

Investigating the translation capabilities of Large Language Models trained on parallel data only

Anonymous ACL submission

Abstract

In recent years, Large Language Models (LLMs) have demonstrated exceptional proficiency across a broad spectrum of Natural Language Processing (NLP) tasks, including Machine Translation. However, previous methods predominantly relied on iterative processes such as instruction fine-tuning or continual pre-training, leaving unexplored the challenges of training LLMs solely on parallel data. In this work, we introduce PLUME (Parallel Language Model), a collection of three 2B LLMs featuring varying vocabulary sizes (32k, 128k, and 256k) trained exclusively on Catalan-centric parallel examples. These models perform comparably to previous encoder-decoder architectures on 16 supervised translation directions and 56 zero-shot ones. Utilizing this set of models, we conduct a thorough investigation into the translation capabilities of LLMs, probing their performance, the impact of the different elements of the prompt, and their cross-lingual representation space. We will make our models publicly available¹.

1 Introduction

Neural Machine Translation (NMT) has traditionally relied on encoder-decoder architectures, where an encoder processes the source sentence and a decoder generates the target sentence based on the encoder’s output. However, recent advancements have moved away from this paradigm, with the introduction of decoder-only Large Language Models (LLMs). In these models, the source sentence acts as a prompt, eliminating the need for a conventional encoder.

With the rise of LLMs, research has increasingly focused on adapting these models for translation tasks by using techniques such as prompt-tuning (Zhang et al., 2023), instruction-finetuning (Xu

et al., 2024), or continual pretraining (Rei et al., 2022a). While these methods have shown impressive results, they open new questions about the performance of LLMs when trained exclusively on parallel data, and therefore, the possibility of having models that are trained directly on the task of machine translation. Additionally, the majority of these models are trained predominantly on English-centric corpora.

To address these questions, our paper proposes a new approach consisting of training LLMs solely on parallel corpora to evaluate their efficacy in machine translation (MT). Our investigation revolves around questions such as: How does an LLM trained exclusively on parallel data perform? And how does the model leverage prompt information to ensure accurate translations?

Our contributions are twofold: Firstly, we introduce PLUME (Parallel Language Model), an innovative ensemble comprising three multilingual 2B LLMs, trained from scratch on Catalan-centric parallel data. Each model has a different vocabulary size (32k, 128k and 256k). All models are proficient in 16 supervised translation directions, as well as 56 zero-shot translation directions. Results show comparable results to previous encoder-decoder architectures of similar size.

Secondly, to understand how these models work, we study how they utilize contextual information across different layers to execute translation tasks effectively. Our experiments show distinctive attention patterns associated with the different parts of the prompt, and how they vary through the different attention blocks. We also observe how languages use the source tag information differently, leading to a large performance variability when this token is missing. As a byproduct, we propose a strategy to remove attention heads with minimal performance loss. Finally, we study the cross-lingual space learned by the models and how it progresses through the model’s attention blocks.

¹We release anonymous code at https://anonymous.4open.science/r/Plume_fork-69D1

2 Related work

Neural Machine Translation (NMT) has predominantly relied on encoder-decoder architectures (Cho et al., 2014; Bahdanau et al., 2015; Sutskever et al., 2014). These methods have proven effective by conditioning language models to generate translations that accurately retain the meaning of the source sentence. Moreover, these systems are easily extendable to multilingual scenarios, enabling zero-shot translation between language pairs that have not been seen together during training (Firat et al., 2016; Wu et al., 2016).

Over the years, some approaches to NMT have dropped the traditional encoder-decoder setup to adopt decoder-only architectures (Fonollosa et al., 2019; He et al., 2018). Although these methods showed promise, they did not become the standard due to issues with context loss and hallucinations (Fu et al., 2023).

Recent advancements in training Large Language Models (LLMs) (Touvron et al., 2023; Jiang et al., 2023; Gemma Team et al., 2024; Abdin et al., 2024), including techniques like scaling and Rotary Embeddings (Su et al., 2024), have significantly enhanced the ability of decoder-only architectures to handle long contexts of hundreds or even thousands of tokens. Consequently, several studies have proposed leveraging pretrained LLMs for NMT through continual pretraining and instruction tuning (Alves et al., 2024; Xu et al., 2024; Yang et al., 2023). These methods have demonstrated results comparable to traditional encoder-decoder systems, while also supporting multiple translation directions.

However, training and adapting these systems to various languages remains challenging (Ali et al., 2024). Creating a vocabulary that accurately represents all supported languages can lead to performance disparities of up to 68% on some downstream tasks. Additionally, interpretability methods have gained popularity in order to understand better how models utilize provided information and to guide further improvements (Voita et al., 2019b,a; Ferrando et al., 2024).

3 Methodology

3.1 Catalan-Centric Dataset

In order to study zero-shot translation using a decoder-only architecture, we employ a Catalan-centric dataset. This dataset pairs Catalan sentences

with their counterparts in one of eight other languages: Spanish, French, Italian, Portuguese, Galician, German, English, and Basque. Specifically, for each language, we include translation directions both to Catalan ($xx \rightarrow ca$) and from Catalan ($ca \rightarrow xx$). The dataset consists of 783.6M sentences and 30.9 billion words.

Data preprocessing All data is first filtered using LaBSE (Feng et al., 2022). Off-target translations are removed using the *Lingua*² library. Following the filtering process, the data undergoes deduplication and punctuation normalization using the *Bifixer* library (Ramírez-Sánchez et al., 2020). Further details about the dataset are provided in Appendix A.

3.2 Tokenizer

Prior studies have shown that vocabulary overlap plays a crucial role in zero-shot translation for encoder-decoder architectures (Stap et al., 2023; Tan and Monz, 2023). More related to our work concerning tokenizer size in decoder-only architectures is the study by Ali et al. (2024), who found that larger vocabulary sizes lead to improved downstream performance in multilingual settings. The main difference is that our focus is in Multilingual Neural Machine Translation (MNMT) while Ali et al. (2024) focused on more general multilingual tasks (Natural language inference, Question Answering, etc.).

To investigate the impact of vocabulary sharing on zero-shot MNMT for decoder-only architectures, we train 3 tokenizers using BPE (Sennrich et al., 2016) from the *Huggingface tokenizer* library (Moi and Patry, 2023) with different vocabulary sizes; 32k, 128k, and 256k. Regarding the training data used to train the tokenizer, recent work has shown that while NMT performance is relatively robust to language imbalance, better performance is often achieved when languages are more equally represented in the training data (Zhang et al., 2022). In this work, we equally sample Romance languages and we oversample English, Basque, and German to avoid underrepresenting these languages and to achieve near parity (Petrov et al., 2024) and fertility among all language pairs. More details about tokenizer experiments can be found in Appendix B.

²<https://github.com/pemistahl/lingua-py>

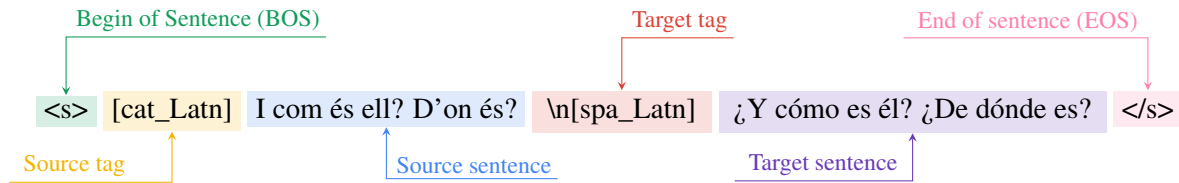


Figure 1: Prompt strategy used to train PLUME.

3.3 Model

We trained one model for each of our three tokenizers using the same architecture as GEMMA 2B³ (Gemma Team et al., 2024) to train a 2 billion parameter, transformer-based, decoder-only model. Following the scaling law proposed by (Hoffmann et al., 2022), each model was trained on 30.9 billion words, corresponding to 54.7, 46.8, and 44.6 billion tokens for vocabularies of 32k, 128k, and 256k respectively. Details about the model size and model architecture can be found in Appendix C.

3.4 Training

We train all PLUME models with a context window of 2048 tokens, utilizing the Adam optimizer (Kingma and Ba, 2015) and the causal language modeling objective. Note that the main focus of this study is to understand how LLMs perform translation. Thus, PLUME models are not trained for state-of-the-art performance on MNMT. A more detailed description of the training configuration can be found in Appendix D.

Formatting Figure 1 presents an example of a formatted sentence for the Catalan to Spanish translation direction. During batching, we concatenate formatted sentences up to a context length of 2048 tokens, mixing different translation directions within a single batch. Padding is added to fill out the remainder of the sequence.

3.5 Evaluation

To compute reference-based translation quality we use COMET-22 (Rei et al., 2022a) and BLEU (Papineni et al., 2002) metrics on the FLORES-200 devtest (NLLB Team et al., 2022) and NTREX-101 (Federmann et al., 2022) datasets. We additionally report CHRf (Popović, 2015) and COMET-KIWI-22 (Rei et al., 2022b) in appendix I. We use TOWEREVAL⁴ (Alves et al., 2024) to compute all the evaluation metrics. For inference, we use beam

³<https://huggingface.co/google/gemma-2b>

⁴TOWEREVAL uses the sacreBLEU implementation to compute BLEU and CHRf metrics.

search decoding with a beam size of 5 and limiting the translation length to 512 tokens.

We compare PLUME models with the following bilingual and multilingual models.

- NLLB (NLLB Team et al., 2022): A transformer encoder-decoder model that supports 202 languages. We use the 600 million, the 1.3 billion, and the 3.3 billion parameter variants.
- Bilingual models BSC: Transformer encoder-decoder models, trained from scratch on language pairs that include Catalan. These models were developed as part of the Aina Project⁵.

It is important to note that NLLB has seen parallel data for our zero-shot directions, therefore zero-shot only describes the condition in PLUME models. Our setup is designed to study the potential of a decoder-only architecture to perform zero-shot translation, specifically using Catalan as the pivot language.

4 Results

Table 1 shows results for all PLUME models aggregated by supervised and zero-shot directions. The PLUME 32k, 128k and 256k variants perform equally well in supervised directions, achieving similar BLEU and COMET scores for both NTREX and FLORES-200 datasets. In supervised directions, PLUME models demonstrate competitive performance, matching the COMET scores of the Bilingual BSC models and achieving scores comparable to the NLLB variants.

In zero-shot directions, the PLUME models exhibit a decline in performance compared to supervised directions. However, the decline is more pronounced in the BLEU scores than in the COMET scores, indicating that the overall quality remains relatively robust. Specifically, the PLUME 256k variant achieves a COMET score of 0.84 on the

⁵<https://huggingface.co/projecte-aina>

	Supervised directions				Zero-shot directions			
	FLORES-200		NTREX		FLORES-200		NTREX	
	BLEU	COMET	BLEU	COMET	BLEU	COMET	BLEU	COMET
NLLB-3.3B	32.02	0.87	30.48	0.85	28.97	0.86	28.74	0.84
NLLB-1.3B	31.02	0.86	29.68	0.85	28.48	0.86	28.37	0.84
NLLB-600M	29.24	0.85	28.37	0.84	27.04	0.85	27.25	0.84
Bilinguals BSC	31.93	0.86	29.77	0.84	-	-	-	-
PLUME 32k	30.44	0.86	28.46	0.84	23.25	0.83	23.03	0.80
PLUME 128k	30.81	0.86	28.78	0.84	23.97	0.83	23.53	0.81
PLUME 256k	30.72	0.86	28.87	0.84	24.42	0.84	23.81	0.81

Table 1: Averaged BLEU and COMET scores on supervised and zero-shot directions for FLORES-200 devtest and NTREX.

FLORES-200 dataset and 0.81 on the NTREX dataset, which, although lower than its supervised performance, still demonstrates its zero-shot translation capabilities when training using only Catalan as the bridge language.

Larger vocabulary sizes improve zero-shot translation. The results in Table 1 show that higher vocabulary sizes consistently yield better zero-shot capabilities. Specifically, the PLUME 256k variant outperforms the 32k and 128k variants in zero-shot scenarios for both FLORES-200 and NTREX datasets.

To further understand the influence of the vocabulary size in zero-shot translation quality, we calculated the vocabulary overlap (Tan and Monz, 2023) for each zero-shot translation direction as follows:

$$Overlap = \frac{|V_{src} \cap V_{tgt}|}{|V_{tgt}|} \quad (1)$$

where V_{src}, V_{tgt} are the set of unique words in the source and target language vocabulary respectively. We show the correlation between vocabulary overlap and both BLEU and COMET for zero-shot directions in table 2. On average there is a positive correlation between the vocabulary overlap and the translation quality of 0.3 for BLEU and 0.57 for COMET, which diminishes as vocabulary size increases. This suggests that vocabulary overlap between the source and target languages further helps explain zero-shot performance, particularly for smaller vocabulary sizes.

	PLUME 32k	PLUME 128k	PLUME 256k
BLEU	0.351	0.280	0.255
COMET	0.593	0.588	0.538

Table 2: Correlation between vocabulary overlap and BLEU, COMET metrics for different vocabulary sizes in zero-shot directions.

4.1 Understanding translation with an LLM

Our goal is to understand how an LLM performs translation. We start by examining which parts of the prompt the model focuses on. This helps us determine the most important attention heads for each section of the prompt. Then, we study the model’s cross-lingual representation space by extracting contextualized token embeddings.

4.2 Attention

For each attention head, we assess its importance by calculating coverage as defined by (Tu et al., 2016). Originally, coverage was proposed for encoder-decoder attention and refers to the total attention a source token receives from target tokens. We adapt coverage for masked-self attention. Given a set of prompt’s tokens I , the coverage formula for a single sentence is defined as:

$$cov_I(\text{head}) = \sum_{j \in J} \left(\sum_{i \in I} \alpha_{i,j} \right)^2 \quad (2)$$

where $\alpha_{i,j}$ denotes the attention weight from token i to token j and J represent the set of the decoded (target) tokens.

Each coverage metric is computed and averaged




Figure 2: Coverage evaluating on FLORES-200 devtest using PLUME 32k. Each heatmap for each part of the prompt shows the coverage scores for each layer (vertical axis) and for each head (horizontal axis) in the model.

over the FLORES-200 devtest for each head in the model and for each translation direction. To understand which part of the prompt the model is focusing on in each head we study coverage separately for different parts of the prompt: BOS, source tag, source sentence and target tag. Figure 3 shows a graphical illustration of the regions in the attention matrix that are used to compute coverage based on the part of the prompt.

In Figure 2, we show the average coverage across all translation directions for each part of the prompt, employing PLUME 32k. We note that heads within the same layer generally exhibit similar coverage patterns. Future work may investigate how these patterns arise and how they are related to the usage of Multi-Query attention⁶ (Shazeer, 2019).

We find that source tag is the part of the prompt with least coverage. However, BOS, source sentence and target tag tokens exhibit varying degrees of coverage with some coverage spikes in specific layers and heads. Interestingly, layers 5, 6, 10 and 11 show coverage uniquely for the BOS token which suggests that all attention mass is given to the BOS token, leaving the residual stream unchanged. This patterns have recently been observed in autoregressive language models and are named attention sink mechanisms (Xiao et al., 2024; Ferrando and Voita, 2024; Ferrando et al., 2024; Cancedda,

2024). For instance, Cancedda (2024) demonstrates that in Llama 2, the feed-forward blocks embed crucial information into the residual stream of the BOS token, enabling the attention sink mechanism to happen in subsequent layers. We show in appendix E the coverage heatmaps for PLUME 128k and 256k.

Source tag importance As previously pointed out, the source tag receives less attention than the other parts of the prompt. Specifically, it has




Figure 3: Illustration of the regions in the attention matrix used to compute coverage for each part of the prompt. We show the cross-attention regions between decoded tokens and the BOS, source tag, source sentence and target tag tokens in green, yellow, blue, and red, respectively.

⁶When we use Multi-Query attention with *num_kv_heads* set to 1, the keys and values are shared across all heads from a specific layer and is only the query that differs which may hinder the specialization of the heads.

	PLUME 32k	PLUME 128k	PLUME 256k
ca→xx	-1.80	-0.54	-0.83
es→xx	-0.43	0.23	-0.33
pt→xx	-8.13	-6.01	-5.54
gl→xx	-6.52	-4.18	-4.92
it→xx	-6.57	-10.79	-5.03
fr→xx	-13.16	-19.90	-17.63
de→xx	-7.54	-2.73	-6.73
en→xx	-19.83	-25.52	-20.03
eu→xx	-16.73	-11.03	-13.23
Avg.	-8.97	-8.94	-8.25

Table 3: Relative BLEU change with respect to PLUME models after ignoring the source tag. We label languages according to their BCP-47 language code (see Table 6 from Appendix A).

an average coverage of 0.56 which is 3.7 times less coverage than the target token or 18.5 times less coverage than the BOS token. This motivates our next experiments which consist of evaluating PLUME models without indicating the source language. Specifically, we replace the source tag with another BOS token to maintain the same learned positional encodings and evaluate the model’s performance on FLORES-200 devtest using BLEU. Table 3 shows the relative BLEU change with respect to the original model aggregated by language pair. Results show varying impacts across different language pairs when the source tag is omitted. For languages like English, French and Basque, the drop in BLEU scores is particularly significant. However, for other translation directions like Spanish and Catalan, the decrease in BLEU scores is negligible. This suggests that the model is more reliant on the source tag to represent certain languages, particularly those which are less related to the bridge language or those that the model has seen less during training.

Regarding the vocabulary size, the model with a 256k vocabulary shows the smallest average decrease in BLEU scores, suggesting that a larger vocabulary may improve the model’s representation of the source language.

Redundant heads Previous work on MNMT has shown that coverage is a good indicator for pruning cross attention heads in encoder-decoder architectures (Kim et al., 2021). Following Kim et al. (2021), we study whether coverage can be used to prune heads in a decoder-only architecture without

sacrificing the model’s performance. Specifically, we mask all attention heads within a specific layer that fall below a predetermined coverage threshold. We compute coverage per layer for a specific direction as follows:

$$\text{COV}_l = \phi\left(\sum_{i=1}^H \sum_{j \in \text{Pr}} \text{cov}_j(\text{head}_{l,i})\right) \quad (3)$$

$\text{Pr} = \{\text{BOS}, \text{Source tag}, \text{Source sentence}, \text{Target tag}\}$

where COV_l represents the coverage of layer l , H is the total number of attention heads in the model, and Pr is a set that contains sets of tokens for each part of the prompt. Finally, ϕ is a MinMax Scaler used to normalize the metric between 0 and 1.

We use FLORES-200 devtest to evaluate the impact of masking heads per layer based on the coverage criterion (Equation 3). Figure 4 (left) illustrates the evolution of BLEU scores as we mask heads in PLUME 32k for the Spanish to Catalan direction (supervised). The right axis indicates the number of heads that are masked. We find that up to 64 heads can be masked without degrading the model’s performance using a threshold of 0.2, representing 47.05% of the model’s total heads. In Figure 4 (right), we show the cumulative coverage for the different parts of the prompt. We observe that for a threshold of 0.2, the masked heads represent 9.05%, 2.61%, 36% and 58.9% total coverage for the BOS, source tag, source sentence and target tag tokens respectively. This indicates that the majority of the masked heads are paying attention to the target tag token and to a lesser extent to the source sentence tokens. This suggests that these heads are less critical for maintaining translation quality. Specifically, when masking these 64 heads we are only using heads from layers 5, 6, 8, 9, 10, 11, 15, and 16 which are the layers with higher coverage for the BOS, source tag and source sentence tokens (see Figure 2). Regarding the source tag, we find that even though it is the part of the prompt with the lowest coverage, it is still useful for maintaining the translation quality. This observation aligns with our previous findings from section 4.2.

In table 14 from appendix F we report the number of heads that we can mask without losing more than 2 BLEU points for different translation directions and for different vocabulary sizes. We find that for larger vocabulary sizes we can mask a higher number of heads without hurting the model’s

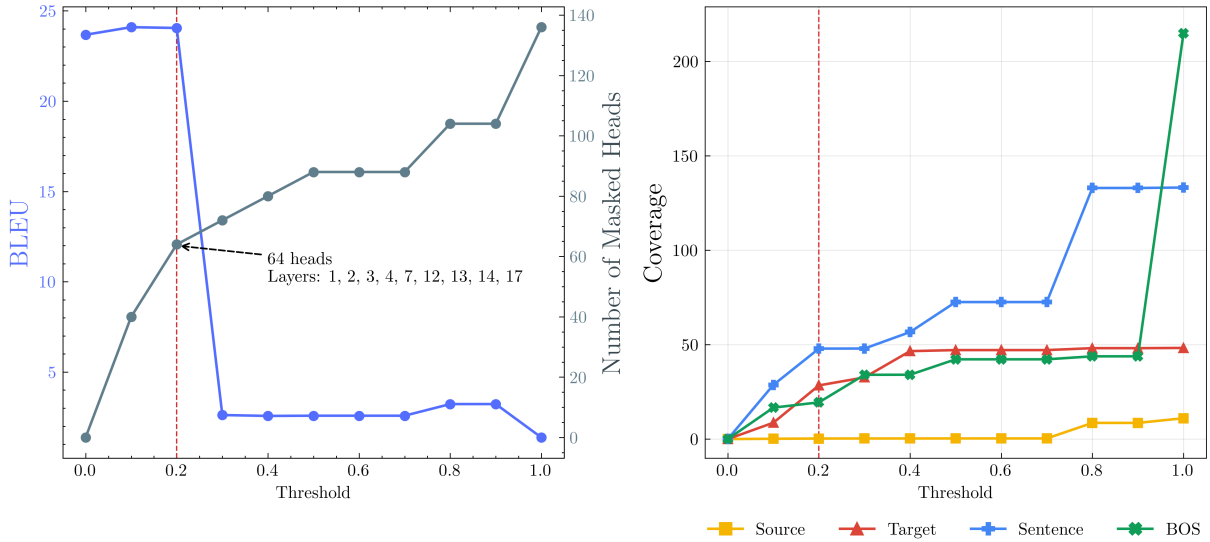


Figure 4: Impact of masking on BLEU score and number of masked heads across different coverage thresholds (left). Accumulated coverage of masked heads for source tag, target tag, source sentence, and BOS (right). Experiments are evaluated on the Spanish to Catalan direction.

performance. Specifically, with PLUME 256k, we can mask 88 heads on average, which represents 64.7% of the total number of heads. Future work may investigate how this can be used for model pruning.

4.3 Language subspaces

Subspace distances We first extract sub-word tokens output by each layer in the Transformer. Specifically, we use the first 300 sentences from FLORES-200 devtest for each source language, denoted as s . These sentences are used to create translation prompts from s to each target language ($300 * 8 = 2,400$ prompts). For each prompt, we extract the token embeddings from each layer of the model and concatenate the consecutive tokens to form H_l^s . Then, we apply singular value decomposition (SVD) on H_l^s after subtracting the mean. We calculate pairwise distances among the 9 languages using the affine subspace for each language computed by the SVD, utilizing the Riemannian metric on the space of positive definite matrices described in (Chang et al., 2022), which is both symmetric and invariant to affine transformations.

Figure 5 shows the mean distance between language subspaces in each layer. As we can see, the distance between language subspaces decreases with model depth. Initially, from the embeddings layer to layer 0 we can observe a significant decrease of approximately 5.07%, and from layer 0 to layer 1, a further reduction of 7.23%. In middle layers (layers 3 to 11), distances are relatively stable




Figure 5: Mean distance between language subspaces grouped by vocabulary size. Additional plots grouped by languages and vocabulary sizes are included in Appendix G.

and show minimal variations. This suggests that the model applies only minimal transformations to the representations along these layers. Interestingly, in layer 8 we can observe a small decrease in the distance of 0.05% which we hypothesize may be due to the model’s attention focusing more on the source token at this layer (see Figure 2). As we move to the deeper layers, the distances continue to decrease, with a significant drop of approximately 8.88% from layer 11 to layer 12, a trend that persists through layers 12 to 16. However, in the last layer, there is a notable increase in distance by approximately 23.06%. These results align with previous work on encoder-only models,




Figure 6: From left to right: Representations at the token embeddings in the embedding layer, the penultimate layer, and the last layer of PLUME 32k grouped by source language. See Appendix H for more additional plots.

468 which suggest that in intermediate layers the model
 469 representations diverge more from the embedding
 470 layer representation and from the final layer. Both
 471 the embedding layer and the final layer are highly
 472 language-sensitive (Chang et al., 2022; Libovický
 473 et al., 2020; Pires et al., 2019).

474 Regarding the vocabulary size, as shown in Fig-
 475 ure 5, we observe that for PLUME 32k the distance
 476 between embeddings are higher than PLUME 128k
 477 or PLUME 256k until layer 12, where distances be-
 478 come similar. This can be attributed to the higher
 479 vocabulary overlap between languages in PLUME
 480 32k, where each embedding represents a more di-
 481 verse concept, limiting its ability to learn language-
 482 agnostic representations.

483 **Visualization** In the previous subsection, we
 484 found that the distances between embeddings ini-
 485 tially decrease, then in the last layer, the embed-
 486 dings spread out. To understand this phenomenon,
 487 we visualize the token embeddings using Uniform
 488 Manifold Approximation and Projection (UMAP)
 489 (McInnes et al., 2018). We construct prompts from
 490 each source language to Galician. Token embed-
 491 dings per layer are concatenated to form \mathbf{P}_l^s , then
 492 we apply UMAP to reduce the dimensionality of
 493 the representations.

494 Figure 6 shows the UMAP visualizations for
 495 token embeddings in the embedding layer and the
 496 two last layers of the model coloured by source
 497 language. As we can see, token embeddings remain
 498 language-neutral as they pass through the model
 499 until the last layer, where token embeddings group
 500 by source language. This suggests that the model
 501 must align embeddings cross-linguistically until
 502 reaching the last layer where it clusters by source
 503 language. This explains the distance of the last

504 layer (see Figure 5). See Appendix H for more
 505 additional plots⁷.

506 5 Conclusions

507 This work demonstrates the successful training of
 508 an LLM-based machine translation system from
 509 scratch using only parallel data. The achieved re-
 510 sults are comparable to those of existing encod-
 511 er-decoder architectures for supervised translation
 512 tasks. We identified that larger vocabulary sizes
 513 consistently improve translation quality across
 514 zero-shot directions, suggesting the potential ben-
 515 efits of experimenting with even larger or language-
 516 specific vocabularies.

517 Further analysis revealed that different LLM lay-
 518 ers focus on distinct aspects of the prompt, par-
 519 ticularly the source language tag, which exhibits
 520 significant language variation. By leveraging this
 521 insight and employing an appropriate criterion, we
 522 achieved a performance reduction of less than 2
 523 BLEU score while removing over 47% of attention
 524 heads. Additionally, our exploration of the learned
 525 cross-lingual space demonstrates that languages
 526 get closer in the cross-lingual space as they get to
 527 deeper layers and highlight the layers with the most
 528 significant impact on the learned space.

529 This research opens doors for further investi-
 530 gation. We identified "sink heads" that primarily
 531 focus on the BOS token. Exploring their utility
 532 and relationship to the learned cross-lingual rep-
 533 resentations presents an opportunity for future work.
 534 Additionally, further research into the optimization
 535 of vocabulary size along model size could also lead
 536 to better NMT models.

⁷Additionally, we include UMAP Spherical Voronoi dia-
 grams as supplementary materials (see Appendix H.1).

6 Limitations

This study focused on understanding the capabilities of an LLM trained solely on parallel data, without aiming to achieve state-of-the-art translation quality or extensive language support. Here are some key limitations to consider when interpreting the results:

Data Scope: The experiment employed non-English centric data with a focus on Western, Latin-script languages. This approach aimed to isolate the impact of vocabulary size and overlap, but limits generalizability to languages with different scripts or historical connections. However, the inclusion of Basque, a non-Indo-European Subject-Object-Verb (SOV) language, provides valuable insights into the model’s handling of structural variations.

Scalability: The study did not explore the impact of model scale and data availability on translation across diverse languages and scripts. Further research is necessary to understand how these factors influence performance in more complex settings.

These two main aspects will be considered as future work by studying the scalability of these architectures on both model size and translation directions.

References

Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Qin Cai, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Yen-Chun Chen, Yi-Ling Chen, Parul Chopra, Xiyang Dai, Allie Del Giorno, Gustavo de Rosa, Matthew Dixon, Ronen Eldan, Victor Fragoso, Dan Iter, Mei Gao, Min Gao, Jianfeng Gao, Amit Garg, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Jamie Huynh, Mojan Javaheripi, Xin Jin, Piero Kauffmann, Nikos Karampatziakis, Dongwoo Kim, Mahoud Khademi, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Ce Liu, Mengchen Liu, Weishung Liu, Eric Lin, Zeqi Lin, Chong Luo, Piyush Madan, Matt Mazzola, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norrick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma,

Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Xin Wang, Lijuan Wang, Chunyu Wang, Yu Wang, Rachel Ward, Guanhua Wang, Philipp Witte, Haiping Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Sonali Yadav, Fan Yang, Jianwei Yang, Ziyi Yang, Yifan Yang, Donghan Yu, Lu Yuan, Chengruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *Preprint*, arXiv:2404.14219.

Mehdi Ali, Michael Fromm, Klaudia Thellmann, Richard Rutmann, Max Lübbering, Johannes Leveling, Katrin Klug, Jan Ebert, Niclas Doll, Jasper Schulze Buschhoff, Charvi Jain, Alexander Arno Weber, Lena Jurkschat, Hammam Abdelwahab, Chelsea John, Pedro Ortiz Suarez, Malte Ostendorff, Samuel Weinbach, Rafet Sifa, Stefan Kesselheim, and Nicolas Flores-Herr. 2024. [Tokenizer choice for llm training: Negligible or crucial?](#) *Preprint*, arXiv:2310.08754.

Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, 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](#). *Preprint*, arXiv:2402.17733.

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*.

Nicola Cancedda. 2024. [Spectral filters, dark signals, and attention sinks](#). *Preprint*, arXiv:2402.09221.

Tyler Chang, Zhuowen Tu, and Benjamin Bergen. 2022. [The geometry of multilingual language model representations](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 119–136, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the properties of neural machine translation: Encoder-decoder approaches](#). In *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Doha, Qatar, 25 October 2014*, pages 103–111. Association for Computational Linguistics.

Christian Federmann, Tom Kocmi, and Ying Xin. 2022. [NTREX-128 – news test references for MT evaluation of 128 languages](#). In *Proceedings of the First Workshop on Scaling Up Multilingual Evaluation*, pages 21–24, Online. Association for Computational Linguistics.

Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic](#)

647	BERT sentence embedding . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 878–891, Dublin, Ireland. Association for Computational Linguistics.	
648		
649		
650		
651		
652	Javier Ferrando, Gabriele Sarti, Arianna Bisazza, and Marta R. Costa-jussà. 2024. A primer on the inner workings of transformer-based language models . <i>Preprint</i> , arXiv:2405.00208.	
653		
654		
655		
656	Javier Ferrando and Elena Voita. 2024. Information flow routes: Automatically interpreting language models at scale . <i>Preprint</i> , arXiv:2403.00824.	
657		
658		
659	Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. Multi-way, multilingual neural machine translation with a shared attention mechanism . In <i>NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016</i> , pages 866–875. The Association for Computational Linguistics.	
660		
661		
662		
663		
664		
665		
666		
667	José A. R. Fonollosa, Noe Casas, and Marta R. Costa-jussà. 2019. Joint source-target self attention with locality constraints . <i>Preprint</i> , arXiv:1905.06596.	
668		
669		
670	Zihao Fu, Wai Lam, Qian Yu, Anthony Man-Cho So, Shengding Hu, Zhiyuan Liu, and Nigel Collier. 2023. Decoder-only or encoder-decoder? interpreting language model as a regularized encoder-decoder . <i>Preprint</i> , arXiv:2304.04052.	
671		
672		
673		
674		
675	Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimentko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray	
676		
677		
678		
679		
680		
681		
682		
683		
684		
685		
686		
687		
688		
689		
690		
691		
692		
693		
694		
695		
696		
697		
698		
699		
700		
701		
702		
703		
704		
705		
	Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. Gemma: Open models based on gemini research and technology . <i>Preprint</i> , arXiv:2403.08295.	706
		707
		708
		709
		710
		711
	Tianyu He, Xu Tan, Yingce Xia, Di He, Tao Qin, Zhibo Chen, and Tie-Yan Liu. 2018. Layer-wise coordination between encoder and decoder for neural machine translation . In <i>Advances in Neural Information Processing Systems</i> , volume 31. Curran Associates, Inc.	712
		713
		714
		715
		716
	Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Thomas Hennigan, Eric Noland, Katherine Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karén Simonyan, Erich Elsen, Oriol Vinyals, Jack Rae, and Laurent Sifre. 2022. An empirical analysis of compute-optimal large language model training . In <i>Advances in Neural Information Processing Systems</i> , volume 35, pages 30016–30030. Curran Associates, Inc.	717
		718
		719
		720
		721
		722
		723
		724
		725
		726
		727
		728
	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 . <i>Preprint</i> , arXiv:2310.06825.	729
		730
		731
		732
		733
		734
		735
		736
	Zae Myung Kim, Laurent Besacier, Vassilina Nikoulina, and Didier Schwab. 2021. Do multilingual neural machine translation models contain language pair specific attention heads? In <i>Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021</i> , pages 2832–2841, Online. Association for Computational Linguistics.	737
		738
		739
		740
		741
		742
		743
	Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization . In <i>International Conference on Learning Representations (ICLR)</i> , San Diego, CA, USA.	744
		745
		746
		747
	Jindřich Libovický, Rudolf Rosa, and Alexander Fraser. 2020. On the language neutrality of pre-trained multilingual representations . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 1663–1674, Online. Association for Computational Linguistics.	748
		749
		750
		751
		752
		753
	Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. Umap: Uniform manifold approximation and projection . <i>The Journal of Open Source Software</i> , 3(29):861.	754
		755
		756
		757
	Anthony Moi and Nicolas Patry. 2023. Huggingface’s tokenizers .	758
		759
	NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht,	760
		761
		762

763	Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic 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 . <i>Preprint</i> , arXiv:2207.04672.	
776	Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation . In <i>Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics</i> , pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.	
783	Aleksandar Petrov, Emanuele La Malfa, Philip Torr, and Adel Bibi. 2024. Language model tokenizers introduce unfairness between languages. <i>Advances in Neural Information Processing Systems</i> , 36.	
787	Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual BERT? In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 4996–5001, Florence, Italy. Association for Computational Linguistics.	
793	Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation . In <i>Proceedings of the Tenth Workshop on Statistical Machine Translation</i> , pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.	
798	Gema Ramírez-Sánchez, Jaime Zaragoza-Bernabeu, Marta Bañón, and Sergio Ortiz-Rojas. 2020. Bifixer and bicleaner: two open-source tools to clean your parallel data. In <i>Proceedings of the 22nd Annual Conference of the European Association for Machine Translation</i> , pages 291–298, Lisboa, Portugal. European Association for Machine Translation.	
805	Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022a. COMET-22: Unbabel-IST 2022 submission for the metrics shared task . In <i>Proceedings of the Seventh Conference on Machine Translation (WMT)</i> , pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.	
813	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 <i>Proceedings of the Seventh Conference on Machine Translation (WMT)</i> , pages 634–645, Abu Dhabi,	
	United Arab Emirates (Hybrid). Association for Computational Linguistics.	821 822
	Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units . In <i>Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.	823 824 825 826 827 828 829
	Noam Shazeer. 2019. Fast transformer decoding: One write-head is all you need . <i>Preprint</i> , arXiv:1911.02150.	830 831 832
	David Stap, Vlad Niculae, and Christof Monz. 2023. Viewing knowledge transfer in multilingual machine translation through a representational lens . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 14973–14987, Singapore. Association for Computational Linguistics.	833 834 835 836 837 838
	Jianlin Su, Murtadha H. M. Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding . <i>Neurocomputing</i> , 568:127063.	839 840 841 842
	Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks . In <i>Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada</i> , pages 3104–3112.	843 844 845 846 847 848
	Shaomu Tan and Christof Monz. 2023. Towards a better understanding of variations in zero-shot neural machine translation performance . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing</i> , pages 13553–13568, Singapore. Association for Computational Linguistics.	849 850 851 852 853 854
	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models . <i>Preprint</i> , arXiv:2307.09288.	855 856 857 858 859 860 861 862 863 864 865 866 867 868 869 870 871 872 873 874 875 876 877

878	Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation . In <i>Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 76–85, Berlin, Germany. Association for Computational Linguistics.	936
879		937
880		938
881		939
882		940
883		941
884		942
885	Elena Voita, Rico Sennrich, and Ivan Titov. 2019a. The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019</i> , pages 4395–4405. Association for Computational Linguistics.	
886		
887		
888		
889		
890		
891		
892		
893		
894		
895	Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019b. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned . In <i>Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers</i> , pages 5797–5808. Association for Computational Linguistics.	
896		
897		
898		
899		
900		
901		
902		
903	Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation . <i>Preprint</i> , arXiv:1609.08144.	
904		
905		
906		
907		
908		
909		
910		
911		
912		
913		
914		
915	Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. Efficient streaming language models with attention sinks . <i>Preprint</i> , arXiv:2309.17453.	
916		
917		
918		
919	Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. 2024. A paradigm shift in machine translation: Boosting translation performance of large language models . <i>Preprint</i> , arXiv:2309.11674.	
920		
921		
922		
923	Wen Yang, Chong Li, Jiajun Zhang, and Chengqing Zong. 2023. Bigtranslate: Augmenting large language models with multilingual translation capability over 100 languages . <i>Preprint</i> , arXiv:2305.18098.	
924		
925		
926		
927	Biao Zhang, Barry Haddow, and Alexandra Birch. 2023. Prompting large language model for machine translation: A case study . In <i>International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA</i> , volume 202 of <i>Proceedings of Machine Learning Research</i> , pages 41092–41110. PMLR.	
928		
929		
930		
931		
932		
933		
934	Shiyue Zhang, Vishrav Chaudhary, Naman Goyal, James Cross, Guillaume Wenzek, Mohit Bansal, and	
935		

A Dataset

943

Table 4 shows the number of sentences and number of words per language pair in the created Catalan-Centric dataset.

944

945

Pair	N sentences	N words	Dataset	Language	Id
ca _{SYN} ↔ de	187,483,456	6,847,140,698	Aina-ca-en-Parallel-Corpus	Catalan	ca
ca ↔ de	12,516,544	603,121,312	CCAligned	German	de
ca _{SYN} ↔ it	181,034,146	6,526,304,128	Covost2	English	en
ca ↔ it	18,965,862	577,243,404	DOGC	Spanish	es
ca ↔ es	171,907,026	8,252,262,032	EUBookshop	Basque	eu
ca _{SYN} ↔ pt	62,858,532	2,429,548,286	Europarl	Italian	it
ca ↔ pt	12,319,262	504,959,082	Globalvoices	Galician	gl
ca ↔ en	60,046,068	2,429,961,320	Gnome	French	fr
ca ↔ fr	37,269,716	1,114,635,790	HLPT	Portuguese	pt
ca _{SYN} ↔ eu	17,998,782	749,042,034	KDE4		
ca ↔ eu	2,091,356	61,237,122	MultiCCAligned		
ca _{SYN} ↔ gl	11,434,180	531,773,730	NLLB		
ca ↔ gl	7,713,022	263,280,596	OpenSubtitles		
Total	783,637,952	30,890,509,534	ParaCrawl		
			Tatoeba		
			TildeModel		
			Ubuntu		
			Wikimatrix		
			Wikimedia		
			XLEnt		

Table 4: Number of sentences and words for each language pair. We label languages with their BCP-47 language code. SYN means synthetic data generated on the source side for the ca-xx direction.

Table 6: List of BCP-47 language codes.

Table 5: Data sources.

B Tokenizer

946

In our experiments, we utilized the BPE algorithm (Sennrich et al., 2016) from the *Huggingface Tokenizer* library (Moi and Patry, 2023). The settings used for training the tokenizer are detailed in Table 7. Every language tag is represented by a BCP-47 tag sequence where the base subtag is a three-letter ISO 639-3 code, followed by ISO 15924 script subtags.

947

948

949

950

Hyper-Parameter	Value(s)
model_type	BPE
vocab_size	32k & 128k & 256k
nfkd_normalizer	True
lowercase_normalizer	False
pre_tokenizer	ByteLevel
add_prefix_space	False
special_tokens	<s>, </s>, <pad>, <mask>, [deu_Latn], [eng_Latn], [eus_Latn], [fra_Latn], [glg_Latn], [ita_Latn], [por_Latn], [spa_Latn], [cat_Latn]

Table 7: BPE tokenizer configuration.

We trained various tokenizers employing two distinct sampling strategies for each vocabulary size, then we evaluated them on fertility and parity (Petrov et al., 2024) metrics on FLORES-200 devtest. For a

951

952

953 given tokenizer T and a set of sentences S, fertility is determined by dividing the total number of tokens
 954 generated from S (using T) by the total number of words in S. Parity is defined as achieving a balanced
 955 tokenization ratio between two languages. Specifically, a tokenizer T achieves parity for language A
 956 with respect to language B if the ratio $\frac{|T(s_A)|}{|T(s_B)|} \approx 1$, where s_A and s_B denote the sets of all sentences for
 957 languages A and B, respectively.

958 We experimented with both unigram and BPE implementations from the *Huggingface Tokenizer* library.
 959 We tested two sampling strategies: one involving the sampling of 1 million sentences from all languages,
 960 and another involving the equal sampling of 1 million sentences from Romance languages, with an
 961 oversampling of 3 million sentences for English, Basque, and German. Figure 7 presents the fertility
 962 metrics on English, Basque, and German. Given the results, we decided to use the BPE algorithm with the
 963 oversampling strategy for our final experiments. We also report obtained parity metrics by vocabulary
 964 size in figure 8 and average fertility (average of fertility per each language) per vocabulary size as well as
 965 the number of tokens in the dataset in Table 8⁸.




Figure 7: From left to right: fertility evaluated on Basque, English and German. Fertility is in the vertical axis, and vocabulary size is in the horizontal axis.




Figure 8: Parity for the different vocabulary sizes.

Vocabulary size	Average Fertility	Number of tokens
PLUME 32k	1.77	54.7B
PLUME 128k	1.52	46.8B
PLUME 256k	1.44	44.6B

Table 8: Fertility and number of tokens in the dataset grouped by vocabulary size.

⁸We compute the number of tokens as Average Fertility * Number of words in the dataset. The number of words is 30,890,509,534.

C Model Architecture

966

Table 9 summarizes the architecture used for PLUME models.

967

Hyper-Parameter	Value
Hidden Dimension	2048
Layers	18
Intermediate Size (in MLPs)	16384
Attention-Heads	8
Head size	256
Num KV Heads	1
Max Seq Length	2048
Position Embeddings	Rotary
Rope Theta	10000
Precision	float-32
RMSNorm ϵ	1e-06

Table 9: Model architecture.

D Training

968

For training, the learning rate is warmed up from 1×10^{-7} to a maximum of 3×10^{-4} over the first 2000 steps. We apply a weight decay of 0.1 and a gradient clipping of 1.0. During training, we set an effective batch size of 81,920 tokens per gradient step distributed over 40 NVIDIA H100-64GB GPUs. We use DeepSpeed with full float32 training.

969

970

971

972

Hyper-Parameter	
Batch size	40
Number of Epochs	1
Optimizer	Adam
Adam- β_1	0.9
Adam- β_2	0.999
Adam- ϵ	1e-08
Learning rate	3e-04
LR Scheduler	Linear
Warmup Steps	2000

Table 10: Model training hyper-parameters

Num examples	26,301,993
Num tokens = Num examples * 2048 (considering pad tokens)	53,866,481,664
Num Epochs	1
Instantaneous batch size per device	1
Total train batch size (w. parallel, distributed & accumulation)	40
Gradient Accumulation steps	1
Total optimization steps	657,550
Number of trainable parameters	2,047,420,416

Table 11: Training and performance information for PLUME 32k.

Num examples	23,093,719
Num tokens = Num examples * 2048 (considering pad tokens)	47,295,936,512
Num Epochs	1
Instantaneous batch size per device	1
Total train batch size (w. parallel, distributed & accumulation)	40
Gradient Accumulation steps	1
Total optimization steps	577,343
Number of trainable parameters	2,244,028,416

Table 12: Training and performance information for PLUME 128k.

Num examples	22,213,825
Num tokens = Num examples * 2048 (considering pad tokens)	45,493,913,600
Num Epochs	1
Instantaneous batch size per device	1
Total train batch size (w. parallel, distributed & accumulation)	40
Gradient Accumulation steps	1
Total optimization steps	555,346
Number of trainable parameters	2,506,172,416

Table 13: Training and performance information for PLUME 256k.




Figure 9: Training loss.

E Coverage metrics

We show in Figure 11 and Figure 12 the coverage heatmaps for PLUME 32k, 128k and 256k respectively. In Figure 13 we show the average coverage per layer for the different vocabulary sizes. We notice that PLUME 32k, 128k and 256k exhibit a similar coverage pattern across layers.

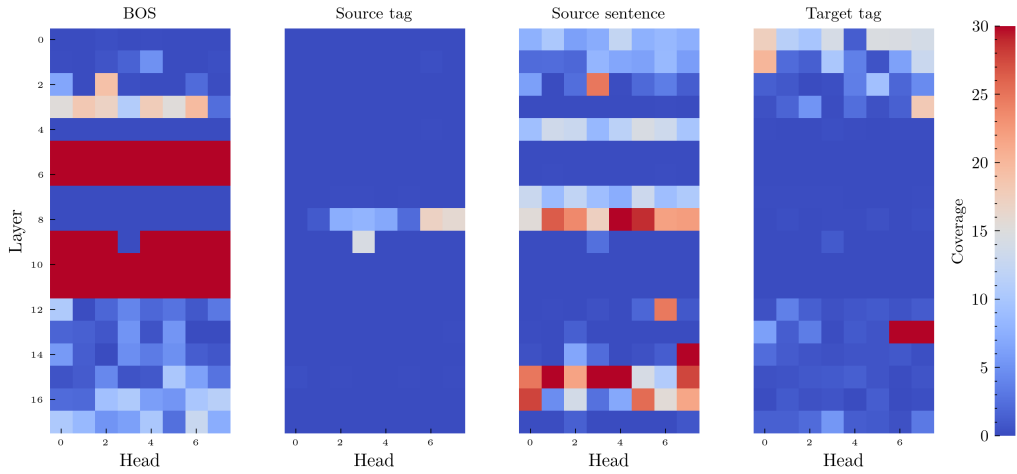


Figure 10: Coverage evaluating on FLORES-200 devtest using PLUME 32k. Each heatmap for each part of the prompt shows the coverage scores for each layer (vertical axis) and for each head (horizontal axis) in the model.




Figure 11: Coverage evaluating on FLORES-200 devtest using PLUME 128k. Each heatmap for each part of the prompt shows the coverage scores for each layer (vertical axis) and for each head (horizontal axis) in the model.




Figure 12: Coverage evaluating on FLORES-200 devtest using PLUME 256k. Each heatmap for each part of the prompt shows the coverage scores for each layer (vertical axis) and for each head (horizontal axis) in the model.

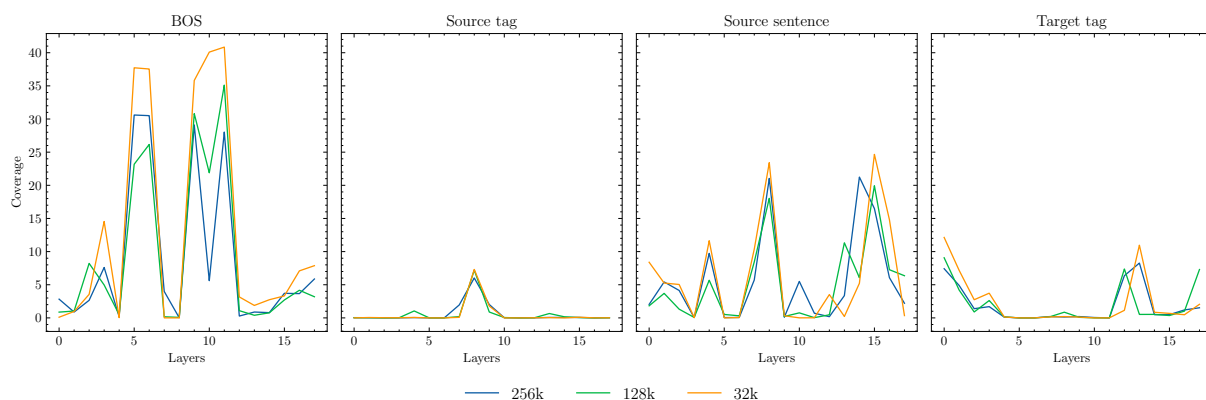


Figure 13: Average coverage per layer for each part of the prompt across various vocabulary sizes.

E.1 Attention matrices

An attention sink mechanism occurs when all the attention mass is given to some special tokens. We visualize the attention matrices for the first head of layer 9 and layer 17 (last layer) in Figure 14. We observe that in layer 9, the model is giving all the attention mass to the BOS token⁹ which allows the model to keep the residual stream of the network unchanged.




Figure 14: Attention weights for head 1 in layer 9 (left) and head 1 in layer 17 (right).

F Redundant Heads

Table 14 shows the number of heads that can be masked in various directions for different vocabulary sizes without losing more than 2 BLEU points. We observe that as vocabulary size increases, we can mask more heads. Specifically, on average we can mask 41.56%, 49.41% and 64.7% of the model’s heads for PLUME 32k, PLUME 128k and PLUME 256k respectively.

⁹There is a special token created by Huggingface BPE implementation, which is positioned between the BOS and the source tag tokens. We consider this special token as part of the BOS token.

	PLUME 32k	PLUME 128k	PLUME 256k
de→ca	64	64	88
de→en	32	72	88
de→pt	64	64	88
es→ca	64	104	88
es→en	64	72	88
es→pt	64	104	88
fr→ca	64	64	88
fr→en	24	72	88
fr→pt	64	0	88
gl→ca	64	104	88
gl→en	24	72	88
gl→pt	64	64	88
it→ca	64	80	88
it→en	64	72	88
it→pt	64	0	88
Avg.	56.53	67.2	88

Table 14: Number of masked heads across different language pairs and vocabulary sizes such that BLEU drop is less than 2 points.

G Subspace distances

We show in Figure 15 the distances between language subspaces computed using the Riemannian metric on the space of positive definite matrices as detailed in (Chang et al., 2022) grouped by language and for each vocabulary size. We observe that for all the vocabulary sizes, Basque’s subspace is further from the rest of the languages subspaces which could explain why model’s performance on Basque is lower compared to other languages.




Figure 15: Mean distance between language subspaces grouped by languages and vocabulary sizes.

H UMAP Plots

Below we show the token representations¹⁰ using Uniform Manifold Approximation and Projection (UMAP) (McInnes et al., 2018) for all the layers in PLUME 32k, 128k and 256k. We employ the cosine distance and we set the number of neighbours to 8 for computing UMAP’s embeddings.

¹⁰We use the first sentence from FLORES-200 devtest in each source language to construct the prompts: "We now have 4-month-old mice that are non-diabetic that used to be diabetic," he added.

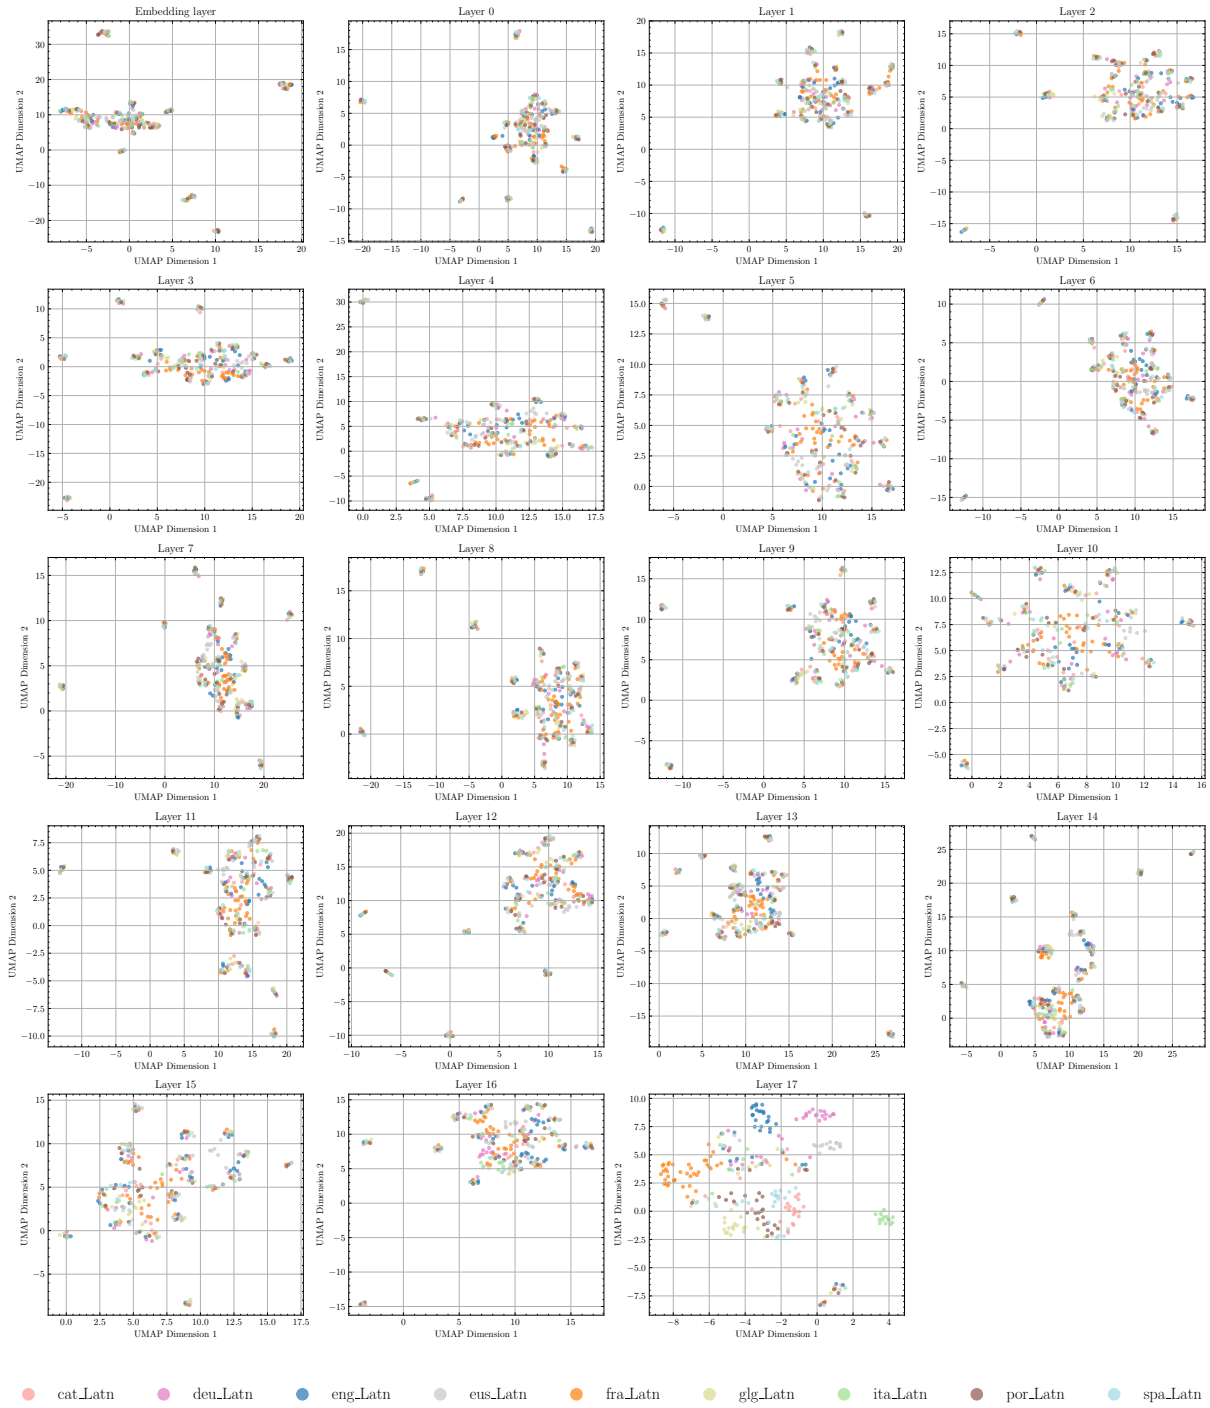


Figure 16: UMAP representations at the token embeddings in each layer grouped by source language using PLUME 32k.

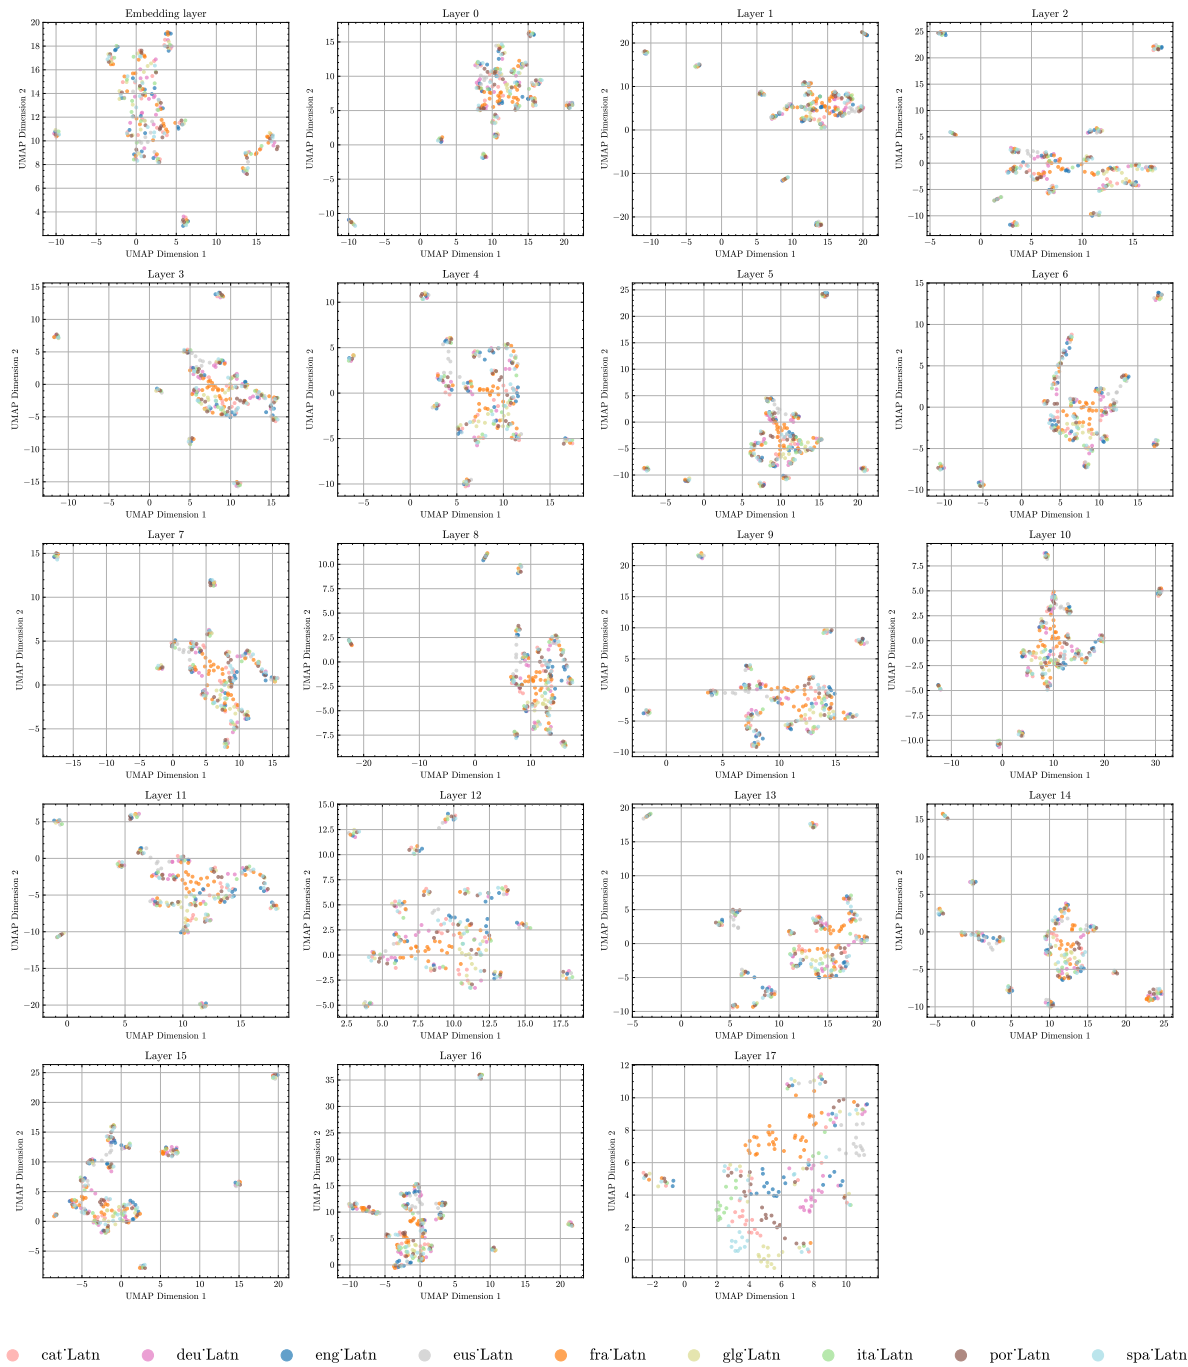


Figure 17: UMAP representations at the token embeddings in each layer grouped by source language using PLUME 128k.




Figure 18: UMAP representations at the token embeddings in each layer grouped by source language using PLUME 256k.

997 H.1 Spherical Voronoi diagrams

998 To better visualize high-dimensional token embeddings in PLUME models, we used spherical voronoi
 999 diagrams. Specifically, we reduced the embeddings to a 2D space, optimizing for cosine similarity using
 1000 UMAP. Then, the 2D UMAP embeddings were projected onto a unit sphere. Specifically, each 2D point
 1001 (x, y) was mapped to 3D coordinates (X, Y, Z) as follows:

$$1002 \quad X = \sin(x) \cdot \cos(y) \quad Y = \sin(x) \cdot \sin(y) \quad Z = \cos(x) \quad (4)$$

1003 Then, for each language, we calculated the centroid of its corresponding tokens on the sphere and using
 1004 these centroids, we computed Voronoi regions (where each region contains all the closest points to a

specific centroid). We add as supplementary material the spherical voronoi diagrams for each layer in PLUME 32k.

I Detailed results

We report in the following tables the results of PLUME models for each translation direction. We also provide comparisons for TOWERBASE 7B (Alves et al., 2024) in those directions that PLUME and TOWERBASE 7B share, as well as comparisons with NLLB 3.3B (NLLB Team et al., 2022).

Table 15: Results for ca→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
ca-de	BSC Bilinguals	33.30	61.12	0.85	0.84	25.04	55.00	0.83	0.83
	NLLB 3.3B	31.19	58.41	0.85	0.84	21.72	53.41	0.81	0.82
	PLUME 128k	28.00	57.53	0.83	0.82	21.98	53.36	0.80	0.81
	PLUME 256k	28.55	57.63	0.83	0.82	21.39	52.72	0.80	0.81
	PLUME 32k	27.81	57.00	0.83	0.82	27.79	56.66	0.83	0.84
ca-en	BSC Bilinguals	46.29	70.44	0.88	0.86	41.20	66.57	0.87	0.86
	NLLB 3.3B	49.65	71.68	0.89	0.86	33.22	62.82	0.85	0.85
	PLUME 128k	42.91	68.69	0.88	0.86	33.73	63.07	0.85	0.85
	PLUME 256k	42.47	68.47	0.88	0.85	32.82	62.14	0.85	0.84
	PLUME 32k	41.92	68.15	0.87	0.85	37.61	64.98	0.87	0.85
ca-es	BSC Bilinguals	24.70	53.42	0.86	0.86	36.89	61.83	0.86	0.85
	NLLB 3.3B	25.62	53.73	0.86	0.86	35.44	61.27	0.86	0.85
	PLUME 128k	24.66	53.44	0.86	0.86	35.66	61.23	0.86	0.85
	PLUME 256k	24.59	53.37	0.86	0.85	35.70	61.24	0.86	0.85
	PLUME 32k	24.50	53.37	0.86	0.86	35.97	61.40	0.86	0.85
ca-eu	BSC Bilinguals	18.26	57.03	0.86	0.81	9.83	46.47	0.80	0.74
	NLLB 3.3B	13.13	50.47	0.83	0.75	12.40	49.99	0.82	0.78
	PLUME 128k	14.88	53.41	0.84	0.79	12.09	49.96	0.82	0.78
	PLUME 256k	14.97	53.75	0.84	0.78	12.17	49.58	0.81	0.77
	PLUME 32k	14.38	53.29	0.84	0.78	14.08	52.70	0.84	0.81
ca-fr	BSC Bilinguals	38.25	63.23	0.85	0.84	27.60	56.73	0.84	0.85
	NLLB 3.3B	39.89	64.05	0.86	0.85	25.20	54.13	0.81	0.82
	PLUME 128k	35.46	61.08	0.84	0.83	25.48	54.16	0.81	0.82
	PLUME 256k	35.72	61.18	0.84	0.83	24.94	53.76	0.81	0.82
	PLUME 32k	34.32	60.68	0.83	0.82	27.71	55.53	0.82	0.83
ca-gl	BSC Bilinguals	31.96	59.66	0.87	0.84	34.07	60.52	0.86	0.84
	NLLB 3.3B	32.78	59.25	0.87	0.85	33.23	60.22	0.86	0.84
	PLUME 128k	32.22	59.73	0.87	0.84	33.37	60.24	0.86	0.83
	PLUME 256k	32.07	59.51	0.87	0.84	33.23	60.27	0.86	0.84
	PLUME 32k	32.21	59.73	0.87	0.85	32.59	59.76	0.85	0.82
ca-it	BSC Bilinguals	26.92	56.55	0.87	0.85	29.46	58.00	0.87	0.85
	NLLB 3.3B	26.38	55.66	0.88	0.86	27.91	57.43	0.86	0.84
	PLUME 128k	25.77	55.78	0.87	0.85	28.11	57.62	0.86	0.84
	PLUME 256k	25.76	55.94	0.87	0.85	27.80	57.33	0.85	0.84
	PLUME 32k	25.45	55.51	0.87	0.85	29.07	57.95	0.86	0.84
ca-pt	BSC Bilinguals	37.18	62.73	0.88	0.84	31.46	57.67	0.86	0.84
	NLLB 3.3B	36.68	61.97	0.88	0.85	27.79	55.97	0.85	0.83
	PLUME 128k	36.27	62.12	0.88	0.84	28.50	56.29	0.85	0.83
	PLUME 256k	35.76	61.88	0.88	0.84	27.92	55.91	0.85	0.83
	PLUME 32k	35.81	61.67	0.88	0.84	28.19	56.17	0.85	0.83

Table 16: Results for de→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
de-ca	BSC Bilinguals	30.15	57.65	0.83	0.82	28.24	55.02	0.83	0.84
	NLLB 3.3B	31.45	57.99	0.86	0.85	28.34	55.03	0.82	0.82
	PLUME 128k	32.23	59.02	0.85	0.83	28.13	54.66	0.82	0.82
	PLUME 256k	31.76	58.73	0.85	0.83	27.94	54.58	0.82	0.81
	PLUME 32k	31.76	58.56	0.85	0.83	24.49	53.60	0.78	0.80
de-en	NLLB 3.3B	46.02	69.30	0.90	0.85	41.01	66.16	0.88	0.84
	TOWERBASE 7B	43.69	68.56	0.89	0.84	41.01	66.16	0.88	0.84
	PLUME 128k	36.17	63.49	0.86	0.82	29.73	59.26	0.84	0.81
	PLUME 256k	36.99	64.04	0.87	0.83	29.80	59.39	0.84	0.81
	PLUME 32k	34.12	62.13	0.86	0.81	28.73	58.11	0.83	0.80
de-es	NLLB 3.3B	23.86	51.39	0.84	0.86	31.13	57.36	0.84	0.85
	TOWERBASE 7B	21.66	50.94	0.83	0.85	31.13	57.36	0.84	0.85
	PLUME 128k	22.00	50.41	0.82	0.83	28.41	54.92	0.81	0.82
	PLUME 256k	22.35	50.80	0.82	0.83	28.76	54.89	0.81	0.82
	PLUME 32k	20.90	49.74	0.82	0.82	27.83	54.18	0.81	0.81
de-eu	NLLB 3.3B	9.83	45.23	0.78	0.71	7.83	41.70	0.76	0.69
	PLUME 128k	9.91	46.23	0.78	0.73	8.18	42.65	0.75	0.72
	PLUME 256k	11.48	47.52	0.79	0.74	8.93	43.59	0.76	0.73
	PLUME 32k	10.77	46.22	0.77	0.72	8.46	42.39	0.74	0.71
de-fr	NLLB 3.3B	37.62	62.60	0.86	0.85	28.06	56.03	0.83	0.85
	TOWERBASE 7B	34.84	61.23	0.85	0.85	28.06	56.03	0.83	0.85
	PLUME 128k	28.50	56.32	0.80	0.80	20.26	49.16	0.77	0.78
	PLUME 256k	29.01	56.15	0.80	0.79	20.84	49.13	0.77	0.78
	PLUME 32k	27.13	54.89	0.79	0.78	20.37	48.30	0.75	0.76
de-gl	NLLB 3.3B	28.87	55.70	0.85	0.85	29.17	56.21	0.84	0.84
	PLUME 128k	26.01	54.15	0.83	0.83	24.55	52.87	0.81	0.81
	PLUME 256k	25.20	53.46	0.83	0.82	24.87	52.86	0.81	0.81
	PLUME 32k	25.31	53.11	0.82	0.82	24.11	51.92	0.80	0.80
de-it	NLLB 3.3B	25.88	54.95	0.87	0.86	27.84	56.12	0.86	0.85
	TOWERBASE 7B	24.73	54.26	0.86	0.85	27.84	56.12	0.86	0.85
	PLUME 128k	22.47	52.44	0.84	0.83	22.77	52.04	0.82	0.82
	PLUME 256k	22.74	52.34	0.85	0.83	23.12	52.16	0.82	0.82
	PLUME 32k	21.36	51.19	0.84	0.82	22.39	51.53	0.81	0.81
de-pt	NLLB 3.3B	33.42	59.32	0.87	0.85	29.42	55.97	0.85	0.85
	TOWERBASE 7B	30.94	58.48	0.86	0.85	29.42	55.97	0.85	0.85
	PLUME 128k	30.02	57.17	0.85	0.83	24.09	51.90	0.82	0.82
	PLUME 256k	30.36	57.46	0.85	0.83	24.06	51.90	0.82	0.82
	PLUME 32k	29.19	55.98	0.84	0.81	23.00	51.09	0.80	0.80

Table 17: Results for en→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
en-ca	BSC Bilinguals	44.05	67.95	0.88	0.85	37.49	62.38	0.85	0.83
	NLLB 3.3B	42.33	65.97	0.88	0.85	35.80	61.29	0.83	0.81
	PLUME 128k	42.29	66.44	0.87	0.84	35.95	61.30	0.83	0.81
	PLUME 256k	42.64	66.59	0.87	0.84	35.05	60.72	0.82	0.80
	PLUME 32k	42.32	66.39	0.86	0.84	37.93	63.19	0.84	0.82
en-de	NLLB 3.3B	39.88	65.14	0.88	0.84	32.46	60.93	0.85	0.84
	TOWERBASE 7B	37.53	64.47	0.87	0.84	32.46	60.93	0.85	0.84
	PLUME 128k	31.27	59.30	0.82	0.80	24.31	54.33	0.78	0.77
	PLUME 256k	31.81	60.17	0.83	0.81	24.94	55.13	0.79	0.78
	PLUME 32k	29.86	58.22	0.82	0.79	23.46	53.42	0.77	0.75
en-es	NLLB 3.3B	28.14	55.85	0.86	0.86	39.33	63.79	0.85	0.84
	TOWERBASE 7B	26.38	55.02	0.86	0.86	39.33	63.79	0.85	0.84
	PLUME 128k	24.34	53.01	0.83	0.84	35.62	60.75	0.81	0.80
	PLUME 256k	25.00	53.43	0.84	0.84	36.42	61.36	0.82	0.81
	PLUME 32k	23.47	52.61	0.83	0.83	34.86	60.10	0.81	0.79
en-eu	NLLB 3.3B	15.71	53.25	0.85	0.82	11.62	47.74	0.81	0.79
	PLUME 128k	13.02	48.69	0.81	0.78	10.51	44.21	0.76	0.75
	PLUME 256k	12.95	50.05	0.81	0.79	10.96	45.41	0.77	0.75
	PLUME 32k	13.03	48.89	0.80	0.78	10.73	44.79	0.75	0.74
en-fr	NLLB 3.3B	50.90	71.70	0.88	0.87	34.77	61.69	0.84	0.85
	TOWERBASE 7B	49.28	70.83	0.88	0.87	34.77	61.69	0.84	0.85
	PLUME 128k	36.49	62.25	0.82	0.82	26.36	54.27	0.77	0.79
	PLUME 256k	38.27	63.03	0.83	0.83	27.20	54.95	0.77	0.79
	PLUME 32k	36.11	61.92	0.81	0.81	26.36	54.15	0.76	0.78
en-gl	NLLB 3.3B	35.98	61.55	0.87	0.85	39.01	63.75	0.85	0.83
	PLUME 128k	32.26	59.64	0.85	0.83	33.28	59.53	0.81	0.79
	PLUME 256k	32.61	59.66	0.85	0.83	33.13	59.59	0.81	0.79
	PLUME 32k	31.16	58.92	0.84	0.82	31.88	58.48	0.80	0.77
en-it	NLLB 3.3B	30.63	59.52	0.88	0.87	37.68	63.84	0.87	0.85
	TOWERBASE 7B	29.64	59.13	0.88	0.87	37.68	63.84	0.87	0.85
	PLUME 128k	25.58	55.15	0.84	0.84	28.84	57.37	0.82	0.81
	PLUME 256k	25.64	55.75	0.85	0.85	30.73	58.42	0.82	0.81
	PLUME 32k	24.51	54.69	0.84	0.84	29.55	57.32	0.81	0.80
en-pt	NLLB 3.3B	49.45	70.54	0.90	0.85	37.37	62.46	0.87	0.84
	TOWERBASE 7B	49.67	71.36	0.90	0.85	37.37	62.46	0.87	0.84
	PLUME 128k	40.94	65.75	0.87	0.83	30.59	57.41	0.82	0.79
	PLUME 256k	42.62	66.47	0.87	0.83	31.27	57.81	0.82	0.79
	PLUME 32k	40.57	65.13	0.86	0.82	30.13	56.87	0.81	0.78

Table 18: Results for es→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-K _{IWI}	BLEU	CHRf	COMET	COMET-K _{IWI}
es-ca	BSC Bilinguals	23.34	53.98	0.86	0.84	34.47	60.52	0.86	0.84
	NLLB 3.3B	25.70	55.24	0.86	0.84	33.16	60.59	0.86	0.83
	PLUME 128k	23.43	54.22	0.86	0.84	33.41	60.49	0.86	0.83
	PLUME 256k	23.42	54.20	0.86	0.84	33.23	60.60	0.86	0.83
	PLUME 32k	23.55	54.30	0.86	0.84	34.14	60.73	0.86	0.83
es-de	NLLB 3.3B	22.88	53.27	0.84	0.84	24.63	55.15	0.83	0.84
	TOWERBASE 7B	18.86	51.44	0.82	0.84	24.63	55.15	0.83	0.84
	PLUME 128k	17.69	50.73	0.80	0.81	19.90	52.08	0.79	0.81
	PLUME 256k	18.06	51.26	0.81	0.82	20.41	52.30	0.80	0.81
	PLUME 32k	17.63	50.19	0.80	0.80	19.47	51.49	0.78	0.80
es-en	NLLB 3.3B	32.93	61.52	0.88	0.86	41.88	67.47	0.88	0.86
	TOWERBASE 7B	30.47	60.37	0.87	0.86	41.88	67.47	0.88	0.86
	PLUME 128k	24.74	56.76	0.85	0.85	31.64	62.07	0.85	0.84
	PLUME 256k	24.91	57.16	0.85	0.85	31.53	62.24	0.85	0.84
	PLUME 32k	23.79	56.29	0.84	0.85	31.05	61.38	0.85	0.84
es-eu	NLLB 3.3B	11.31	49.93	0.84	0.81	11.13	47.56	0.81	0.77
	PLUME 128k	10.39	49.12	0.82	0.81	11.45	48.54	0.81	0.79
	PLUME 256k	11.22	49.59	0.83	0.81	11.29	48.92	0.81	0.79
	PLUME 32k	11.26	49.16	0.82	0.79	11.31	47.79	0.80	0.78
es-fr	NLLB 3.3B	29.97	58.18	0.85	0.86	27.92	56.77	0.84	0.85
	TOWERBASE 7B	25.16	55.84	0.84	0.85	27.92	56.77	0.84	0.85
	PLUME 128k	21.91	52.76	0.81	0.82	23.99	52.86	0.80	0.81
	PLUME 256k	22.15	52.87	0.81	0.82	23.85	52.99	0.80	0.81
	PLUME 32k	21.96	52.78	0.81	0.82	24.39	53.10	0.79	0.81
es-gl	NLLB 3.3B	24.64	53.77	0.87	0.84	34.92	61.24	0.87	0.83
	PLUME 128k	21.47	52.69	0.87	0.84	33.34	60.71	0.86	0.83
	PLUME 256k	21.59	52.54	0.86	0.84	33.63	60.81	0.86	0.82
	PLUME 32k	21.29	52.51	0.86	0.84	33.08	60.63	0.86	0.83
es-it	NLLB 3.3B	22.77	52.86	0.87	0.86	29.60	58.19	0.87	0.85
	TOWERBASE 7B	19.95	51.18	0.86	0.86	29.60	58.19	0.87	0.85
	PLUME 128k	18.76	50.27	0.85	0.85	25.08	55.31	0.84	0.83
	PLUME 256k	18.86	50.53	0.85	0.84	25.42	55.57	0.85	0.84
	PLUME 32k	19.29	50.45	0.85	0.84	25.14	55.55	0.84	0.83
es-pt	NLLB 3.3B	26.18	55.23	0.87	0.85	32.30	58.24	0.87	0.84
	TOWERBASE 7B	23.11	53.87	0.87	0.85	32.30	58.24	0.87	0.84
	PLUME 128k	21.16	52.25	0.86	0.84	25.82	54.84	0.85	0.83
	PLUME 256k	21.84	52.70	0.86	0.84	27.27	55.53	0.85	0.83
	PLUME 32k	21.65	52.74	0.86	0.84	27.00	55.35	0.85	0.83

Table 19: Results for eu→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
eu-ca	BSC Bilinguals	26.18	54.14	0.85	0.82	24.56	51.56	0.83	0.81
	NLLB 3.3B	26.70	53.97	0.86	0.82	22.29	49.79	0.81	0.79
	PLUME 128k	24.33	51.85	0.84	0.80	21.70	49.48	0.81	0.78
	PLUME 256k	24.02	51.67	0.84	0.80	20.19	48.69	0.80	0.77
	PLUME 32k	22.92	50.69	0.83	0.79	24.29	51.84	0.82	0.81
eu-de	NLLB 3.3B	22.71	51.75	0.83	0.80	18.96	48.84	0.81	0.79
	PLUME 128k	13.64	44.72	0.76	0.72	11.38	41.74	0.73	0.72
	PLUME 256k	13.58	44.77	0.76	0.72	10.74	41.78	0.73	0.72
	PLUME 32k	10.62	40.74	0.72	0.69	9.30	38.93	0.69	0.69
eu-en	NLLB 3.3B	33.44	60.57	0.87	0.86	29.59	57.37	0.85	0.85
	PLUME 128k	21.49	51.65	0.82	0.81	16.70	48.58	0.79	0.80
	PLUME 256k	22.12	52.31	0.82	0.82	16.41	48.54	0.79	0.80
	PLUME 32k	17.52	48.60	0.79	0.78	13.84	45.54	0.77	0.77
eu-es	NLLB 3.3B	20.50	48.29	0.84	0.84	27.50	53.84	0.84	0.83
	PLUME 128k	17.74	45.98	0.81	0.81	20.71	48.75	0.79	0.79
	PLUME 256k	17.94	45.41	0.81	0.81	20.58	48.54	0.79	0.79
	PLUME 32k	15.61	43.47	0.79	0.79	18.76	47.03	0.78	0.78
eu-fr	NLLB 3.3B	29.05	56.00	0.84	0.83	22.63	50.58	0.81	0.82
	PLUME 128k	18.58	46.77	0.75	0.75	14.90	42.94	0.73	0.73
	PLUME 256k	18.39	46.08	0.75	0.74	14.73	42.58	0.72	0.72
	PLUME 32k	15.77	44.00	0.71	0.71	12.58	40.59	0.69	0.70
eu-gl	NLLB 3.3B	25.16	52.52	0.86	0.83	24.18	52.15	0.83	0.82
	PLUME 128k	19.24	47.58	0.82	0.78	18.04	46.91	0.79	0.77
	PLUME 256k	18.53	46.92	0.81	0.78	18.23	46.74	0.79	0.76
	PLUME 32k	15.91	45.11	0.79	0.75	16.13	44.99	0.77	0.75
eu-it	NLLB 3.3B	21.27	51.07	0.86	0.84	22.45	51.13	0.84	0.83
	PLUME 128k	16.39	45.65	0.81	0.80	16.82	46.45	0.79	0.79
	PLUME 256k	16.46	45.76	0.81	0.80	15.96	46.05	0.79	0.78
	PLUME 32k	14.01	43.52	0.79	0.77	14.34	44.19	0.77	0.76
eu-pt	NLLB 3.3B	27.79	54.65	0.86	0.84	23.93	50.72	0.83	0.82
	PLUME 128k	20.12	48.58	0.82	0.80	16.11	44.79	0.79	0.78
	PLUME 256k	20.89	48.87	0.81	0.80	16.80	45.27	0.79	0.78
	PLUME 32k	17.64	46.34	0.79	0.77	14.05	42.96	0.76	0.76

Table 20: Results for fr→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-K _{IWI}	BLEU	CHRf	COMET	COMET-K _{IWI}
fr-ca	BSC Bilinguals	34.44	60.10	0.86	0.83	29.22	55.76	0.84	0.83
	NLLB 3.3B	34.00	59.82	0.87	0.84	27.30	54.40	0.83	0.82
	PLUME 128k	34.35	60.24	0.86	0.83	27.57	54.40	0.83	0.81
	PLUME 256k	33.63	59.83	0.86	0.83	27.00	54.18	0.83	0.81
	PLUME 32k	34.28	60.16	0.86	0.83	27.03	54.04	0.83	0.81
fr-de	NLLB 3.3B	29.96	57.73	0.85	0.84	23.82	53.55	0.83	0.84
	TOWERBASE 7B	25.48	56.02	0.82	0.84	23.82	53.55	0.83	0.84
	PLUME 128k	24.63	54.96	0.81	0.80	19.07	49.59	0.78	0.78
	PLUME 256k	23.85	54.54	0.82	0.80	18.18	49.18	0.78	0.78
	PLUME 32k	22.45	53.56	0.81	0.78	18.35	48.80	0.77	0.77
fr-en	NLLB 3.3B	48.38	70.72	0.90	0.86	40.30	64.78	0.87	0.86
	TOWERBASE 7B	45.48	69.54	0.89	0.86	40.30	64.78	0.87	0.86
	PLUME 128k	37.37	64.47	0.87	0.85	28.95	58.15	0.84	0.84
	PLUME 256k	37.74	64.80	0.87	0.85	29.11	58.37	0.84	0.84
	PLUME 32k	34.87	63.11	0.86	0.84	28.36	57.38	0.83	0.83
fr-es	NLLB 3.3B	24.45	52.39	0.86	0.86	32.28	57.85	0.85	0.85
	TOWERBASE 7B	22.02	51.42	0.84	0.85	32.28	57.85	0.85	0.85
	PLUME 128k	21.65	50.63	0.84	0.84	27.18	54.18	0.82	0.83
	PLUME 256k	21.80	50.74	0.84	0.84	27.30	54.22	0.82	0.83
	PLUME 32k	21.60	50.66	0.84	0.84	27.23	54.00	0.82	0.82
fr-eu	NLLB 3.3B	10.73	46.16	0.80	0.73	7.79	41.10	0.76	0.69
	PLUME 128k	10.79	48.17	0.80	0.76	9.32	44.51	0.78	0.75
	PLUME 256k	11.78	48.71	0.80	0.77	9.43	44.37	0.78	0.75
	PLUME 32k	11.59	48.08	0.79	0.75	8.65	43.30	0.76	0.72
fr-gl	NLLB 3.3B	30.59	57.45	0.86	0.85	29.61	56.42	0.85	0.84
	PLUME 128k	27.95	55.92	0.85	0.84	24.65	52.84	0.81	0.81
	PLUME 256k	28.49	55.94	0.85	0.84	24.57	52.94	0.82	0.81
	PLUME 32k	27.69	55.65	0.85	0.83	24.11	52.42	0.81	0.81
fr-it	NLLB 3.3B	27.06	56.27	0.88	0.86	28.22	56.47	0.86	0.86
	TOWERBASE 7B	25.14	55.00	0.87	0.86	28.22	56.47	0.86	0.86
	PLUME 128k	24.45	53.92	0.86	0.84	24.25	53.18	0.84	0.83
	PLUME 256k	24.27	53.92	0.86	0.84	24.45	53.22	0.84	0.83
	PLUME 32k	23.98	53.72	0.86	0.84	23.84	53.05	0.83	0.82
fr-pt	NLLB 3.3B	36.18	61.28	0.88	0.85	29.11	55.64	0.85	0.84
	TOWERBASE 7B	33.03	60.10	0.87	0.85	29.11	55.64	0.85	0.84
	PLUME 128k	32.15	59.00	0.86	0.83	24.59	52.51	0.83	0.82
	PLUME 256k	32.86	59.22	0.86	0.83	24.85	52.21	0.82	0.81
	PLUME 32k	31.72	58.70	0.86	0.82	24.33	52.19	0.82	0.81

Table 21: Results for it→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
it-ca	BSC Bilinguals	27.68	56.63	0.86	0.84	31.87	57.96	0.86	0.84
	NLLB 3.3B	27.77	56.56	0.87	0.86	31.18	57.64	0.85	0.83
	PLUME 128k	27.92	57.34	0.87	0.85	31.00	57.62	0.85	0.83
	PLUME 256k	27.86	57.25	0.87	0.85	30.69	57.35	0.85	0.83
	PLUME 32k	27.48	57.19	0.86	0.85	30.67	57.08	0.84	0.82
it-de	NLLB 3.3B	25.33	55.23	0.85	0.86	26.76	56.82	0.84	0.85
	TOWERBASE 7B	18.14	49.13	0.82	0.86	26.76	56.82	0.84	0.85
	PLUME 128k	20.84	52.75	0.82	0.83	20.84	51.69	0.79	0.82
	PLUME 256k	21.05	53.04	0.82	0.83	21.06	52.07	0.80	0.82
	PLUME 32k	19.77	51.78	0.81	0.82	20.28	51.35	0.79	0.80
it-en	NLLB 3.3B	36.33	64.25	0.88	0.87	43.96	67.59	0.88	0.86
	TOWERBASE 7B	32.95	62.57	0.88	0.86	43.96	67.59	0.88	0.86
	PLUME 128k	27.80	58.98	0.86	0.85	33.76	62.30	0.85	0.84
	PLUME 256k	28.91	59.82	0.86	0.86	34.76	62.75	0.85	0.85
	PLUME 32k	27.43	58.75	0.85	0.85	32.90	61.49	0.84	0.84
it-es	NLLB 3.3B	22.70	51.45	0.86	0.87	34.15	59.45	0.86	0.86
	TOWERBASE 7B	20.71	50.87	0.85	0.87	34.15	59.45	0.86	0.86
	PLUME 128k	20.91	50.70	0.85	0.86	30.30	56.88	0.84	0.85
	PLUME 256k	21.35	51.04	0.85	0.86	30.62	56.96	0.84	0.85
	PLUME 32k	20.99	50.72	0.85	0.86	30.06	56.70	0.84	0.85
it-eu	NLLB 3.3B	7.65	43.50	0.79	0.73	8.09	41.63	0.76	0.70
	PLUME 128k	9.77	47.74	0.81	0.79	10.07	45.74	0.79	0.76
	PLUME 256k	11.33	49.20	0.82	0.80	10.82	46.47	0.79	0.77
	PLUME 32k	10.69	48.55	0.81	0.78	10.44	45.82	0.78	0.76
it-fr	NLLB 3.3B	33.24	60.44	0.87	0.87	29.23	57.43	0.84	0.86
	TOWERBASE 7B	29.16	58.49	0.85	0.87	29.23	57.43	0.84	0.86
	PLUME 128k	27.21	56.24	0.83	0.84	23.92	52.66	0.81	0.82
	PLUME 256k	27.89	56.11	0.83	0.84	24.39	52.83	0.80	0.82
	PLUME 32k	26.35	55.67	0.82	0.83	24.04	52.53	0.80	0.81
it-gl	NLLB 3.3B	25.72	54.62	0.87	0.86	32.39	58.86	0.86	0.84
	PLUME 128k	23.80	54.06	0.86	0.85	29.04	56.66	0.84	0.83
	PLUME 256k	23.79	53.94	0.86	0.84	29.34	56.60	0.84	0.82
	PLUME 32k	23.59	53.88	0.85	0.84	28.20	55.97	0.84	0.82
it-pt	NLLB 3.3B	28.17	56.94	0.88	0.86	33.41	58.86	0.87	0.85
	TOWERBASE 7B	24.49	55.37	0.86	0.85	33.41	58.86	0.87	0.85
	PLUME 128k	26.64	56.24	0.87	0.84	28.48	55.43	0.85	0.83
	PLUME 256k	27.10	56.52	0.87	0.85	28.33	55.31	0.84	0.83
	PLUME 32k	25.86	55.58	0.86	0.84	28.03	55.24	0.84	0.82

Table 22: Results for gl→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
gl-ca	BSC Bilinguals	32.43	60.50	0.87	0.84	34.27	60.27	0.86	0.84
	NLLB 3.3B	34.43	60.88	0.87	0.85	34.25	60.34	0.86	0.83
	PLUME 128k	32.77	60.71	0.87	0.84	34.28	60.55	0.86	0.83
	PLUME 256k	33.00	60.85	0.88	0.84	34.10	60.42	0.86	0.83
	PLUME 32k	32.75	60.76	0.87	0.84	33.72	60.27	0.86	0.83
gl-de	NLLB 3.3B	29.57	57.53	0.85	0.84	25.13	55.12	0.83	0.83
	PLUME 128k	23.05	54.44	0.81	0.81	20.23	51.72	0.79	0.80
	PLUME 256k	24.25	55.47	0.82	0.82	20.35	52.31	0.79	0.80
	PLUME 32k	22.89	54.11	0.80	0.80	19.75	51.46	0.78	0.79
gl-en	NLLB 3.3B	44.14	68.60	0.89	0.86	43.52	67.80	0.88	0.85
	PLUME 128k	35.47	64.50	0.86	0.85	33.40	62.42	0.85	0.84
	PLUME 256k	34.74	64.17	0.86	0.84	32.56	62.21	0.85	0.84
	PLUME 32k	34.15	63.48	0.86	0.84	30.76	61.22	0.84	0.83
gl-es	NLLB 3.3B	25.59	53.47	0.87	0.85	36.99	61.92	0.87	0.84
	PLUME 128k	23.67	52.86	0.86	0.85	35.18	61.04	0.86	0.84
	PLUME 256k	23.79	52.87	0.86	0.85	35.84	61.32	0.86	0.84
	PLUME 32k	23.59	52.83	0.86	0.85	35.48	61.15	0.86	0.84
gl-eu	NLLB 3.3B	12.37	48.45	0.82	0.73	9.06	43.94	0.78	0.70
	PLUME 128k	13.23	51.10	0.83	0.77	11.89	48.13	0.81	0.76
	PLUME 256k	13.68	51.27	0.83	0.77	11.28	48.44	0.81	0.76
	PLUME 32k	12.78	50.05	0.82	0.75	10.94	47.31	0.80	0.74
gl-fr	NLLB 3.3B	38.37	63.38	0.86	0.85	29.03	56.98	0.84	0.84
	PLUME 128k	29.14	57.49	0.82	0.82	23.19	52.26	0.79	0.81
	PLUME 256k	30.24	57.82	0.82	0.82	23.80	52.55	0.79	0.80
	PLUME 32k	29.84	57.65	0.81	0.81	23.56	52.22	0.79	0.80
gl-it	NLLB 3.3B	26.14	55.52	0.88	0.85	30.79	58.39	0.87	0.84
	PLUME 128k	22.73	53.29	0.86	0.84	26.47	55.68	0.84	0.83
	PLUME 256k	23.20	53.77	0.86	0.84	27.00	56.19	0.84	0.83
	PLUME 32k	22.45	53.22	0.86	0.84	26.36	55.84	0.84	0.83
gl-pt	NLLB 3.3B	34.42	60.37	0.88	0.83	31.87	58.16	0.87	0.83
	PLUME 128k	28.42	57.24	0.87	0.83	26.36	54.81	0.85	0.81
	PLUME 256k	29.11	57.70	0.87	0.83	27.82	55.65	0.85	0.81
	PLUME 32k	29.23	57.83	0.87	0.83	27.50	55.41	0.85	0.81

Table 23: Results for pt→xx.

Pair	Model	FLORES-200				NTREX			
		BLEU	CHRf	COMET	COMET-KIWI	BLEU	CHRf	COMET	COMET-KIWI
pt-ca	BSC Bilinguals	35.75	61.22	0.87	0.84	32.04	58.28	0.86	0.83
	NLLB 3.3B	34.64	60.68	0.87	0.84	31.17	57.91	0.85	0.83
	PLUME 128k	35.50	61.41	0.87	0.84	31.05	57.84	0.85	0.83
	PLUME 256k	35.38	60.95	0.87	0.83	31.12	57.84	0.85	0.83
	PLUME 32k	35.50	61.26	0.87	0.83	30.95	57.66	0.85	0.82
pt-de	NLLB 3.3B	31.27	58.75	0.85	0.85	25.56	55.62	0.84	0.84
	TOWERBASE 7B	25.48	56.02	0.82	0.84	25.56	55.62	0.84	0.84
	PLUME 128k	25.45	55.44	0.82	0.82	19.99	51.73	0.80	0.80
	PLUME 256k	26.51	55.90	0.83	0.82	20.03	51.96	0.80	0.81
	PLUME 32k	25.01	54.48	0.81	0.81	20.48	51.29	0.79	0.79
pt-en	NLLB 3.3B	52.50	73.31	0.90	0.85	43.94	68.11	0.88	0.85
	TOWERBASE 7B	50.16	72.76	0.90	0.85	43.94	68.11	0.88	0.85
	PLUME 128k	42.71	68.42	0.88	0.84	33.21	62.26	0.85	0.83
	PLUME 256k	43.31	68.95	0.88	0.84	33.50	62.46	0.86	0.83
	PLUME 32k	41.73	67.58	0.87	0.83	32.87	61.63	0.85	0.82
pt-es	NLLB 3.3B	25.76	53.31	0.86	0.86	34.85	60.45	0.86	0.85
	TOWERBASE 7B	22.82	51.90	0.85	0.85	34.85	60.45	0.86	0.85
	PLUME 128k	22.97	51.85	0.85	0.85	30.89	57.40	0.85	0.84
	PLUME 256k	23.04	51.82	0.85	0.84	31.32	57.66	0.85	0.84
	PLUME 32k	22.72	51.74	0.85	0.84	30.84	57.25	0.85	0.84
pt-eu	NLLB 3.3B	10.38	45.45	0.79	0.72	8.14	41.30	0.76	0.69
	PLUME 128k	11.18	49.09	0.82	0.79	9.93	46.18	0.80	0.77
	PLUME 256k	13.37	50.70	0.82	0.79	10.26	46.86	0.80	0.77
	PLUME 32k	12.68	49.77	0.81	0.78	10.50	46.72	0.79	0.76
pt-fr	NLLB 3.3B	40.85	64.94	0.87	0.86	29.39	57.41	0.84	0.85
	TOWERBASE 7B	36.52	62.44	0.85	0.85	29.39	57.41	0.84	0.85
	PLUME 128k	33.25	59.78	0.83	0.83	23.91	52.93	0.80	0.81
	PLUME 256k	33.80	59.69	0.83	0.82	24.72	53.34	0.81	0.81
	PLUME 32k	32.60	58.97	0.82	0.82	24.11	52.80	0.80	0.80
pt-gl	NLLB 3.3B	31.12	57.92	0.88	0.83	32.55	59.00	0.87	0.82
	PLUME 128k	28.83	56.91	0.87	0.82	28.27	56.48	0.85	0.81
	PLUME 256k	28.58	56.52	0.87	0.82	28.54	56.57	0.85	0.81
	PLUME 32k	28.64	56.61	0.87	0.82	28.01	56.32	0.85	0.81
pt-it	NLLB 3.3B	26.42	55.44	0.88	0.85	31.19	59.11	0.87	0.85
	TOWERBASE 7B	22.31	52.69	0.85	0.85	31.19	59.11	0.87	0.85
	PLUME 128k	24.06	53.75	0.86	0.84	26.97	56.30	0.85	0.83
	PLUME 256k	24.24	53.75	0.86	0.84	27.46	56.52	0.85	0.83
	PLUME 32k	23.67	53.46	0.85	0.83	27.60	56.49	0.85	0.83