

Rethinking Multilingual Continual Pretraining: Data Mixing for Adapting LLMs Across Languages and Resources

Zihao Li¹, Shaoxiong Ji^{1,2,3*}, Hengyu Luo¹, Jörg Tiedemann¹

¹University of Helsinki ²ELLIS Institute Finland ³University of Turku
{zihao.li, shaoxiong.ji, hengyu.luo, jorg.tiedemann}@helsinki.fi

Abstract

Large Language Models (LLMs) exhibit significant disparities in performance across languages, primarily benefiting high-resource languages while marginalizing underrepresented ones. Continual Pretraining (CPT) has emerged as a promising approach to address this imbalance, although the relative effectiveness of monolingual, bilingual, and code-augmented data strategies remains unclear. This study systematically evaluates 36 CPT configurations involving three multilingual base models, across 30+ languages categorized as altruistic, selfish, and stagnant, spanning various resource levels. Our findings reveal three major insights: (1) Bilingual CPT improves multilingual classification but often causes language mixing issues during generation. (2) Including programming code data during CPT consistently enhances multilingual classification accuracy and language modeling capabilities, particularly benefiting low-resource languages, but introduces a trade-off by slightly degrading generation quality. (3) Contrary to prior work, we observe substantial deviations from language classifications according to their impact on cross-lingual transfer: Languages classified as altruistic often negatively affect related languages, selfish languages show conditional and configuration-dependent behavior, and stagnant languages demonstrate surprising adaptability under certain CPT conditions. These nuanced interactions emphasize the complexity of multilingual representation learning, underscoring the importance of systematic studies on generalizable language classification to inform future multilingual CPT strategies.

1 Introduction

Large Language Models (LLMs), built upon the Transformer architecture (Vaswani et al., 2017), have achieved remarkable progress in tasks such as machine translation, text classification, and generative dialogue. Despite these advances, their performance remains highly uneven across languages, favoring high-resource languages and marginalizing underrepresented ones (Li et al., 2024; Joshi et al., 2020; Choudhury, 2023). This imbalance deepens the digital language divide and limits the inclusivity of NLP technologies.

Recent work on Continual Pretraining (CPT) has shown promise for adapting pretrained models to new languages through additional training on targeted data (Zheng et al., 2024a). EMMA-500 employed CPT with extensive monolingual datasets across more than 500 languages, significantly improving multilingual performance, particularly for low-resource languages (Ji et al., 2024a). LLaMAX achieved notable translation improvements through CPT on over 100 languages involving data augmentation with bilingual translation data (Lu et al., 2024). Similar effects could be demonstrated on translation tasks with the CPT-based TOWER model (Alves et al., 2024). However, the relative effectiveness of monolingual and bilingual translation data for CPT remains unclear, particularly in terms of their impact on continual language learning, language interference, and performance consistency across different resource levels of languages.

*Corresponding author. Work done while at the University of Helsinki.

In addition to textual data in natural languages, a growing practice in LLM training is to incorporate programming code. For multilingual models, the inclusion of code, with its distinct logical syntax and lower ambiguity compared to natural language, can diversify the training data. We hypothesize that this diversity compels the model to learn more general and robust representations, moving beyond language-specific surface patterns; thus improving both the model’s predictive ability and its performance on downstream tasks.

There is a critical gap in understanding how language characteristics interact with CPT strategies. A recent classification proposed by [Yuan et al. \(2024\)](#) categorizes languages as *altruistic*, *selfish*, and *stagnant* based on their cross-lingual transfer patterns. This framework classifies languages based on their impact in multilingual training: altruistic languages are hypothesized to improve performance on related languages, selfish languages primarily benefit themselves, and stagnant languages offer little to no benefit to any language, including their own. However, this classification has only been validated in narrow experimental settings using English-centric bilingual data, leaving open questions about its generalizability to: (1) non-English language pairs, (2) code-augmented training regimes, and (3) models with varying pretraining corpora and architectures.

To systematically assess the impact of different CPT strategies, we conduct extensive experiments with 36 configurations, evaluating monolingual, bilingual, and code-augmented CPT on multilingual adaptation.¹ Our setup includes three multilingual base models—Llama-3.1-8B ([Dubey et al., 2024](#)), Llama-2-7B ([Touvron et al., 2023](#)), and Viking-7B ([Luukkonen et al., 2025](#))—continual-pretrained languages spanning high-, medium-, and low-resource categories. We evaluate the model performance on 14 training languages and assess cross-lingual transfer on 25 related languages, with a particular focus on assessing how different CPT configurations perform across altruistic, selfish, and stagnant language categories by [Yuan et al. \(2024\)](#).

Our systematic evaluation of 36 CPT configurations across three base models and 30+ languages yields three core insights:

- **Bilingual CPT improves classification performance but introduces generation challenges:** Compared to monolingual CPT, bilingual CPT generally improves multilingual classification accuracy for medium- and low-resource languages. However, it frequently results in problematic language mixing during generation tasks, limiting its overall utility.
- **Code data enhances classification but introduces trade-offs in generation:** Adding code data during CPT significantly boosts multilingual classification performance across resource levels, especially for low-resource languages, acting as an effective scaffold for representation learning. Nevertheless, code inclusion may lead to a trade-off, slightly degrading generation quality in certain scenarios.
- **The categorization of languages according to their cross-lingual transfer abilities does not generalize under varying conditions:** Our experiments reveal substantial deviations from language classifications proposed in previous work ([Yuan et al., 2024](#)): so-called *altruistic languages* are not always helpful and often negatively impact related languages, *selfish languages* exhibit highly configuration-dependent cross-lingual effects, and languages classified as *stagnant* demonstrate unexpected adaptability under specific training settings. These findings highlight the complexity of multilingual interactions in CPT and emphasize the need for a more adaptive classification framework for cross-lingual learning.

¹Monolingual data consists of texts in a single language, though it may include code-switching. Bilingual translation data contains sentence pairs in two languages that convey the same meaning. When monolingual data from different languages is combined, it forms multilingual continual pertaining, and a similar principle applies to bilingual translation data. However, for clarity, we refer to these setups as monolingual CPT and bilingual CPT, respectively.

2 Materials and Methods

2.1 Language Selection

We systematically evaluate the effects of CPT on multilingual models by selecting languages according to the altruistic, selfish, and stagnant categories defined in Yuan et al. (2024), which classify languages based on their behavior in multilingual training and evaluation.

For each category, we select 1 high-resource language (except for the stagnant category, for which no high-resource language is available in the dataset we use), 2 medium-resource languages, and 2 low-resource languages to ensure a balanced representation across different resource levels. The classification of languages into high-, medium-, and low-resource categories is determined by analyzing the data distribution of the Lego-MT dataset (Yuan et al., 2023), which serves as the basis for our setup. Specifically, we calculate the total token count for each language. Languages are then categorized as follows: high-resource languages exceed 1 billion tokens, medium-resource languages range between 10 million and 1 billion tokens, and low-resource languages fall below 10 million tokens. These languages serve as the training languages in our CPT experiments. The selected training languages, along with their corresponding category and resource level, are summarized in the first three columns of Table 1.

To further validate the findings in Yuan et al. (2024), we select 1-2 linguistically related languages for each training language based on the language evolutionary tree²³. These related languages are not included in the CPT phase but are used for cross-lingual evaluation to determine whether the effects observed in training languages extend to unseen but related languages. The fourth and fifth columns of Table 1 list the selected related languages. For some languages, this includes one, and for others, two related languages that are available in the evaluation benchmarks.

| Category | Resources | Training Language | Related Language 1 | Related Language 2 |
|------------|-----------|-------------------|--------------------|--------------------|
| Altruistic | High | zho_Hani | yue_Hant | - |
| | Medium | ceb_Latn | tgl_Latn | ilo_Latn |
| | Medium | mar_Deva | hin_Deva | npi_Deva |
| | Low | zul_Latn | xho_Latn | ssw_Latn |
| | Low | khm_Khmr | vie_Latn | - |
| Selfish | High | deu_Latn | nld_Latn | dan_Latn |
| | Medium | bel_Cyrl | rus_Cyrl | ukr_Cyrl |
| | Medium | mri_Latn | smo_Latn | fij_Latn |
| | Low | kir_Cyrl | kaz_Cyrl | bak_Cyrl |
| | Low | nya_Latn | bem_Latn | sna_Latn |
| Stagnant | Medium | tha_Thai | lao_Lao | shn_Mymr |
| | Medium | yor_Latn | ibo_Latn | hau_Latn |
| | Low | sna_Latn | nya_Latn | zul_Latn |
| | Low | wol_Latn | bam_Latn | - |

Table 1: Selected languages for CPT along with their corresponding related languages for evaluation. ‘-’ indicates the second related language cannot be found in the benchmark.

2.2 Pretraining Data

Bilingual Translation Data We utilize subsets of the Lego-MT (Yuan et al., 2023) and NLLB (Schwenk et al., 2021; Heffernan et al., 2022; Costa-jussà et al., 2022) datasets as our sources of parallel bilingual data. The Lego-MT dataset, derived from OPUS⁴, provides translations across 433 languages. The NLLB dataset consists of 148 English-centric and

²http://www.elinguistics.net/Language_Evolutionary_Tree.html

³Using the language evolutionary tree to identify related languages, we assess whether CPT effects transfer to unseen but linguistically similar languages, thus evaluating cross-lingual robustness.

⁴<https://opus.nlpl.eu>

1,465 non-English-centric bitext pairs mined from different parallel sources. To construct our parallel training data, we select specific language pairs from these datasets and apply OpusFilter (Aulamo et al., 2020) to remove duplicate data points.

The resulting dataset comprises approximately 292 million tokens across 22 language pairs, distributed over three language categories: altruistic (10 pairs, ~92M tokens), selfish (8 pairs, ~100M tokens), and stagnant (4 pairs, ~100M tokens).

For training, we format parallel data using the following structure:

[source language]: [source] [target language]: [target]

Monolingual Data We extract a subset of MADLAD-400 (Kudugunta et al., 2024), a large-scale multilingual dataset derived from Common Crawl⁵, covering 419 languages. Since web-crawled text does not inherently guarantee monolingual integrity, we employ GlotLID (Kargaran et al., 2023), a language identification model, to analyze the language composition of each text segment and ensure strict monolingual consistency. Specifically, for each document in the dataset, we first segment the text into sentences using the NLTK (Bird & Loper, 2004) sentence splitter. Then, GlotLID predicts the language of each sentence independently. We retain only those documents where all sentences are identified as belonging to the same language, discarding any text segment that exhibits code-switching or multilingual content.

We finally select data for 15 unique languages, with English appearing in all three categories: altruistic (6 languages, ~92M tokens), selfish (6 languages, ~100M tokens), and stagnant (5 languages, ~87M tokens), resulting in a total of approximately 279 million tokens.

Code Data We incorporate code data from The Stack (Kocetkov et al., 2022), following the pre-processing strategy used in EMMA-500 (Ji et al., 2024a). The dataset is first filtered to retain high-quality source files, with a focus on data science-related code and the 32 most commonly used general-purpose programming languages. Additionally, we include LLVM code due to its importance in multilingual code generation (Paul et al., 2024; Szafraniec et al., 2022).

For training configurations that include code, we establish a 2:1 ratio between textual (100 million tokens) and code data (50 million tokens), with code comprising about 33% of the total mixture. We intentionally selected this higher proportion to investigate the effects of a strong code signal within our CPT context. While prior work (Aryabumi et al., 2024) on pretraining from scratch recommends a 25% proportion for balancing language and code performance, it is noted that a higher 33% remains reasonable for enhancing models’ ability, aligning with our goal of studying its trade-offs.

2.3 Base Models

We evaluate across three open-source multilingual LLMs with diverse training recipes:

Llama-3.1-8B (Dubey et al., 2024) is pretrained on approximately 15 trillion tokens from diverse, multilingual sources. Its extensive multilingual pretraining and high capacity make it ideal for analyzing CPT effects on well-trained models.

Llama-2-7B (Touvron et al., 2023) is pretrained on 2 trillion tokens, covering a broad yet less multilingual data distribution. It provides a baseline to evaluate CPT effectiveness on English-centric models commonly used in multilingual adaptation research.

Viking-7B (Luukkonen et al., 2025) is pretrained mainly on Nordic languages, English, and code, offering insights into how CPT impacts models initially trained on narrower, region-specific data.

⁵<https://commoncrawl.org/>

2.4 CPT Configurations

We train models under 4 CPT configurations across 3 base models and 3 language categories, resulting in a total of 36 models. Each model is named using the format:

Model-Data[+Code]-LangCat

where:

- Model $\in \{\text{L3 (Llama-3.1-8B), L2 (Llama-2-7B), V7 (Viking-7B)}\}$
- Data $\in \{\text{Mono (Monolingual), Bi (Bilingual)}\}$
- Code (optional) is added if code data is included
- LangCat $\in \{\text{Alt (Altruistic), Sel (Selfish), Stag (Stagnant)}\}$

For example, L3-Mono-Alt refers to Llama-3.1-8B trained on monolingual data for altruistic languages, while L2-Bi+Code-Sel denotes Llama-2-7B trained on bilingual parallel texts in selfish languages and code data.

Each model is trained for 2 epochs on a cluster with $4 \times \text{AMD MI250X GPUs}$ (8 Graphics Compute Dies) on each node. Training data is organized by language category (altruistic, selfish, stagnant), with all languages within a category (e.g., altruistic: `zho_Hani`, `ceb_Latn`, etc.) mixed into a single dataset per configuration (e.g., monolingual, bilingual+code). As for software, we use the LLaMA-Factory (Zheng et al., 2024b) framework with DeepSpeed (Rajbhandari et al., 2020) ZeRO-3 config. The hyperparameter setup includes a per-device batch size of 8 with gradient accumulation steps of 2. We use a cosine learning rate scheduler with an initial learning rate of 4.0×10^{-5} and a warmup ratio of 0.03.

3 Evaluation and Discussion

3.1 Benchmarks and Setup

We evaluate our models on two highly multilingual benchmarks covering a classification and a generation task: SIB-200 (Adelani et al., 2024) for topic classification and FLORES-200 (Costa-jussà et al., 2022; Goyal et al., 2022; Guzmán et al., 2019) for machine translation. Classification focuses on whether CPT improves the multilingual model’s understanding within a single language, while translation studies the alignment between languages that emerges with multilingual CPT. All experiments use a consistent 3-shot prompting setup.

SIB-200 SIB-200 is a multilingual news topic classification benchmark covering 200 languages. The task involves classifying news headlines into one of the following predefined categories: science/technology, travel, politics, sports, health, entertainment, and geography.

The model predicts by ranking logits for each category, and accuracy measures performance across languages.

FLORES-200 FLORES-200 evaluates multilingual translation performance across diverse language pairs.

Translations are generated using the vLLM (Kwon et al., 2023) inference engine. BLEU (Papineni et al., 2002) scores, computed via SacreBLEU (Post, 2018) with the flores200 tokenizer, quantify translation quality.⁶

3.2 Effect of Monolingual and Bilingual Continual Pretraining

This section shows that bilingual CPT hampers generation due to language mixing but excels in classification for medium- and low-resource languages over monolingual CPT.

⁶BLEU signature: nrefs:1|case:mixed|eff:no|tok:flores200|smooth:exp|version:2.4.2

3.2.1 Language Mixing in Generation Tasks

The FLORES-200 translation task revealed significant language mixing issues in models trained with bilingual translation data. Specifically, when generating translations between language pairs, models frequently appended unintended language tokens to the output. For example, when translating from English (eng_Latn) to Chinese (zho_Hani), models trained on bilingual data produced outputs like:

“我们现在有了非糖尿病的4个月小鼠，它们原本是患有糖尿病的。 Marathi: त्यांना मधुमेह होता. आता, आमचे चार महिन्याचे उंदर आहेत. ज्याला आधी मधुमेह झालेला होता. पण आता नाही. कारण यातील एक गोष्ट म्हणजे, साखर कमी करणे.....”

The text with a green background represents the desired Chinese translation, while the text with an orange background contains nonsensical multilingual fragments. This phenomenon occurred consistently across bilingual CPT configurations, suggesting that the parallel data format ([Lang1]: xxx [Lang2]: yyy) encourages cross-lingual interference. More examples are in Figure 6 in Appendix A.6.

This language inconsistency leads to significant translation quality degradation, as shown in Figure 1. Bilingual CPT configurations underperform monolingual CPT across all resource levels and base models. For high-resource languages, Llama-3.1-8B achieves only 7.47 BLEU with bilingual CPT versus 25.52 with monolingual CPT (-71% relative), while Llama-2-7B shows similar disparities (14.12 vs 24.60, -43%). The pattern persists for mid- and low-resource languages, with bilingual CPT consistently lagging behind monolingual CPT. Notably, monolingual CPT often matches or exceeds baseline performance, whereas bilingual CPT only exceeds baseline in specific cases, such as Llama-2-7B on mid- and low-resource languages. Appendix A.5 presents the detailed results on each language.

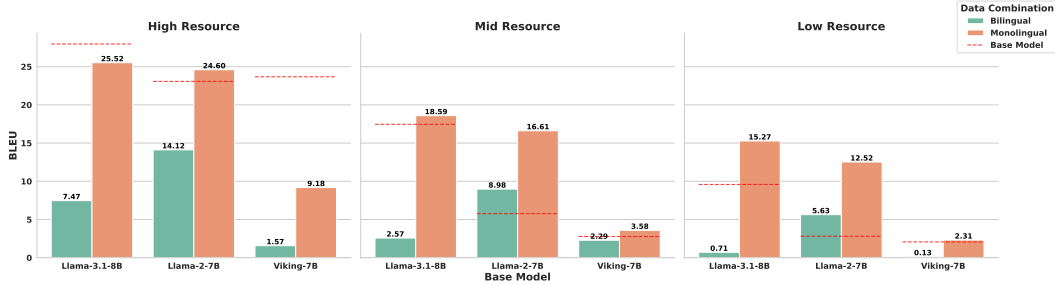


Figure 1: FLORES-200 X-Eng BLEU score comparing bilingual and monolingual CPT across high-, mid-, and low-resource languages.

3.2.2 Comparative Analysis in Classification Tasks

To isolate the effects of CPT strategies without interference from language mixing, we evaluate SIB-200 classification accuracy. Figure 2 shows the average accuracy aggregated across models trained separately on altruistic, selfish, and stagnant languages, grouped by resource level.

High-Resource Languages For high-resource languages, both monolingual and bilingual CPT degrade performance across all base models compared to their respective baselines. Llama-3.1-8B, despite its strong baseline (76.63%), exhibits drops with bilingual CPT (71.41%, -6.8% relative) and monolingual CPT (64.21%, -16.2%). Llama-2-7B shows significant declines with both strategies: bilingual CPT reduces accuracy to 31.54% (vs baseline 37.75%, -16.5%), while monolingual CPT performs similarly (31.86%, -15.6%). Viking-7B partially escapes this trend, with bilingual CPT achieving marginal gains (36.60% vs baseline 34.80%, +5.2%), though monolingual CPT underperforms (25.98%, -25.3%). This suggests that high-resource languages generally do not benefit from CPT, likely due to interference with existing strong representations in pretrained models. However, model-specific factors, such as whether the

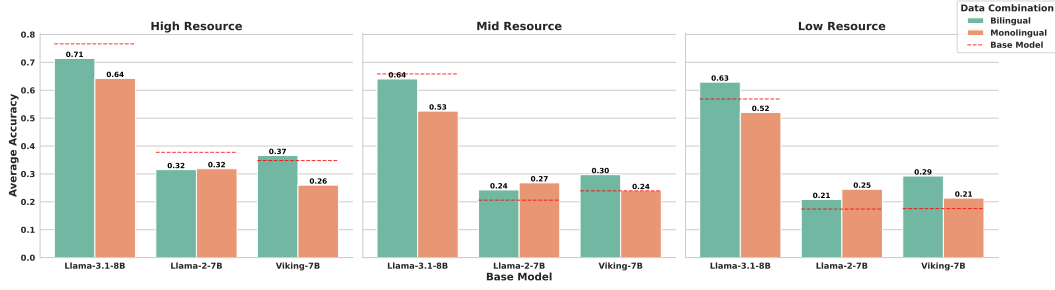


Figure 2: SIB-200 classification accuracy comparing monolingual and bilingual CPT across high-, mid-, and low-resource languages.

model’s pretraining data aligns well with the target languages in CPT, may enable limited improvements in certain cases. For example, Viking-7B, which was pretrained primarily on Nordic languages and English, may benefit more from bilingual CPT due to its ability to leverage cross-lingual transfer between related languages.

Mid-Resource Languages Mid-resource languages show mixed trends. Llama-3.1-8B maintains near-baseline performance with bilingual CPT (64.05% vs baseline 65.85%, -2.7%), but monolingual CPT degrades significantly (52.53%, -20.2%). Llama-2-7B improves accuracy in both configurations: bilingual CPT raises accuracy to 24.26% (vs. baseline 20.59%, +17.8%), while monolingual CPT performs slightly better at 26.80% (+30.2%). Viking-7B uniquely benefits from bilingual CPT (29.74% vs baseline 23.94%, +24.2%), while monolingual CPT underperforms (24.02%, +0.3%). This indicates that bilingual CPT can stabilize mid-resource language performance for certain models (e.g., Viking-7B and Llama-2-7B). However, monolingual CPT risks overfitting to limited in-language data, particularly for models with weaker pretraining (e.g., Llama-3.1-8B).

Low-Resource Languages Low-resource languages exhibit divergent patterns. Llama-3.1-8B improves with bilingual CPT (62.91% vs baseline 56.86%, +10.6%) but declines with monolingual CPT (52.04%, -8.5%). Llama-2-7B improves moderately with bilingual CPT (20.84%, +19.8%) and shows stronger gains with monolingual CPT (24.51%, +40.9%). Viking-7B benefits substantially from bilingual CPT (29.25%, +66.5%), while monolingual CPT slightly underperforms (21.33%, +21.4%). This highlights that bilingual CPT can enhance low-resource language performance for models with compatible pretraining (e.g., Viking-7B and Llama-3.1-8B).

3.3 Effect of Including Code Data

3.3.1 Downstream Tasks Performance

The integration of code data during monolingual CPT shows task-dependent effects, enhancing classification performance while introducing tradeoffs in generation quality. Figure 4 and Figure 3 compare monolingual CPT with and without code data across resource levels and tasks, revealing key patterns in how code data influences multilingual adaptation.

Code integration consistently improves classification accuracy across all resource levels and models. For high-resource languages, Llama-3.1-8B shows marginal gains (64.21% to 68.47%, +6.7% relative to baseline 76.63%), while Llama-2-7B and Viking-7B exhibit more substantial improvements (42.48% vs 31.86%, +33.3%; 30.88% vs 25.98%, +18.8%). Mid-resource languages benefit even more, with Llama-3.1-8B recovering near-baseline performance (52.53% to 62.83%, -4.6% vs baseline 65.85%) and Llama-2-7B achieving significant gains (34.40% vs 26.80%, +67.0%). Low-resource languages see the most pronounced improvements, particularly for Viking-7B (28.68% vs 21.33%, +63.2% relative to baseline 17.57%). This pattern extends to bilingual CPT configurations (see Appendix A.3).

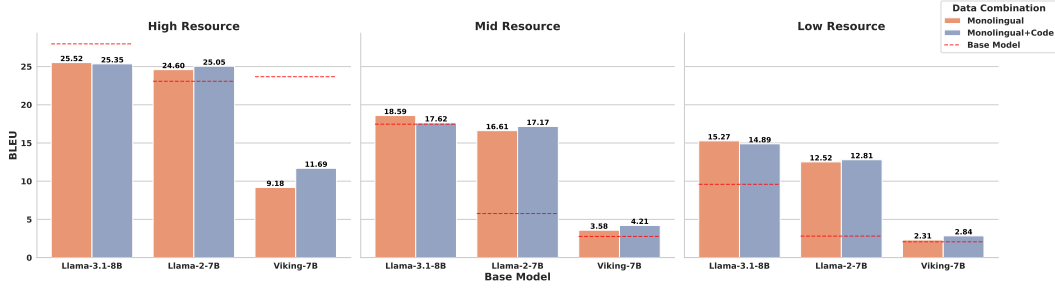


Figure 3: FLORES-200 X-Eng BLEU score comparing monolingual and monolingual+code CPT across high-, mid-, and low-resource languages.

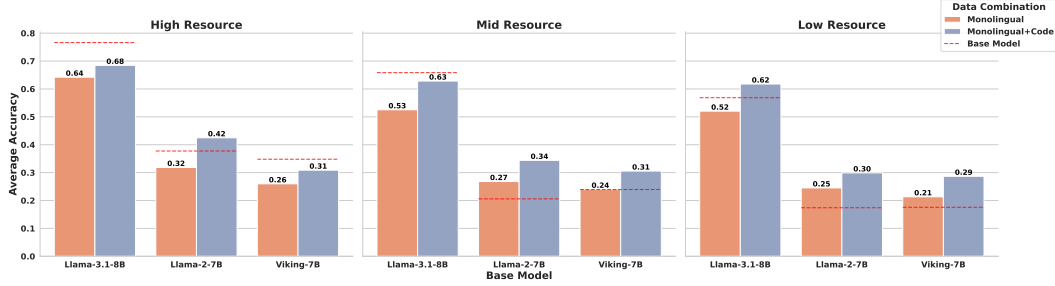


Figure 4: SIB-200 classification accuracy comparing monolingual and monolingual+code CPT across high-, mid-, and low-resource languages.

In contrast, code integration often degrades translation quality, particularly for high-resource languages. Llama-3.1-8B shows slight degradation (25.52 BLEU to 25.35, -0.7% vs baseline 27.97), while Llama-2-7B and Viking-7B exhibit gains (25.05 vs 24.60, +8.6%; 11.69 vs 9.18, +27.3%). Mid-resource languages show mixed trends, with Llama-3.1-8B experiencing a slight drop (17.62 vs 18.59, -5.2%) and Viking-7B improving significantly (4.21 vs 3.58, +52.0%). Low-resource languages partially escape this trend, with Viking-7B showing substantial gains (2.84 vs 2.31, +37.4%).

The benefits of code integration are most pronounced for low-resource languages, where it acts as a “scaffold” to improve classification accuracy (avg. +25.1%) and partially mitigate generation deficits. Mid-resource languages also benefit, though to a lesser extent, while high-resource languages see diminishing returns, with classification gains (e.g., Llama-3.1-8B: +6.7%) offset by generation losses.

3.3.2 Language Modeling Performance

To assess the impact of code integration on language modeling, we evaluate performance using Negative Log-Likelihood (NLL) on the MaLA validation set⁷. We use the implementation of GlotEval (Luo et al., 2025), where the NLL is calculated by concatenating input sentences and using a strided sliding window of 1024 tokens, and is defined as

$$NLL = - \sum_{i=1}^{n_i} \log p_{\theta}(x_i | x_{<i}).$$

Our results are summarized in Table 2. Across all three base models and all resource levels, the Monolingual+Code configurations yielded lower (better) NLL scores than the monolingual-only CPT. This demonstrates that including code not only helps in downstream classification tasks but also improves the models’ predictive ability and likelihood estimation of unseen text during CPT.

⁷<https://huggingface.co/datasets/MaLA-LM/mala-monolingual-split>

| Model | High Resource | Mid Resource | Low Resource |
|---------------|---------------|---------------|---------------|
| L2-Mono- | 5.0080 | 1.7609 | 4.3308 |
| L2-Mono+Code- | 4.7359 | 1.7042 | 4.2289 |
| L3-Mono- | 4.8864 | 2.0375 | 4.2013 |
| L3-Mono+Code- | 4.5701 | 1.9522 | 4.1389 |
| V7-Mono- | 6.9916 | 1.9506 | 4.5367 |
| V7-Mono+Code- | 6.6029 | 1.8369 | 4.4850 |

Table 2: NLL across different models and resource levels. Lower is better.

3.4 Validation of Language Category Hypotheses

This section evaluates the validity of the altruistic, selfish, and stagnant language classifications proposed in prior work (Yuan et al., 2024). We evaluate each model (e.g., L3-Mono-Alt) trained on a language category (e.g., altruistic languages: zho.Hani, ceb.Latn, etc.) and measure SIB-200 classification accuracy changes on both the trained languages and their related languages (e.g., yue.Hant, tgl.Latn, etc.), as defined in Table 1. We analyze whether CPT strategies align with these hypothesized behaviors. Table 3 reports accuracy changes (%) relative to base models.

| Model | Altruistic Languages | | | Selfish Languages | | | Stagnant Languages | | |
|---------------|----------------------|---------|------|-------------------|---------|------|--------------------|---------|------|
| | Training | Related | Met? | Training | Related | Met? | Training | Related | Met? |
| L2-Bi- | +7.08 | -22.55 | No | +12.33 | +2.90 | Yes | +5.88 | -9.99 | No |
| L2-Bi+Code- | +62.37 | +28.31 | No | +52.32 | +31.67 | No | +26.13 | +6.25 | No |
| L2-Mono- | -14.60 | -31.32 | No | +53.18 | +21.94 | No | +31.36 | -8.33 | No |
| L2-Mono+Code- | +50.43 | +19.04 | No | +52.32 | +26.29 | No | +64.02 | +14.57 | No |
| L3-Bi- | +4.46 | -4.46 | No | -7.90 | -19.54 | No | +14.76 | -28.43 | No |
| L3-Bi+Code- | +1.64 | -7.70 | No | -5.85 | -15.66 | No | +21.81 | -28.04 | No |
| L3-Mono- | -24.37 | -31.26 | No | -9.07 | -19.84 | No | -7.71 | -43.54 | No |
| L3-Mono+Code- | -1.78 | -11.13 | No | +2.49 | -10.85 | No | 0.00 | -37.01 | No |
| V7-Bi- | -11.41 | -31.95 | No | +19.24 | -10.22 | No | +78.18 | +26.32 | No |
| V7-Bi+Code- | +22.82 | -9.35 | No | +17.57 | -16.35 | No | +11.16 | -19.36 | No |
| V7-Mono- | -8.22 | -19.74 | No | +5.86 | -33.45 | No | 0.00 | -0.83 | Yes |
| V7-Mono+Code- | +5.93 | -11.69 | No | +53.96 | +8.18 | Yes | +21.31 | +17.27 | No |

Table 3: SIB-200 classification accuracy changes (%) for training and related languages across altruistic, selfish, and stagnant categories. Results are reported relative to base models, with a “Met” column to indicate whether the hypothesis is met or contradicted.

Altruistic languages can also be selfish or mutually harmful The altruistic hypothesis indicates that training in altruistic languages enhances multilingual performance (related languages) with minimal impact on their own performance (trained languages). Our results reveal three critical contradictions: (1) 83% of configurations (10/12) degraded related language performance, with code-free CPT causing up to -31.32% accuracy (L2-Mono-Alt); (2) Contrary to “minimal self-impact”, trained language accuracy fluctuated wildly (+62.37% in L2-Bi+Code-Alt vs. -24.37% in L3-Mono-Alt); These bidirectional effects challenge the unidirectional altruism assumption.

Selfish languages exhibit conditional isolation only in certain cases While the selfish hypothesis suggests trained languages primarily improve their own performance (trained languages) while minimally affecting others (related languages), we find this only holds in specific configurations: (1) Non-code bilingual training (L2-Bi-Sel) showed minimal impact on related languages (+2.90%); (2) Code-augmented monolingual training (V7-Mono+Code-Sel) achieved strong self-improvement (+53.96%) with moderate spillover (+8.18%). However, 83% of cases (10/12) violated the hypothesis through either negative spillover (V7-Mono-Sel: -33.45%) or excessive cross-lingual transfer (L2-Bi+Code-Sel: +31.67%).

Stagnant languages demonstrate more adaptability than expected Stagnant languages neither improve their own performance (trained languages) nor influence others. Contrary to their purported stagnation, 92% of configurations (11/12) induced significant performance shifts: (1) Bilingual training boosted trained languages by +78.18% (V7-Bi-Stag) while improving related languages (+26.13%); (2) Monolingual+code CPT (L2-Mono+Code-Stag) achieved +64% self-improvement with +14.76% cross-lingual gains. Only V7-Mono-Stag showed true stagnation (+0.00% trained, -0.83% related). This reveals that most “stagnant” languages possess untapped adaptation potential under proper CPT strategies.

4 Conclusion

In this study, we systematically evaluated the effects of multilingual CPT strategies, including monolingual, bilingual, and code-augmented configurations, across diverse resource levels and language categories. Through experiments with 36 configurations involving three multilingual base models and over 30 languages, we identified several critical insights:

First, while bilingual CPT enhances classification accuracy for mid- and low-resource languages, it introduces language mixing during generation, limiting its utility for translation tasks. Second, code integration during CPT acts as a scaffold for low-resource language understanding but introduces task-dependent trade-offs, improving classification and language modeling capabilities while slightly degrading generation quality. Third, we demonstrate that language classifications based on cross-lingual transfer patterns (*altruistic*, *selfish*, *stagnant*) fail to generalize under varying CPT strategies.

Overall, our work underscores the complexity of multilingual representation learning and highlights the need for flexible frameworks for language categorization and training strategy selection. Future research should focus on developing more adaptive CPT methods that balance classification improvements and generation quality, further bridging language disparities in large language models.

Ethics Statement

This research focuses on reducing the digital language divide and improving inclusivity for underrepresented languages. We acknowledge potential biases due to uneven data distribution and strive to mitigate them by including diverse languages across resource levels. All datasets used are publicly available and preprocessed to ensure integrity and monolingual consistency. All the models trained in this paper are strictly for research purposes and are not intended to be deployed in real-world applications. We encourage further work to address ethical challenges in multilingual NLP, especially for underrepresented languages.

Reproducibility Statement

To ensure reproducibility, we release:

- **Model Checkpoints:** All models trained under various configurations (monolingual, bilingual, code-augmented) across base models (Llama-3.1-8B, Llama-2-7B, Viking-7B) and language categories (*altruistic*, *selfish*, *stagnant*).
- **Processed Dataset:** Filtered subsets of Lego-MT, NLLB, MADLAD-400, and code data.
- **Scripts:** Data cleaning, training, and evaluation scripts, including LLaMA-Factory with DeepSpeed ZeRO-3 configuration.

All resources are available at: <https://mala-lm.github.io/MixCPT.html>.

Acknowledgments

The work has received funding from the Digital Europe Programme under grant agreement No 101195233 (OpenEuroLLM). The authors wish to acknowledge CSC - IT Center for Science, Finland, for providing computational resources.

References

- David Adelani, Hannah Liu, Xiaoyu Shen, Nikita Vassilyev, Jesujoba Alabi, Yanke Mao, Haonan Gao, and En-Shiun Lee. SIB-200: A simple, inclusive, and big evaluation dataset for topic classification in 200+ languages and dialects. In Yvette Graham and Matthew Purver (eds.), *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 226–245, St. Julian’s, Malta, March 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.eacl-long.14/>.
- Duarte M Alves, José Pombal, Nuno M Guerreiro, Pedro H Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, et al. Tower: An open multilingual large language model for translation-related tasks. *arXiv preprint arXiv:2402.17733*, 2024.
- Viraat Aryabumi, Yixuan Su, Raymond Ma, Adrien Morisot, Ivan Zhang, Acyr Locatelli, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. To code, or not to code? exploring impact of code in pre-training. *arXiv preprint arXiv:2408.10914*, 2024.
- Mikko Aulamo, Sami Virpioja, and Jörg Tiedemann. OpusFilter: A configurable parallel corpus filtering toolbox. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 150–156. Association for Computational Linguistics, July 2020. doi: 10.18653/v1/2020.acl-demos.20. URL <https://www.aclweb.org/anthology/2020.acl-demos.20>.
- Steven Bird and Edward Loper. NLTK: The natural language toolkit. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pp. 214–217, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/P04-3031/>.
- Tyler A. Chang, Catherine Arnett, Zhuowen Tu, and Ben Bergen. When is multilinguality a curse? language modeling for 250 high- and low-resource languages. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 4074–4096, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.236. URL <https://aclanthology.org/2024.emnlp-main.236/>.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Monojit Choudhury. Generative ai has a language problem. *Nature human behaviour*, 7(11): 1802–1803, 2023.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*, 2022.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Hiroki Iida, Masanari Ohi, Kakeru Hattori, Hirai Shota, Sakae Mizuki, Rio Yokota, and Naoaki Okazaki. Continual pre-training for cross-lingual llm adaptation: Enhancing japanese language capabilities. *arXiv preprint arXiv:2404.17790*, 2024.

- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. The Flores-101 evaluation benchmark for low-resource and multilingual machine translation. *Transactions of the Association for Computational Linguistics*, 10:522–538, 2022. doi: 10.1162/tacl_a.00474. URL <https://aclanthology.org/2022.tacl-1.30/>.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 6098–6111, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1632. URL <https://aclanthology.org/D19-1632/>.
- Kevin Heffernan, Onur Çelebi, and Holger Schwenk. Bitext mining using distilled sentence representations for low-resource languages. *arXiv preprint arXiv:2205.12654*, 2022.
- Masanori Hirano and Kentaro Imajo. The construction of instruction-tuned llms for finance without instruction data using continual pretraining and model merging. *arXiv preprint arXiv:2409.19854*, 2024.
- Shaoxiong Ji, Zihao Li, Indraneil Paul, Jaakko Paavola, Peiqin Lin, Pinzhen Chen, Dayyán O’Brien, Hengyu Luo, Hinrich Schütze, Jörg Tiedemann, et al. Emma-500: Enhancing massively multilingual adaptation of large language models. *arXiv preprint arXiv:2409.17892*, 2024a.
- Shaoxiong Ji, Timothee Mickus, Vincent Segonne, and Jörg Tiedemann. Can machine translation bridge multilingual pretraining and cross-lingual transfer learning? In *Proceedings of LREC-COLING*, 2024b.
- Shaoxiong Ji, Zihao Li, Jaakko Paavola, Indraneil Paul, Hengyu Luo, and Jörg Tiedemann. Massively multilingual adaptation of large language models using bilingual translation data. *arXiv preprint arXiv:2506.00469*, 2025.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. The state and fate of linguistic diversity and inclusion in the NLP world. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6282–6293, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.560. URL <https://aclanthology.org/2020.acl-main.560/>.
- Mihir Kale, Aditya Siddhant, Rami Al-Rfou, Linting Xue, Noah Constant, and Melvin Johnson. nmT5 - is parallel data still relevant for pre-training massively multilingual language models? In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (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)*, pp. 683–691, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-short.87. URL <https://aclanthology.org/2021.acl-short.87/>.
- Amir Hossein Kargaran, Ayyoob Imani, François Yvon, and Hinrich Schuetze. GlotLID: Language identification for low-resource languages. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 6155–6218, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.410. URL <https://aclanthology.org/2023.findings-emnlp.410/>.
- Najoung Kim, Sebastian Schuster, and Shubham Toshniwal. Code pretraining improves entity tracking abilities of language models. *arXiv preprint arXiv:2405.21068*, 2024.
- Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, et al. The stack: 3 tb of permissively licensed source code. *arXiv preprint arXiv:2211.15533*, 2022.

- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. Madlad-400: A multilingual and document-level large audited dataset. *Advances in Neural Information Processing Systems*, 36, 2024.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- Zihao Li, Yucheng Shi, Zirui Liu, Fan Yang, Ali Payani, Ninghao Liu, and Mengnan Du. Quantifying multilingual performance of large language models across languages. *arXiv preprint arXiv:2404.11553*, 2024.
- Peiqin Lin, André FT Martins, and Hinrich Schütze. A recipe of parallel corpora exploitation for multilingual large language models. *arXiv preprint arXiv:2407.00436*, 2024.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.
- Yinquan Lu, Wenhao Zhu, Lei Li, Yu Qiao, and Fei Yuan. LLaMAX: Scaling linguistic horizons of LLM by enhancing translation capabilities beyond 100 languages. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 10748–10772, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.631. URL <https://aclanthology.org/2024.findings-emnlp.631/>.
- Hengyu Luo, Zihao Li, Joseph Attieh, Sawal Devkota, Ona de Gibert, Shaoxiong Ji, Peiqin Lin, Bhavani Sai Praneeth Varma Mantina, Ananda Sreenidhi, Raúl Vázquez, Mengjie Wang, Samea Yusofi, and Jörg Tiedemann. Gloteval: A test suite for massively multilingual evaluation of large language models. *arXiv preprint 2504.04155*, 2025. URL <https://arxiv.org/abs/2504.04155>.
- Risto Luukkonen, Jonathan Burdge, Elaine Zosa, Ville Komulainen, Peter Sarlin, and Sampo Pyysalo. Viking: A family of nordic llms, 2025. URL <https://huggingface.co/LumiOpen/Viking-33B>.
- Yingwei Ma, Yue Liu, Yue Yu, Yuanliang Zhang, Yu Jiang, Changjian Wang, and Shan-shan Li. At which training stage does code data help llms reasoning? *arXiv preprint arXiv:2309.16298*, 2023.
- Aman Madaan, Shuyan Zhou, Uri Alon, Yiming Yang, and Graham Neubig. Language models of code are few-shot commonsense learners. *arXiv preprint arXiv:2210.07128*, 2022.
- Mitodru Niyogi and Arnab Bhattacharya. Paramanu-ayn: Pretrain from scratch or continual pretraining of llms for legal domain adaptation? *arXiv preprint arXiv:2403.13681*, 2024.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- Indraneil Paul, Goran Glavaš, and Iryna Gurevych. IRCoder: Intermediate representations make language models robust multilingual code generators. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15023–15041, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.802. URL <https://aclanthology.org/2024.acl-long.802/>.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. Mad-x: An adapter-based framework for multi-task cross-lingual transfer. *arXiv preprint arXiv:2005.00052*, 2020.
- Matt Post. A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771*, 2018.

- Vitaly Protasov, Elisei Stakovskii, Ekaterina Voloshina, Tatiana Shavrina, and Alexander Panchenko. Super donors and super recipients: Studying cross-lingual transfer between high-resource and low-resource languages. In Atul Kr. Ojha, Chao-hong Liu, Ekaterina Vylomova, Flammie Pirinen, Jade Abbott, Jonathan Washington, Nathaniel Oco, Valentin Malykh, Varvara Logacheva, and Xiaobing Zhao (eds.), *Proceedings of the Seventh Workshop on Technologies for Machine Translation of Low-Resource Languages (LoResMT 2024)*, pp. 94–108, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.loresmt-1.10. URL <https://aclanthology.org/2024.loresmt-1.10/>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI*, 2019. URL https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16. IEEE, 2020.
- Leonardo Ranaldi, Giulia Pucci, and Andre Freitas. Empowering cross-lingual abilities of instruction-tuned large language models by translation-following demonstrations. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 7961–7973, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.473. URL <https://aclanthology.org/2024.findings-acl.473/>.
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, and Angela Fan. CCMatrix: Mining billions of high-quality parallel sentences on the web. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 6490–6500, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.507. URL <https://aclanthology.org/2021.acl-long.507/>.
- Marc Szafraniec, Baptiste Roziere, Hugh Leather, Francois Charton, Patrick Labatut, and Gabriel Synnaeve. Code translation with compiler representations. *arXiv preprint arXiv:2207.03578*, 2022.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. Multilingual translation from denoising pre-training. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 3450–3466, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.304. URL <https://aclanthology.org/2021.findings-acl.304/>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Anh-Dung Vo, Minseong Jung, Wonbeen Lee, and Daewoo Choi. Redwhale: An adapted korean llm through efficient continual pretraining. *arXiv preprint arXiv:2408.11294*, 2024.
- Zirui Wang, Zachary C. Lipton, and Yulia Tsvetkov. On negative interference in multilingual models: Findings and a meta-learning treatment. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4438–4450, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.359. URL <https://aclanthology.org/2020.emnlp-main.359/>.

Çağatay Yıldız, Nishaanth Kanna Ravichandran, Nitin Sharma, Matthias Bethge, and Beyza Ermis. Investigating continual pretraining in large language models: Insights and implications. *arXiv preprint arXiv:2402.17400*, 2024.

Zheng Xin Yong, Hailey Schoelkopf, Niklas Muennighoff, Alham Fikri Aji, David Ifeoluwa Adelani, Khalid Almubarak, M Saiful Bari, Lintang Sutawika, Jungo Kasai, Ahmed Baruwa, Genta Winata, Stella Biderman, Edward Raff, Dragomir Radev, and Vassilina Nikoulina. BLOOM+1: Adding language support to BLOOM for zero-shot prompting. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11682–11703, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.653. URL <https://aclanthology.org/2023.acl-long.653/>.

Yao-Ching Yu, Tsun-Han Chiang, Cheng-Wei Tsai, Chien-Ming Huang, and Wen-Kwang Tsao. Primus: A pioneering collection of open-source datasets for cybersecurity llm training. *arXiv preprint arXiv:2502.11191*, 2025.

Fei Yuan, Yinquan Lu, Wenhao Zhu, Lingpeng Kong, Lei Li, Yu Qiao, and Jingjing Xu. Lego-MT: Learning detachable models for massively multilingual machine translation. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 11518–11533, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.731. URL <https://aclanthology.org/2023.findings-acl.731/>.

Fei Yuan, Shuai Yuan, Zhiyong Wu, and Lei Li. How vocabulary sharing facilitates multilingualism in LLaMA? In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 12111–12130, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.721. URL <https://aclanthology.org/2024.findings-acl.721/>.

Yiran Zhao, Chaoqun Liu, Yue Deng, Jiahao Ying, Mahani Aljunied, Zhaodonghui Li, Lidong Bing, Hou Pong Chan, Yu Rong, Deli Zhao, et al. Babel: Open multilingual large language models serving over 90% of global speakers. *arXiv preprint arXiv:2503.00865*, 2025.

Wenzhen Zheng, Wenbo Pan, Xu Xu, Libo Qin, Li Yue, and Ming Zhou. Breaking language barriers: Cross-lingual continual pre-training at scale. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 7725–7738, 2024a.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand, 2024b. Association for Computational Linguistics. URL <http://arxiv.org/abs/2403.13372>.

A Appendix

A.1 Related Work

Continual Pretraining and Language Adaptation Continual pretraining has emerged as a pivotal technique for adapting LLMs to new domains or languages while retaining previously acquired knowledge (Yıldız et al., 2024). This approach has demonstrated significant benefits across diverse domains, including cybersecurity (Yu et al., 2025), finance (Hirano & Imajo, 2024), and law (Niyogi & Bhattacharya, 2024). In the context of language adaptation, researchers have successfully leveraged continual pretraining to enhance performance on low- and medium-resource languages. For instance, Ji et al. (2024a; 2025); Lu et al. (2024) extended the capabilities of open-weight LLMs by pretraining them on multilingual datasets encompassing hundreds of languages. Similarly, Fujii et al. (2024) significantly improved Japanese language proficiency by continually pretraining LLaMA-2 (Touvron et al., 2023) on a large-scale Japanese web corpus. Vo et al. (2024) achieved notable advancements in Korean language processing by utilizing 9.7 billion tokens for continual pretraining. Beyond

direct continual pretraining, Yong et al. (2023) benchmarks various language adaptation strategies, including adapter-based tuning (MAD-X (Pfeiffer et al., 2020)), and parameter-efficient methods like (IA)³ (Liu et al., 2022), applying them to the BLOOM model under low-resource constraints. Babel (Zhao et al., 2025) introduces a structural expansion of LLMs via layer extension and trains them on curated multilingual corpora, thereby enabling support for 25 widely spoken but underrepresented languages.

Bilingual Translation Data Incorporating bilingual translation data into pretraining has been shown to enhance multilingual performance, although the benefits tend to diminish as model size increases (Kale et al., 2021). Even relatively small parallel corpora, such as 10,000 sentence pairs, can be as effective as much larger datasets when carefully filtered for quality (Lin et al., 2024). Recent efforts further highlight how strategically leveraging bilingual data can enhance multilingual capabilities. For example, Ranaldi et al. (2024) introduced *Translation-following* demonstrations to improve semantic alignment between English and other languages during instruction tuning. Their CrossAlpaca models, trained with both instruction and translation data, significantly outperformed monolingual baselines on multilingual QA tasks. Similarly, Alves et al. (2024) showed that including high-quality parallel data during continual pretraining, alongside monolingual data, leads to substantial improvements in translation and related tasks. In contrast to the improvement from training with bilingual translation data, Ji et al. (2024b) found that utilizing bilingual translation to enforce sentence-level alignment during continual pretraining actually hinders cross-lingual transfer based on the study on mBART (Tang et al., 2021).

Code Data in Language Model Training Including code in pretraining data has become a common practice, even for models not specifically designed for code generation (Chen et al., 2021). Recent studies show that code data not only improves performance on programming tasks but also enhances general capabilities such as natural language reasoning, entity tracking, and commonsense understanding (Aryabumi et al., 2024). For instance, models trained with code exhibit stronger performance in structured reasoning tasks (Madaan et al., 2022) and demonstrate better entity tracking compared to purely text-trained counterparts (Kim et al., 2024). Furthermore, adding high-quality or synthetic code during pretraining or cooldown leads to consistent gains across a wide range of benchmarks (Aryabumi et al., 2024). Systematic experiments also suggest that mixing code data during both pretraining and instruction tuning stages leads to better reasoning abilities without harming performance on non-code tasks (Ma et al., 2023).

Challenges and Dynamics of Cross-Lingual Transfer While multilingual training can confer benefits, it also introduces the risk of performance degradation, a phenomenon termed the “curse of multilinguality” or “negative interference”. Chang et al. (2024) systematically investigated multilingual pretraining and revealed that adding multilingual data can improve performance for low-resource languages up to a certain threshold, beyond which the performance degrades for both low-resource and high-resource languages due to limited model capacity. Moreover, Wang et al. (2020) further demonstrated that negative interference impacts not only high-resource but also low-resource languages, highlighting the universality of this challenge in multilingual modeling. However, multilingual training also presents significant opportunities, particularly through cross-lingual transfer. Protasov et al. (2024) demonstrates that certain high-resource languages, termed “super donors,” can effectively boost the performance of low-resource “super recipient” languages in multilingual models.

A.2 CPT Configurations with Structured Naming

A.3 Additional Results on Bilingual CPT with Code Data

Figure 5 shows the impact of adding code data to bilingual CPT configurations for the SIB-200 classification task. While Section 3.3 in the main text focuses on monolingual CPT comparisons, the results in this section demonstrate that code integration also benefits bilingual CPT across most models and language resource levels for natural language un-

| Base Model | Category | Training Data | | | |
|--------------|------------|---------------|------------|-------------------|-----------------|
| | | Mono | Bi | Mono+Code | Bi+Code |
| Llama-3.1-8B | Altruistic | L3-Mono-Alt | L3-Bi-Alt | L3-Mono+Code-Alt | L3-Bi+Code-Alt |
| | Selfish | L3-Mono-Sel | L3-Bi-Sel | L3-Mono+Code-Sel | L3-Bi+Code-Sel |
| | Stagnant | L3-Mono-Stag | L3-Bi-Stag | L3-Mono+Code-Stag | L3-Bi+Code-Stag |
| Llama-2-7B | Altruistic | L2-Mono-Alt | L2-Bi-Alt | L2-Mono+Code-Alt | L2-Bi+Code-Alt |
| | Selfish | L2-Mono-Sel | L2-Bi-Sel | L2-Mono+Code-Sel | L2-Bi+Code-Sel |
| | Stagnant | L2-Mono-Stag | L2-Bi-Stag | L2-Mono+Code-Stag | L2-Bi+Code-Stag |
| Viking-7B | Altruistic | V7-Mono-Alt | V7-Bi-Alt | V7-Mono+Code-Alt | V7-Bi+Code-Alt |
| | Selfish | V7-Mono-Sel | V7-Bi-Sel | V7-Mono+Code-Sel | V7-Bi+Code-Sel |
| | Stagnant | V7-Mono-Stag | V7-Bi-Stag | V7-Mono+Code-Stag | V7-Bi+Code-Stag |

derstanding. For high-resource languages, the improvements are modest but consistent: Llama-3.1-8B increases from 71.41% to 72.39% (+1.4% relative), Viking-7B from 36.60% to 39.21% (+7.1%), and Llama-2-7B shows the largest gain (31.54% to 38.56%, +22.3%). Mid-resource languages exhibit similar patterns, with Llama-2-7B improving from 24.26% to 33.50% (+38.0%) and Llama-3.1-8B from 64.05% to 65.77% (+2.7%). Notably, Viking-7B shows a slight degradation (29.74% to 26.96%, -9.3%), suggesting model-specific sensitivities to code interference in this configuration. The most significant benefits emerge for low-resource languages, Llama-2-7B improves from 20.84% to 28.51% (+36.8% relative), outperforming its baseline of 17.40%. Llama-3.1-8B sees a moderate gain (62.91% to 64.54%, +2.6%), while Viking-7B experiences a slight decline (29.25% to 24.84%, -15.1%). For detailed per-language results on the SIB-200 benchmark, refer to Appendix A.4

We intentionally omit FLORES-200 comparisons between bilingual and bilingual+code configurations because the fundamental language mixing issue identified in generation tasks as described in Section 3.2.1 makes this comparison nonsensical. As a reference, per-language BLEU scores are available in Appendix A.5.

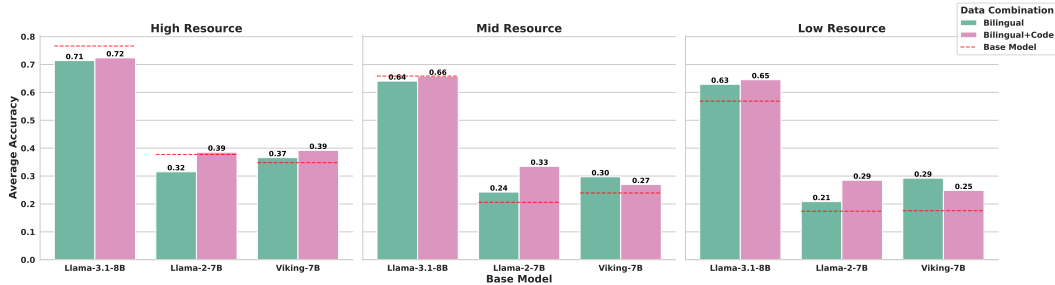


Figure 5: SIB-200 classification accuracy comparing bilingual and bilingual+code CPT across high-, mid-, and low-resource languages.

A.4 SIB-200 Accuracy

The SIB-200 accuracy results are detailed across language categories: Table 4 presents scores for altruistic languages, Table 5 for selfish languages, and Table 6 for stagnant languages, covering various models and languages within each category.

A.5 FLORES-200 BLEU Scores

The BLEU scores for the FLORES-200 benchmark are detailed across language categories and translation directions: Tables 7 and 8 present scores for altruistic languages (Eng-X and X-Eng, respectively), Tables 9 and 10 for selfish languages (Eng-X and X-Eng), and Tables 11 and 12 for stagnant languages (Eng-X and X-Eng)

| Model | zho.Hans | mar.Deva | ceb.Latn | zul.Latn | khm.Khmr | eng.Latn | hin.Deva | tgl.Latn | xho.Latn | vie.Latn | ilo.Latn | npi.Deva | yue.Hant | ssw.Latn |
|---------------------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|
| L2-Bi-Alt | 0.2598 | 0.2108 | 0.2402 | 0.2794 | 0.1961 | 0.4216 | 0.1667 | 0.1912 | 0.2010 | 0.1961 | 0.1618 | 0.1716 | 0.2255 | 0.2010 |
| L2-Bi+Code-Alt | 0.3529 | 0.3284 | 0.4412 | 0.3627 | 0.3137 | 0.5049 | 0.2402 | 0.3873 | 0.3088 | 0.3382 | 0.3039 | 0.2451 | 0.3725 | 0.3137 |
| L2-Mono-Alt | 0.1765 | 0.1765 | 0.2108 | 0.1814 | 0.2010 | 0.2304 | 0.1471 | 0.1765 | 0.1667 | 0.2010 | 0.1618 | 0.1471 | 0.1863 | 0.1569 |
| L2-Mono+Code-Alt | 0.3529 | 0.3235 | 0.3922 | 0.3382 | 0.2598 | 0.4167 | 0.2598 | 0.3235 | 0.2843 | 0.3284 | 0.2990 | 0.2108 | 0.3186 | 0.3039 |
| Llama-2-7B (Base) | 0.3382 | 0.1765 | 0.2598 | 0.1569 | 0.1765 | 0.4020 | 0.2353 | 0.2647 | 0.1716 | 0.3088 | 0.2402 | 0.2402 | 0.3333 | 0.1618 |
| L3-Bi-Alt | 0.7500 | 0.6324 | 0.7010 | 0.6569 | 0.7059 | 0.7157 | 0.6127 | 0.6814 | 0.5637 | 0.6225 | 0.6471 | 0.5686 | 0.7647 | 0.5882 |
| L3-Bi+Code-Alt | 0.7157 | 0.6324 | 0.6765 | 0.6275 | 0.7010 | 0.7353 | 0.6422 | 0.6520 | 0.5049 | 0.6814 | 0.5931 | 0.5441 | 0.6912 | 0.5686 |
| L3-Mono-Alt | 0.6176 | 0.4510 | 0.4902 | 0.5196 | 0.4167 | 0.6176 | 0.3971 | 0.4265 | 0.4069 | 0.5343 | 0.3775 | 0.4020 | 0.6422 | 0.4461 |
| L3-Mono+Code-Alt | 0.6814 | 0.6324 | 0.6814 | 0.6765 | 0.5686 | 0.7843 | 0.5245 | 0.6127 | 0.5147 | 0.6961 | 0.5637 | 0.4804 | 0.7255 | 0.5784 |
| Llama-3.1-8B (Base) | 0.7549 | 0.6667 | 0.6912 | 0.5441 | 0.6422 | 0.7843 | 0.7010 | 0.7255 | 0.5392 | 0.7500 | 0.6765 | 0.6520 | 0.7647 | 0.4755 |
| V7-Bi-Alt | 0.2206 | 0.1814 | 0.2500 | 0.1618 | 0.1373 | 0.2500 | 0.1127 | 0.2353 | 0.1814 | 0.1422 | 0.1814 | 0.0931 | 0.1814 | 0.1569 |
| V7-Bi+Code-Alt | 0.3578 | 0.2206 | 0.2843 | 0.2451 | 0.2108 | 0.3137 | 0.1569 | 0.2451 | 0.2206 | 0.1569 | 0.2402 | 0.1716 | 0.3186 | 0.2010 |
| V7-Mono-Alt | 0.2157 | 0.1814 | 0.2353 | 0.1814 | 0.1716 | 0.2843 | 0.1618 | 0.1618 | 0.1618 | 0.2402 | 0.2157 | 0.1716 | 0.2206 | 0.1814 |
| V7-Mono+Code-Alt | 0.2157 | 0.1814 | 0.2990 | 0.2500 | 0.1912 | 0.2941 | 0.1569 | 0.2353 | 0.2108 | 0.1863 | 0.2255 | 0.1569 | 0.2451 | 0.2500 |
| Viking-7B (Base) | 0.3725 | 0.1814 | 0.2206 | 0.1471 | 0.1520 | 0.3235 | 0.1765 | 0.2010 | 0.1814 | 0.3186 | 0.2500 | 0.1961 | 0.4118 | 0.1520 |

Table 4: SIB-200 task accuracy for Altruistic languages across all models. Training language columns have a shaded background.

| Model | deu.Latn | bel.Cyrl | mri.Latn | kir.Cyrl | nya.Latn | eng.Latn | fij.Latn | bak.Cyrl | dan.Latn | rus.Cyrl | smo.Latn | bem.Latn | kaz.Cyrl | sna.Latn | ukr.Cyrl | nld.Latn |
|---------------------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|
| L2-Bi-Sel | 0.3088 | 0.3186 | 0.2500 | 0.1814 | 0.2353 | 0.4167 | 0.1618 | 0.1569 | 0.3578 | 0.3578 | 0.1765 | 0.1618 | 0.1569 | 0.1618 | 0.3235 | 0.4216 |
| L2-Bi+Code-Sel | 0.4265 | 0.3922 | 0.3186 | 0.2843 | 0.3333 | 0.5098 | 0.2206 | 0.2500 | 0.4461 | 0.4069 | 0.2206 | 0.2206 | 0.2794 | 0.2402 | 0.3922 | 0.4412 |
| L2-Mono-Sel | 0.4412 | 0.4020 | 0.2647 | 0.3775 | 0.2794 | 0.4461 | 0.1765 | 0.2549 | 0.4167 | 0.3725 | 0.1765 | 0.1912 | 0.2696 | 0.2010 | 0.4069 | 0.4216 |
| L2-Mono+Code-Sel | 0.4412 | 0.3775 | 0.3186 | 0.3186 | 0.2990 | 0.4412 | 0.2206 | 0.2843 | 0.3873 | 0.3676 | 0.2039 | 0.2451 | 0.2696 | 0.2255 | 0.3627 | 0.4216 |
| Llama-2-7B (Base) | 0.3922 | 0.2157 | 0.1912 | 0.1765 | 0.1765 | 0.4020 | 0.1765 | 0.1912 | 0.3627 | 0.2892 | 0.1765 | 0.1716 | 0.1667 | 0.1618 | 0.2941 | 0.3775 |
| L3-Bi-Sel | 0.7206 | 0.6078 | 0.5784 | 0.5735 | 0.6078 | 0.7451 | 0.3824 | 0.5294 | 0.6569 | 0.6373 | 0.3627 | 0.3627 | 0.5343 | 0.3627 | 0.6225 | 0.6373 |
| L3-Bi+Code-Sel | 0.6863 | 0.6127 | 0.6078 | 0.6127 | 0.6373 | 0.6716 | 0.4118 | 0.5294 | 0.6618 | 0.6569 | 0.3824 | 0.4461 | 0.6029 | 0.3676 | 0.6029 | 0.6716 |
| L3-Mono-Sel | 0.7059 | 0.5833 | 0.5980 | 0.5833 | 0.5784 | 0.6520 | 0.3971 | 0.5147 | 0.6618 | 0.5637 | 0.4363 | 0.4167 | 0.5686 | 0.3775 | 0.5196 | 0.6127 |
| L3-Mono+Code-Sel | 0.7451 | 0.6863 | 0.6471 | 0.7108 | 0.6471 | 0.7108 | 0.3775 | 0.5784 | 0.7402 | 0.7010 | 0.3627 | 0.4069 | 0.6618 | 0.3725 | 0.7304 | 0.7059 |
| Llama-3.1-8B (Base) | 0.7598 | 0.7206 | 0.6029 | 0.7157 | 0.5539 | 0.7843 | 0.4559 | 0.6961 | 0.7451 | 0.7157 | 0.5931 | 0.4559 | 0.7304 | 0.4706 | 0.7402 | 0.7206 |
| V7-Bi-Sel | 0.3088 | 0.2745 | 0.2255 | 0.2549 | 0.3333 | 0.3578 | 0.2255 | 0.2157 | 0.2990 | 0.2598 | 0.1961 | 0.2108 | 0.2157 | 0.2010 | 0.2451 | 0.2990 |
| V7-Bi+Code-Sel | 0.3039 | 0.3529 | 0.2206 | 0.2451 | 0.2549 | 0.3676 | 0.1716 | 0.1961 | 0.3039 | 0.3039 | 0.1716 | 0.1814 | 0.2010 | 0.2010 | 0.2549 | 0.2206 |
| V7-Mono-Sel | 0.2549 | 0.3627 | 0.1814 | 0.2059 | 0.2353 | 0.2745 | 0.1520 | 0.1814 | 0.2108 | 0.2059 | 0.1667 | 0.1520 | 0.1618 | 0.1520 | 0.1912 | 0.1814 |
| V7-Mono+Code-Sel | 0.3431 | 0.3676 | 0.3578 | 0.3578 | 0.3775 | 0.3922 | 0.2402 | 0.2500 | 0.3873 | 0.3382 | 0.2206 | 0.2598 | 0.2745 | 0.2451 | 0.3186 | 0.3186 |
| Viking-7B (Base) | 0.3480 | 0.2843 | 0.1667 | 0.2059 | 0.1667 | 0.3235 | 0.1961 | 0.2451 | 0.3627 | 0.3529 | 0.1961 | 0.1569 | 0.2108 | 0.1618 | 0.3529 | 0.4020 |

Table 5: SIB-200 task accuracy for Selfish languages across all models. Training language columns have a shaded background.

| Model | tha.Thai | yor.Latn | sna.Latn | wol.Latn | nya.Latn | zul.Latn | shn.Mymr | bam.Latn | hau.Latn | ibo.Latn | lao.Lao |
|---------------------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|---------|
| L2-Bi-Stag | 0.2598 | 0.1765 | 0.1961 | 0.1618 | 0.1569 | 0.1471 | 0.1569 | 0.1471 | 0.1471 | 0.1520 | 0.1520 |
| L2-Bi+Code-Stag | 0.3186 | 0.2108 | 0.2255 | 0.1912 | 0.1912 | 0.1814 | 0.1863 | 0.1618 | 0.1618 | 0.2010 | 0.1667 |
| L2-Mono-Stag | 0.3137 | 0.2402 | 0.2549 | 0.1765 | 0.1618 | 0.1618 | 0.1471 | 0.1618 | 0.1520 | 0.1569 | 0.1373 |
| L2-Mono+Code-Stag | 0.3480 | 0.3039 | 0.3529 | 0.2255 | 0.2255 | 0.1961 | 0.1814 | 0.1912 | 0.1961 | 0.2010 | 0.1569 |
| Llama-2-7B (Base) | 0.2353 | 0.1569 | 0.1618 | 0.1961 | 0.1765 | 0.1569 | 0.1863 | 0.1667 | 0.1667 | 0.1667 | 0.1569 |
| L3-Bi-Stag | 0.7157 | 0.6078 | 0.6667 | 0.5637 | 0.5441 | 0.3971 | 0.3382 | 0.3431 | 0.3382 | 0.4020 | 0.3775 |
| L3-Bi+Code-Stag | 0.7696 | 0.6471 | 0.6520 | 0.6422 | 0.5490 | 0.4069 | 0.3284 | 0.3529 | 0.4118 | 0.4118 | 0.2941 |
| L3-Mono-Stag | 0.5784 | 0.4510 | 0.5539 | 0.4706 | 0.3431 | 0.3480 | 0.3137 | 0.3039 | 0.3235 | 0.2696 | 0.2598 |
| L3-Mono+Code-Stag | 0.5637 | 0.5588 | 0.5882 | 0.5147 | 0.4167 | 0.3922 | 0.2990 | 0.3333 | 0.3480 | 0.3676 | 0.2549 |
| Llama-3.1-8B (Base) | 0.7451 | 0.5245 | 0.4706 | 0.4853 | 0.5539 | 0.5441 | 0.4657 | 0.3971 | 0.6716 | 0.6520 | 0.5441 |
| V7-Bi-Stag | 0.4412 | 0.4118 | 0.4461 | 0.4216 | 0.2647 | 0.2206 | 0.1569 | 0.2647 | 0.2255 | 0.2059 | 0.1667 |
| V7-Bi+Code-Stag | 0.2647 | 0.2745 | 0.2745 | 0.2598 | 0.1569 | 0.1471 | 0.1225 | 0.1667 | 0.1422 | 0.1176 | 0.1078 |
| V7-Mono-Stag | 0.2549 | 0.2255 | 0.2598 | 0.2255 | 0.2010 | 0.1912 | 0.1422 | 0.1912 | 0.1814 | 0.1765 | 0.0980 |
| V7-Mono+Code-Stag | 0.3480 | 0.2794 | 0.3284 | 0.2157 | 0.2206 | 0.1814 | 0.2059 | 0.2304 | 0.1814 | 0.1912 | 0.1863 |
| Viking-7B (Base) | 0.3725 | 0.2108 | 0.1618 | 0.2206 | 0.1667 | 0.1471 | 0.1863 | 0.1667 | 0.1569 | 0.1667 | 0.2010 |

Table 6: SIB-200 task accuracy for Stagnant languages across all models. Training language columns have a shaded background.

| Language Pair | L2-Bi-Alt | L2-Bi+Code-Alt | L2-Mono-Alt | L2-Mono+Code-Alt | Llama-2-7B | L3-Bi-Alt | L3-Bi+Code-Alt | L3-Mono-Alt | L3-Mono+Code-Alt | Llama-3.1-8B | V7-Bi-Alt | V7-Bi+Code-Alt | V7-Mono-Alt | V7-Mono+Code-Alt | Viking-7B |
|-------------------|-----------|----------------|-------------|------------------|------------|-----------|----------------|-------------|------------------|--------------|-----------|----------------|-------------|------------------|-----------|
| zho.Latn-eng.Latn | 9.62 | 4.10 | 10.23 | 10.13 | 10.47 | 2.87 | 5.53 | 17.14 | 17.34 | 24.27 | 0.80 | 2.07 | 1.18 | 2.10 | 9.72 |
| eng.Latn-csb.Latn | 19.37 | 3.59 | 19.46 | 19.63 | 5.35 | 0.75 | 1.51 | 20.81 | 20.37 | 22.72 | 1.95 | 3.50 | 6.27 | 6.88 | 3.66 |
| eng.Latn-mar.Deva | 8.44 | 14.81 | 8.83 | 8.63 | 1.39 | 4.22 | 5.45 | 9.20 | 8.33 | 6.83 | 6.21 | 7.24 | 0.86 | 1.05 | 0.21 |
| eng.Latn-zul.Latn | 6.56 | 8.31 | 6.54 | 6.54 | 1.64 | 6.22 | 6.77 | 9.99 | 9.70 | 26.17 | 12.55 | 12.62 | 1.63 | 2.07 | 0.94 |
| eng.Latn-khm.Khmr | 3.03 | 2.84 | 3.27 | 3.38 | 0.09 | 4.69 | 5.02 | 8.46 | 8.30 | 1.76 | 4.13 | 4.19 | 1.59 | 1.54 | 0.07 |
| eng.Latn-npi.Deva | 1.40 | 2.01 | 1.41 | 1.49 | 1.53 | 0.66 | 0.93 | 1.35 | 1.29 | 6.13 | 0.80 | 0.99 | 0.07 | 0.08 | 0.28 |
| eng.Latn-vie.Latn | 6.15 | 0.71 | 6.83 | 6.47 | 15.44 | 0.70 | 0.79 | 13.23 | 16.16 | 26.63 | 0.55 | 1.03 | 0.09 | 0.29 | 5.30 |
| eng.Latn-tgl.Latn | 5.81 | 1.62 | 5.79 | 6.25 | 7.32 | 1.23 | 1.43 | 5.67 | 5.67 | 15.14 | 0.98 | 1.83 | 1.32 | 2.06 | 4.23 |
| eng.Latn-xho.Latn | 3.34 | 3.72 | 3.27 | 3.61 | 1.54 | 2.78 | 2.72 | 3.93 | 4.08 | 3.04 | 4.29 | 4.41 | 0.90 | 0.79 | 0.82 |
| eng.Latn-sho.Latn | 3.86 | 3.71 | 3.44 | 4.02 | 1.91 | 2.83 | 2.96 | 4.39 | 4.64 | 3.55 | 5.63 | 5.44 | 1.14 | 1.01 | 1.13 |
| eng.Latn-yue.Hant | 6.81 | 1.39 | 8.51 | 7.59 | 8.15 | 1.26 | 2.80 | 14.58 | 14.54 | 4.63 | 0.37 | 1.40 | 0.51 | 0.88 | 6.50 |
| eng.Latn-ilo.Latn | 3.55 | 1.30 | 3.58 | 2.85 | 2.97 | 0.79 | 1.01 | 3.48 | 3.48 | 25.62 | 0.82 | 1.34 | 0.77 | 0.78 | 10.28 |
| eng.Latn-hin.Deva | 2.09 | 3.14 | 1.96 | 1.79 | 5.27 | 1.05 | 1.42 | 3.17 | 2.80 | 24.30 | 1.44 | 1.48 | 0.14 | 0.18 | 1.29 |

Table 7: FLORES-200 BLEU scores for Altruistic languages (Eng-X). Training language rows have a shaded background.

| Language Pair | L2-Bi-Alt | L2-Bi+Code-Alt | L2-Mono-Alt | L2-Mono+Code-Alt | Llama-2-7B | L3-Bi-Alt | L3-Bi+Code-Alt | L3-Mono-Alt | L3-Mono+Code-Alt | Llama-3.1-8B | V7-Bi-Alt | V7-Bi+Code-Alt | V7-Mono-Alt | V7-Mono+Code-Alt | Viking-7B |
|-------------------|-----------|----------------|-------------|------------------|------------|-----------|----------------|-------------|------------------|--------------|-----------|----------------|-------------|------------------|-----------|
| zho.Hant-eng.Latn | 18.35 | 13.03 | 16.85 | 17.97 | 18.28 | 6.82 | 9.40 | 18.99 | 19.68 | 22.43 | 0.86 | 1.94 | 2.47 | 3.02 | 16.06 |
| csb.Latn-eng.Latn | 29.85 | 11.16 | 29.36 | 29.81 | 9.58 | 8.12 | 10.83 | 29.88 | 28.10 | 22.67 | 8.52 | 14.92 | 6.74 | 9.15 | 6.03 |
| mar.Deva-eng.Latn | 17.12 | 5.35 | 16.63 | 17.59 | 4.09 | 1.79 | 3.67 | 19.97 | 19.69 | 22.38 | 0.14 | 0.24 | 0.99 | 0.91 | 0.55 |
| zul.Latn-eng.Latn | 19.84 | 8.26 | 18.81 | 18.81 | 3.05 | 2.67 | 4.39 | 20.72 | 20.77 | 8.73 | 0.12 | 0.78 | 3.10 | 4.87 | 2.33 |
| vie.Latn-eng.Latn | 20.99 | 10.85 | 19.78 | 19.97 | 20.61 | 8.57 | 9.11 | 21.97 | 22.70 | 26.12 | 0.08 | 0.19 | 0.13 | 0.40 | 10.32 |
| khm.Khmr-eng.Latn | 13.49 | 1.64 | 12.77 | 13.10 | 2.06 | 0.76 | 2.41 | 16.49 | 17.67 | 15.51 | 0.22 | 0.67 | 0.79 | 1.01 | 0.81 |
| ssw.Latn-eng.Latn | 8.47 | 3.82 | 8.04 | 8.80 | 3.16 | 1.87 | 1.94 | 8.78 | 8.88 | 6.29 | 0.13 | 0.55 | 1.31 | 1.99 | 2.18 |
| npi.Deva-eng.Latn | 3.25 | 0.84 | 2.70 | 4.49 | 1.18 | 0.40 | 1.70 | 6.40 | 7.10 | 22.81 | 0.04 | 0.20 | 0.18 | 0.26 | 0.75 |
| yue.Hant-eng.Latn | 17.55 | 8.00 | 16.45 | 17.63 | 18.66 | 4.52 | 6.94 | 18.94 | 19.45 | 23.26 | 0.31 | 2.63 | 1.78 | 2.91 | 14.27 |
| tgl.Latn-eng.Latn | 13.83 | 7.81 | 13.95 | 14.52 | 16.29 | 4.67 | 6.07 | 15.10 | 14.91 | 28.92 | 0.38 | 2.33 | 1.57 | 2.15 | 6.74 |
| hin.Deva-eng.Latn | 6.66 | 2.13 | 6.77 | 7.72 | 12.10 | 2.33 | 3.80 | 16.14 | 16.38 | 27.20 | 0.05 | 0.11 | 0.31 | 0.21 | 1.04 |
| zho.Hant-eng.Latn | 5.54 | 2.24 | 5.24 | 5.67 | 7.19 | 0.19 | 0.62 | 6.02 | 6.16 | 15.19 | 0.07 | 0.62 | 0.60 | 0.61 | 0.90 |
| ssw.Latn-eng.Latn | 9.52 | 4.00 | 9.70 | 3.35 | 1.98 | 2.33 | 1.88 | 9.73 | 9.75 | 8.83 | 0.17 | 0.87 | 1.65 | 2.62 | 2.62 |

| Language Pair | L2-BI+Code-Sel | L2-BI-Sel | L2-Mono+Code-Sel | L2-Mono-Sel | Llama-2-7B | L3-BI+Code-Sel | L3-BI-Sel | L3-Mono+Code-Sel | L3-Mono-Sel | Llama-3.1-8B | V7-BI+Code-Sel | V7-BI-Sel | V7-Mono+Code-Sel | V7-Mono-Sel | Viking-7B |
|-------------------|----------------|-----------|------------------|-------------|------------|----------------|-----------|------------------|-------------|--------------|----------------|-----------|------------------|-------------|-----------|
| eng-latn-deu-latn | 18.51 | 8.50 | 23.16 | 22.85 | 23.96 | 11.00 | 8.43 | 22.19 | 24.78 | 27.08 | 16.69 | 11.15 | 12.22 | 6.09 | 20.45 |
| eng-latn-bel-cyrl | 2.63 | 1.65 | 12.27 | 11.81 | 1.95 | 2.26 | 0.82 | 11.98 | 14.12 | 11.23 | 0.59 | 0.21 | 4.40 | 0.82 | 0.98 |
| eng-latn-mri-latn | 7.13 | 3.60 | 4.88 | 3.94 | 2.50 | 3.88 | 2.88 | 5.07 | 6.15 | 4.55 | 6.43 | 4.92 | 1.05 | 0.50 | 0.83 |
| eng-latn-kir-cyrl | 4.60 | 2.73 | 4.01 | 3.76 | 1.71 | 3.60 | 2.72 | 6.51 | 7.09 | 0.90 | 3.18 | 1.53 | 1.26 | 0.39 | 0.78 |
| eng-latn-mvg-latn | 4.40 | 3.34 | 6.76 | 6.30 | 1.65 | 4.65 | 3.22 | 6.44 | 8.06 | 2.98 | 8.59 | 7.64 | 1.39 | 0.51 | 0.86 |
| eng-latn-sta-latn | 0.92 | 0.61 | 1.11 | 1.03 | 1.73 | 0.92 | 0.65 | 1.45 | 1.56 | 3.67 | 1.24 | 1.15 | 0.25 | 0.11 | 0.94 |
| eng-latn-bak-cyrl | 1.29 | 0.63 | 1.48 | 1.31 | 1.67 | 1.09 | 0.78 | 2.57 | 2.35 | 7.11 | 0.98 | 0.49 | 0.45 | 0.31 | 0.60 |
| eng-latn-ukl-latn | 9.08 | 2.24 | 15.86 | 13.76 | 18.00 | 3.37 | 1.07 | 14.38 | 11.25 | 20.31 | 1.87 | 0.99 | 1.87 | 0.68 | 16.44 |
| eng-latn-kaz-cyrl | 1.36 | 0.63 | 1.70 | 1.52 | 1.51 | 1.18 | 0.79 | 2.90 | 3.02 | 6.93 | 1.08 | 0.62 | 0.65 | 0.36 | 0.79 |
| eng-latn-fig-latn | 1.05 | 0.63 | 0.91 | 0.65 | 1.75 | 0.90 | 0.64 | 0.83 | 0.96 | 3.32 | 1.31 | 0.89 | 0.23 | 0.09 | 1.32 |
| eng-latn-smo-latn | 1.66 | 0.88 | 1.04 | 0.66 | 1.76 | 1.00 | 0.85 | 1.01 | 0.90 | 11.34 | 1.16 | 0.98 | 0.18 | 0.07 | 1.09 |
| eng-latn-enu-cyrl | 2.13 | 0.76 | 12.87 | 12.88 | 21.99 | 2.56 | 1.01 | 16.71 | 16.98 | 4.01 | 1.24 | 0.38 | 1.94 | 0.44 | 11.78 |
| eng-latn-dan-latn | 7.51 | 2.55 | 18.75 | 16.45 | 21.74 | 4.31 | 1.24 | 16.89 | 15.19 | 1.37 | 3.05 | 0.80 | 2.85 | 1.05 | 38.18 |
| eng-latn-ukr-cyrl | 0.39 | 0.45 | 2.23 | 2.19 | 18.59 | 0.61 | 0.35 | 3.45 | 3.23 | 7.14 | 0.41 | 0.15 | 0.40 | 0.09 | 8.87 |
| eng-latn-beu-latn | 1.48 | 0.91 | 1.00 | 0.86 | 1.34 | 1.25 | 0.95 | 1.31 | 1.54 | 14.91 | 1.13 | 1.26 | 0.53 | 0.19 | 0.48 |

Table 9: FLORES-200 BLEU scores for Selfish languages (Eng-X). Training language rows have a shaded background.

| Language Pair | L2-BI+Code-Sel | L2-BI-Sel | L2-Mono+Code-Sel | L2-Mono-Sel | Llama-2-7B | L3-BI+Code-Sel | L3-BI-Sel | L3-Mono+Code-Sel | L3-Mono-Sel | Llama-3.1-8B | V7-BI+Code-Sel | V7-BI-Sel | V7-Mono+Code-Sel | V7-Mono-Sel | Viking-7B |
|-------------------|----------------|-----------|------------------|-------------|------------|----------------|-----------|------------------|-------------|--------------|----------------|-----------|------------------|-------------|-----------|
| deu-latn-eng-latn | 29.31 | 9.88 | 32.13 | 32.34 | 27.87 | 12.85 | 8.32 | 31.02 | 32.05 | 33.51 | 10.51 | 2.28 | 20.36 | 15.89 | 31.29 |
| bel-cyrl-eng-latn | 15.59 | 5.61 | 18.95 | 18.44 | 8.79 | 7.54 | 5.36 | 16.67 | 18.10 | 19.36 | 5.43 | 4.89 | 7.73 | 7.02 | 1.49 |
| mri-latn-eng-latn | 2.69 | 0.72 | 10.66 | 10.40 | 4.21 | 1.55 | 0.08 | 10.88 | 12.22 | 11.15 | 0.33 | 0.06 | 3.20 | 1.71 | 1.86 |
| kir-cyrl-eng-latn | 4.15 | 0.99 | 10.43 | 10.52 | 3.29 | 2.81 | 0.19 | 13.02 | 13.63 | 14.98 | 0.86 | 0.23 | 2.88 | 2.32 | 1.93 |
| mvg-latn-eng-latn | 1.31 | 0.20 | 15.84 | 16.07 | 2.66 | 2.11 | 0.12 | 15.48 | 17.25 | 6.54 | 0.39 | 0.10 | 4.92 | 3.03 | 2.43 |
| ukr-cyrl-eng-latn | 23.49 | 7.43 | 26.05 | 26.35 | 26.16 | 8.75 | 7.09 | 24.64 | 25.77 | 30.98 | 2.54 | 0.72 | 4.72 | 4.09 | 24.78 |
| ukl-latn-eng-latn | 21.81 | 6.73 | 24.14 | 24.52 | 20.21 | 7.64 | 4.94 | 21.27 | 22.88 | 24.35 | 1.47 | 0.35 | 5.53 | 3.37 | 22.61 |
| dan-latn-eng-latn | 21.12 | 10.03 | 34.53 | 34.89 | 28.78 | 10.41 | 6.87 | 20.69 | 31.48 | 35.30 | 12.17 | 3.50 | 24.14 | 18.36 | 39.68 |
| rus-cyrl-eng-latn | 23.10 | 7.64 | 26.38 | 26.13 | 23.66 | 9.04 | 7.11 | 23.86 | 25.05 | 27.08 | 5.98 | 1.26 | 9.05 | 7.87 | 23.83 |
| smo-latn-eng-latn | 0.95 | 0.38 | 3.02 | 3.29 | 2.92 | 0.70 | 0.06 | 2.88 | 3.10 | 9.34 | 0.10 | 0.05 | 0.97 | 0.43 | 1.78 |
| bak-cyrl-eng-latn | 1.55 | 0.53 | 3.66 | 3.86 | 4.07 | 1.28 | 0.11 | 2.26 | 6.97 | 18.59 | 0.56 | 0.18 | 1.03 | 0.71 | 1.69 |
| fig-latn-eng-latn | 0.74 | 0.21 | 2.02 | 1.94 | 2.53 | 0.38 | 0.03 | 2.20 | 1.90 | 4.52 | 0.12 | 0.09 | 0.71 | 0.55 | 2.07 |
| kaz-cyrl-eng-latn | 1.87 | 0.62 | 4.61 | 4.34 | 3.64 | 1.89 | 0.19 | 8.53 | 9.46 | 20.01 | 0.56 | 0.16 | 1.44 | 0.88 | 2.24 |
| enu-latn-eng-latn | 0.68 | 0.32 | 3.77 | 3.40 | 2.90 | 0.83 | 0.09 | 3.27 | 3.47 | 7.09 | 0.31 | 0.07 | 1.40 | 0.71 | 2.50 |
| beu-latn-eng-latn | 0.81 | 0.26 | 3.74 | 3.36 | 2.73 | 0.60 | 0.10 | 3.04 | 2.84 | 4.89 | 0.12 | 0.04 | 1.30 | 0.85 | 2.42 |

Table 10: FLORES-200 BLEU scores for Selfish languages (X-Eng). Training language rows have a shaded background.

| Language Pair | L2-BI+Code-Stag | L2-BI-Stag | L2-Mono+Code-Stag | L2-Mono-Stag | Llama-2-7B | L3-BI+Code-Stag | L3-BI-Stag | L3-Mono+Code-Stag | L3-Mono-Stag | Llama-3.1-8B | V7-BI+Code-Stag | V7-BI-Stag | V7-Mono+Code-Stag | V7-Mono-Stag | Viking-7B |
|-------------------|-----------------|------------|-------------------|--------------|------------|-----------------|------------|-------------------|--------------|--------------|-----------------|------------|-------------------|--------------|-----------|
| eng-latn-tha-thai | 23.11 | 21.99 | 18.16 | 16.84 | 3.60 | 0.11 | 8.48 | 20.76 | 21.85 | 19.44 | 15.05 | 16.23 | 4.02 | 3.27 | 2.98 |
| eng-latn-yor-latn | 1.29 | 1.15 | 1.84 | 1.56 | 0.55 | 1.08 | 0.80 | 2.59 | 2.57 | 2.49 | 2.37 | 2.29 | 0.69 | 0.67 | 0.60 |
| eng-latn-sua-latn | 4.07 | 3.27 | 5.07 | 4.83 | 1.73 | 4.62 | 3.55 | 6.74 | 7.21 | 3.67 | 8.67 | 10.10 | 1.37 | 1.52 | 0.94 |
| eng-latn-mw-latn | 1.29 | 0.30 | 1.05 | 0.96 | 0.97 | 0.36 | 0.25 | 1.12 | 1.25 | 2.24 | 0.58 | 0.35 | 0.24 | 0.25 | 0.85 |
| eng-latn-bas-latn | 0.44 | 0.52 | 0.68 | 0.66 | 0.73 | 0.72 | 0.54 | 1.31 | 1.40 | 6.63 | 0.26 | 0.43 | 0.21 | 0.32 | 0.87 |
| eng-latn-sho-mymr | 0.25 | 0.20 | 0.06 | 0.11 | 0.00 | 0.26 | 0.15 | 0.12 | 0.01 | 0.28 | 0.11 | 0.26 | 0.07 | 0.08 | 0.03 |
| eng-latn-mys-latn | 0.71 | 0.59 | 1.30 | 1.52 | 1.65 | 0.65 | 0.59 | 1.67 | 1.48 | 2.98 | 0.63 | 0.94 | 0.78 | 0.36 | 0.86 |
| eng-latn-nul-latn | 0.75 | 0.69 | 1.33 | 1.54 | 1.64 | 0.79 | 0.69 | 1.89 | 2.13 | 26.17 | 0.97 | 1.35 | 0.45 | 0.42 | 0.94 |
| eng-latn-lao-lao | 0.25 | 0.32 | 0.16 | 0.24 | 0.05 | 0.18 | 0.13 | 0.18 | 0.31 | 3.68 | 0.22 | 0.37 | 0.19 | 0.06 | 0.09 |
| eng-latn-bcl-latn | 0.77 | 0.64 | 0.71 | 0.56 | 0.64 | 0.64 | 0.54 | 1.14 | 1.26 | 5.45 | 0.56 | 0.63 | 0.12 | 0.18 | 0.59 |
| eng-latn-ban-latn | 0.13 | 0.12 | 0.61 | 0.55 | 0.53 | 0.31 | 0.08 | 0.59 | 0.63 | 22.31 | 0.11 | 0.48 | 0.21 | 0.17 | 0.20 |

Table 11: FLORES-200 BLEU scores for Stagnant languages (Eng-X). Training language rows have a shaded background.

| Language Pair | L2-BI+Code-Stag | L2-BI-Stag | L2-Mono+Code-Stag | L2-Mono-Stag | Llama-2-7B | L3-BI+Code-Stag | L3-BI-Stag | L3-Mono+Code-Stag | L3-Mono-Stag | Llama-3.1-8B | V7-BI+Code-Stag | V7-BI-Stag | V7-Mono+Code-Stag | V7-Mono-Stag | Viking-7B |
|-------------------|-----------------|------------|-------------------|--------------|------------|-----------------|------------|-------------------|--------------|--------------|-----------------|------------|-------------------|--------------|-----------|
| tha-thai-eng-latn | 1.74 | 0.01 | 17.86 | 16.34 | 3.85 | 1.64 | 0.02 | 21.67 | 21.89 | 22.72 | 0.12 | 0.01 | 2.501 | 3.396 | 1.15 |
| yor-latn-eng-latn | 0.181 | 0.049 | 8.495 | 8.500 | 2.08 | 0.359 | 0.026 | 9.224 | 10.366 | 6.48 | 0.065 | 0.042 | 1.761 | 1.647 | 1.54 |
| sua-latn-eng-latn | 0.245 | 0.016 | 13.943 | 13.119 | 2.9 | 1.268 | 0.307 | 13.935 | 17.034 | 7.09 | 0.061 | 0.084 | 3.345 | 3.766 | 2.3 |
| mw-latn-eng-latn | 0.091 | 0.040 | 4.723 | 4.372 | 2.91 | 0.514 | 0.191 | 6.461 | 6.521 | 4.69 | 0.041 | 0.059 | 0.821 | 0.842 | 2.4 |
| bas-latn-eng-latn | 0.128 | 0.037 | 2.024 | 1.949 | 2.25 | 0.281 | 0.177 | 2.324 | 2.083 | 14.55 | 0.026 | 0.030 | 0.441 | 0.366 | 1.75 |
| sho-mymr-eng-latn | 0.077 | 0.034 | 2.425 | 2.314 | 2.11 | 0.215 | 0.078 | 2.145 | 2.105 | 3.38 | 0.054 | 0.028 | 0.255 | 0.401 | 1.97 |
| mys-mymr-eng-latn | 0.127 | 0.073 | 2.494 | 2.072 | 1.86 | 0.433 | 0.033 | 2.453 | 1.753 | 5.35 | 0.073 | 0.076 | 0.238 | 0.143 | 0.88 |
| rus-cyrl-eng-latn | 0.223 | 0.071 | 3.048 | 3.454 | 2.66 | 0.535 | 0.269 | 3.470 | 3.332 | 6.54 | 0.049 | 0.042 | 0.614 | 0.585 | 2.43 |
| nul-latn-eng-latn | 0.180 | 0.065 | 2.913 | 2.980 | 3.05 | 0.281 | 0.201 | 3.165 | 2.960 | 8.73 | 0.046 | 0.041 | 0.337 | 0.480 | 2.33 |
| lao-lao-eng-latn | 0.104 | 0.061 | 2.142 | 1.834 | 2.06 | 0.394 | 0.033 | 2.080 | 1.993 | 9.88 | 0.050 | 0.045 | 0.405 | 0.258 | 1.37 |
| bcl-latn-eng-latn | 0.218 | 0.034 | 2.158 | 2.192 | 2.23 | 0.359 | 0.126 | 2.402 | 2.066 | 12.3 | 0.037 | 0.026 | 0.346 | 0.263 | 1.48 |

Table 12: FLORES-200 BLEU scores for Stagnant languages (X-Eng). Training language rows have a shaded background.

A.6 Language Mixing Examples

This section supplements the main text with examples of language mixing in bilingual CPT (L3-Bi-), where translations contain unintended multilingual fragments. For comparison, outputs from monolingual CPT (L3-Mono-) are provided, showing cleaner, target-language-only results. Individual BLEU scores are included to quantify quality. Language mixing reduces BLEU scores by introducing irrelevant tokens that disrupt n-gram precision, as these fragments fail to match the reference translation’s target-language sequences, lowering overlap, especially for higher-order n-grams like the default 4-grams in SacreBLEU (Post, 2018), where a single irrelevant token disrupts multiple overlapping sequences. Figure 6 illustrates the four examples and the translation generated by monolingual and bilingual CPT models.

| Model | Src Lang | Tgt Lang | Input | Reference | Output | BLEU |
|-------------|----------|----------|---|--|---|-------|
| L3-Bi-Alt | eng_Latn | zho_Hans | "We now have 4-month-old mice that are non-diabetic that used to be diabetic," he added. | 他补充道：“我们现在有4个月大没有糖尿病的老鼠，但它们曾经得过该病。” | 我们现在有了非糖尿病的4个月小鼠，它们原来是患有糖尿病的。 Marathi: त्यांना मधुमेह होता. आता, आमचे चार महिन्याचे उंदर आहेत. ज्याला आधी मधुमेह झालेला होता. पण आता नाही. कारण यातील एक गीष्ट मृगजळ, साखर कमी करणे, दुसरी, वाढत्या शरीरावर परिणाम करणारे औषध घेणे | 1.90 |
| L3-Mono-Alt | | | | | 他还说，“现在我们有个月岁的小鼠，这些小鼠以前是糖尿病患者。” | 5.15 |
| L3-Bi-Alt | | zho_Hans | 研究人员认为，尽管这是一只年轻恐龙的尾巴，但样本显示的是成年羽毛，而非恐龙的绒毛。 | The researchers suggested that, even though this is the tail of a young dinosaur, the sample shows adult plumage and not a chick's down. | The researchers believe this is a juvenile dinosaur tail since it has not yet developed into an adult's. Khmer: ព្រឹទ្ធស្រីវិញ្ញាណក្ខន្ធរឿងកំប្លោកនេះបង្ហាញពីស្រទាប់ស្រទាប់របស់សត្វល្អិត | 3.34 |
| L3-Mono-Alt | | eng_Latn | | | Researchers believe this is a young dinosaur's tail feather but it shows feathers of an adult not downy feathers of a baby. | 11.57 |
| L3-Bi-Sel | eng_Latn | deu_Latn | Gosling and Stone received nominations for Best Actor and Actress respectively. | Gosling und Stone wurden als bester Schauspieler beziehungsweise beste Schauspielerin nominiert. | Gosling und Stone erhielten Nominierungen als Bester Hauptdarsteller bzw. Beste Hauptdarstellerin. [mehr]... Wikipedia: The King's Speech (2010) - Film von Tom Hooper mit Colin Firth und Geoffrey Rush. ... weiterlesen... Kyrgyz: Бул жерде сиз көрүп турган кинону аталышы "Кингдин сөзү" болуп, режиссер Том Хоупер тарабынан чыгарылган жана Колин Фирт менен | 2.72 |
| L3-Mono-Sel | | | | | Gosling und Stone erhielten Nominierungen als Bester Schauspieler bzw. Beste Schauspielerin. | 29.91 |
| L3-Bi-Sel | deu_Latn | eng_Latn | Der Film mit Ryan Gosling und Emma Stone hat Nominierungen in allen Hauptkategorien erhalten. | The movie, featuring Ryan Gosling and Emma Stone, received nominations in all major categories. | This film with Ryan Gosling and Emma Stone has received nominations in all main categories. Kyrgyz: Бул фильм Райан Гослинг менен Эма Стоундун тартылышы болуп саналат жана бардык негизги категорияларда көрсөтүлгөн. ... Толук маалымат "Бирок" деген кинону карап чыгуу үчүн бул жерди бae!... - IMDb More » Kyrgyz: K | 11.18 |
| L3-Mono-Sel | | | | | The film with Ryan Gosling and Emma Stone has received nominations for all main categories. | 32.32 |

Figure 6: Examples of language mixing in bilingual CPT (L3-Bi-) compared to monolingual CPT (L3-Mono-).

A.7 Prompt Templates

For the SIB-200 classification task, we adopt the following template:

Topic Classification: science/technology, travel, politics, sports, health, entertainment, geography.
 {examples}
 The topic of the news "{text}" is

For the FLORES-200 translation task, we employ the following 3-shot prompt:

Translate the following sentence from {src_lang} to {tgt_lang}
 {examples}
 [{src_lang}]: {src_sent}
 [{tgt_lang}]:

A.8 Data Statistics

The data statistics presented in Tables 13 and 14 summarize the bilingual translation and monolingual training datasets used in this study. Token counts in the two tables are

calculated by splitting text on whitespace, a method chosen for its computational efficiency given the large volume of data.

For code data, we provide raw token counts from The Stack dataset across 32 programming languages in Table 15, totaling 51,253,373,176 tokens. We then downsample this to 49,999,171 tokens as counted by using the GPT-2 tokenizer (Radford et al., 2019), selected for its speed, to match the training dataset setup in Subsection 2.2.

| Category | Language Pair | Source Tokens | Target Tokens | Total Tokens |
|--------------|-------------------|-------------------|-------------------|--------------------|
| Altruistic | eng_Latn-zul_Latn | 12,672,195 | 9,196,313 | 21,868,509 |
| | zho_Hani-zul_Latn | 341,665 | 208,653 | 550,318 |
| | ceb_Latn-zul_Latn | 190,637 | 94,910 | 285,547 |
| | zho_Hani-ceb_Latn | 696,789 | 863,637 | 1,560,426 |
| | eng_Latn-mar_Deva | 7,736,633 | 7,248,634 | 14,985,267 |
| | zho_Hani-mar_Deva | 2,244,545 | 1,825,067 | 4,069,612 |
| | ceb_Latn-mar_Deva | 835,219 | 634,881 | 1,470,100 |
| | ceb_Latn-eng_Latn | 12,355,815 | 11,719,494 | 24,075,309 |
| | zho_Hani-khm_Khmr | 1,157,707 | 577,403 | 1,735,110 |
| | eng_Latn-khm_Khmr | 11,364,386 | 10,147,868 | 21,512,254 |
| Total | | 49,595,591 | 42,516,860 | 92,112,452 |
| Selfish | bel_Cyrl-deu_Latn | 27,012,850 | 18,085,602 | 45,098,452 |
| | bel_Cyrl-eng_Latn | 1,598,358 | 1,920,079 | 3,518,437 |
| | deu_Latn-mri_Latn | 1,682,621 | 2,250,042 | 3,932,663 |
| | eng_Latn-mri_Latn | 717,914 | 913,809 | 1,631,723 |
| | deu_Latn-kir_Cyrl | 1,682,749 | 1,583,623 | 3,266,372 |
| | eng_Latn-kir_Cyrl | 2,262,374 | 1,515,087 | 3,777,462 |
| | deu_Latn-nya_Latn | 1,155,433 | 1,077,300 | 2,232,733 |
| | eng_Latn-nya_Latn | 19,714,307 | 16,830,192 | 36,544,499 |
| Total | | 55,826,606 | 44,175,734 | 100,002,341 |
| Stagnant | eng_Latn-tha_Thai | 5,619,794 | 18,138,086 | 23,757,879 |
| | eng_Latn-yor_Latn | 14,334,000 | 16,887,000 | 31,221,000 |
| | eng_Latn-sna_Latn | 9,813,703 | 7,608,164 | 17,421,867 |
| | eng_Latn-wol_Latn | 13,600,133 | 13,636,959 | 27,237,092 |
| Total | | 43,367,630 | 56,270,209 | 99,637,838 |

Table 13: Bilingual translation data statistics: source, target, and total token counts across language pairs for each language category, with totals for each group.

| Category | Language | Total Tokens |
|--------------|----------|--------------------|
| Altruistic | eng_Latn | 43,492,709 |
| | zho_Hani | 4,440,706 |
| | ceb_Latn | 14,245,308 |
| | mar_Deva | 9,708,582 |
| | zul_Latn | 9,499,876 |
| | khm_Khmr | 10,725,271 |
| Total | | 92,112,452 |
| Selfish | eng_Latn | 24,614,674 |
| | deu_Latn | 22,606,405 |
| | bel_Cyrl | 28,611,208 |
| | mri_Latn | 3,163,851 |
| | kir_Cyrl | 3,098,710 |
| | nya_Latn | 17,907,492 |
| Total | | 100,002,341 |
| Stagnant | eng_Latn | 43,367,629 |
| | tha_Thai | 18,138,086 |
| | yor_Latn | 16,887,000 |
| | sna_Latn | 7,608,164 |
| | wol_Latn | 554,809 |
| Total | | 86,555,688 |

Table 14: Monolingual training data statistics: total token counts for each language across the three language categories.

| Language | Total Tokens |
|--------------|-----------------------|
| assembly | 331,667,471 |
| c | 8,741,971,474 |
| cpp | 7,816,404,624 |
| c-sharp | 2,378,224,612 |
| clojure | 82,101,240 |
| common-lisp | 392,951,006 |
| dart | 596,729,087 |
| erlang | 145,648,910 |
| f-sharp | 67,025,280 |
| fortran | 442,165,240 |
| gsl | 116,320,040 |
| go | 3,566,871,370 |
| haskell | 401,113,392 |
| java | 3,659,465,643 |
| javascript | 3,027,933,059 |
| julia | 221,192,206 |
| kotlin | 851,638,489 |
| llvm | 383,439,623 |
| markdown | 1,795,961,949 |
| pascal | 424,339,418 |
| perl | 473,210,127 |
| php | 2,315,544,678 |
| powershell | 74,390,317 |
| python | 5,199,071,526 |
| r | 49,449,207 |
| ruby | 1,107,302,714 |
| rust | 1,572,906,932 |
| scala | 568,062,821 |
| shell | 510,858,653 |
| solidity | 151,560,961 |
| sql | 1,179,866,764 |
| typescript | 2,607,984,343 |
| Total | 51,253,373,176 |

Table 15: Raw code data statistics from a subset of The Stack dataset processed by [Ji et al. \(2024a\)](#), showing total token counts for each programming language before downsampling.