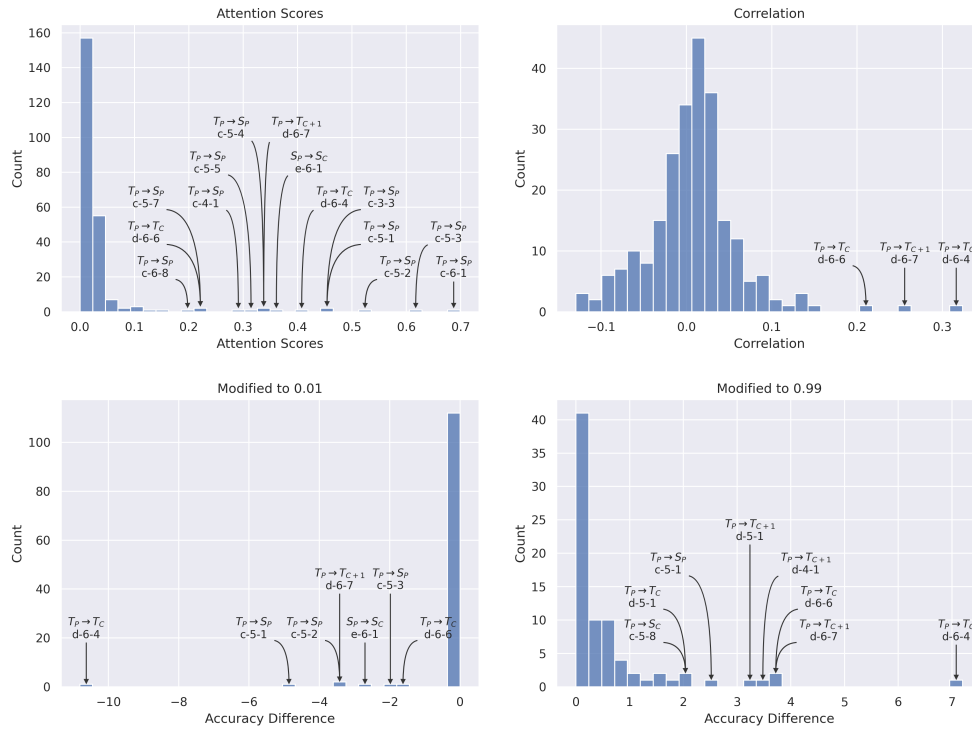


Figure 19: Histograms of average attention scores, correlations between attention scores and accuracy on ContraPro dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **context-aware OpusMT en-de** model with the **context size of three** with the values of noticeable heads annotated with arrows.

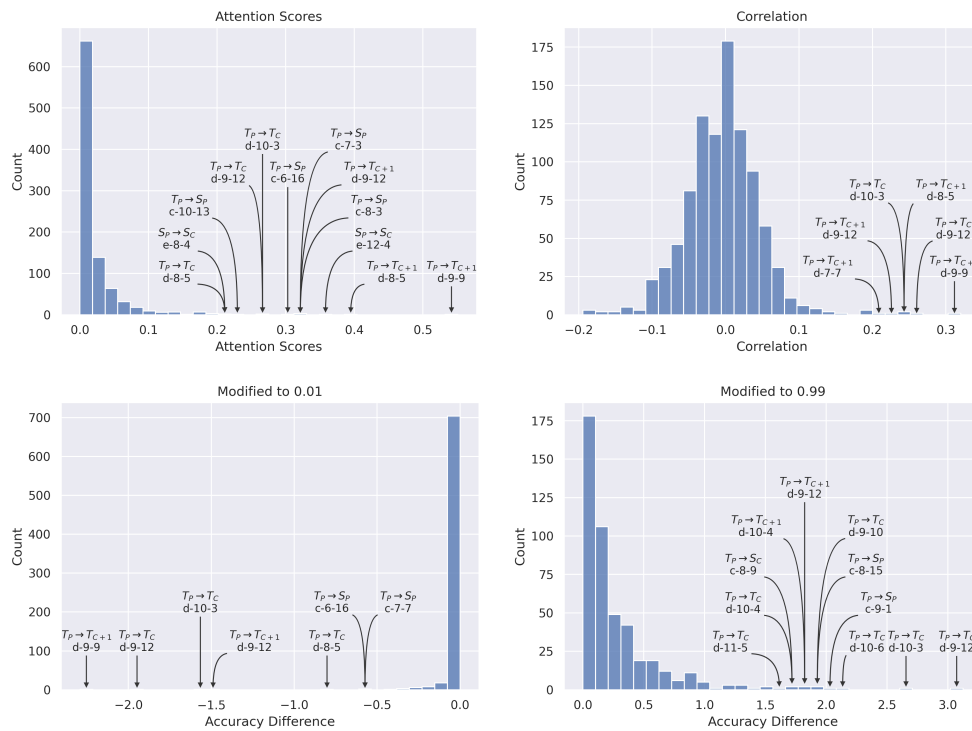


Figure 20: Histograms of average attention scores, correlations between attention scores and accuracy on **ContraPro** (English-to-German) dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **sentence-level NLLB-200** model with the values of noticeable heads annotated with arrows.

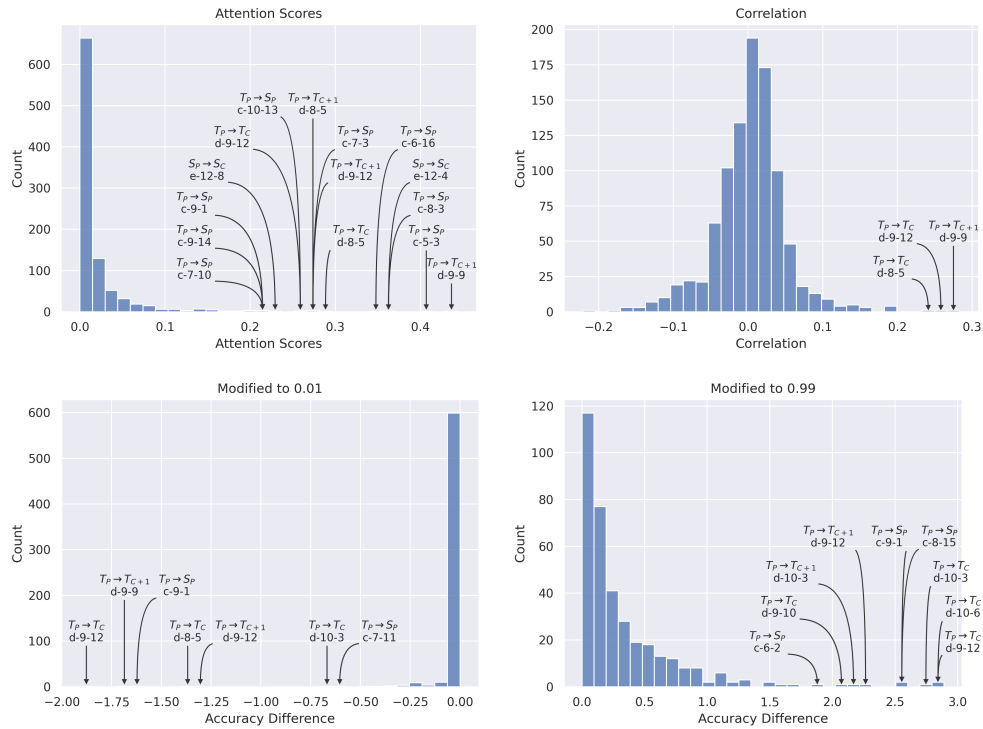


Figure 21: Histograms of average attention scores, correlations between attention scores and accuracy on **ContraPro** (English-to-German) dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **context-aware NLLB-200** model with the values of noticeable heads annotated with arrows.

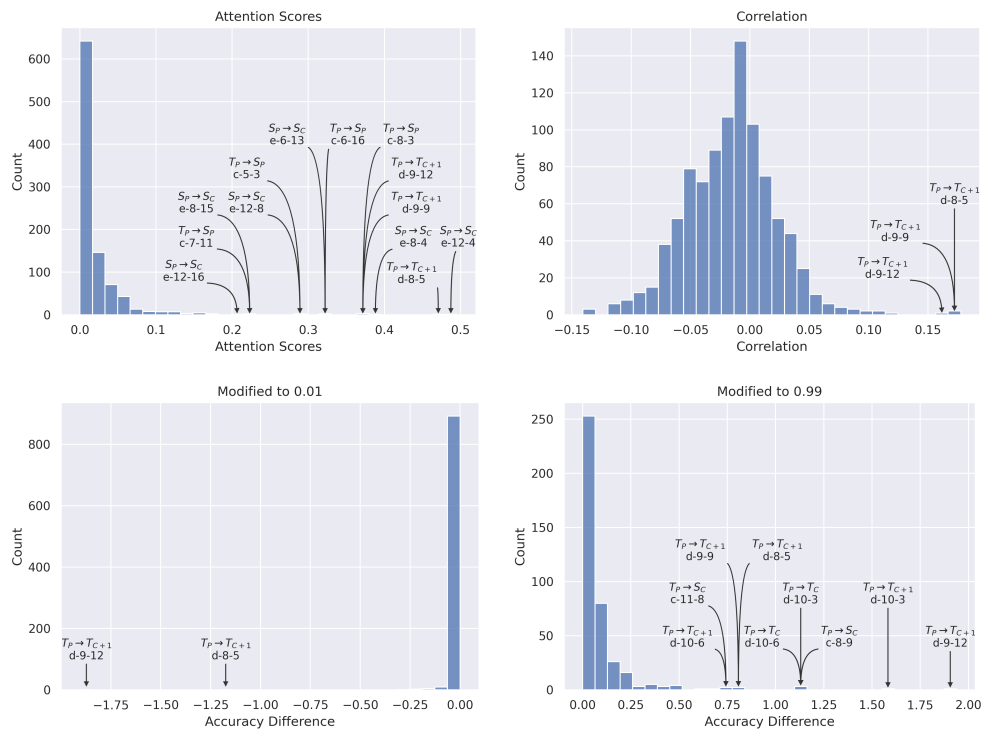


Figure 22: Histograms of average attention scores, correlations between attention scores and accuracy on **LCPT** (English-to-French) dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **sentence-level NLLB-200** model with the values of noticeable heads annotated with arrows.




Figure 23: Histograms of average attention scores, correlations between attention scores and accuracy on LCPT (English-to-French) dataset, and the differences in accuracy for: Modifying Heads to 0.01 (excluding positive values) and Modifying Heads to 0.99 (excluding negative values) of each head of the **context-aware NLLB-200** model with the values of noticeable heads annotated with arrows.