

# Do Not Change Me: On Transferring Entities Without Modification in Neural Machine Translation – a Multilingual Perspective

Dawid Wisniewski<sup>1,2</sup>, Mikolaj Pokrywka<sup>1,3</sup>, Zofia Rostek<sup>1</sup>

<sup>1</sup>Laniqo.com, <sup>2</sup>Poznan University of Technology, Poland <sup>3</sup>Adam Mickiewicz University, Poland,

Correspondence: dawid.wisniewski@laniqo.com

## Abstract

Current machine translation models provide us with high-quality outputs in most scenarios. However, they still face some specific problems, such as detecting which entities should not be changed during translation. In this paper, we explore the abilities of popular NMT models, including models from the OPUS project, Google Translate, MADLAD, and EuroLLM, to preserve entities such as URL addresses, IBAN numbers, or emails when producing translations between four languages: English, German, Polish, and Ukrainian. We investigate the quality of popular NMT models in terms of accuracy, discuss errors made by the models, and examine the reasons for errors. Our analysis highlights specific categories, such as emojis, that pose significant challenges for many models considered. In addition to the analysis, we propose a new multilingual synthetic dataset of 36,000 sentences that can help assess the quality of entity transfer across nine categories and four aforementioned languages<sup>1</sup>.

## 1 Introduction

Machine translation is one of the oldest branches of Natural Language Processing (NLP), which, thanks to the Transformer (Vaswani et al., 2017) architecture incorporating the (self-)attention (Bahdanau et al., 2015) mechanism, achieves human-like quality in many translation directions, especially in the case of sentence-level translations (Läubli et al., 2018).

Although the overall quality of recent neural models is very high both in terms of human perception and metrics such as BLEU (Papineni et al., 2002) or COMET (Rei et al., 2022a), there are still some weaknesses that need to be addressed.

© 2025 The authors. This article is licensed under a Creative Commons 4.0 licence, no derivative works, attribution, CC-BY-ND.

<sup>1</sup>Project supported by grant no. 0311/SBAD/0763 - Młoda Kadra financed by Poznan University of Technology.

One of them is the problem of identifying text fragments that should be copied without modification into the target sentence. Many of these entities: phone numbers, email addresses, or company names, represent categories that occur frequently in texts. However, the possible instantiations of these categories are so numerous (e.g., all possible phone numbers or email addresses) that models must rely on contextual information to detect them, rather than memorizing these entities.

In this paper, we show that many popular NMT models have problems with transferring such entities without modification. To analyze this issue, we focus on translations between four languages: English, German, Polish, and Ukrainian, using sentences containing entities from 9 categories: alphanumeric sequences, emails, emojis, IBANs, IP numbers, ISBNs, phone numbers, social handlers, and URLs.

The main contributions of this paper are: (i) a new multilingual dataset consisting of 36,000 sentences, each of which contains entities that should not be modified in the translation process, (ii) an analysis of eight popular NMT models, and (iii) a discussion of the causes of errors made by the models considered. This paper addresses three research questions: **RQ1**: Are there categories of entities that popular NMT models cannot transfer without modification? **RQ2**: What are the causes of the problems with transferring entities without modification? **RQ3**: Which models offer the highest quality solutions to this problem?

## 2 Related work

**Translation errors** In recent years, much attention has been paid to the evaluation of neural machine translation (NMT) models. Although the Transformer-era NMT models generally perform very well, some of them are prone to small but sometimes critical errors. The problem analyzed in

this paper has so far been considered in the context of the sensitivity of evaluation metrics to critical errors, among which unexpected input sequence modification has often been selected as one of the critical error types.

ACES (Amrhein et al., 2022) – the set of translation accuracy challenges, used to assess the sensitivity of metrics to critical errors, is a data set consisting of approximately 36,500 translated sentences. These sentences are expressed in 146 languages and describe 68 error phenomena, including hallucinations, erroneous unit conversions, adding unnecessary information, producing nonsense words, or translating entities that should not be translated. The error categories in ACES follow the Multidimensional Quality Metrics (MQM) ontology, providing a taxonomy of 108 translation problems defined at multiple levels of granularity (Lommel et al., 2014).

Another similar dataset, DEMETR (Karpinska et al., 2022), was developed to analyze machine translation metrics. This dataset contains examples of 35 perturbations that include: deleting parts of speech, using hyperonyms, replacing words, misspelling, using wrong capitalization, or repeating parts of text.

SMAUG (Alves et al., 2022), a sentence-level multilingual augmentation project focuses on generating translations that include critical errors and may be used to evaluate the robustness of MT metrics. SMAUG is a tool for introducing perturbations in existing sentences and is focused on categories such as deviation from named entities, numbers or meaning, or modification of content.

However, our goal is to analyze the problem from a different perspective. Instead of perturbing existing translations, we observe which models are more likely to generate perturbations. Furthermore, we increase the number of examples per category, which helps us to draw statistically significant conclusions. In our work, we provide 1,000 examples per category and language pair, while in other works, e.g., ACES, 36,500 examples cover 146 languages and 68 phenomena, which gives an average of about 3.5 examples per language and category combination. Moreover, instead of focusing only on translations from a given language to English, as in the DEMETR dataset, we analyze all possible translation directions between the four languages under consideration.

**Recent neural translation models** Since the advent of the Transformer (Vaswani et al., 2017) architecture, the quality of machine translation models has increased by a large margin (Stahlberg, 2020). For this reason, multiple Transformer-based MT models have been proposed in recent years. One of the well-known projects in this area is OPUS (Tiedemann et al., 2024), which provides datasets and MT models trained using MARIAN (Junczys-Dowmunt et al., 2018) and specialized in translation between various pairs of languages. Other similar models try to explore the multilingual abilities of Transformers, e.g., MBART (Tang et al., 2020) supporting translation in 50 languages, No Language Left Behind (NLLB) (Costa-jussà et al., 2022) supporting 200 languages, or MADLAD (Kudugunta et al., 2023) with the support of more than 400 languages. These models are trained using multilingual corpora, which help to exploit relationships between languages. In addition to scaling models in terms of language number, other approaches try to add different modalities as additional sources of knowledge. A popular example is SeamlessM4T (Barraut et al., 2023), which is able to process textual and audio data.

An interesting avenue in machine translation is the use Large Language Models (LLMs) to translate between languages (Wu and Hu, 2023). Due to the multilingualism of large web corpora and the variety of tasks observed in these corpora that can help to better understand the language, these general models are an interesting alternative to specialized translation models (Jiang et al., 2023; Dubey et al., 2024; Rivière et al., 2024; Chu et al., 2024). Some LLMs are trained with a substantial amount of parallel corpora, for example, TowerLLM (Alves et al., 2024), which supports 10 languages and, in addition to machine translation, can perform grammatical correction of texts, identify named entities, or post-edit texts. EuroLLM (Martins et al., 2024) is another attempt to use LLMs for machine translation. This model supports 35 languages and is trained on various sources of data, including good-quality parallel translation corpora, mathematical equations, code, and general web data.

### 3 Dataset

The ACES, DEMETR, and SMAUG have different scopes than this research; thus, even if they consider entities that should not be translated (hereafter

referred to as *no-translate entities*), they use them to evaluate metrics. They also provide relatively small numbers of examples regarding no-translate entities. For example, ACES introduces only 100 examples representing the no-translate category, all of them representing English to German translations. DEMETR focuses on translating from one of 10 languages to English, providing 100 examples per a given perturbation and source language pair. SMAUG operates on existing sentences and can be used to introduce perturbations (e.g., punctuation removal, random sequence injection, or named entity modification) into provided sentences.

For this reason, we decided to create a new dataset, entirely focused on entities that should not be modified during translation. This dataset can be used to measure how well different models transfer no-translate entities without modification.

We selected 9 entity types that are simple to identify using regular expressions, these are: *e-mail addresses, URLs, phone numbers, emojis, social handlers* (e.g., Instagram or TikTok user identifiers), *IP addresses, alphanumeric sequences* – artificial identifiers represented as mixtures of letters and numbers, International Bank Account Numbers (*IBANs*), and International Standard Book Numbers (*ISBNs*). We did not include dates and numbers as their formats may vary depending on the language (e.g., DD/MM/YYYY vs. MM/DD/YYYY dates or 1,000.00 vs. 1.000,00 as representations of "one thousand" in different languages).

In addition to analyzing individual categories, we wanted to investigate the transferability of entities across languages. We selected 4 languages: English, German, Polish, and Ukrainian — with the goal of selecting popular languages (English, German) and less popular ones that have interesting features, e.g. strongly inflected languages (Polish) or non-Latin alphabets (Ukrainian with Cyrillic). These 9 categories and 4 languages were used for the subsequent generation of the dataset.

### 3.1 Sentences generation

We asked Gemma 2 (Rivière et al., 2024) 9B instruction-following model<sup>2</sup> to generate 20,000 sentences expressed in each of the languages considered (English, German, Polish, Ukrainian) and provide examples for each category considered. The prompts used to generate the examples are listed in Appendix B.

<sup>2</sup><https://huggingface.co/google/gemma-2-9b-it>

To generate sentences, we used vLLM (Kwon et al., 2023) library setting the following sampling parameters: `temperature=1.2`, `top_p=0.95`, `length_penalty=1.0`, `min_tokens=16`, `max_tokens=512`, `max_model_len=4096`, `gpu_memory_utilization=0.95`.

The result of this step was a set of 720,000 generated sentences (20,000 sentences  $\times$  4 languages  $\times$  9 categories).

### 3.2 Representative sample selection

As LLM outputs may be of varying quality, we performed a sampling procedure to retain only the most representative sentences.

We applied the following procedure for this purpose: (i) we discarded additional information occasionally generated by Gemma, which frequently consists of the English translation of the sentence, comments about entities, etc. These additional texts were added mainly due to the length constraint `min_tokens=16`, so that for short sentences generated, the model continued the generation process to meet the criterion (e.g., providing English translation for languages other than English). Gemma frequently places these additional remarks after double newlines so they can be easily removed, (ii) we discarded examples expressed in a language other than the expected one using the `langdetect`<sup>3</sup> library, (iii) for a given language and category pair, we grouped sentences of similar size together by sorting them by length and placing in 20 buckets of equal size, (iv) we sampled 50 sentences from each bucket ensuring that each sentence has exactly one entity of the expected category (using regular expressions listed in Appendix A) and has no grammar errors according to `language tool` (Miłkowski, 2010)<sup>4</sup>. The buckets provide a diverse length representation, (v) Finally, we combined the selected samples forming a set of 1,000 examples (50 examples  $\times$  20 buckets) for each language and category pair.

This procedure applied to each language and category created a high-quality dataset of 36,000 examples (1,000 examples  $\times$  9 categories  $\times$  4 languages). The dataset is published online<sup>5</sup>.

<sup>3</sup><https://pypi.org/project/langdetect/>

<sup>4</sup><https://pypi.org/project/language-tool-python/>

<sup>5</sup><https://github.com/laniqo-public/do-not-change-me>

## 4 Explorative data analysis

The average sentence in our dataset consists of  $18.71 (\pm 6.81 \text{ std. dev.})$  tokens<sup>6</sup>. As can be seen in Figure 1, Polish, Ukrainian, and German sentences tend to be similar in length – with an average of almost 17 tokens. Meanwhile English sentences are longer, with the average of 23.99 tokens. This discrepancy arises from Gemma’s tendency to include English translations for short inputs and non-English target languages due to the `min_tokens=16` constraint. In these cases, our filtering procedure described in Section 3.2 is applied, making the average sentence shorter.




Figure 1: Number of NLTK (Bird and Loper, 2004) tokens in sentences expressed in a given language.

There is a greater variability in the number of tokens in sentences when analyzing texts by category. These values, presented in Appendix D, Figure 7, show that sentences introducing phone numbers are on average the longest with  $23.21 \pm 6.68$  tokens. The shortest sentences are related to emojis, with  $13.06 \pm 3.94$  tokens.

Considering the lengths of the entities in characters, we see that the average entity consists of  $17.91 \pm 9.99$  characters. Contrary to general sentence lengths, per-language entity length averages are very similar to each other, with English entities having  $18.13 \pm 10.48$  characters, German –  $18.01 \pm 9.64$ , Polish –  $17.93 \pm 9.81$ , and Ukrainian –  $17.58 \pm 10.01$ . The shortest entities are emojis, usually consisting of only 1 character, while the longest are URLs, emails, IBANs, and phone numbers. The longest entity is an English URL address consisting of more than 100 characters.

<sup>6</sup>Tokens were generated using the NLTK’s (Bird and Loper, 2004) `word_tokenize` function

Table 1: Models selected for experiments.

Model	# of Params
OPUS <sup>7</sup>	12×75M
mBART <sup>8</sup>	611M
NLLB <sup>9</sup>	3,300M
M2M100 <sup>10</sup>	1,200M
EuroLLM 9B <sup>11</sup>	9,000M
MADLAD 7B <sup>12</sup>	7,000M
SeamlessM4T <sup>13</sup>	2,300M
Google Transl.	no data

## 5 Methodology

To understand the quality of the entity transfer, we applied the following methodology: (i) we chose a set of NMT models to translate between supported languages, (ii) we used each selected model to translate each sentence from our dataset to all languages considered, and (iii) we extracted entities from source and target sentences and compared them in terms of the Levenshtein distance. The following subsections describe these steps in detail.

### 5.1 Models selection

We selected the most popular NMT models from the Huggingface repository, searching for models that support all the languages considered (English, German, Polish, and Ukrainian), and those that fit consumer GPUs (not bigger than 9B parameters). The list of the selected models is given in Table 1.

Since OPUS-MT models are unidirectional, with specialized models translating between a given language pair, we selected 12 OPUS models that support all possible language pairs considered.

Moreover, we selected Google Translate as a reference as it is one of the leading commercial translation services.

### 5.2 Translation procedure

For each model, language, and category, we translated sentences from a given language to all other supported languages. We used the Huggingface Transformers<sup>14</sup> and vLLM<sup>15</sup> libraries for inference. In this way, we generated 108,000 translations ( $3 \times 36,000$ ).

For Google Translate, we used the Google Docs translation feature to translate sentences in batches. For each language and category, we concatenated all sentences related to a given language and category into one document using double newlines to

<sup>14</sup><https://huggingface.co/>

<sup>15</sup><https://github.com/vllm-project/vllm>




Figure 2: Distribution of the number of characters for each entity category.

separate sentences and translated these documents. In this way, we reduced the number of requests.

The generated translations are published online<sup>16</sup>.

### 5.3 Source and target sentence entity comparison

For each source and translation (target), we used the hand-crafted regular expressions listed in Appendix A to detect if the desired entity is transferred correctly. Due to the selection process, each source sentence contains exactly one entity of a given type. Thanks to regular expressions, we checked if entities of the same type are detected in both source and target, and measure differences between them if needed, using Levenshtein distance. To add an additional perspective on the general quality of the translations, we used the CometKiwi<sup>17</sup> (Rei et al., 2022b) model.

## 6 Results

We analyzed the results from various perspectives, as described in the following paragraphs.

**Per-category accuracy** For each category considered, we examined each model’s entity transfer accuracy averaging scores over all language directions with equal weights. In this experiment, a true positive is generated when a regular expression related to the expected category detects exactly the same sequence of characters on the source and target sides.

Table 2 summarizes this scenario and shows that the best-performing models are **EuroLLM**

**9B**, **Google Translate**, and **MADLAD**. The worst-performing models are **OPUS**, **SeamlessM4T**, and **NLLB**. Emoji was the category where most models struggled. Five models: **OPUS**, **MBART**, **SeamlessM4T**, **NLLB** and **M2M100** achieved an accuracy of less than 5.5% in this category. The three easiest-to-transfer categories for all models are: IP (average over models: 91.79), Phone (91.29), and ISBN (90.69). These scores suggest that all those models are able to transfer sequences of numbers relatively well.

**Per-direction accuracy** Similarly to per-category analysis, we also grouped the results by translation direction, macro-averaging the scores over all categories considered. The reason for this analysis is to understand whether all the language pairs considered are handled with similar quality.

The corresponding statistics are presented in Table 3. The highest quality overall is observed for the en → uk direction, with an average accuracy of 98.85. The lowest accuracy (14.60) is observed for de → uk for the **OPUS** model. The average accuracy over all models and directions is equal to 78.65 (± 19.3 of std. dev.). Averaging over all models, the best quality directions are: en → pl (88.32 ± 9.06), en → de (87.99 ± 8.18), de → pl (81.6 ± 11.25) and de → pl (81.27 ± 11.13). On average, the worst quality directions are pl → uk (70.67 ± 23.38), pl → en (74.76 ± 21.6), and de → uk (75.13 ± 25.27). However, it should be noted that the weak average scores involving Ukrainian language are due to **OPUS** and **SeamlessM4T** models, which handle this language relatively poorly in the context of entity transfer. While, according to Table 3, among all directions, the average accuracy for en → uk direction is the highest for **Google Translate**, second highest for **MADLAD**,

<sup>16</sup><https://github.com/laniqo-public/do-not-change-me>

<sup>17</sup>Python3.12.7/Comet2.2.4/fp32/Unbabel/wmt22-cometkiwi-dair0

Table 2: Accuracy averaged over all language pairs for each category. Values in bold are the highest, while underlined are the worst.

Category	Google T..	OPUS	MADLAD	MBART	SeamlessM4T	NLLB	EuroLLM 9B	M2M100
alphanumeric	87.39	<u>71.67</u>	81.68	86.65	78.82	83.31	87.77	<b>88.54</b>
email	85.66	<u>30.31</u>	85.52	89.15	82.78	84.72	<b>92.12</b>	90.14
emoji	<b>98.59</b>	<u>0.02</u>	81.64	2.1	0.67	5.3	96.78	4.18
iban	95.45	<u>40.45</u>	79.53	84.87	80.47	88.96	<b>98.55</b>	94.86
ip	98.61	<u>58.12</u>	94.49	97.82	91.28	95.62	<b>99.58</b>	98.78
isbn	<b>99.16</b>	<u>69.91</u>	89.22	98.51	76.54	97.05	98.41	96.7
phone	98.38	<u>68.08</u>	90.52	96.45	83.22	96.24	98.36	<b>99.07</b>
social	88.08	<u>34.51</u>	84.37	88.79	77.50	75.04	<b>96.23</b>	72.00
url	88.52	<u>38.09</u>	81.96	90.43	50.92	79.11	<b>95.20</b>	86.09
macro avg.	93.32	<u>45.68</u>	85.44	81.64	69.13	78.37	<b>95.89</b>	81.15

Table 3: Accuracy macro-averaged over categories for each direction. Values in bold are the highest, while underlined are the worst.

Direction	Google T.	OPUS	MADLAD	MBART	SeamlessM4T	NLLB	EuroLLM 9B	M2M100
de → en	90.31	68.74	92.84	82.91	<u>66.06</u>	74.56	<b>95.75</b>	78.98
de → pl	89.52	<u>58.33</u>	88.92	82.25	84.20	75.42	<b>94.92</b>	79.27
de → uk	90.09	<u>14.60</u>	85.61	80.52	83.04	76.25	<b>94.62</b>	76.27
en → de	96.92	81.27	95.36	87.14	<u>74.01</u>	85.41	<b>97.18</b>	86.66
en → pl	97.04	77.58	<b>97.88</b>	88.07	<u>74.23</u>	86.31	97.75	87.70
en → uk	<b>98.85</b>	<u>21.42</u>	96.07	86.85	38.45	85.26	97.08	86.60
pl → de	90.29	<u>54.56</u>	88.37	77.54	83.17	69.98	<b>94.70</b>	74.35
pl → en	91.70	59.44	94.69	81.35	<u>31.89</u>	68.85	<b>95.49</b>	74.67
pl → uk	90.86	<u>25.65</u>	47.01	79.33	<u>82.57</u>	70.95	<b>95.17</b>	74.60
uk → de	93.22	<u>21.61</u>	86.23	76.23	84.66	82.59	<b>95.21</b>	84.34
uk → en	<b>96.60</b>	<u>33.85</u>	92.24	80.78	42.43	82.32	96.48	84.96
uk → pl	94.41	<u>31.15</u>	60.02	76.69	84.91	82.58	<b>96.29</b>	85.45
macro avg.	93.32	<u>45.68</u>	85.44	81.64	69.13	78.37	<b>95.89</b>	81.15

and third highest on **MBART**, **NLLB**, **EuroLLM**, and **M2M100**, it is ranked 11th out of 12 possible places for **OPUS** and **SeamlessM4T**.

**Error distribution** The accuracy scores analyzed in the previous paragraphs give only partial information on what happens to entities during translations, as the accuracy only checks whether exactly the same entity is detected on the source and target sides. To address this problem, we analyzed the distribution of errors in terms of the Levenshtein distance. For each model, we checked what percentage of errors are due to the lack of an entity of a given type detected on the target side (no-match case), or matching an entity of a given type but differing in the sequences detected. In that case, we measured the edit distance between the expected and the observed entities (between the source and target entities). The summary of this experiment is presented in Figure 3. For 6 models (**OPUS**, **MADLAD**, **MBART**, **SeamlessM4T**, **NLLB**, and **M2M100**), the no-match category is by far the most dominant. This is mainly due to the emoji category, which according to Table 2 has very low scores assigned to exactly these models (except **MADLAD**). The low scores for emojis are

due to the lack of model’s support for byte-level tokenization.

A deeper analysis reveals that these six models indeed produce most of the no-match errors from emojis: 85.84% of **MBART** no-match errors are due to emojis, 85.58% for **M2M100**, 68.99% for **NLLB**, 60.78% for **SeamlessM4T**, 41.16% for **OPUS**, and 25.37% for **MADLAD**. Models that are not dominated by no-match errors (**Google Translate** and **EuroLLM**, have a relatively low percentage of emojis in the no-match category: 12.5% and 26.73%, respectively. For comparison, Figure 8 in Appendix D represents the error distribution without the Emoji category. In every scenario, small differences (e.g., 1 or 2 characters) are more likely than larger ones (e.g., 3 or 4). The highest chance of one-character errors is observed for **Google Translate**. Models, such as **EuroLLM**, **MADLAD**, or **Google Translate** have a relatively high number of large differences (Levenshtein distance > 5).

Table 6 provides information on the category that each model has the most problems with, listing the categories with the highest number of errors per each model considered. This analysis is performed separately for edit distances equal to one,

Table 4: Prompts used for EuroLLM.

Prompt id	Text
GENERIC	<lim_start>system You are a professional {source_language} to {target_language} translator. Your goal is to accurately convey the meaning and nuances of the original {source_language} text while adhering to {target_language} grammar, vocabulary, and cultural sensitivities. Return only translation without any prefixes and explanations. <lim_end><lim_start>user Translate the following {source_language} source text to {target_language}: {source} {target_language}: <lim_end><lim_start>assistant
FOCUSED	<lim_start>system You are a professional {source_language} to {target_language} translator. Your goal is to accurately convey the meaning and nuances of the original {source_language} text while adhering to {target_language} grammar, vocabulary, and cultural sensitivities. <b>Translate the provided text while ensuring that all non-translatable elements, such as numbers, email addresses, alphanumeric strings, emoticons, and similar elements, remain unchanged in the translated text.</b> Return only translation without any prefixes and explanations. <lim_end><lim_start>user Translate the following {source_language} source text to {target_language}: {source} {target_language}: <lim_end><lim_start>assistant

Table 5: Average change when using large models and smaller ones. For EuroLLM, 2 prompts are evaluated.

Category	Reference	Focused prompt		Generic prompt		MADLAD 3.3B	MADLAD 7B
	Google T..	Euro 1.7B	Euro 9B	Euro 1.7B	Euro 9B		
alphanum	87.49	83.03	<b>87.77</b>	<u>82.56</u>	86.42	81.25	81.68
email	85.66	71.56	<b>92.12</b>	<u>71.06</u>	90.61	86.47	85.52
emoji	<b>98.59</b>	89.43	96.78	<u>84.00</u>	93.19	75.95	81.64
iban	95.45	92.64	<b>98.55</b>	<u>91.80</u>	98.12	80.53	79.53
ip	<u>98.61</u>	98.86	<b>99.58</b>	<u>98.77</u>	99.35	96.72	94.49
isbn	<b>99.16</b>	95.73	98.41	<u>95.06</u>	98.19	88.17	89.22
phone	<b>98.38</b>	97.35	98.36	<u>97.25</u>	98.08	88.48	90.52
social	88.08	<u>84.74</u>	<b>96.23</b>	<u>85.27</u>	95.29	81.41	84.37
url	88.52	88.19	<b>95.20</b>	<u>87.93</u>	94.26	85.23	81.96
macro avg.	93.32	89.06	<b>95.89</b>	<u>88.19</u>	94.83	84.91	85.44

two, and larger than 5. As can be seen, the most frequent differences by one character can be observed for IBANs and social handlers (2/8 models struggled with each of those categories). Larger differences, with edit distance = 2, are more common among IBANs and ISBNs (three models struggled with IBANs and other three with ISBNs). When considering edit distances larger than 5, social handlers are the most problematic for 3 models, and alphanumeric sequences for 2 models. Some recurring problems that are observed among translations are: repeating the same phrase over and over, dropping a subset of characters, translating fragments of entities, and omitting an entity. Table 10 presents a subset of problems observed for **Google Translate**.

**Prompt selection for EuroLLM** We experimented with two types of prompts for inference using **EuroLLM** models: generic and focused ones, as described in Table 4. The focused prompt differs from the generic one by adding a sentence requesting that the entities representing categories considered should remain unchanged in translated sentences. The results, presented in Table 5, indicate that the focused prompt significantly improves the handling of entities for both versions of **EuroLLM** (1.7B and 9B). Additionally, the results of CometKiwi, shown in Table 9, demonstrate that inference with the focused prompt also enhances translation quality. Consequently, we reported the results of **EuroLLM** using the focused prompt in all experiments.

**Model size vs. accuracy** The models analyzed differ both in size and the number of supported language directions. As bigger models are frequently linked to higher quality in various NLP tasks, it may be interesting to evaluate the relationship between the model size and the number of parameters in comparison to the average accuracy. With the details on model sizes and parameters provided in Table 8, we can observe that the biggest models, namely **EuroLLM (9B)** and **MADLAD (7B)** are indeed of the best quality and the smallest ones, the **OPUS** family, is the worst. Interestingly, those large models are of the highest quality despite the large number of language directions supported. **OPUS**, with the highest number of parameters per a direction (75M) is the worst-quality model, which may indicate that multilingual models do not lose their entity transfer abilities with more languages. We also directly compared two of the largest and best performing models – **EuroLLM (9B)** and **MADLAD (7B)** with their smaller counterparts – **EuroLLM (1.7B)** and **MADLAD (3B)**. As presented in Table 5, the larger versions perform better. **EuroLLM 9B** beats the 1.7B version in the case of every category, with the average accuracy more than 6 percentage points higher in the case of the 9B model. Larger **MADLAD**, on the other hand, beats the smaller version in 5 out of 9 categories, with similar scores in 3 other categories, and a drop of more than 3 percentage points compared to the 3B model on URL addresses.

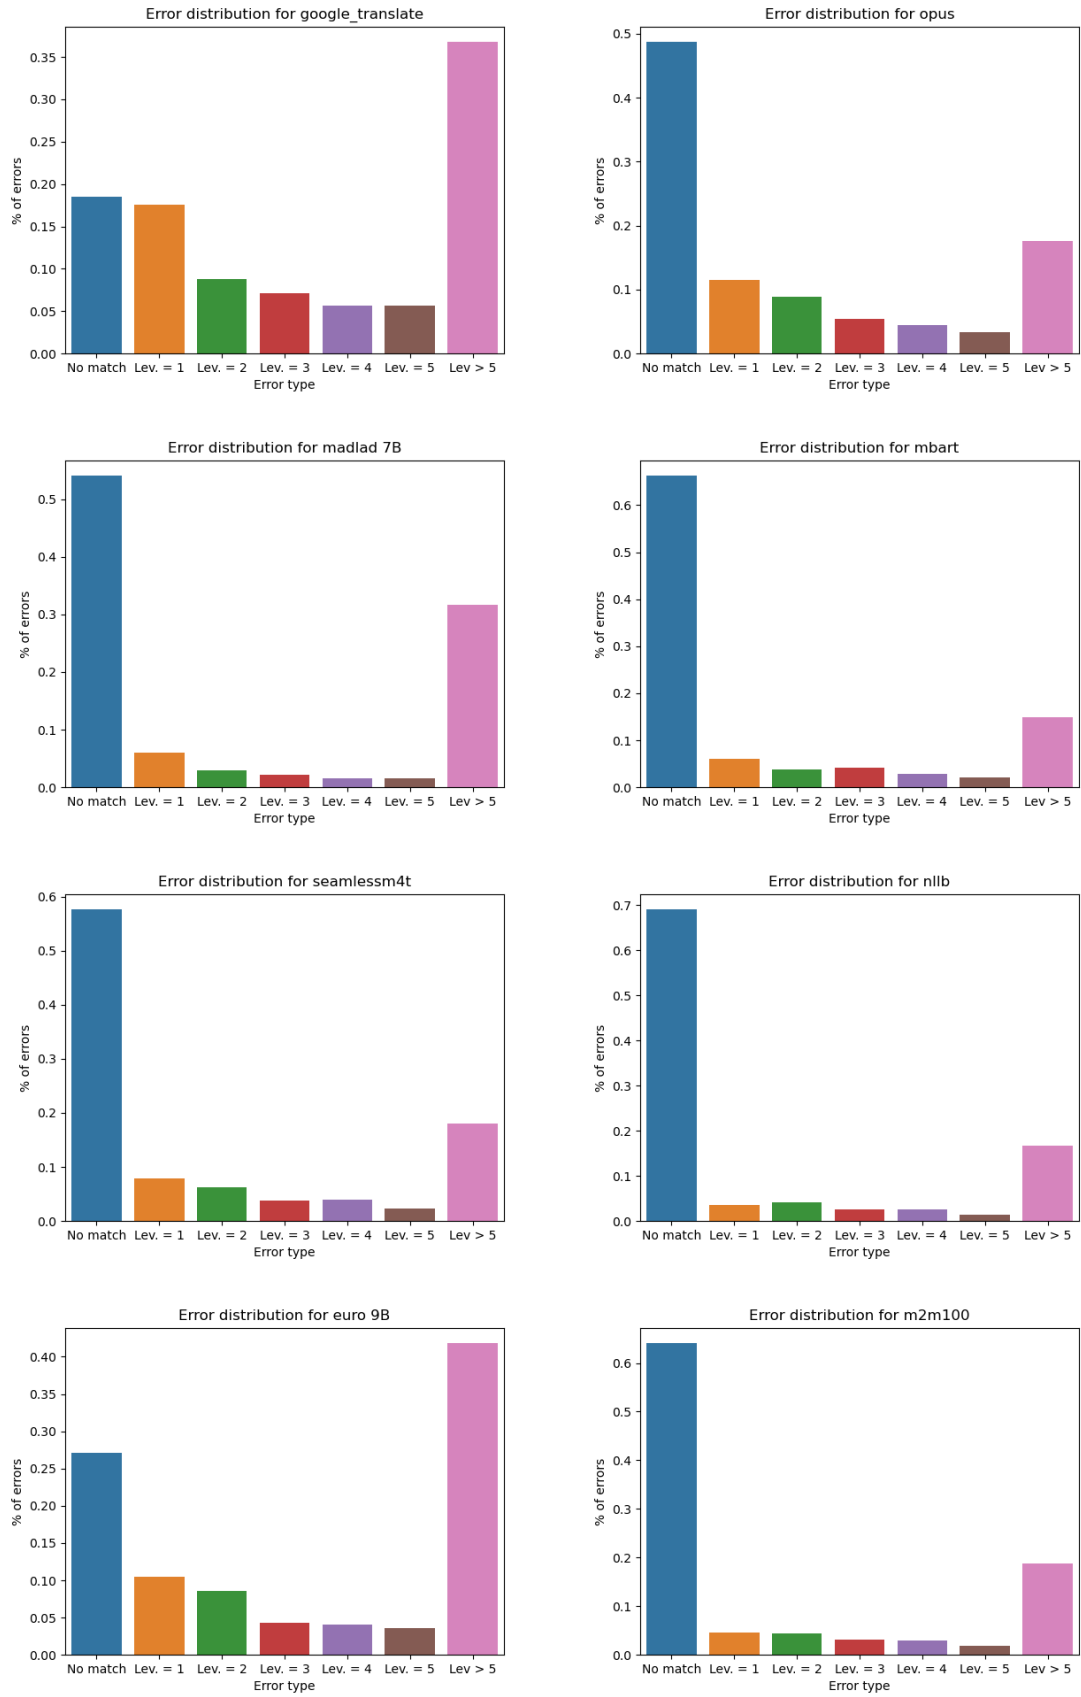


Figure 3: Distribution of the size of errors among models



Table 6: Categories generating highest number of errors in relation to the edit distance considered.

Model	Lev. = 1		Lev. = 2		Lev. >5	
	category	# errors	category	# errors	category	# errors
Google Translate	IBAN	339	e-mail	195	alphanumeric	736
OPUS	phone	1074	ISBN	1479	social	2275
MADLAD	IBAN	381	IBAN	166	phone	955
MBART	alphanumeric	288	IBAN	178	IBAN	829
SeamlessM4T	ISBN	677	ISBN	774	URL	3239
NLLB	social	303	IBAN	290	social	1574
EuroLLM	e-mail	94	ISBN	120	alphanumeric	786
M2M100	social	314	URL	239	social	1364

**Relationship between accuracy and average entity length** Our intuition tells us that the probability of errors may depend on the number of (sub)tokens a given entity is tokenized into by a given model. As different models use different tokenizers, they can split entities in a different manner. In Table 8, we put together the accuracies of models with the average entity lengths measured in tokens generated with a tokenizer of a given model. The Spearman’s rank correlation measured between these values tells us that there is a non-significant medium positive relationship between the accuracy and average entity (sub)tokens number ( $r_s = 0.3929$ ,  $p = 0.395$ ). The same conclusion is reached for Pearson’s correlation with  $r = 0.4317$  and  $p = 0.334$ . This observation may indicate that models splitting entities into a longer sequence of smaller (sub)tokens deal better with the transfer task, whereas models that chunk entities into bigger portions struggle in that context. For example, **EuroLLM** tokenizes an example ISBN number 0176 7890 1234 5678 9012 3456 into 30 tokens: "\_", "0", "1", "7", "6", "\_", "7", "8", "9", "0", "\_", "1", "2", "3", "4", "\_", "5", "6", "7", "8", "\_", "9", "0", "1", "2", "\_", "3", "4", "5", "6", while **OPUS** generates only 12 tokens: "\_01", "76", "\_78", "90", "\_12", "34", "\_56", "78", "\_90", "12", "\_34", "56". As can be seen, **OPUS**, instead of relating sequences of digits to ISBNs (as **EuroLLM** did), needs to consider possible pairs of digits, which are more numerous.

Based on the token length distribution presented in Figure 2, we selected categories with high average length and analyzed the correlation between the number of entity’s (sub)tokens and the likelihood of observing an error in an entity. The results presented in Table 7 show that there is a positive correlation between URL, alphanumeric, and e-mail lengths and error likelihood, ranging from very small values (**EuroLLM** and URL category – Pearson’s coeff. = 0.179 with p-value =

0.303) and very high ones (**OPUS** with e-mail and **URL** categories – 0.822 with p-value = 4.9e-06 and 0.724 with p-value = 5.94e-08, respectively). This analysis shows that some models, e.g., **OPUS** and **MBART** are more prone to making errors in longer sequences, whereas, e.g., **EuroLLM** model is much more robust.

**Context length vs. accuracy** Each entity is mentioned within a sentence, which can be considered its context. To understand whether longer sentences are more beneficial, we checked if the number of (sub)tokens in the sentence influences the accuracy of transfer. The intuition is that longer sentences should make entities less ambiguous, leading the model to realize that those entities should be transferred without changes. To measure the degree of this relationship, the following procedure was applied: we sorted source sentences according to lengths and split them into five bins of equal size so that texts of similar lengths are in the same bin. Then, for each bin, we measured the average accuracy of a given model, and the ratios of situations where entities are modified (Levenshtein > 0) or are not transferred at all (no-match errors). Figure 4 shows the percentage of correctly transferred entities for each bin. Figures 5 and 6 in Appendix D show distributions of non-matched and modified entities in target sentences, respectively. As can be seen, there is no visible relationship between the length of the sentence and the accuracy of the entity transfer. For some models, e.g., **OPUS**, the quality even decreases with increasing sentence length.

**General translation quality** The CometKiwi translation evaluation is presented in Table 9 and described in Appendix C. The scores range from 0.7 to 0.73 in most scenarios, with **Google Translate** – the top-rated model – assigned a score of 0.76. This means that, in general, the overall translation quality was good for all models considered.




Figure 4: Percentage of correctly transferred entities among different sentence lengths.

Table 7: Pearson’s correlation coefficient between the number of subtokens in an entity and the likelihood of observing an error in an entity

Model	Alphanumeric		email		URL		Phone	
	Pearson’s c.	p-value	Pearson’s c.	p-value	Pearson’s c.	p-value	Pearson’s c.	p-value
EuroLLM	0.344	0.1	0.401	0.111	0.179	0.303	-0.07	0.665
MADLAD	0.578	0.003	0.453	0.059	0.179	0.29	-0.187	0.241
OPUS	0.61	0.002	0.822	4.9e-06	0.724	5.94e-08	0.624	0.001
M2M100	0.479	0.024	0.572	0.01	0.046	0.785	-0.277	0.224
SeamlessM4T	0.402	0.071	0.696	0.001	0.397	0.018	0.289	0.192
NLLB	0.473	0.026	0.47	0.049	0.194	0.263	0.192	0.418
MBART	0.814	1.27e-05	0.603	0.01	0.39	0.023	-0.351	0.13

Table 8: Relationship between the accuracy of a model and selected model characteristics.

Model	Num. of params (M)	Supported languages	params per language (M)	avg. tokenized entity length	Average accuracy
EuroLLM 9B	9,000	1,225 (35 <sup>2</sup> )	7.35	13.55	95.89
MADLAD	7,000	175,561 (419 <sup>2</sup> )	0.04	13.42	85.44
OPUS	75	1 (1 <sup>1</sup> )	75	10.31	45.68
M2M100	1,200	10,000 (100 <sup>2</sup> )	0.12	9.64	81.15
SeamlessM4T	2,300	9,216 (96 <sup>2</sup> )	0.25	9.55	69.13
NLLB	3,300	40,000 (200 <sup>2</sup> )	0.083	9.51	78.37
MBART	611	2500 (50 <sup>2</sup> )	0.244	8.59	81.64

## 7 Conclusions

In this paper, we find that modern medium-sized LLMs such as **EuroLLM**, despite being only partially trained on parallel corpora, may excel in terms of entity transfer quality and can reach the quality similar to **Google Translate**. As the best performing model is **EuroLLM 9B**, it is the answer to **RQ3**. In contrast, the smallest **OPUS**-related models are scored as the worst. This may be due to two factors: on the one hand, bigger models may learn more about the world, but also they may learn better token representations using fine-grained tokenization as discussed in Section 6. We showed that increasing the length of the sentence does not correlate with the transfer accuracy, thus, the context information does not necessarily help to guide

a model to transfer a given sequence without modifications. Considering **RQ2**, the manual analysis of entities transferred by models with high Levenshtein distances reveals that entities are frequently partially translated, differing in some numbers, or repeated in fragments. Also, for some models (e.g., **OPUS**, **MBART**), the longer entities in terms of subtokens increase the probability of generating an error. This is in line with the intuition that it may be harder to transfer long sequences without errors as compared to the short ones. While most of the categories considered are decently transferred by most of the models considered, emoji is a big struggle – **OPUS**, **MBART**, **SeamlessM4T**, **NLLB**, and **M2M100** cannot transfer this category correctly (which answers **RQ1**).

## References

- Duarte Alves, Ricardo Rei, Ana C Farinha, José G. C. de Souza, and André F. T. Martins. 2022. [Robust MT evaluation with sentence-level multilingual augmentation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 469–478, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Duarte M. Alves, José Pombal, Nuno Miguel Guerreiro, Pedro Henrique Martins, João Alves, M. Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024. [Tower: An open multilingual large language model for translation-related tasks](#). *CoRR*, abs/2402.17733.
- Chantal Amrhein, Nikita Moghe, and Liane Guillou. 2022. [ACES: Translation accuracy challenge sets for evaluating machine translation metrics](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 479–513, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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*.
- Loïc Barrault et al. 2023. [Seamlessm4t-massively multilingual & multimodal machine translation](#). *CoRR*, abs/2308.11596.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. 2024. [Qwen2-audio technical report](#). *CoRR*, abs/2407.10759.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Y. Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley 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](#). *CoRR*, abs/2207.04672.
- Abhimanyu Dubey et al. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Marzena Karpinska, Nishant Raj, Katherine Thai, Yixiao Song, Ankita Gupta, and Mohit Iyyer. 2022. [Demetr: Diagnosing evaluation metrics for translation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2023. [MADLAD-400: A multilingual and document-level large audited dataset](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*, pages 611–626. ACM.
- Samuel Läubli, Rico Sennrich, and Martin Volk. 2018. [Has machine translation achieved human parity? A case for document-level evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4791–4796. Association for Computational Linguistics.
- Arle Lommel, Hans Uszkoreit, and Aljoscha Burchardt. 2014. [Multidimensional quality metrics \(mqm\): A framework for declaring and describing translation quality metrics](#). *Tradumática*, (12):0455–463.
- Pedro Henrique Martins, Patrick Fernandes, João Alves, Nuno Miguel Guerreiro, Ricardo Rei, Duarte M. Alves, José Pombal, M. Amin Farajian, Manuel Faysse, Mateusz Klimaszewski, Pierre Colombo, Barry Haddow, José G. C. de Souza, Alexandra Birch, and André F. T. Martins. 2024. [Eurollm:](#)

- Multilingual language models for europe. *CoRR*, abs/2409.16235.
- Marcin Miłkowski. 2010. Developing an open-source, rule-based proofreading tool. *Softw. Pract. Exper.*, 40(7):543–566.
- 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, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Ricardo Rei, José G. C. de Souza, Duarte M. Alves, Chrysoula Zerva, Ana C. Farinha, Taisiya Glushkova, Alon Lavie, Luísa Coheur, and André F. T. Martins. 2022a. [COMET-22: unbabel-ist 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation, WMT 2022, Abu Dhabi, United Arab Emirates (Hybrid), December 7-8, 2022*, pages 578–585. Association for Computational Linguistics.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C. Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022b. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Morgane Rivi re et al. 2024. [Gemma 2: Improving open language models at a practical size](#). *CoRR*, abs/2408.00118.
- Felix Stahlberg. 2020. [Neural machine translation: A review](#). *J. Artif. Intell. Res.*, 69:343–418.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#). *CoRR*, abs/2008.00401.
- J rg Tiedemann, Mikko Aulamo, Daria Bakshandaeva, Michele Boggia, Stig-Arne Gr nroos, Tommi Nieminen, Alessandro Raganato, Yves Scherrer, Ra l V zquez, and Sami Virpioja. 2024. [Democratizing neural machine translation with OPUS-MT](#). *Lang. Resour. Evaluation*, 58(2):713–755.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Yangjian Wu and Gang Hu. 2023. [Exploring prompt engineering with GPT language models for document-level machine translation: Insights and findings](#). In *Proceedings of the Eighth Conference on Machine Translation, WMT 2023, Singapore, December 6-7, 2023*, pages 166–169. Association for Computational Linguistics.
- Sustainability statement** The experiments were performed on a single PC with one GeForce RTX 4090 GPU. The total time for the experiments was equal to approximately 12 hours of translation time, which is related to 1.56 kg of CO<sub>2</sub>eq according to <https://mlco2.github.io/impact/>.
- A Regular expressions for category detection**
- List of regular expressions utilized for identifying non-translatable units:
- Alphanumeric:**
- ```
\b[\p{N}\p{L}][\p{N}\p{L}\p{P}]*
(\p{L}\p{N}|\p{N}\p{L})
[\p{N}\p{L}\p{P}]*[\p{N}\p{L}]\b
```
- E-mail:**
- ```
\b[\p{L}\p{N}._%+-]+@[\p{L}\p{N}.-]+\
.[\p{L}]{2,}\b
```
- IBAN:**
- ```
\b([A-Z]{2})[\ -]?([0-9]{2})[\ -]?
([A-Z0-9]{9,30})\b
```
- IP:**
- ```
\b\d{1,3}\.\d{1,3}\
.\d{1,3}\.\d{1,3}\b
```
- ISBN:**
- ```
\b(?:ISBN(?:-13)??:?\ )?(?=[0-9]{13}$|
(?:?:[0-9]+[-\ ]){4})[-\ 0-9]{17}$)
97[89][-\ ]?[0-9]{1,5}[-\ ]?[0-9]+
[-\ ]?[0-9]+[-\ ]?[0-9]\b
```
- Phone:**
- ```
\b[\d\+\-\=\%\^\(\)\[\]\{\}\]
[\d\. , \+:\-\*\/\=\%\^\(\)\[\]\{\}\]
{2,}
[\d\+\:\-\*\/\=\%\^\(\)\[\]\{\}\]\b
```
- Social handler:**
- ```
@[\0-9_.\p{L}]{2,24}[\0-9_.\p{L}]\b
```
- URL:**
- ```
\b((imap|s3|file|ftp|https?):\//\
[\p{L}\p{N}_-]+
(\.[-_/?=\p{L}\p{N}]{1,15}|
\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}|
www\.[\p{L}\p{N}_-]+
(\.[-_/?=\p{L}\p{N}]{1,15}))\b
```

## B Prompts used to generate samples

- **Alphanum:** Write me a random and creative sentence in [LANGUAGE] that includes a sequence consisting of multiple digits and letters longer than 5 characters.
- **E-mail:** Write me a random and creative sentence in [LANGUAGE] that includes a random email address.
- **Emoji:** Write me a random and creative sentence in [LANGUAGE] that includes a random emoji.
- **IBAN:** Write me a random and creative sentence in [LANGUAGE] with a sequence including an artificial IBAN number in IBAN format.
- **ISBN:** Write me a random and creative sentence in [LANGUAGE] with a sequence including an artificial ISBN number in ISBN format.
- **IP:** Write me a random and creative sentence in [LANGUAGE] that includes a random IP number.
- **Phone:** Write me a random and creative sentence in [LANGUAGE] that includes a long random phone number.
- **Social:** Write me a random and creative sentence in [LANGUAGE] that includes a social media handler starting with the @ sign (e.g., Twitter, Instagram).
- **URL:** Write me a random and creative sentence in [LANGUAGE] that includes a random URL address.

## C Detailed CometKiwi evaluation

Table 9 summarizes translation performance across models, language pairs, and input types. Google Translate leads with a macro-average score of 0.7598, performing consistently well in both high-resource pairs (e.g., de-en: 0.795) and low-resource pairs (e.g., uk-pl: 0.7245). EuroLLM models, particularly the 9B version with focused prompts, achieve strong results with a macro-average of 0.7202, making them the closest competitors. Performance generally declines for low-resource pairs and some specialized input types like email and url. Overall, larger models and prompt optimization (as

seen with EuroLLM) significantly enhance performance. It is important to note that this test set is synthetic, and the evaluation process is intended primarily as a sanity check to assess machine translation quality.

## D Detailed Figures and data analysis

Table 9: CometKiwi evaluation of translation accuracy across various models, language pairs, and data categories, showcasing macro-averaged performance by language and prompt-specific configurations.

	Google T.	OPUS	MADLAD 3.3B	MADLAD 7B	MBART	SeamlessM4T	NLLB	M2M100	EuroLLM 1.7B generic prompt	EuroLLM 1.7B focused prompt	EuroLLM 9B generic prompt	EuroLLM 9B focused prompt	macro avg. by lang and category
de-en	0.795	0.7861	0.7852	0.7863	0.7839	0.7899	0.787	0.7775	0.7787	0.7796	0.781	0.7822	0.7844
de-pl	0.7474	0.6853	0.7042	0.71	0.6966	0.6859	0.6968	0.7047	0.7033	0.7022	0.7056	0.7085	0.7042
de-uk	0.7529	0.6511	0.6816	0.6918	0.6676	0.6745	0.6742	0.6882	0.6895	0.6906	0.6962	0.701	0.6883
en-de	0.8113	0.7876	0.786	0.7888	0.782	0.8021	0.7911	0.7769	0.7775	0.7779	0.7805	0.7826	0.7870
en-pl	0.7793	0.7442	0.738	0.7426	0.7323	0.7665	0.7491	0.7318	0.731	0.7303	0.7337	0.7365	0.7429
en-uk	0.7711	0.7107	0.7193	0.7248	0.7144	0.7516	0.7199	0.7157	0.7145	0.7135	0.718	0.7216	0.7246
pl-de	0.7489	0.665	0.6772	0.6861	0.6652	0.6556	0.6808	0.6841	0.6853	0.6861	0.6914	0.6961	0.6852
pl-en	0.7562	0.7303	0.7343	0.7365	0.7319	0.7326	0.7345	0.73	0.7314	0.7325	0.7344	0.7361	0.7351
pl-uk	0.7476	0.6776	0.6651	0.6542	0.6775	0.7031	0.6914	0.6584	0.6647	0.6696	0.6773	0.6838	0.6809
uk-de	0.7242	0.6277	0.6496	0.6603	0.6347	0.6404	0.6488	0.6595	0.6623	0.6644	0.6699	0.6745	0.6597
uk-en	0.7595	0.7129	0.724	0.7281	0.7185	0.7277	0.7221	0.7231	0.7257	0.7278	0.7305	0.7329	0.7277
uk-pl	0.7245	0.6629	0.6458	0.6416	0.6546	0.6845	0.6746	0.6455	0.6511	0.6555	0.6622	0.6679	0.6642
alphanumeric	0.7507	0.699	0.7037	0.707	0.7005	0.7079	0.7089	0.7085	0.7093	0.71	0.7134	0.7164	0.7113
email	0.7338	0.6719	0.6821	0.6869	0.6761	0.6925	0.6847	0.6883	0.688	0.6878	0.6918	0.6952	0.6899
emoji	0.7606	0.7045	0.7041	0.7079	0.6989	0.7207	0.7121	0.7069	0.709	0.7107	0.715	0.7186	0.7141
iban	0.7616	0.7043	0.7089	0.709	0.7072	0.7194	0.7154	0.7113	0.7125	0.7134	0.7176	0.7211	0.7168
ip	0.7777	0.7278	0.7321	0.735	0.7278	0.7382	0.7378	0.7357	0.7368	0.7375	0.741	0.7439	0.7393
isbn	0.772	0.7178	0.7238	0.7276	0.7205	0.7291	0.7292	0.7286	0.7295	0.7301	0.7338	0.7369	0.7316
phone	0.766	0.7127	0.7147	0.7175	0.7119	0.7236	0.7228	0.7183	0.7193	0.72	0.7237	0.7269	0.7231
social	0.769	0.7082	0.7171	0.722	0.7113	0.7285	0.7194	0.7215	0.7229	0.7239	0.7279	0.7313	0.7253
url	0.7471	0.6849	0.6963	0.7005	0.6903	0.701	0.6973	0.7014	0.7025	0.7034	0.7074	0.7107	0.7036
macro avg. by model	0.7598	0.7035	0.7092	0.7126	0.7049	0.7179	0.7142	0.7103	0.7117	0.7127	0.7168	0.7202	




Figure 5: Percentage of entities not matched in target sentences among different lengths.




Figure 6: Percentage of modified entities among different lengths.




Figure 7: Number of NLTK (Bird and Loper, 2004) tokens per category. Sentences expressed in all languages were collected together for this analysis.

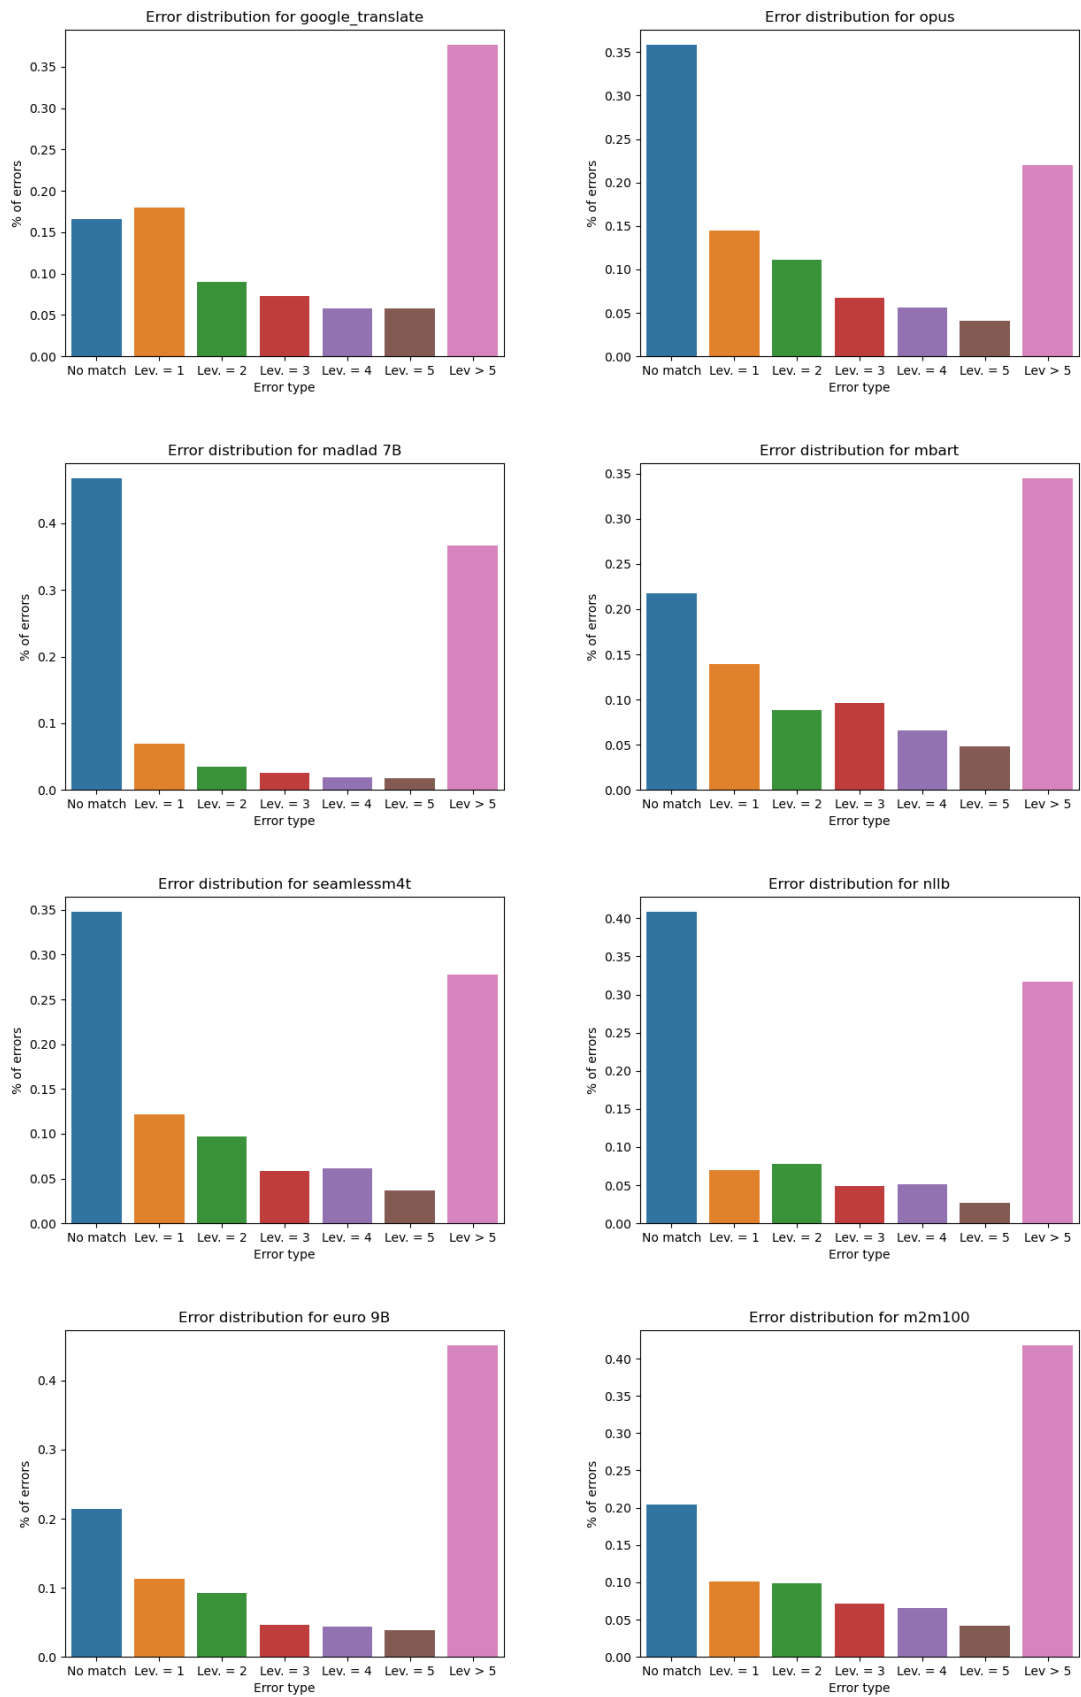


Figure 8: Errors distributions without emoji category



Table 10: Examples of errors made by Google Translate (verified on 01.01.2025)

comment srcEntity tgtEntity srcText tgtText	<b>IBAN modified</b> , pl → de <b>PL60109010000000000000000000</b> <b>PL601090100000000000000000</b> Panie profesorze Janie Kowalski, może przesłać mi nową listę prac do sprawdzenia na konto o numerze PL601090100000000000000000? Professor Jan Kowalski, können Sie mir eine neue Liste der zu prüfenden Werke an die Kontonummer PL601090100000000000000000 senden?
comment srcEntity tgtEntity srcText tgtText	<b>Translated social handler</b> , en → de <b>@klimatyzacja</b> <b>@airconditioning</b> Przechodząc przez park, nagle usłyszałam @klimatyzacja śpiewającą piosenkę o letnim słońcu. Als ich durch den Park spazierte, hörte ich plötzlich @airconditioning ein Lied über die Sommersonne singen.
comment srcEntity tgtEntity srcText tgtText	<b>Changed characters in alphanumeric</b> , en → de <b>tenotypic123CBSprk</b> <b>tenotypisch123CBSprk</b> Gniazdo pająka, o symbolu tenotypic123CBSprk, wisiało pod niebem usianym szumami. Das Spinnennest, Symbol tenotypisch123CBSprk, hing unter einem mit Lärm übersäten Himmel.
comment srcEntity tgtEntity srcText tgtText	<b>Changed characters in URL</b> , de → en <b>www.irgendwohin.com</b> <b>www.somewhere.com</b> Die Katze las www.irgendwohin.com vor dem Frühstück und kraulte verschmitzt um Aufmerksamkeit. The cat read www.somewhere.com before breakfast and playfully scratched for attention.
comment srcEntity tgtEntity srcText tgtText	<b>Modified e-mail</b> de → en <b>liebevollchenpinguin@aya.at</b> <b>lovingchenpinguin@aya.at</b> Die sprechende Mandarin geschickt ein Bild an liebevollchenpinguin@aya.at. The talking mandarin sent a picture to lovingchenpinguin@aya.at.
comment srcEntity tgtEntity srcText tgtText	<b>IP removed</b> , pl → de <b>192.168.1.108</b> <b>NO ENTITY DETECTED</b> Wiatr gonił pozółkłe liście, aż mu się zadało z 192.168.1.108 i spróbowało wziąć pod nie chwytem. Der Wind verfolgte die vergilbten Blätter, bis er müde wurde und versuchte, sie zu ergreifen.
comment srcEntity tgtEntity srcText tgtText	<b>Dropped Phone number</b> , de → pl <b>49 030 1234567890</b> <b>NO ENTITY DETECTED</b> Die alten Ratten spielten Karten und diskutierten leidenschaftlich laut vor Telefonnummer +49 030 1234567890 verband 123politisch. Stare szczury grały w karty i głośno i namiętnie dyskutowały o polityce.
comment srcEntity tgtEntity srcText tgtText	<b>ISBN dropped</b> pl → de <b>978-83-12-34567-8</b> <b>NO ENTITY DETECTED</b> Podczas lekcji astronomii, Paweł natknął się na książkę o istocie czasoprzestrzeni, której ISBN 978-83-12-34567-8 zdradził tajemnicę kosmicznej harmonii. Während einer Astronomiestunde stieß Paweł auf ein Buch über die Natur der Raumzeit, das das Geheimnis der kosmischen Harmonie enthüllte.