

Tradutor: Building a Variety Specific Translation Model

Hugo Sousa^{1,2*}, Satya Almasian^{3*}, Ricardo Campos^{2,4,5}, Alípio Jorge^{1,2}

¹Faculty of Sciences, University of Porto, Porto, Portugal

²INESC TEC, Porto, Portugal

³Institute of Computer Science, Heidelberg University, Germany

⁴Department of Informatics, University of Beira Interior, Covilhã, Portugal

⁵Ci2 - Smart Cities Research Center, Tomar, Portugal

hugo.o.sousa@inesctec.pt, almasian@informatik.uni-heidelberg.de

Abstract

Language models have become foundational to many widely used systems. However, these seemingly advantageous models are double-edged swords. While they excel in tasks related to resource-rich languages like English, they often lose the fine nuances of language forms, dialects, and varieties that are inherent to languages spoken in multiple regions of the world. Languages like European Portuguese are neglected in favor of their more popular counterpart, Brazilian Portuguese, leading to suboptimal performance in various linguistic tasks. To address this gap, we introduce the first open-source translation model specifically tailored for European Portuguese, along with a novel dataset specifically designed for this task. Results from automatic evaluations on two benchmark datasets demonstrate that our best model surpasses existing open-source translation systems for Portuguese and approaches the performance of industry-leading closed-source systems for European Portuguese. By making our dataset, models, and code publicly available, we aim to support and encourage further research, fostering advancements in the representation of underrepresented language varieties.

1 Introduction

In an era, where Language Models (LMs) form the foundation of numerous tools and systems, a significant concern arises: how can we ensure these tools are equally beneficial to all communities? One major hurdle in this direction is *inclusiveness*, particularly when it comes to scarce language varieties and the lack of models that take into account the diversity of dialects and culturally significant nuances of a given language (Team 2022). For example, while Portuguese is not considered a low-resource language, most of the content available is in Brazilian Portuguese. This is due to the Brazilian population representing more than 70% of all native Portuguese speakers¹. Consequently, a LM or a translation system trained on a Portuguese corpus is inclined to produce text containing phrases and grammatical structures of Brazilian Portuguese. As a result, countries like Portugal and Mozambique, which use different varieties of Por-

tuguese, could find themselves at a disadvantage, particularly in deploying LM-based systems in critical areas such as healthcare and judiciary, where the grammar and lexicons of the language are of great importance (Scherre and Duarte 2016; Kato and Martins 2016; Brito and Lopes 2016).

One solution would be to create LMs specific to a certain language variety. However, this presents its own set of challenges for training and evaluation (Armengol-Estapé et al. 2021; Rodrigues et al. 2023). Training a LM, be it from scratch or fine-tuning, requires a large, carefully curated corpus for training and several benchmarks for evaluation (Albalak et al. 2024), which are either nonexistent or contaminated by the dominant language varieties. Another way to overcome these challenges is to create machine translation (MT) models dedicated to a specific language variety (Zbib et al. 2012; Sennrich, Haddow, and Birch 2016; Riley et al. 2023). With a sufficiently powerful MT model for the low-resource variety, we can take the first step toward resource creation and inclusion of these languages. Such models can be used to translate training and evaluation benchmarks, which are predominantly in English, to the low-resource variety. Additionally, they can serve as an off-the-shelf intermediary between widely-used LMs and systems in low-resource settings or even be used to artificially generate data in the desired language variety.

In this paper, we present a novel methodology for developing a neural machine translation (NMT) model tailored to low-resource language varieties, where manually annotated data is scarce or unavailable. Our approach leverages a retro-translation technique: we first gather texts in the low-resource language variety and translate them into a resource-rich language. This newly created parallel corpus is then used to fine-tune a pre-trained language model. We validate our methodology using European Portuguese as a case study, though the same steps can be applied to other languages with similar challenges. As part of this work, we have created and publicly released a meticulously curated parallel corpus for European Portuguese, comprising 1,719,002 documents – the largest of its kind to date. We evaluate our model against existing open and close source translation systems for Portuguese, as well as zero-shot language models, demonstrating the effectiveness of our approach. Finally, to ensure our translations remain true to the intended language variety, we use a language variety classification model to quantify if the

*These authors contributed equally.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Statistic derived from Wikipedia https://en.wikipedia.org/wiki/Portuguese_language#Lusophone_countries

texts produced by our translation model are effectively European Portuguese.

In summary, our contributions are the following:

1. We propose a methodology to create a parallel corpus for a low-resource language variety using an on-the-shelf translation model. To this end, we provide the community with the largest translation dataset for European Portuguese and English, named PTradutor.
2. We provide the first open-source translation models from English to European Portuguese which outperform the generic Portuguese open-source translation systems and close the gap to state-of-the-art close-source translations systems for European Portuguese.
3. We offer a comprehensive evaluation of our models, emphasizing not only translation quality but also linguistic alignment to the desired language variety.

The remainder of this manuscript is organized as follows: Section 2 offers a comprehensive overview of the development of our corpus, PTradutor, which serves as the foundation for training our models. Sections 3 and 4 describe the experimental setup and present the results of our experiments, respectively. In Section 5, we position our research within the context of existing approaches. Finally, Section 6 summarizes our findings and suggests directions for future research.

2 PTradutor

As noted previously, a significant challenge in training a translation model for a specific language variety is the scarcity of datasets tailored for this purpose. To overcome this obstacle, we devised a three-step approach to automatically generate a parallel corpus. Industry-grade translation systems typically achieve higher translation quality when translating *into* a resource-rich language, such as English, rather than a low-resource language, such as European Portuguese. We exploit this concept to transform a monolingual European Portuguese corpus into a high-quality parallel corpus, using the following method:

1. **Mono-lingual Corpus:** We compiled a large monolingual collection of texts written in European Portuguese.
2. **Translation:** Using a translation system, we created a parallel corpus from the mono-lingual corpus, by translating from European Portuguese to English.
3. **Filtering:** We conducted filtering and quality checks to ensure the integrity of the dataset.

This process creates a corpus with parallel data pairs for English and European Portuguese, enabling the training of MT systems. In the following sections, we describe each step in detail.

European Portuguese Corpus Collection

For the collection of texts in European Portuguese, we used two sources: the DSL-TL (Zampieri et al. 2024) and the PtBrVid corpus (Sousa et al. 2025). The DSL-TL corpus com-

prises a total of 4,458 news articles² written in Portuguese, manually annotated as “European Portuguese”, “Brazilian Portuguese”, or “Both”. For our purposes, we use the train partition of this dataset and keep the texts labeled as “European Portuguese” and “Both” as they are both valid texts in European Portuguese, resulting in a total of 1,734 documents.

The larger chunk of our data comes from the PtBrVid corpus, which was originally created to train a Portuguese variety identifier that discriminates between European and Brazilian Portuguese. This corpus was constructed by compiling several datasets for each variety across six domains – journalism, web, social media, literature, legal, and politics – using metadata from the original datasets to label the variety of its texts. For example, the political subset uses the CETEM Público (Rocha and Santos 2000) corpus for European Portuguese and the CETEM Folha³ for Brazilian Portuguese. For this research, we kept all entries that were labeled as “European Portuguese”.

Translation

Monolingual data was translated into English using Google Translate with the Python library `deep_translator`⁴. The reason for this choice is the accessibility and relatively good quality of Google Translate, however, this step can be achieved by other translation systems. When this dataset was created, Google Translate did not distinguish between European and Brazilian Portuguese. Consequently, using it as the translation engine assumes that the translation from European Portuguese to English is unaffected by this limitation.

To verify this hypothesis, we conducted an experiment using one of our test datasets, namely FRMT (Riley et al. 2023), which contains 2,616 examples in English, Brazilian Portuguese, and European Portuguese. We used Google Translate to translate the European Portuguese and Brazilian Portuguese texts into English and then compared the English translations produced from the different varieties. We found that Google Translate produced exactly the same translation for approximately 87.2% of the dataset and achieved a BLEU score of 96.8%. This confirms that the quality of the dataset is only mildly affected by Google’s lack of differentiation between varieties. Nonetheless, this assumption is language-specific and might not hold for different language varieties and different translation systems.

Filtering

The aggregation of the mentioned resources results in a total of 3,966,538 documents. As previous research focused on training and fine-tuning LMs has emphasized the importance of data quality (Wenzek et al. 2020; Penedo et al. 2024), we apply extensive filtering and quality checks to our dataset.

Specifically, our filtering pipeline begins by using the

²Although the authors report 4,953 articles in their paper, the files shared on the project repository <https://github.com/LanguageTechnologyLab/DSL-TL> contain 4,458 rows.

³https://linguateca.pt/cetenfolha/index_info.html

⁴<https://github.com/nidhaloff/deep-translator>

justText⁵ library – designed to clean boilerplate content by classifying textual blocks based on various features such as length and stop word density – to remove entries classified as low-quality text. As depicted in Figure 1, this is the most stringent filter in our process, eliminating approximately 1.9 million documents from the collection. A closer examination reveals that most of the discarded documents originated from the Social Media partition of the PtBrVid corpus, which initially contained 2,014,752 documents, of which only 260,315 remained after applying this filter.

To further refine the dataset, we removed all entries for which the Portuguese text were duplicated, contained ASCII characters that are not used in the Portuguese language (like “ø” or “ž”), or included repetitive templates that were over-represented in the dataset (steps “Duplicates”, “Invalid Chars” and “Patterns” in Figure 1). For example, we identified 1,132 documents that began with “Lista de alterações recentes” (“List of recent changes” in English) from a badly scrapped page and 221 documents that started with “Filtrar por” (“Filter by”), referring to search engine filters.

Additionally, we filtered out all documents exceeding 900 tokens in the combined source and target texts⁶. This ensures that, regardless of the prompt template used during training, no entry exceeds 1024 tokens. The data loss from this step was minimal, as the dataset primarily consists of single-paragraph documents without line breaks. In fact, documents exceeding 900 tokens accounted for only 0.54% of the entire dataset. However, removing these documents significantly improved training speed by enabling a larger batch size. When deploying this model, it is crucial to be aware of this limitation and implement strategies to handle larger inputs, such as sentence splitting.




Figure 1: Number of documents (in millions) remaining after each step of our filtering pipeline.

The final step of our filtering pipeline involves a series of miscellaneous checks, such as ensuring that each text contains a matching number of opening and closing brackets

⁵<https://github.com/miso-belica/justText>

⁶Token count was determined using the LLaMA3 tokenizer.

and quotation marks. For detailed information about the operations performed, we refer to our code repository.

We refer to the final parallel corpus as PTradutor. Table 1 presents relevant statistics of the dataset after filtering, including the number of tokens and the domains covered. The code to replicate the dataset is available as open-source⁷, and the final corpus is publicly accessible on HuggingFace⁸.

3 Experimental Setup

In this section, we describe the training paradigm and model variations used in this study. We primarily focus on fine-tuning small-sized LMs, since smaller models are easier for the community to deploy and use.

Models

To align with the pre-training objective of language models (LMs), we model the translation task as a causal language modeling problem, conditioned on both the translation prompt and the input English text. Specifically, given a translation prompt template $prompt = T(text_{en})$ (with $text_{en}$ being the text in English) and a parallel corpus of Portuguese and English, we fine-tune a pre-trained model $LM(prompt)$ to generate the translated version of the English text in European Portuguese, $text_{pt}$. The language models we use are as follows:

Gemma-2 The latest installment in Google’s model series, Gemma-2, includes models with 2, 9, and 27 billion parameters (Team 2024). These models are recognized for their strong performance in reasoning and problem-solving tasks. We use the 2 billion parameter model from this series.

Phi-3-mini Developed by Microsoft (Abdin et al. 2024), the Phi-3 models are compact transformer decoder architectures (Vaswani et al. 2017) with a great percentage of the training corpus being composed of syntactic data. We employ the model with 3.8 billion parameters.

LLaMA-3 This is the third iteration of Meta’s large language model series (AI 2024). LLaMA-3 models are noted for their enhanced reasoning abilities, logical consistency, and reduced hallucination compared to earlier versions. We use the smallest model from this series, which has 8 billion parameters.

Since the goal is to fine-tune the models to follow a specific instruction (translate the English text to European Portuguese), we use the instruct-tuned version of all models as the starting checkpoint. The prompt templates for each model are designed by following the models’ guidelines, with the system message of “You are a translator from English to European Portuguese” (in case the model supports a system message) and the user prompt “Translate this text from English to European Portuguese: $text_{en}$ ”.

⁷<https://github.com/hmosousa/ptradutor>

⁸<https://huggingface.co/datasets/hugosousa/PTradutor>

Domain	# Docs	# Tokens PT				# Tokens EN			
		Min	Max	Mean	Total	Min	Max	Avg	Total
Journalistic	1,250,982	45	511	202.9	253,767,361	25	433	150.3	188,072,054
Literature	12,082	51	510	121.0	1,461,651	37	360	89.8	1,085,296
Web	9,006	44	555	224.7	2,024,062	28	416	167.1	1,504,751
Politics	477	53	524	244.9	116,836	36	380	171.5	81,801
Legal	282,870	44	451	87.1	24,635,676	25	385	64.9	18,346,240
Social Media	163,585	41	129	71.0	11,622,673	26	121	55.2	9,025,327
DSL-TL (news)	1,734	14	135	63.6	110,334	10	108	47.2	81,821
All	1,719,002	14	555	170.8	293,628,259	10	433	126.9	218,115,469

Table 1: Statistics describing the PTradutor corpus, including the number of documents (# Docs) and the minimum, maximum, average, and total number of tokens per domain.

Training Approaches

We utilize two types of instruction fine-tuning:

Full fine-tuning This is similar to the pre-training paradigm, in which given a prompt as instruction, the LM is trained end-to-end, updating all the parameters during optimization. This method although providing better results is often costly, due to the large size of current LMs.

Parameter efficient fine-tuning (PEFT) To decrease the computational and storage costs methods for PEFT (Xu et al. 2023) enable efficient adaptation of LMs by only fine-tuning a small number of extra model parameters instead of all the model’s parameters. LoRA (Hu et al. 2022; Xu et al. 2024) is one of the most popular methods in this domain, which achieves this by injecting trainable low-rank matrices into each layer of the transformer architecture. In addition to full fine-tuning, we also train the LoRA variant of each model.

4 Evaluation

In this section, we provide a detailed evaluation of our system against various baselines, on two European Portuguese benchmarks. The code for training and evaluation, as well as the trained checkpoints of our models, is available in our repository⁹.

Metrics

For a comprehensive evaluation, we include both classical and embedding-based metrics.

N-gram based metrics Classical machine translation metrics (Popovic 2015; Papineni et al. 2002; Lin 2004; Banerjee and Lavie 2005) usually focus on n-gram overlap between the reference translation and the translation generated by the system. Although by design this metrics fail to recognize semantic similarity beyond the lexical level, they remain widely adopted in academic research. In this work, we incorporate the two most common metrics: BLEU (Papineni et al. 2002) and ROUGE (Lin 2004).

Learnable metrics These methods focus on directly learning human judgment through training. For this family, we include COMET (Rei et al. 2020), which leverages a pre-trained multilingual model. Although COMET can function as a reference-less metric, in this work, we report results exclusively for its direct assessment variant.

Language variety metric To assess if the text produced by our translation system is indeed in European Portuguese, we use a Portuguese language variant classifier (Sousa et al. 2025), which distinguishes between Brazilian and European Portuguese. After translation of the benchmark data, we employ the classifier to label all generated texts and to compute the percentage of documents that are labeled as European Portuguese. Since the classifier might contain intrinsic errors and bias, we also compute the percentage of documents labeled as European Portuguese in the reference translations. This step is crucial for handling cases where the translated text for the two variants might be identical. In such scenarios, the classifier might incorrectly classify the variety as the wrong one due to the lack of distinguishing features in the text. By comparing the results of our system with the reference translations, we can correct for this potential bias and obtain a more accurate assessment of how well our system produces European Portuguese. We refer to the ratio of these two percentages as the VID score, which serves as a measure of the system’s effectiveness in generating European Portuguese text.

Test Benchmarks

As a low-resource language variant, the number of benchmarks that include European Portuguese is limited. In this study, we use two high-quality publicly available datasets that feature this variant:

- **FRMT**: This dataset is specifically designed to contain regional variants of Portuguese and Chinese (Riley et al. 2023), containing human translations of sentences from English Wikipedia articles that were manually translated to European and Brazilian Portuguese.
- **NTrex**: The dataset consists of high-quality translations

⁹<https://github.com/hmosousa/tradutor>

by speakers who are bilingual in English and in one of the 128 target languages, including 123 documents and 1,997 sentences for each language (Federmann, Kocmi, and Xin 2022).

Baselines

We compare our models to three sets of baselines:

Closed Baselines We include industry-standard systems for Portuguese translation, including Google Translate and DeepL. Recently, Google Translate introduced a model specifically designed for European Portuguese, referred to as `Googlept`, which we include in our evaluation alongside the original model that does not distinguish between Portuguese varieties, referred to as `Googlebr`.

Open Baselines For open-source systems, we evaluate ArgosTranslate¹⁰, which uses OpenNMT (Klein et al. 2017) as its backend. Although Portuguese is listed as a supported language, the specific variety is not indicated. We also consider the Opus-MT project (Tiedemann and Thottingal 2020), another open-source system that provides a model for translating from English to Portuguese. However, like ArgosTranslate, this system does not differentiate between Portuguese varieties.

Zero-shot Additionally, we assess the zero-shot capabilities of language models without applying our task-specific fine-tuning to demonstrate the effectiveness of the fine-tuning process.

Implementation Details

All models were trained and evaluated on a server with six A-100 GPUs, each with 40GB of memory. The batch size and training duration varied depending on the memory requirements of each model. In our repository, we provide training and evaluation scripts compatible with the two libraries used to train the language models: `torchtune`¹¹ and `transformers`¹². While the `transformers` library was chosen for its practicality, we found it limiting when training larger models. In that scenario, `torchtune` presented as a reliable alternative with significantly better memory management. Training runs we executed with early stopping – using the test set the DSL-TL corpus as validation set – with patience of 3,000 steps. As a result, the number of training steps varied across models. All LoRA variants were trained with an alpha of 128 and a rank of 256. Detailed training configurations can be found in our repository.

The parameter setup is as follows:

- **Phi-3:** For both the LoRA and full fine-tuning of Phi-3 models, we used a batch size of 512, a learning rate of $2e-5$, a weight decay of 0.1, and a warm-up of 1,000 steps.
- **Gemma-2:** For both variants, the learning rate was set to $2e-5$ with a weight decay of 0.1. The full fine-tuned model was trained with a 1,000-step warm-up and a batch

size of 512, while the LoRA variant had 500 warm-up steps and a batch size of 256.

- **LLaMA-3:** Both variants were trained with a batch size of 256 and a learning rate of $2e-5$. The LoRA variant additionally includes a warm-up of 100 steps and a weight decay of 0.1 on the learning rate.

Results

The result of our evaluation is shown in Table 2, where the best overall values are marked with an underline, and the most effective open-source systems are marked in bold. The general trend on both datasets is quite similar and high-performing models maintain a stable performance across both benchmarks. Yet, values for NTrex are slightly below FRMT for all systems, indicating a harder benchmark. In the following, we describe our main findings and highlight avenues for future research.

LoRA models: These variants effectively learn the vocabulary of European Portuguese, but their overall translation quality remains subpar. This discrepancy is evident as they score highly on the VID metric, nevertheless, the BLEU, ROUGE-L, and COMET metrics suggest that they struggle with generating high-quality text. Upon closer inspection of the generated translations, we found that both the Phi-3 and LLaMA-3 models, when trained with LoRA, tend to enter a repetition loop, where the same token is generated repeatedly until the process is interrupted. This suggests that for a medium size language model translation is a complex task, and simply adding adapter parameters is insufficient to fully capture its nuances. This issue, combined with the fact that the early stopping criteria were reached quickly during training (training loss had plateaued while the evaluation was increasing), suggests that the models may require increased capacity (by adjusting the alpha and rank parameters) to better learn from the training data.

Full fine-tuning (FFT): The fully fine-tuned models sacrifice their mastery of European Portuguese nuances in favor of high-quality, coherent translations. These models yield more moderate scores for the VID metric while achieving significantly higher text quality metrics. In particular, the fine-tuned LLaMA-3 model beats all open source software on all metrics and produces results comparable to `Googlebr` in terms of BLEU and ROUGE-L on both benchmarks, only falling short on the COMET metric.

It is important to note that the COMET metric, as a learnable metric, is subject to the same biases that affect any trainable neural model. Since the training data does not differentiate between European and Brazilian Portuguese and most open-source resources are skewed toward Brazilian Portuguese, this bias may have influenced the scores produced by this metric. The slight difference between `Googlebr` and fine-tuned LLaMA-3 models suggests a potential bias in the COMET score toward Brazilian Portuguese.

The overall performance of the fully fine-tuned models has a direct correlation with model size, where our largest model, LLaMA-3, with 8 billion parameters outperforms the smaller models of Phi-3 (3.8 billion parameters) and

¹⁰<https://github.com/argosopentech/argos-translate>

¹¹<https://github.com/pytorch/torchtune/>

¹²<https://huggingface.co/>

	Model	FRMT				NTrex			
		BLEU	ROUGE-L	COMET	VID	BLEU	ROUGE-L	COMET	VID
Close Baselines	Google _{br}	43.20	68.43	87.44±0.25	0.445	35.80	63.44	86.88±0.32	0.361
	Google _{pt}	47.81	71.66	87.87±0.25	0.956	39.92	66.72	87.17±0.33	0.900
	DeepL	<u>49.77</u>	<u>72.44</u>	<u>88.48±0.23</u>	0.999	<u>44.76</u>	<u>67.77</u>	<u>87.96±0.31</u>	<u>0.997</u>
Open Baselines	Argos	38.39	65.07	83.99±0.35	0.511	30.44	58.72	80.30±0.54	0.446
	Opus-MT	40.41	66.25	85.67±0.31	0.413	32.99	60.24	83.81±0.46	0.229
Zero-shot	Gemma-2	25.37	49.56	75.66±0.51	0.807	17.47	41.34	69.46±0.66	0.858
	Phi-3	17.59	43.99	57.90±0.56	0.942	13.16	38.38	54.62±0.66	1.003
	LLaMA-3	31.47	60.61	82.95±0.40	0.811	21.09	52.34	78.70±0.60	0.805
LoRA	Gemma-2	19.83	56.87	79.62±0.64	0.530	14.41	49.42	76.18±0.81	0.514
	Phi-3	24.70	53.34	72.19±0.58	1.178	20.10	48.38	67.67±0.73	1.203
	LLaMA-3	25.42	51.51	74.06±0.56	1.092	17.74	41.24	67.96±0.73	1.140
FFT	Gemma-2	33.76	66.41	85.25±0.35	1.066	18.35	59.75	82.62±0.54	1.049
	Phi-3	38.16	66.31	85.35±0.34	1.055	27.89	60.18	82.91±0.49	1.031
	LLaMA-3	41.12	66.92	86.12±0.28	0.968	35.76	62.02	84.42±0.42	0.933

Table 2: Model effectiveness was assessed using the FRMT and NTrex benchmarks. Confidence intervals for the COMET metric were computed using a t -distribution with a 95% confidence level. The best results among open-source systems are highlighted in bold, while the best overall results are underlined.

Gemma-2 (2 billion parameters). This behavior is typical for large language models and indicates that increasing their size can lead to better results. This indicates that by applying this methodology to even larger models, it may be possible to achieve results that are on par with industry-standard systems.

Since our focus is on translation specific to European Portuguese, it is important to examine the VID scores across both benchmarks. Our best model, LLaMA-3, once again beats all open source baselines and achieves scores significantly higher than Google_{br} and is comparable to Google_{pt} and DeepL. This suggests that the proposed methodology is effective in achieving the targeted goal of producing text specific to a language variety. Even our smaller-size models, perform comparable to open-source baselines on translation metrics, dramatically improving on the VID score.

It is true that in terms of text quality metrics, the LLaMA-3 model still lags behind the European Portuguese-specific industry models, namely Google_{pt} and DeepL. However, it is important to emphasize that our goal was not to beat the specialized model from industry, but to propose a computationally efficient, adaptable, and resource-efficient method for adapting small language models to translate specific language varieties. Surpassing the current open-source software and achieving a score close to industry-level models, which benefit from dedicated teams of experts and annotators for each language, is a significant accomplishment.

5 Related Work

Despite the availability of industry-level translation systems like Google Translate¹³ and DeepL¹⁴ and a handful of open-source software with unclear language variant definition (Klein et al. 2017; Tiedemann and Thottingal 2020), there are, to the best of our knowledge, no other translation models specifically dedicated to European Portuguese in the literature. Like many other industry systems, these models lack transparency, as neither their internal workings nor the data used for training are publicly accessible.

Since this work focuses on Portuguese, we begin by reviewing relevant research in this area. Lakew, Erofeeva, and Federico (2018) explore NMT from English into four pairs of language varieties, including European and Brazilian Portuguese, unlike our focus on Portuguese to English. They train a transformer model using transcripts of movies and TED talks but do not open-source their models for comparison. The work primarily highlights the challenges of low-resource settings and the initial steps to address these issues.

Another work in this direction is from Team (2022), which aims to train a translation model between standard national varieties of the same language, namely between Brazilian and European Portuguese. They provide a taxonomy of distinctive characteristics between these two language varieties, which we refer curious readers to for a deeper understanding of the language varieties discussed in this study.

Other previous work in this direction explored a variety-

¹³<https://translate.google.com/>

¹⁴<https://deepl.com/>

targeted MT system in other languages, which contains varieties or dialects specific to a region. It is worth noting that although some of these works are similar to textual style transfer in methodology, they are different in the task definition (Jhamtani et al. 2017).

One of the earlier works in this area focuses on the development of translation systems for Arabic dialects (Zbib et al. 2012; Sajjad, Darwish, and Belinkov 2013). Efforts for Arabic took off with the introduction of AraBench benchmark (Sajjad et al. 2020), an evaluation suite for dialectal Arabic to English MT.

Similar efforts have been undertaken for the low-resource language family of Swiss German, which is widely spoken in Switzerland, but rarely written (Honnet et al. 2018; Scherrer and Ljubecic 2016). These systems focus on normalizing Swiss German to the standard variant of High German.

Similarly to Arabic and German other languages, such as Chinese, Russian, Hindi and Turkish, are also slowly finding their way into this study (Wan et al. 2020; Kumar et al. 2021; Nguyen and Chiang 2017; Durrani et al. 2010).

Another area of research loosely related to ours involves adapting language models to low-resource language varieties. As previously mentioned, the effectiveness of large language models in this domain is often limited by the scarcity of representative datasets (Alam et al. 2024). Recent efforts have focused on creating specialized language models for specific dialects (Ondrejová and Šuppa 2024; Faisal and Anastasopoulos 2024; Nguyen et al. 2024; Lopes, Magalhães, and Semedo 2024), mainly through data augmentation or the introduction of new datasets. Although these are not translation systems and therefore are not directly comparable to our work, our methodology and provided resources can also be used to provide artificial training data.

6 Conclusion & Future Work

In this paper, we present a methodology for creating a parallel corpus and training a translation model tailored to a low-resource language variety. Specifically, we developed and open-sourced the largest European Portuguese-English parallel corpus, along with European Portuguese-specific translation models. Our extensive evaluation demonstrates the effectiveness of our approach and the fidelity of the generated translations to the desired language variety. Thanks to the proposed methodology we managed to achieve performance on par with industry-level translation systems with minimal resources and limited computation.

In future work, it would be interesting to investigate the impact of different generation configurations on the translations produced by our model. In this study, we used greedy decoding, but other generation techniques, such as beam search, could yield better results. Another promising direction is to explore prompt optimization both before and after model training. This approach has been shown to improve outcomes in other studies (Soylu, Potts, and Khattab 2024) and might enhance our system’s performance as well. Finally, we plan to conduct a human evaluation with linguists to identify areas where our model falls short compared to other systems.

Acknowledgements

This research is supported by national funding from the Portuguese Foundation for Science and Technology (FCT) under the project with DOI 10.54499/LA/P/0063/2020. The authors also acknowledge the support of the StorySense project (DOI 10.54499/2022.09312.PTDC) and the advanced computing project PTicola (ID CPCA-IAC/AV/594794/2023). Hugo Sousa further acknowledges FCT for funding his PhD grant (ID 2022.14691.BD).

References

- Abdin, M.; Aneja, J.; Awadalla, H.; Awadallah, A.; Awan, A. A.; Bach, N.; Bahree, A.; Bakhtiari, A.; Bao, J.; Behl, H.; Benhaim, A.; Bilenko, M.; Bjorck, J.; Bubeck, S.; Cai, M.; Cai, Q.; Chaudhary, V.; Chen, D.; Chen, D.; Chen, W.; Chen, Y.-C.; Chen, Y.-L.; Cheng, H.; Chopra, P.; Dai, X.; Dixon, M.; Eldan, R.; Fragoso, V.; Gao, J.; Gao, M.; Gao, M.; Garg, A.; Giorno, A. D.; Goswami, A.; Gunasekar, S.; Haider, E.; Hao, J.; Hewett, R. J.; Hu, W.; Huynh, J.; Iter, D.; Jacobs, S. A.; Javaheripi, M.; Jin, X.; Karampatziakis, N.; Kauffmann, P.; Khademi, M.; Kim, D.; Kim, Y. J.; Kurilenko, L.; Lee, J. R.; Lee, Y. T.; Li, Y.; Li, Y.; Liang, C.; Liden, L.; Lin, X.; Lin, Z.; Liu, C.; Liu, L.; Liu, M.; Liu, W.; Liu, X.; Luo, C.; Madan, P.; Mahmoudzadeh, A.; Majercak, D.; Mazzola, M.; Mendes, C. C. T.; Mitra, A.; Modi, H.; Nguyen, A.; Norick, B.; Patra, B.; Perez-Becker, D.; Portet, T.; Pryzant, R.; Qin, H.; Radmilac, M.; Ren, L.; de Rosa, G.; Rosset, C.; Roy, S.; Ruwase, O.; Saarikivi, O.; Saied, A.; Salim, A.; Santacroce, M.; Shah, S.; Shang, N.; Sharma, H.; Shen, Y.; Shukla, S.; Song, X.; Tanaka, M.; Tupini, A.; Vaddamanu, P.; Wang, C.; Wang, G.; Wang, L.; Wang, S.; Wang, X.; Wang, Y.; Ward, R.; Wen, W.; Witte, P.; Wu, H.; Wu, X.; Wyatt, M.; Xiao, B.; Xu, C.; Xu, J.; Xu, W.; Xue, J.; Yadav, S.; Yang, F.; Yang, J.; Yang, Y.; Yang, Z.; Yu, D.; Yuan, L.; Zhang, C.; Zhang, C.; Zhang, J.; Zhang, L. L.; Zhang, Y.; Zhang, Y.; Zhang, Y.; and Zhou, X. 2024. Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone.
- AI, M. 2024. Introducing Meta Llama 3: The most capable openly available LLM to date.
- Alam, F.; Chowdhury, S. A.; Boughorbel, S.; and Hasanain, M. 2024. LLMs for Low Resource Languages in Multilingual, Multimodal and Dialectal Settings. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Tutorial Abstracts*, 27–33.
- Albalak, A.; Elazar, Y.; Xie, S. M.; Longpre, S.; Lambert, N.; Wang, X.; Muennighoff, N.; Hou, B.; Pan, L.; Jeong, H.; Raffel, C.; Chang, S.; Hashimoto, T.; and Wang, W. Y. 2024. A Survey on Data Selection for Language Models.
- Armengol-Estapé, J.; Carrino, C. P.; Rodriguez-Penagos, C.; de Gibert Bonet, O.; Armentano-Oller, C.; Gonzalez-Agirre, A.; Melero, M.; and Villegas, M. 2021. Are Multilingual Models the Best Choice for Moderately Under-resourced Languages? A Comprehensive Assessment for Catalan. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds.,

- Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 4933–4946. Association for Computational Linguistics.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In Goldstein, J.; Lavie, A.; Lin, C.-Y.; and Voss, C., eds., *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 65–72. Association for Computational Linguistics.
- Brito, A. M.; and Lopes, R. E. 2016. *The Structure of DPs*, 254–274. Wiley Blackwell, 1st edition.
- Durrani, N.; Sajjad, H.; Fraser, A.; and Schmid, H. 2010. Hindi-to-Urdu Machine Translation through Transliteration. In Hajič, J.; Carberry, S.; Clark, S.; and Nivre, J., eds., *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, 465–474. Association for Computational Linguistics.
- Faisal, F.; and Anastasopoulos, A. 2024. Data-Augmentation-Based Dialectal Adaptation for LLMs. In *Proceedings of the Eleventh Workshop on NLP for Similar Languages, Varieties, and Dialects (VarDial 2024)*, 197–208. Association for Computational Linguistics.
- Federmann, C.; Kocmi, T.; and Xin, Y. 2022. NTREX-128 – News Test References for MT Evaluation of 128 Languages. In *Proceedings of the First Workshop on Scaling Up Multilingual Evaluation*, 21–24. Association for Computational Linguistics.
- Honnet, P.-E.; Popescu-Belis, A.; Musat, C.; and Baeriswyl, M. 2018. Machine Translation of Low-Resource Spoken Dialects: Strategies for Normalizing Swiss German. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jhamtani, H.; Gangal, V.; Hovy, E.; and Nyberg, E. 2017. Shakespearizing Modern Language Using Copy-Enriched Sequence to Sequence Models. In *Proceedings of the Workshop on Stylistic Variation*, volume 6, 10–19. Association for Computational Linguistics.
- Kato, M. A.; and Martins, A. M. 2016. *European Portuguese and Brazilian Portuguese: An Overview on Word Order*, 15–40. Wiley Blackwell, 1st edition.
- Klein, G.; Kim, Y.; Deng, Y.; Senellart, J.; and Rush, A. 2017. OpenNMT: Open-Source Toolkit for Neural Machine Translation. In *Proceedings of ACL 2017, System Demonstrations*, 67–72. Association for Computational Linguistics.
- Kumar, S.; Anastasopoulos, A.; Wintner, S.; and Tsvetkov, Y. 2021. Machine Translation into Low-resource Language Varieties. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 110–121. Association for Computational Linguistics.
- Lakew, S. M.; Erofeeva, A.; and Federico, M. 2018. Neural Machine Translation into Language Varieties. In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, 156–164. Association for Computational Linguistics.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lopes, R.; Magalhães, J.; and Semedo, D. 2024. GlóRIA: A Generative and Open Large Language Model for Portuguese. In *Proceedings of the 16th International Conference on Computational Processing of Portuguese, PROPOR 2024, Santiago de Compostela, Galicia/Spain, 12-15 March, 2024*, 441–453. Association for Computational Linguistics.
- Nguyen, T. Q.; and Chiang, D. 2017. Transfer Learning across Low-Resource, Related Languages for Neural Machine Translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017, Volume 2: Short Papers*, 296–301. Asian Federation of Natural Language Processing.
- Nguyen, X.-P.; Zhang, W.; Li, X.; Aljunied, M.; Hu, Z.; Shen, C.; Chia, Y. K.; Li, X.; Wang, J.; Tan, Q.; Cheng, L.; Chen, G.; Deng, Y.; Yang, S.; Liu, C.; Zhang, H.; and Bing, L. 2024. SeaLLMs - Large Language Models for Southeast Asia. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, volume abs/2312.00738, 294–304. Association for Computational Linguistics.
- Ondrejová, V.; and Šuppa, M. 2024. Can LLMs Handle Low-Resource Dialects? A Case Study on Translation and Common Sense Reasoning in Šariš. In *Proceedings of the Eleventh Workshop on NLP for Similar Languages, Varieties, and Dialects (VarDial 2024)*, 130–139. Association for Computational Linguistics.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, 311–318. ACL.
- Penedo, G.; Malartic, Q.; Hesslow, D.; Cojocaru, R.; Alobeidli, H.; Cappelli, A.; Pannier, B.; Almazrouei, E.; and Lounay, J. 2024. The RefinedWeb dataset for falcon LLM: outperforming curated corpora with web data only. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. Curran Associates Inc.
- Popovic, M. 2015. ChrF: Character n-gram F-score for Automatic MT Evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation, WMT@EMNLP 2015, 17-18 September 2015, Lisbon, Portugal*, 392–395. The Association for Computer Linguistics.
- Rei, R.; Stewart, C.; Farinha, A. C.; and Lavie, A. 2020. COMET: A Neural Framework for MT Evaluation. In Web-

- ber, B.; Cohn, T.; He, Y.; and Liu, Y., eds., *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2685–2702. Association for Computational Linguistics.
- Riley, P.; Dozat, T.; Botha, J. A.; Garcia, X.; Garrette, D.; Riesa, J.; Firat, O.; and Constant, N. 2023. FRMT: A Benchmark for Few-Shot Region-Aware Machine Translation. *Transactions of the Association for Computational Linguistics*, 11: 671–685.
- Rocha, P. G. M.; and Santos, D. 2000. CETEMPública: Um corpus de grandes dimensões de linguagem jornalística portuguesa.
- Rodrigues, J.; Gomes, L.; Silva, J.; Branco, A.; Santos, R.; Cardoso, H. L.; and Osório, T. 2023. *Advancing Neural Encoding of Portuguese with Transformer Albertina PT-**, 441–453. Springer Nature Switzerland. ISBN 9783031490088.
- Sajjad, H.; Abdelali, A.; Durrani, N.; and Dalvi, F. 2020. AraBench: Benchmarking Dialectal Arabic-English Machine Translation. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, 5094–5107. International Committee on Computational Linguistics.
- Sajjad, H.; Darwish, K.; and Belinkov, Y. 2013. Translating Dialectal Arabic to English. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 2: Short Papers*, 1–6. The Association for Computer Linguistics.
- Scherre, M. M. P.; and Duarte, M. E. L. 2016. Main current processes of morphosyntactic variation. *The Handbook of Portuguese Linguistics*, 526–544.
- Scherrer, Y.; and Ljubesic, N. 2016. Automatic normalisation of the Swiss German ArchiMob corpus Using Character-level Machine Translation. In *Proceedings of the 13th Conference on Natural Language Processing, KONVENS 2016, Bochum, Germany, September 19-21, 2016*, volume 16.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Controlling Politeness in Neural Machine Translation via Side Constraints. In Knight, K.; Nenkova, A.; and Rambow, O., eds., *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 35–40. Association for Computational Linguistics.
- Sousa, H.; Almeida, R.; Silvano, P.; Cantante, I.; Campos, R.; and Jorge, A. 2025. Enhancing Portuguese Variety Identification with Cross-Domain Approaches. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39.
- Soylu, D.; Potts, C.; and Khattab, O. 2024. Fine-Tuning and Prompt Optimization: Two Great Steps that Work Better Together. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 10696–10710. Association for Computational Linguistics.
- Team, G. 2024. Gemma 2: Improving Open Language Models at a Practical Size.
- Team, N. 2022. No Language Left Behind: Scaling Human-Centered Machine Translation.
- Tiedemann, J.; and Thottingal, S. 2020. OPUS-MT – Building open translation services for the World. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, 479–480. European Association for Machine Translation.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 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*, 5998–6008. ISBN 9781510860964.
- Wan, Y.; Yang, B.; Wong, D. F.; Chao, L. S.; Du, H.; and Ao, B. C. H. 2020. Unsupervised Neural Dialect Translation with Commonality and Diversity Modeling. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34: 9130–9137.
- Wenzek, G.; Lachaux, M.-A.; Conneau, A.; Chaudhary, V.; Guzmán, F.; Joulin, A.; and Grave, E. 2020. CCNet: Extracting High Quality Monolingual Datasets from Web Crawl Data. In Calzolari, N.; Béchet, F.; Blache, P.; Choukri, K.; Cieri, C.; Declerck, T.; Goggi, S.; Isahara, H.; Maegaard, B.; Mariani, J.; Mazo, H.; Moreno, A.; Odijk, J.; and Piperidis, S., eds., *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 4003–4012. European Language Resources Association. ISBN 979-10-95546-34-4.
- Xu, L.; Xie, H.; Qin, S.-Z. J.; Tao, X.; and Wang, F. L. 2023. Parameter-Efficient Fine-Tuning Methods for Pre-trained Language Models: A Critical Review and Assessment.
- Xu, Y.; Xie, L.; Gu, X.; Chen, X.; Chang, H.; Zhang, H.; Chen, Z.; Zhang, X.; and Tian, Q. 2024. QA-LoRA: Quantization-Aware Low-Rank Adaptation of Large Language Models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Zampieri, M.; North, K.; Jauhiainen, T.; Felice, M.; Kumari, N.; Nair, N.; and Banger, Y. M. 2024. Language Variety Identification with True Labels. In Calzolari, N.; Kan, M.-Y.; Hoste, V.; Lenci, A.; Sakti, S.; and Xue, N., eds., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 10100–10109. ELRA and ICCL.
- Zbib, R.; Malchiodi, E.; Devlin, J.; Stallard, D.; Matsoukas, S.; Schwartz, R.; Makhoul, J.; Zaidan, O. F.; and Callison-Burch, C. 2012. Machine Translation of Arabic Dialects. In Fosler-Lussier, E.; Riloff, E.; and Bangalore, S., eds., *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 49–59. Association for Computational Linguistics.