

On Bilingual Lexicon Induction with Large Language Models

Yaoyiran Li Anna Korhonen Ivan Vulić

Language Technology Lab, TAL, University of Cambridge
{y1711, alk23, iv250}@cam.ac.uk

Abstract

Bilingual Lexicon Induction (BLI) is a core task in multilingual NLP that still, to a large extent, relies on calculating cross-lingual word representations. Inspired by the global paradigm shift in NLP towards Large Language Models (LLMs), we examine the potential of the latest generation of LLMs for the development of bilingual lexicons. We ask the following research question: *Is it possible to prompt and fine-tune multilingual LLMs (mLLMs) for BLI, and how does this approach compare against and complement current BLI approaches?* To this end, we systematically study **1**) zero-shot prompting for unsupervised BLI and **2**) few-shot in-context prompting with a set of seed translation pairs, both without any LLM fine-tuning, as well as **3**) standard BLI-oriented fine-tuning of smaller LLMs. We experiment with 18 open-source text-to-text mLLMs of different sizes (from 0.3B to 13B parameters) on two standard BLI benchmarks covering a range of typologically diverse languages. Our work is the first to demonstrate strong BLI capabilities of text-to-text mLLMs. The results reveal that few-shot prompting with in-context examples from nearest neighbours achieves the best performance, establishing new state-of-the-art BLI scores for many language pairs. We also conduct a series of in-depth analyses and ablation studies, providing more insights on BLI with (m)LLMs, also along with their limitations.

1 Introduction and Motivation

Bilingual Lexicon Induction (BLI), also known as word translation, is a fundamental research topic in multilingual NLP that aims to bridge the lexical gap between languages (Ruder et al., 2019). It has a wide range of applications such as machine translation (Artetxe et al., 2018b; Marchisio et al., 2020; Chronopoulou et al., 2021) and cross-lingual transfer learning, especially for low-resource languages (Sun et al., 2021; Zhou et al., 2021; Wang et al., 2022). Over the past decade, state-of-the-art

(SotA) BLI approaches have been predominantly supported by learning a cross-lingual word embedding (CLWE) space, with which BLI is tackled via nearest neighbour retrieval (Artetxe et al., 2018a; Heyman et al., 2019; Peng et al., 2021; Li et al., 2022a; Marchisio et al., 2022, *inter alia*).

Meanwhile, autoregressive text-to-text large language models (LLMs) have emerged as the cornerstone of cutting-edge NLP research (Raffel et al., 2020; Brown et al., 2020; Ouyang et al., 2022; Chowdhery et al., 2022). For example, multilingual LLMs (mLLMs) have shown (sentence-level) machine translation capabilities (Vilar et al., 2022; Briakou et al., 2023), although they have not been pretrained for machine translation in a supervised manner. Motivated by the recent remarkable success of (m)LLMs, in this work we investigate **1**) the potential of prompting and fine-tuning of mLLMs for BLI and **2**) how their capabilities compare against and complement current BLI approaches. We focus on how to expose word-level bilingual knowledge and elicit word translations from multilingual LLMs. To our best knowledge, we are the *first* to leverage autoregressive mLLMs for BLI.¹

We systematically study zero-shot and few-shot prompting for BLI with off-the-shelf encoder-decoder and decoder-only autoregressive mLLMs (Radford et al., 2019; Raffel et al., 2020; Brown et al., 2020), respectively. In the few-shot scenario, we propose to incorporate in-context examples from nearest neighbours into the prompts to boost the BLI performance. In order to guide the mLLMs’ generation, we hand-craft ‘*mask-filling-style*’ and ‘*GPT-style*’ templates catering to the characteristics of different LLMs and conduct extensive template search for BLI. In addition to pro-

¹We point out that the work of Li et al. (2022a) leverages only the encoder part of mT5 (the decoder is dropped) in one of their experiments to extract CLWEs for BLI. In contrast, our experiments with mT5 utilise its full encoder-decoder structure and generate the target words autoregressively, completely different from CLWE-based approaches.

viding a complete and effective pipeline for BLI via prompting off-the-shelf mLLMs, we also investigate BLI-oriented fine-tuning with the LLMs’ own pretraining objectives, aiming at specialising mLLMs into ‘few-shot word translators’.

We conduct extensive experiments on two standard BLI benchmarks, XLING (Glavaš et al., 2019) and PanLex-BLI (Vulić et al., 2019), investigating the word translation capabilities of off-the-shelf mLLMs (we adopt 18 models from 5 LLM families) in various BLI setups. Our comprehensive comparisons between mLLMs confirm, as expected, that **1)** different LLM families display varying word translation capabilities and **2)** stronger BLI performance tends to be associated with larger model sizes. To demonstrate the effectiveness of our prompt-based approach, we benchmark our method against two SotA CLWE-based baselines. Notably, our approach with LLaMA_{13B} outperforms the CLWE-based SotA on the XLING dataset by a considerable margin, establishing new SotA results on many language pairs in all BLI setups. Meanwhile, we also identify two limitations of BLI with mLLMs: **1)** they are less competitive on the PanLex-BLI benchmark for lower-resource languages; **2)** CLWE-based approaches usually support more languages than mLLMs. Finally, we run a series of insightful ablations and discuss the usefulness of BLI-oriented fine-tuning. In short, our work validates the BLI capabilities of mLLMs and proposes new methodology for BLI. We hope that the combination of our comprehensive analyses and discussions, including on limitations, will pave the way for the development of stronger BLI systems in the future. Our code is publicly available at github.com/cambridgeltl/prompt4bli.

2 Related Work

Bilingual Lexicon Induction. Over the past decade, predominant BLI approaches have relied on the calculation of cross-lingual word embeddings (CLWEs) where, in the most popular BLI variant, two transformation functions are learned to respectively map source and target monolingual static word embedding spaces into a shared cross-lingual space (Xing et al., 2015; Lample et al., 2018; Joulin et al., 2018; Artetxe et al., 2018a; Alvarez-Melis and Jaakkola, 2018; Patra et al., 2019; Mohiuddin et al., 2020; Glavaš and Vulić, 2020; Peng et al., 2021; Li et al., 2022a; Marchisio et al., 2022). Then, relying on the learned CLWE

space, BLI has been conducted via nearest neighbour retrieval. A detailed overview of different BLI principles can be found, e.g., in the work of Ruder et al. (2019).

More recently, researchers have attempted BLI by leveraging *encoder-only* multilingual masked language models (mMLMs) such as mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) whose neural architecture consists of only Transformer encoders (Vaswani et al., 2017). Gonen et al. (2020) prompt mBERT with templates where the target word is replaced with a ‘<mask>’ token, and the language modelling head of mBERT outputs a subword token to fill the mask. This method is theoretically flawed because it cannot address the cases where the target word comprises two or more subword tokens. Therefore, Gonen et al. (2020) only evaluate BLI on a small set of ‘toy’ examples rather than standard BLI datasets. In terms of performance, this method lags far behind traditional BLI approaches. A more successful way of leveraging mMLMs is to extract decontextualised word representations from them (Zhang et al., 2021). The strongest CLWEs for BLI so far are learned via a two-stage contrastive approach combining both static (e.g., fastText) and mMLM-extracted features (Li et al., 2022a).²

Text-to-Text LLMs. Autoregressive LLMs have established new state-of-the-art results on many NLP tasks. The prominent model groups include **1)** *encoder-decoder* LLMs such as BART (Lewis et al., 2020) and T5 (Raffel et al., 2020); **2)** OpenAI’s *decoder-only* GPT series such as GPT-2 (Radford et al., 2019), GPT-3 (Brown et al., 2020), and InstructGPT (Ouyang et al., 2022); **3)** other GPT-like LLMs with specific improvements such as Chinchilla (Hoffmann et al., 2022), PaLM (Chowdhery et al., 2022), and LLaMA LLM series (Touvron et al., 2023).

Our work adopts five families of open-source text-to-text multilingual LLMs for BLI, including mT5 (Xue et al., 2021), mT0 (Muennighoff et al., 2022), XGLM (Lin et al., 2022), mGPT (Shliazhko et al., 2022), and LLaMA (Touvron et al., 2023). We introduce each of these in more detail in §3.1. Unlike the encoder-only MLMs, text-to-text LLMs are theoretically capable of generating words consisting of arbitrary numbers of subword tokens.

²Using mMLMs alone still underperforms purely fastText-based methods since mMLMs are contextualised encoders pretrained for sentence-level tasks (Li et al., 2022a).

3 Methodology

BLI Task: Preliminaries and Terminology. Assuming a bilingual scenario with a source language L^x and a target language L^y with their respective vocabularies denoted as \mathcal{X} and \mathcal{Y} , the BLI task is typically formulated as a standard information retrieval task (Gaussier et al., 2004; Glavaš et al., 2019). The goal is to rank the words from \mathcal{Y} with respect to their similarity to the input source word w^x . The vocabulary size for each language is typically set to 200k (Li et al., 2022a), covering the most frequent 200k word types in each language. A bilingual lexicon then comprises a set of one-to-one source and target word translation pairs (Mikolov et al., 2013), and we denote a word pair as $\pi=(w^x, w^y)$ where $w^x \in \mathcal{X}, w^y \in \mathcal{Y}$.

We assume a set \mathcal{D}_S of N available seed translation pairs, constituting the so-called *seed dictionary*, which are used as the training set. Depending on the number of training pairs, the task is usually referred to as *supervised BLI* (typically, $N \geq 5K$), *semi-supervised BLI* (e.g., $0 < N \leq 1K$), and *unsupervised BLI* ($N = 0$) in the literature (Artetxe et al., 2018a; Zhao et al., 2020; Li et al., 2022a). For convenience, we also refer to the unsupervised setup as *zero-shot BLI* ($N = 0$) and denote the setup with a handful of seed translation pairs as *few-shot BLI* ($N > 0$), corresponding to how we prompt mLLMs for BLI (we describe zero-shot and few-shot prompts for BLI later in §3). A test set \mathcal{D}_T , where $\mathcal{D}_S \cap \mathcal{D}_T = \emptyset$, is used for evaluation.

In some cases, a source word may have more than one ground-truth translation (i.e., there exist two or more word pairs in a BLI dictionary that share the same source word). Following previous work (Lample et al., 2018; Glavaš et al., 2019; Li et al., 2022a), we consider a prediction correct as long as it is any of the ground-truth translations. The BLI scores are reported based on the standard Precision@K (P@K) BLI measure, where K denotes the length of the ranked list.

3.1 Prompting Multilingual LLMs for BLI

This study employs five families of mainstream multilingual text-to-text LLMs (mLLMs): mT5, mT0, XGLM, mGPT, and LLaMA.³ Based on their model structures, we group these models into two

³We also experimented with mBART (Liu et al., 2020). However, we do not report the results with mBART since its BLI performance proved inferior in our preliminary investigations.

categories; in what follows, we briefly introduce each of them and showcase some simple templates used for ‘BLI-prompting’ the LLMs.

The first category includes mT5 and mT0, two *encoder-decoder* LLM families that leverage the full Transformer architecture (Vaswani et al., 2017). Each model family comes in five different sizes, and we evaluate all these ten models.

- **mT5** (Xue et al., 2021) is pretrained on the mC4 dataset covering 101 languages. The LLM leverages a *span-corruption* objective that tries to reconstruct consecutive spans of dropped-out tokens replaced with special mask tokens.

- **mT0** (Muennighoff et al., 2022) is a multitask-finetuned mLLM based on instruction fine-tuning from the original mT5 model. The fine-tuning is conducted with English prompts on mT0’s xP3 dataset spanning 46 languages.⁴

For these two encoder-decoder style mLLMs, we aim to derive prompts such that the first word of the output sequence serves as its guess for w^y . Catering to its span-corruption objective, for mT5 we propose to design *mask-filling-style* English templates where ‘<mask>’ tokens are used as placeholders for the target words. Here is an example template: ‘The L^x word w^x in L^y is <mask>.’, where L^x , L^y , and w^x are placeholders for the source language, target language, and the input source word, respectively.⁵ When a prompt based on this template is fed into mT5, its decoder will then output a sequence to fill the mask. Since mT0 is based on mT5, we found that mask-filling-style prompts are also applicable to mT0. However, unlike for mT5, the instruction-tuned mT0 fits templates without the ‘<mask>’ token.⁶ For simplicity, we will denote all such templates without any ‘<mask>’ tokens as ‘*GPT-style templates*’.

The second model category comprises XGLM, mGPT, and LLaMA as three *decoder-only* LLMs pretrained with causal LM losses. Our experiments involve five XGLM and two LLaMA models whose

⁴Among all the LLMs covered in our work, mT0 is the only one trained for sentence-level translation tasks (machine translation is one of the tasks during its multitask fine-tuning). However, our experimental results reported later indicate that this does not benefit BLI in our prompting setups.

⁵The ‘<mask>’ token for mT5 and mT0 is actually ‘<extra_id_0>’. Therefore, an example of an actual prompt would be ‘The German word *gebouw* in French is <extra_id_0>.’.

⁶For instance, ‘The L^x word w^x in L^y is’ may prompt mT0 to output w^y to complete the input sentence, and ‘How do you say w^x in L^y ?’ would prompt mT0 to generate w^y to answer the question.

model sizes are no larger than 13B parameters, while mGPT only releases one model of size 1.4B. Unlike encoder-decoder LLMs for conditional generation, the decoder-only causal LLMs first repeat the input sequence in their output, and we construct prompts that induce LLMs to produce w^y immediately after the repeated input sequence.

- **XGLM** (Lin et al., 2022) offers multilingual LLMs similar to GPT-3 (Brown et al., 2020) and is reported to outperform GPT-3 of comparable size in a series of tasks. The work builds a CC100-XL dataset based on Conneau et al. (2020) and Wenzek et al. (2020), and XGLM is pretrained with a subset of it covering 30 languages.
- **mGPT** (Shliazhko et al., 2022) reproduces the GPT-3 structure and is trained on 60 languages using Wikipedia and C4 data (Raffel et al., 2020).
- **LLaMA** (Touvron et al., 2023) is a recently released SotA LLM family trained on trillions of tokens exclusively from publicly available datasets; it supports 20 languages. LLaMA also features its efficient implementation, and it adopts a series of recent improvements on normalisation, activation functions, and positional embeddings.

Our decoder-only LLMs solely leverage GPT-style prompts introduced above for mT0, since their tokenisers usually do not support ‘<mask>’ tokens.

3.2 Retrieval-Augmented In-Context Learning

In §3.1, we presented some simple zero-shot prompts (i.e., prompts without in-context examples) for BLI. However, recent work highlights the few-shot capabilities of modern LLMs (Brown et al., 2020). Therefore, we also investigate few-shot templates for improved BLI performance.⁷

We propose to retrieve the nearest neighbours of a source word which we use to construct in-context samples to boost BLI performance. More specifically, given \mathcal{D}_S and an input source word w^x , we extract n word pairs $(w_i^x, w_i^y) \in \mathcal{D}_S, 1 \leq i \leq n$, such that $w_i^x, 1 \leq i \leq n$ are n nearest neighbours of w^x in the auxiliary static monolingual word embedding space of \mathcal{X} . This auxiliary space is based on pretrained fastText word embeddings (Bojanowski et al., 2017)⁸ and we use the cosine

⁷Note again that few-shot in-context learning does not require any actual fine-tuning of LLMs. The word ‘learning’ here only refers to inserting in-context examples into the input prompt sequence.

⁸See Appendix B for more details about the fastText WEs used in our work.

similarity measure for the retrieval.⁹ We again design mask-filling-style and GPT-style few-shot templates for the mLLMs, as discussed in §3.1. Similar to zero-shot prompts, for few-shot prompts we also extract *the first word* after removing special tokens (e.g., start-of-sentence, padding, and ‘<mask>’ tokens) and repeated input sequence (for decoder-only models) as the prediction of w^y .

3.3 Template Design and BLI Inference

Template Design. We hand-craft in total 102 English zero-shot and few-shot templates, respectively listed in Tables 10 and 11 of Appendix C. A small set of basic templates is fully manually designed, and additional variants are then created by modifying or replacing the punctuation (see the tables). For each LLM, we search for its best zero-shot template and best few-shot template on a randomly chosen language pair (German, French) and fix the template choices for experiments on all other language pairs. The best template choices for each LLM are provided in Table 12 (Appendix C).

BLI Inference. At inference, we adopt beam search for both encoder-decoder and decoder-only LLMs and make the generator return the final beam ranked by their sequence scores. For each input prompt corresponding to w^x , we iterate through the returned set of sequences, and for each sequence we extract the word after removing any redundant prefix content, as described in §3.2. The first word extracted that appears in the target vocabulary is returned as our prediction of w^y .

3.4 BLI-Oriented Fine-Tuning

This work predominantly focuses on ‘learningless’ experiments based on zero-shot and few-shot in-context setups with off-the-shelf mLLMs for BLI without any fine-tuning. As a side experiment, we also aim to fine-tune smaller-scale mLLMs, making them specialise into few-shot word translators with our few-shot prompts as input. Our training set is still \mathcal{D}_S , but we now exclude retrieving an input w^x itself as an in-context example. We combine the \mathcal{D}_S of each language pair with which we fine-tune encoder-decoder mLLMs with mT5’s span-corruption loss and fine-tune decoder-only LLMs with the standard causal LM objective.

⁹For the rare cases where an in-context source word retrieved may have more than one translation in \mathcal{D}_S , we only keep the target word with the highest word frequency in fastText’s training data.

4 Experimental Setup

Training and Evaluation Data. Our experiments adopt two standard and publicly available BLI datasets, also used in a body of very recent BLI research (Vulić et al., 2020; Sachidananda et al., 2021; Aboagye et al., 2022; Li et al., 2022a,b; Vulić et al., 2023). **1)** XLING (Glavaš et al., 2019) provides BLI dictionaries covering 8 languages and 56 BLI directions. Among these 8 languages, 5 are supported by all of our mLLMs: English (EN), French (FR), German (DE), Italian (IT), and Russian (RU). Therefore, §5 mainly focuses on and reports results on all the $20 = 5 \times 4$ BLI directions for the 5 languages.¹⁰ For each language pair, XLING provides a test set \mathcal{D}_T of 2K translation pairs. It also provides training sets \mathcal{D}_S of 5K and 1K translation pairs, where the former is the superset of the latter. For brevity, we denote the cases $|\mathcal{D}_S| = 5K$, $|\mathcal{D}_S| = 1K$, and $|\mathcal{D}_S| = 0$ as the 5K setup, 1K setup, and unsupervised setup, respectively.¹¹ **2)** PanLex-BLI (Vulić et al., 2019) offers BLI lexicons spanning 15 lower-resource languages and all 210 BLI directions. We select three languages that are supported by most of our mLLMs: Bulgarian (BG), Catalan (CA), and Hungarian (HU). The test set size of PanLex-BLI is also 2K; under the lower-resource assumption, we only focus on unsupervised and 1K BLI setups.

Main Experiments. In our main experiments, we prompt 18 off-the-shelf models from 5 mLLM families mentioned in §3.1¹² for BLI *without* any fine-tuning¹³ and systematically evaluate their BLI performance in three different BLI setups on XLING and PanLex-BLI datasets introduced above. In 5K and 1K setups, 5-shot in-context learning is adopted for our mLLMs, while in the unsupervised setup, zero-shot prompts are used. We compare the BLI scores between different mLLMs from the perspectives of LLM family and model size, and we also benchmark their performance against two SotA CLWE-based baselines, introduced later. Selected results are summarised in §5.1 while full and

detailed BLI scores are reported in Appendix D.

Side Experiments. We conduct a series of additional experiments to further understand the BLI capabilities of mLLMs. **1)** We investigate how the BLI performance is related to the number of in-context examples (5K and 1K setups). **2)** As an ablation study, we validate the usefulness of (our proposed) in-context samples extracted from nearest neighbours by comparing with randomly sampled in-context examples. **3)** Finally, we fine-tune some of our relatively smaller-scale LLMs, including $mT5_{base}$, $mT5_{large}$, $XGLM_{564M}$, and $XGLM_{1.7B}$ on our 5-shot templated BLI data (XLING) and further study the effectiveness of our BLI-oriented fine-tuning (5K and 1K setups). The training set includes all XLING language pairs, where the 5K and 1K setups have 271, 754 and 55, 228 training instances respectively.

Hyperparameters. We first introduce our hyperparameters for BLI inference. In our main experiments, we adopt $n = 5$ ¹⁴ while in side experiments we further investigate and compare using different numbers of in-context examples n . Concerning the generation of output sequences, we adopt a beam size of 5 for all LLMs, and the maximum sequence length is 5 for encoder-decoder models and 5 plus the input sequence length for decoder-only models which first repeat the input sequence before generating new content. As for encoder-decoder LLMs, we use an evaluation batch size of 100 for smaller models and 8 for larger models as listed in Table 8 (Appendix B). Since the pretraining of decoder-only LLMs usually does not see padding tokens, we adopt a batch size of 1.¹⁵ Following prior work (Li et al., 2022a,b), all our hyperparameters are tuned on (German, French), a randomly selected language pair.

For ‘BLI-oriented’ fine-tuning, we use the XLING data combining all language pairs, and the batch size is 16 for $XGLM_{1.7B}$ and 32 for $mT5_{base,large}$ and $XGLM_{564M}$. We use AdamW (Loshchilov and Hutter, 2019) with $\beta_{tas} = (0.9, 0.98)$ and a weight decay of 0.1. The learning rate is $2e-6$ for $mT5_{base}$, $1e-6$ for $mT5_{large}$, and $5e-8$ for $XGLM_{564M}$; concerning $XGLM_{1.7B}$, $5e-9$ is adopted for the 5K setup and $2e-8$ for the 1K setup. All the hyperparameters are tuned on

¹⁰However, in the appendix, we also provide the results with the remaining 3 languages: Croatian (HR), Finnish (FI), and Turkish (TR).

¹¹Related work often also refers to the 5K and 1K cases as supervised and semi-supervised BLI setups, respectively (Li et al., 2022a).

¹²A summary concerning the detailed information of each mLLM is available in Appendix B.

¹³Experiments on BLI-oriented fine-tuning for a selection of mLLMs are discussed later.

¹⁴Unless otherwise stated, we report 5-shot BLI results throughout our experiments.

¹⁵We found that a larger batch size may cause a drop in the BLI performance.

the same randomly chosen language pair (German, French). Each LLM is fine-tuned for at most 20 epochs capped at 12 hours on 1×80GB A100 GPU.

Baselines. We adopt the following two SotA CLWE-based approaches as our baselines; both are open-source. We follow their original suggested hyperparameter choices respectively for 5K (supervised), 1K (semi-supervised), and unsupervised BLI setups, and we re-verify that the hyperparameters recommended are (near-)optimal. The Cross-domain Similarity Local Scaling (CSLS) retrieval (Lample et al., 2018) is adopted as recommended in the baselines.

- **VECMAP** (Artetxe et al., 2018a) is one of the most representative BLI approaches based on static CLWEs. It induces fastText-based CLWEs in various BLI supervision setups, and is notable for its effective self-learning mechanism, especially in weakly supervised and unsupervised BLI setups.

- **CONTRASTIVEBLI** (Li et al., 2022a) refines CLWEs with a two-stage contrastive learning procedure and reports the currently highest CLWE-based BLI scores on XLING and PanLex-BLI in 5K and 1K BLI setups. We adopt its strongest CLWEs derived with both fastText and mBERT (Devlin et al., 2019). **CONTRASTIVEBLI** does not support unsupervised BLI.

BLI Evaluation. Following previous work, we report the standard Precision@1 (P@1) scores both for our methods and for baseline methods.¹⁶

5 Results and Discussion

5.1 Main Results

Comparison between mLLMs. We compare the average BLI scores on 20 XLING BLI directions derived from all our 18 models from 5 LLM families in Figure 1. In all (5K, 1K, and zero-shot) BLI setups, the same general trends are observed. **1)** As expected, within the same mLLM family, larger models usually present stronger BLI capabilities, although exceptional cases exist (e.g., XGLM_{7.5B} underperforms XGLM_{4.5B}). **2)** For encoder-decoder models, we find that mT5 outperforms mT0, showing that the instruction fine-tuning of mT0 does not benefit BLI in our experimental setups; **3)** LLaMA models achieve the strongest BLI performance among our 5 model families.

¹⁶P@1 is the most authoritative metric for BLI. Other measures such as P@5 and Mean Reciprocal Rank (MRR) show similar trends (Lample et al., 2018; Li et al., 2022a).

[5K Setup]	VECMAP	CONTRASTIVEBLI	mT5 _{xxl}	mT0 _{xxl}	XGLM _{4.5B}	mGPT	LLaMA _{13B}
DE→*	47.65	54.02	49.1	35.73	48.43	37.44	56.44
*→DE	47.34	53.64	46.37	32.66	46.56	32.47	54.78
EN→*	53.54	60.26	59.36	44.2	61.2	45.69	69.0
*→EN	57.38	60.97	57.34	41.48	56.08	45.15	62.35
FR→*	53.36	58.51	53.46	39.7	54.31	44.02	60.81
*→FR	56.74	60.83	57.29	41.9	57.12	43.97	65.18
IT→*	53.52	58.6	52.34	36.73	51.37	40.57	58.66
*→IT	55.61	59.88	55.43	40.67	55.86	47.06	64.4
RU→*	46.74	52.13	49.88	34.09	50.08	38.25	56.26
*→RU	37.74	48.21	47.72	33.75	49.76	37.33	54.45
Avg.	50.96	56.71	52.83	38.09	53.08	41.2	60.23
[1K Setup]	VECMAP	CONTRASTIVEBLI	mT5 _{xxl}	mT0 _{xxl}	XGLM _{4.5B}	mGPT	LLaMA _{13B}
DE→*	44.44	51.79	46.02	33.57	46.9	35.98	53.82
*→DE	44.0	48.92	43.62	30.58	44.06	31.9	53.36
EN→*	47.7	55.11	56.19	40.01	59.3	44.55	68.27
*→EN	55.74	59.66	52.86	39.6	53.51	43.18	57.06
FR→*	48.47	55.33	49.46	37.03	51.08	42.82	58.35
*→FR	54.88	58.65	54.51	38.88	56.74	43.9	63.26
IT→*	49.08	55.92	48.49	33.78	49.97	39.86	56.06
*→IT	53.4	57.08	51.75	38.15	54.19	45.38	62.57
RU→*	43.61	50.17	46.69	32.71	48.67	37.46	52.75
*→RU	25.3	44.02	44.1	29.9	47.42	36.31	53.01
Avg.	46.66	53.66	49.37	35.42	51.18	40.13	57.85
[Zero-Shot]	VECMAP	CONTRASTIVEBLI	mT5 _{xxl}	mT0 _{xxl}	XGLM _{4.5B}	mGPT	LLaMA _{13B}
DE→*	44.44	-	35.76	29.31	38.37	22.65	43.3
*→DE	43.95	-	42.4	31.75	40.76	21.36	47.53
EN→*	47.76	-	50.94	36.48	51.01	28.6	54.44
*→EN	55.8	-	43.16	38.77	44.75	25.78	52.34
FR→*	48.24	-	42.37	35.66	45.24	30.73	51.01
*→FR	54.96	-	46.04	32.16	45.79	32.49	53.87
IT→*	48.97	-	38.84	30.97	39.08	27.4	48.04
*→IT	53.27	-	45.68	37.4	46.81	34.79	53.09
RU→*	43.63	-	42.05	32.19	40.85	21.49	46.82
*→RU	25.08	-	32.69	24.5	36.45	16.46	36.76
Avg.	46.61	-	41.99	32.92	42.91	26.18	48.72

Table 1: Main results on 20 XLING BLI directions in 5K, 1K, and zero-shot (unsupervised) setups. Off-the-shelf mLLMs are used without any fine-tuning. Average P@1×100% scores of each language going to and from other 4 languages are reported. ‘-’: **CONTRASTIVEBLI** does not support unsupervised BLI.

XLING: Main Results. In Table 1, we report the BLI performance of the strongest single model from each LLM family (full results for each of the 18 LLMs are available in Appendix D). Our results on 20 XLING BLI directions reaffirm the leading position of LLaMA where the LLaMA_{13B} variant achieves the highest overall average BLI scores in all BLI setups, also outperforming **CONTRASTIVEBLI**, the previous CLWE-based SotA on the same dataset. We speculate that LLMs are more adept at few-shot learning: LLaMA_{13B} outperforms **VECMAP** by circa 10 P@1 points in few-shot setups, but only by about 2 points in the zero-shot setup. It is also worth mentioning that in 5K and 1K setups, mT5_{xxl} and XGLM_{4.5B} beat **VECMAP** although they underperform **CONTRASTIVEBLI**; however, in the zero-shot setup they still cannot match **VECMAP**.

PanLex-BLI: Main Results. We present our results on lower-resource languages from PanLex-BLI in Table 2. We also provide here only one strongest model of each LLM family, while the full results from all 18 LLMs are available in Appendix D). This time, LLaMA_{13B} still outperforms other LLMs, but we report here XGLM_{7.5B} instead which beats XGLM_{4.5B}. Unlike for XLING, we




Figure 1: Averaged BLI score versus model size (0.3B to 13B): (left) $|\mathcal{D}_S|=5K$; (middle) $|\mathcal{D}_S|=1K$; (right) $|\mathcal{D}_S|=0$.

find that traditional CLWE-based approaches still outperform LLM-elicited BLI in general. This may reveal that current SotA mLLMs (size $\leq 13B$) still lack strong word translation capabilities for a large number of languages and language pairs, even for those they currently cover.

Put simply, while current mLLMs do exhibit strong performance for arguably high-resource languages (from XLING), they still have deficiencies with lower-resource languages as well as with their portability to a much larger number of languages, currently covered by more traditional BLI approaches (Li et al., 2022a). We leave to future work the investigation of larger mLLMs (e.g., LLaMA_{30B}) for BLI with lower-resource languages and languages unseen by the mLLMs.

Statistical Significance. We conduct χ^2 test comparing LLaMA_{13B} against the strongest single baseline in each BLI setup (i.e., CONTRASTIVEBLI in few-shot setups and VECMAP in the zero-shot setup) on the average BLI performance over 20 XLING and 6 PanLex-BLI BLI directions respectively, and we estimate the p -values as follows. **1)** On XLING, p is $2.8e-23$ in the 5K setup, $8.5e-32$ in the 1K setup, and $4.3e-9$ in the zero-shot setup. **2)** For PanLex-BLI, p is $1e-4$ in the 1K setup and $1.9e-35$ in the zero-shot setup. The p -values show that our main findings are clearly statistically significant.¹⁷

5.2 Further Analyses

n -Shot Prompting. To better understand the influence of the number of in-context examples, we pick mT5_{large} (an encoder-decoder LLM) and LLaMA_{13B} (a decoder-only LLM) and run experiments ranging from 0-shot to 10-shot. Figure 2 depicts their average BLI scores on 20 XLING BLI directions in 5K and 1K setups, respectively. The results clearly demonstrate the usefulness of in-context learning. Even when having only one in-context example (one-shot), the same model vari-

¹⁷By convention, $p < 0.05$: statistically significant; $p < 1e-3$: statistically highly significant.

[1K Setup]	VECMAP	CONTRASTIVEBLI	mT5 _{xxl}	mT0 _{xxl}	XGLM _{7.5B}	mGPT	LLaMA _{13B}
BG→CA	39.66	43.93	38.67	31.72	40.19	-	41.71
CA→BG	33.54	40.06	36.2	22.72	40.23	-	41.53
BG→HU	38.77	44.62	36.17	25.46	-	23.61	36.57
HU→BG	36.52	43.03	36.98	24.02	-	28.17	43.2
CA→HU	35.47	41.44	32.43	22.21	-	-	35.3
HU→CA	39.88	47.14	37.68	29.59	-	-	46.04
Avg.	37.31	43.37	36.36	25.95	-	-	40.72

[Zero-Shot]	VECMAP	CONTRASTIVEBLI	mT5 _{xxl}	mT0 _{xxl}	XGLM _{7.5B}	mGPT	LLaMA _{13B}
BG→CA	39.6	-	28.04	28.86	28.5	-	32.77
CA→BG	33.6	-	21.47	16.83	20.17	-	27.03
BG→HU	39.24	-	27.26	24.07	-	7.23	23.61
HU→BG	36.46	-	22.47	16.94	-	9.85	26.5
CA→HU	34.09	-	24.59	22.93	-	-	24.53
HU→CA	37.79	-	25.47	24.48	-	-	38.17
Avg.	36.8	-	24.88	22.35	-	-	28.77

Table 2: Main results on 6 PanLex-BLI BLI directions in 1K and zero-shot (unsupervised) setups. Off-the-shelf mLLMs are used without any fine-tuning. P@1 \times 100% scores are reported. ‘-’: **1)** a language is not supported by the LLM; **2)** CONTRASTIVEBLI does not support unsupervised BLI.




Figure 2: BLI scores averaged over 20 BLI directions from XLING with respect to the number of in-context examples n (0 to 10), with mT5_{large} and LLaMA_{13B} in both 5K and 1K BLI setups.

ant outperforms its zero-shot results by ~ 5 P@1 points. However, with higher values for n (i.e., $n \geq 5$), the gains become saturated.

Ablation Study. One key contribution of our work is that we propose to extract in-context examples from nearest neighbours. To validate the effectiveness of this approach, we conduct an essential ablation study where we use randomly sampled in-context examples instead. As with our main experiments, we present average scores on 20 BLI directions from only one best LLM from each LLM family, and full results on all LLMs are available in Appendix D. Our results in Table 3 demonstrate the following. **1)** The nearest neighbour-based ‘NN (*K)’ scores for every LLM outperform the ‘Ran-

dom (*K)’ scores by a salient margin: this shows the effectiveness of in-context examples from nearest neighbours. **2)** ‘Random (*K)’ outperforms ‘Zero-Shot’, showing that even randomly picked in-context examples can benefit BLI. **3)** ‘NN (5K)’ outperforms ‘NN (1K)’, which means that better in-context examples can be retrieved from a larger (and more varied) 5K seed dictionary.¹⁸ We further show that these findings are statistically significant via χ^2 test and report the p -values in Table 4.

	mT5 _{xxl}	mT0 _{xxl}	XGLM _{4.5B}	mGPT	LLaMA _{13B}
NN (5K)	52.83	38.09	53.08	41.2	60.23
NN (1K)	49.37	35.42	51.18	40.13	57.85
Random (5K)	46.93	34.97	49.85	38.88	56.85
Random (1K)	45.93	33.90	49.92	38.38	56.11
Zero-Shot (Unsupervised)	41.99	32.92	42.91	26.18	48.72

Table 3: Ablation results. Averaged BLI scores (P@1×100%) on 20 XLING BLI directions. Rows 1-2: 5-shot prompting with in-context examples extracted from NN in \mathcal{D}_S of size 5K and 1K. Rows 3-4: 5-shot prompting with random in-context examples in \mathcal{D}_S of size 5K and 1K. Row 5: zero-shot prompting without any in-context examples.

p -value	mT5 _{xxl}	mT0 _{xxl}	XGLM _{4.5B}	mGPT	LLaMA _{13B}
NN (5K) vs. Random (5K)	1.3e-60	1.9e-19	2.4e-19	4.7e-11	1.41e-21
NN (1K) vs. Random (1K)	9.3e-22	9.1e-6	4.6e-4	6.4e-7	1.1e-6
Random (5K) vs. Zero-Shot	1.7e-43	1.8e-9	1.4e-83	1e-311	8.4e-114
Random (1K) vs. Zero-Shot	2.4e-28	3.9e-3	3.2e-85	5.8e-289	2.8e-94
NN (5K) vs. NN (1K)	5.8e-22	1.3e-14	1.2e-7	2.5e-3	1.7e-11

Table 4: Statistical significance associated with Table 3. We conduct χ^2 tests and report p -values.

BLI-Oriented Fine-Tuning. The fine-tuning experiments, due to computational constraints, are conducted on four relatively smaller-scale LLMs. Table 5 reports each model’s average performance on 20 XLING BLI directions before and after fine-tuning. We run fine-tuning experiments three times with different random seeds and report both mean scores and standard deviations. We only observe salient gains on mT5 models and XGLM_{564M} in the 5K setups. For XGLM_{1.7B} in both setups and all models in the 1K setup, the gains are smaller or even non-existent. Even in the 5K setup, the tuned mT5_{base} still cannot match the off-the-shelf mT5_{large}, and the tuned mT5_{large} underperforms mT5_{xl} (cf. Table 18 in Appendix D). This may indicate some of the limitations of our proposed

¹⁸We also observe a slight edge of ‘Random (5K)’ over ‘Random (1K)’, and we speculate this might have to do with how XLING’s test set and seed dictionaries were created. In fact, the 1K seed dictionary contains the most frequent 1K words, and the test set words include less frequent 2K words.

	mT5 _{base}	mT5 _{large}	XGLM _{564M}	XGLM _{1.7B}
Fine-Tuned (5K)	36.11±0.531	46.68±0.058	28.94±0.026	40.92±0.005
Off-The-Shelf (5K)	28.33	43.25	24.51	39.49
Fine-Tuned (1K)	25.61±0.307	38.48±0.193	25.6±0.123	39.59±0.017
Off-The-Shelf (1K)	24.91	38.84	24.18	38.86

Table 5: Comparisons between mLLMs before and after BLI-oriented fine-tuning. Averaged BLI scores (P@1×100%) on 20 XLING BLI directions.

BLI-oriented fine-tuning (i.e., training the mLLMs on BLI data with their own pretraining objectives) and may indicate the following. **1)** With the current training approach and a fixed amount of computational budget, one may prioritise adopting off-the-shelf larger LLMs (with in-context learning) rather than training smaller-scale ones. **2)** In future work, other training objectives and strategies should be investigated for improved BLI performance with mLLMs. As another future research avenue, it is also worth extending the training to larger mLLMs and adopting novel fine-tuning techniques such as prompt tuning (Lester et al., 2021), adapters (Li et al., 2020, 2023) and LoRA (Hu et al., 2022).

Templates. Now, we additionally provide some preliminary findings from our template search.¹⁹ **1)** Models from the same mLLM family may tend to prefer the same template. For example, Table 12 (Appendix C) shows that all five XGLM models prefer the same best zero-shot template and four of them share one best few-shot template. This phenomenon is to some extent seen also on mT5 (zero-shot and few-shot), mT0 (zero-shot and few-shot), and LLaMA (few-shot). This should be due to the same training data, training strategy, and model architecture adopted for all models in the same LLM family. **2)** As already mentioned in §3.1, mT0 is compatible with both mask-filling-style and GPT-style templates: Table 12 shows that some mT0 models prefer templates with ‘<mask>’ and others do not. **3)** Under the ‘GPT-style’ templates, decoder-only models all prefer templates for sentence completion while some of the instruction-tuned mT0 models prefer questions with ‘?’.

5.3 Further Discussion

Few-Shot Learning for BLI. Our main results demonstrate that few-shot learning derives consistent gains over zero-shot prompting. For instance, HR→EN and IT→EN saw 345 and 272 cases in

¹⁹Since we conduct template search only on a random language pair, these findings are yet to be verified by future work for other language pairs.

their test sets respectively where few-shot learning makes the correct prediction but zero-shot learning fails (positive cases). There are only 85 and 87 cases where zero-shot prompting beats few-shot prompting (negative cases). We present 8 positive examples and 4 negative examples for each of HR→EN and IT→EN, comparing five-shot (5K setup) and zero-shot results with LLaMA_{13B} in Table 19 (Appendix E). For instance, ‘gušter (HR) → lizard (EN)’ and ‘sezam (HR) → sesame’ are two positive cases, their in-context examples being five different animal names and five plant names, which may help LLaMA_{13B} to narrow down the scope of the target word to animal and plant names respectively. Similarly, ‘valcer (HR) → waltz (EN)’ (a positive case) is associated with five in-context examples related to either music or dance. However, few-shot learning does not always help. For example, in ‘eventuale (IT) → eventual (EN)’ and ‘scopre (IT) → discovers (EN)’ translation tasks, the LLM seems to make a mistake due to directly copying one of the words provided in the in-context examples, whereas zero-shot prompting predicts the correct answers.

BLI for EN and non-EN Languages. It is noteworthy that the volume of data in English for mLLM pretraining often exceeds that in any other language (e.g., mT5, mT0, and LLaMA), and thus mLLMs may be biased, favouring BLI directions involving EN. However, we did not identify very clear clues indicating that their BLI performance is (heavily) biased. In fact, in the 5K setup (see Table 1), although ‘EN→*’ surpasses ‘non-EN→*’ in absolute BLI scores (for each of our LLMs and also CLWE-based baselines), we meanwhile observe that 1) mT0_{xxl} achieves lower average score in ‘*→EN’ than ‘*→FR’, and 2) for LLaMA_{13B}, ‘*→EN’ lags behind both ‘*→FR’ and ‘*→IT’. Moreover, as an example, the LLaMA_{13B} model supports 6 languages resulting in 30 BLI directions, and 20 of them are between non-EN languages. LLaMA_{13B} outperforms CLWE-based SotA in 16/20 cases and in 18/20 cases respectively in the 5K and 1K setups for the non-EN pairs (cf. Tables 13 and 14). However, we again note that this might hold only for high-resource languages such as the ones covered in XLING.

Impact Statement. Here, we discuss the potential impact of our study on the following two aspects. 1) *On future BLI research.* Our work minimises the technical gap between BLI and prompt-based learn-

ing and opens up new possibilities for BLI research. In fact, LLM prompting provides a generic and straightforward way of leveraging external knowledge for BLI. While we have demonstrated the effectiveness of in-context word translation examples, external information such as word definition, parts-of-speech, spelling, and sentence translation pairs can also be integrated into text prompts. 2) *On NMT and other related fields.* Recent work has incorporated word translation pairs into text templates to prompt LLMs for sentence-level neural machine translation (NMT) and demonstrates that the bilingual lexical ‘hints’ lead to significant gains in NMT (Ghazvininejad et al., 2023; Jones et al., 2023). While a ground-truth bilingual dictionary can be leveraged, BLI is able to provide word translations for language pairs and words not covered in existing bilingual lexica.²⁰ Our work can provide strong word translation pairs for lexicon-enhanced MT, and the improved MT may further benefit, e.g., the field of cross-lingual transfer learning via TRANSLATE-TRAIN/TEST approaches (Conneau et al., 2018; Li et al., 2023).

6 Conclusion

This paper presents the first study on bilingual lexicon induction (BLI) with multilingual text-to-text large language models (mLLMs). We develop the methodology to prompt mLLMs for BLI, conduct extensive template search, and systematically experiment with 5 representative mLLM families (18 models) on a variety of zero-shot and few-shot BLI tasks. Relying on off-the-shelf mLLMs, our experiments on the standard XLING dataset offer strong performance in all BLI setups, where our proposed few-shot prompting with in-context examples from nearest neighbours outperforms the strongest CLWE-based SotA by a considerable margin. However, our study also points out that prompting-based methods still need to be successfully extended to lower-resource languages. Finally, we conduct a series of in-depth analyses covering variants of our few-shot prompting and preliminary investigations on BLI-oriented fine-tuning. Our key findings and comprehensive analyses may pave the way for the development of stronger mLLM-based BLI systems in the future.

²⁰Moreover, given an input word, BLI can offer multiple plausible translations for the downstream NMT to consider, and we speculate this may, to some extent, increase the diversity of MT output.

Limitations

First, most recently released state-of-the-art mLLMs are still unable to support as many languages as static word embeddings, which currently limits their wider portability. For instance, LLaMA supports 20 languages and XGLM supports 30 languages, while fastText provides pretrained static WEs for 294 languages that can be used for the induction of static CLWEs.²¹ Intuitively, this is because training LLMs that support more languages would require higher computational costs (with more training data and typically larger model sizes). We hope that researchers in the future can pretrain and release mLLMs that support a larger set of linguistically diverse languages, which can thus probably extend the success of our approach to more languages and language families.

Second, our work did not investigate open-source LLMs with more than 13B parameters²² due to a large number of experiments conducted combined with our limited computing resources, and we did not evaluate any closed-source LLMs. Quite a few tech companies and AI research labs have been training LLMs with 100+B and even 500+B parameters. We encourage interested readers who have access to adequate computing resources or specific closed-source LLMs to take a step further and investigate if larger LLMs can provide an even stronger BLI performance than reported in this particular work, following the recipe presented in this work.

Third, as also discussed in other BLI work (Li et al., 2022b), existing BLI datasets did not control the synonyms and polysemy well and to a sufficient detail. In fact, when constructing BLI datasets, it is very difficult to collect all correct translations for each source word. Therefore, one limitation of BLI evaluation is that it cannot give credit to correct answers that are not included in the ground-truth translation set, and evaluation is typically conducted out-of-context. Constructing finer-grained BLI datasets with the help of qualified annotators (e.g., linguists, typologists and bilingual speakers) is beyond the scope of this work.

Acknowledgements

We thank the anonymous reviewers and area chairs for their valuable feedback. Yaoyiran Li is sup-

ported by Grace & Thomas C. H. Chan Cambridge International Scholarship. Ivan Vulić is supported by a personal Royal Society University Research Fellowship ‘*Inclusive and Sustainable Language Technology for a Truly Multilingual World*’ (no 221137; 2022–).

References

- Prince Osei Aboagy, Jeff Phillips, Yan Zheng, Junpeng Wang, Chin-Chia Michael Yeh, Wei Zhang, Liang Wang, and Hao Yang. 2022. [Normalization of language embeddings for cross-lingual alignment](#). In *International Conference on Learning Representations*.
- David Alvarez-Melis and Tommi Jaakkola. 2018. [Gromov-Wasserstein alignment of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1881–1890, Brussels, Belgium. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. [A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 789–798, Melbourne, Australia. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018b. [Unsupervised neural machine translation](#). In *International Conference on Learning Representations*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Eleftheria Briakou, Colin Cherry, and George F. Foster. 2023. [Searching for needles in a haystack: On the role of incidental bilingualism in palm’s translation capability](#). *ArXiv*, abs/2305.10266.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul

²¹fasttext.cc/docs/en/pretrained-vectors.html

²²For instance, LLaMA also offers larger 30B and 65B model variants.

- Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam M. Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Benton C. Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier García, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Díaz, Orhan Firat, Michele Catasta, Jason Wei, Kathleen S. Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#). *ArXiv*, abs/2204.02311.
- Alexandra Chronopoulou, Dario Stojanovski, and Alexander Fraser. 2021. [Improving the lexical ability of pretrained language models for unsupervised neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 173–180, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Eric Gaussier, J.M. Renders, I. Matveeva, C. Goutte, and H. Dejean. 2004. [A geometric view on bilingual lexicon extraction from comparable corpora](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, pages 526–533, Barcelona, Spain.
- Marjan Ghazvininejad, Hila Gonen, and Luke Zettlemoyer. 2023. [Dictionary-based phrase-level prompting of large language models for machine translation](#). *ArXiv*, abs/2302.07856.
- Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. [How to \(properly\) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 710–721, Florence, Italy. Association for Computational Linguistics.
- Goran Glavaš and Ivan Vulić. 2020. [Non-linear instance-based cross-lingual mapping for non-isomorphic embedding spaces](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7548–7555, Online. Association for Computational Linguistics.
- Hila Gonen, Shauli Ravfogel, Yanai Elazar, and Yoav Goldberg. 2020. [It’s not Greek to mBERT: Inducing word-level translations from multilingual BERT](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 45–56, Online. Association for Computational Linguistics.
- Geert Heyman, Bregt Verreet, Ivan Vulić, and Marie-Francine Moens. 2019. [Learning unsupervised multilingual word embeddings with incremental multilingual hubs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1890–1902, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and L. Sifre. 2022. [Training compute-optimal large language models](#). *ArXiv*, abs/2203.15556.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Alex Jones, Isaac Caswell, Ishan Saxena, and Orhan Firat. 2023. [Bilex rx: Lexical data augmentation for massively multilingual machine translation](#). *ArXiv*, abs/2303.15265.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. [Loss in translation: Learning bilingual word mapping with a retrieval criterion](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language*

- Processing*, pages 2979–2984, Brussels, Belgium. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Marc’ Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. [Word translation without parallel data](#). In *Proceedings of the International Conference on Learning Representations*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yaoyiran Li, Ching-Yun Chang, Stephen Rawls, Ivan Vulić, and Anna Korhonen. 2023. [Translation-enhanced multilingual text-to-image generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9174–9193, Toronto, Canada. Association for Computational Linguistics.
- Yaoyiran Li, Fangyu Liu, Nigel Collier, Anna Korhonen, and Ivan Vulić. 2022a. [Improving word translation via two-stage contrastive learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4353–4374, Dublin, Ireland. Association for Computational Linguistics.
- Yaoyiran Li, Fangyu Liu, Ivan Vulić, and Anna Korhonen. 2022b. [Improving bilingual lexicon induction with cross-encoder reranking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4100–4116, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yaoyiran Li, Edoardo Maria Ponti, Ivan Vulić, and Anna Korhonen. 2020. [Emergent communication pretraining for few-shot machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4716–4731, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. 2022. [Few-shot learning with multilingual generative language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9019–9052, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *Proceedings of the International Conference on Learning Representations*.
- Sasha Luccioni, Victor Schmidt, Alexandre Lacoste, and Thomas Dandres. 2019. [Quantifying the carbon emissions of machine learning](#). In *NeurIPS 2019 Workshop on Tackling Climate Change with Machine Learning*.
- Kelly Marchisio, Kevin Duh, and Philipp Koehn. 2020. [When does unsupervised machine translation work?](#) In *Proceedings of the Fifth Conference on Machine Translation*, pages 571–583, Online. Association for Computational Linguistics.
- Kelly Marchisio, Neha Verma, Kevin Duh, and Philipp Koehn. 2022. [IsoVec: Controlling the relative isomorphism of word embedding spaces](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6019–6033, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tomás Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. [Exploiting similarities among languages for machine translation](#). *CoRR*, abs/1309.4168.
- Tasnim Mohiuddin, M Saiful Bari, and Shafiq Joty. 2020. [LNMap: Departures from isomorphic assumption in bilingual lexicon induction through non-linear mapping in latent space](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 2712–2723, Online. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Rose Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir R. Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2022. [Crosslingual generalization through multitask finetuning](#). *ArXiv*, abs/2211.01786.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder,

- Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Barun Patra, Joel Ruben Antony Moniz, Sarthak Garg, Matthew R. Gormley, and Graham Neubig. 2019. [Bilingual lexicon induction with semi-supervision in non-isometric embedding spaces](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 184–193, Florence, Italy. Association for Computational Linguistics.
- Xutan Peng, Chenghua Lin, and Mark Stevenson. 2021. [Cross-lingual word embedding refinement by \$\ell_1\$ norm optimisation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2690–2701, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI Blog*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. [A survey of cross-lingual word embedding models](#). *Journal of Artificial Intelligence Research*, 65:569–631.
- Vin Sachidananda, Ziyi Yang, and Chenguang Zhu. 2021. [Filtered inner product projection for crosslingual embedding alignment](#). In *International Conference on Learning Representations*.
- Oleh Shliazhko, Alena Fenogenova, Maria Tikhonova, Vladislav Mikhailov, Anastasia Kozlova, and Tatiana Shavrina. 2022. [mgpt: Few-shot learners go multilingual](#). *ArXiv*, abs/2204.07580.
- Jimin Sun, Hwijee Ahn, Chan Young Park, Yulia Tsvetkov, and David R. Mortensen. 2021. [Cross-cultural similarity features for cross-lingual transfer learning of pragmatically motivated tasks](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2403–2414, Online. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *ArXiv*, abs/2302.13971.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George F. Foster. 2022. [Prompting palm for translation: Assessing strategies and performance](#). *ArXiv*, abs/2211.09102.
- Ivan Vulić, Goran Glavaš, Fangyu Liu, Nigel Collier, Edoardo Maria Ponti, and Anna Korhonen. 2023. [Probing cross-lingual lexical knowledge from multilingual sentence encoders](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2089–2105, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ivan Vulić, Goran Glavaš, Roi Reichart, and Anna Korhonen. 2019. [Do we really need fully unsupervised cross-lingual embeddings?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 4407–4418, Hong Kong, China. Association for Computational Linguistics.
- Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. 2020. [Probing pretrained language models for lexical semantics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 7222–7240, Online. Association for Computational Linguistics.
- Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2022. [Expanding pretrained models to thousands more languages via lexicon-based adaptation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 863–877, Dublin, Ireland. Association for Computational Linguistics.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2020. [CCNet: Extracting high quality monolingual datasets from web crawl data](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4003–4012, Marseille, France. European Language Resources Association.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. [Normalized word embedding and orthogonal transform for bilingual word translation](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011, Denver, Colorado. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and

- Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Jinpeng Zhang, Baijun Ji, Nini Xiao, Xiangyu Duan, Min Zhang, Yangbin Shi, and Weihua Luo. 2021. [Combining static word embeddings and contextual representations for bilingual lexicon induction](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2943–2955, Online. Association for Computational Linguistics.
- Xu Zhao, Zihao Wang, Yong Zhang, and Hao Wu. 2020. [A relaxed matching procedure for unsupervised BLI](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3036–3041, Online. Association for Computational Linguistics.
- Yucheng Zhou, Xiubo Geng, Tao Shen, Wenqiang Zhang, and Daxin Jiang. 2021. [Improving zero-shot cross-lingual transfer for multilingual question answering over knowledge graph](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5822–5834, Online. Association for Computational Linguistics.

A Languages

Family	Language	Code	LLMs
IE:Germanic	English	EN	mT5,mT0,XGLM,mGPT,LLaMA
	German	DE	mT5,mT0,XGLM,mGPT,LLaMA
IE:Romance	Catalan	CA	mT5,mT0,XGLM,LLaMA
	French	FR	mT5,mT0,XGLM,mGPT,LLaMA
	Italian	IT	mT5,mT0,XGLM,mGPT,LLaMA
IE:Slavic	Bulgarian	BG	mT5,mT0,XGLM,mGPT,LLaMA
	Croatian	HR	LLaMA
	Russian	RU	mT5,mT0,XGLM,mGPT,LLaMA
Turkic	Turkish	TR	mT5,mT0,XGLM,mGPT
Uralic	Finnish	FI	mT5,mT0,XGLM,mGPT
	Hungarian	HU	mT5,mT0,mGPT,LLaMA

Table 6: Languages covered in our experiments with their ISO 639-1 codes and the mLLM families that support that language, categorized by language family. IE = Indo-European.

B Reproducibility Checklist

- **BLI Data:** We adopt two publicly available BLI datasets.^{23 24}
- **Static Word Embeddings:** Following the datasets’ own recommendations and other previous work, we use the XLING-preprocessed fastText WEs trained on Wikipedia²⁵ for XLING data and fastText WEs trained on Common Crawl + Wikipedia²⁶ for PanLex-BLI, and the WEs are trimmed to the most frequent 200K words for each language. For fair comparisons, we use the same set of fastText WEs both for the retrieval of nearest neighbours (in-context examples) and for the CLWE-based baselines.
- **Pretrained LLMs and Parameter Counts:** All the LLMs used in our experiments are publicly available from the huggingface.co model hub. We summarise their model identifiers and model sizes in Table 7. Please refer to each LLM’s own copyright and licence before downloading, using, fine-tuning, or redistributing any LLM.
- **Source Code:** Our code is publicly available at <https://github.com/cambridgeltl/prompt4bli>.
- **Computing Infrastructure:** We have run our code on **Wilkes3**, a GPU cluster hosted by **Re-**

²³<https://github.com/codogogo/xling-eval>

²⁴<https://github.com/cambridgeltl/panlex-bli>

²⁵<https://fasttext.cc/docs/en/pretrained-vectors.html>

²⁶<https://fasttext.cc/docs/en/crawl-vectors.html>

search Computing Services at the University of Cambridge, where each run leverages a single Nvidia 80GB A100 GPU and 32× CPU cores.

- **Software:** Slurm 20.11.9, Python 3.9.7, PyTorch 1.10.1+cu113, Transformers 4.28.1.

- **Runtime (Wall Time):** We present the average inference time on one single BLI direction (i.e., circa 2,000 word pairs in an XLING test set; the time required for loading the LLM and the dataset is not included) for each LLM in Table 8. The per-epoch training time for BLI-oriented fine-tuning is provided in Table 9.

- **Hyperparameter Search:** As introduced in §3.3 and §4, our template selection and all our hyperparameter search are conducted on a single randomly chosen language pair (German, French), following previous work (Li et al., 2022a,b). The learning rate for LLM fine-tuning is selected from $[1e-9, 5e-9, 1e-8, 2e-8, 1e-7, 1e-6, 2e-6, 2e-5, 1e-3]$.

- **Significance:** We have discussed the significance of our main results and ablation results in the last paragraph of §5.1 and in Table 4 respectively, which demonstrates that our findings are statistically significant.

- **Randomness:** Our main experiments are completely deterministic since we rely on off-the-shelf LLMs without any fine-tuning, nearest neighbour retrieval for in-context examples (a deterministic retrieval algorithm), and the deterministic beam search. The randomness only exists in two parts of our side analysis. First, we use random in-context examples in our ablation study, and we verify our findings with statistical tests in Table 4. Second, the fine-tuning experiments do have randomness, and we run fine-tuning three times for each model, reporting both average BLI performance and the standard deviation.

- **Carbon Footprint:** All the experiments involved in this project including hyperparameter tuning, template search, BLI inference, and BLI-oriented fine-tuning of our LLMs consume circa 1,650 A100 GPU hours. Based on a publicly available ‘machine learning emissions calculator’ (Luccioni et al., 2019)²⁷ and our computational infrastructure, we estimate that our work causes the emission of circa 200kg CO₂ equivalents.

²⁷<https://mlco2.github.io/impact/#compute>

LLM	Model ID	Number of Parameters
mT5 _{small}	"google/mt5-small"	300,176,768
mT5 _{base}	"google/mt5-base"	582,401,280
mT5 _{large}	"google/mt5-large"	1,229,581,312
mT5 _{xl}	"google/mt5-xl"	3,742,619,648
mT5 _{xxl}	"google/mt5-xxl"	12,921,057,280
mT0 _{small}	"bigscience/mt0-small"	300,176,768
mT0 _{base}	"bigscience/mt0-base"	582,401,280
mT0 _{large}	"bigscience/mt0-large"	1,229,581,312
mT0 _{xl}	"bigscience/mt0-xl"	3,742,619,648
mT0 _{xxl}	"bigscience/mt0-xxl"	12,921,057,280
XGLM _{564M}	"facebook/xglm-564M"	564,463,616
XGLM _{1.7B}	"facebook/xglm-1.7B"	1,732,907,008
XGLM _{2.9B}	"facebook/xglm-2.9B"	2,941,505,536
XGLM _{4.5B}	"facebook/xglm-4.5B"	4,552,511,488
XGLM _{7.5B}	"facebook/xglm-7.5B"	7,492,771,840
mGPT	"sberbank-ai/mGPT"	1,417,596,928
LLaMA _{7B}	"huggyllama/llama-7b"	6,738,415,616
LLaMA _{13B}	"huggyllama/llama-13b"	13,015,864,320

Table 7: LLMs used in our experiments with their huggingface.co model IDs and model sizes.

LLM	Batch Size (Inference)	0-Shot	5-Shot
mT5 _{small}	100	6s	7s
mT5 _{base}	100	7s	8s
mT5 _{large}	100	8s	12s
mT5 _{xl}	8	45s	50s
mT5 _{xxl}	8	52s	83s
mT0 _{small}	100	5s	6s
mT0 _{base}	100	6s	7s
mT0 _{large}	100	8s	14s
mT0 _{xl}	8	46s	50s
mT0 _{xxl}	8	57s	78s
XGLM _{564M}	1	213s	225s
XGLM _{1.7B}	1	228s	237s
XGLM _{2.9B}	1	343s	366s
XGLM _{4.5B}	1	336s	394s
XGLM _{7.5B}	1	382s	461s
mGPT	1	192s	210s
LLaMA _{7B}	1	328s	463s
LLaMA _{13B}	1	434s	636s

Table 8: Inference time (in seconds) of each LLM with 0-Shot and 5-Shot prompts respectively.

LLM	Batch Size (Training)	5K Setup	1K Setup
mT5 _{base}	32	17 min	4 min
mT5 _{large}	32	38 min	8 min
XGLM _{564M}	32	24 min	5 min
XGLM _{1.7B}	16	80 min	15 min

Table 9: Per-epoch training time (in minutes) of each LLM with 5-Shot prompts in 5K and 1K setups respectively.

C Templates

We summarise all our zero-shot templates in Table 10 and few-shot templates in Table 11: these 102 templates constitute our ‘template pool’. Each of Tables 10 and 11 is split into two parts for mask-filling-style and GPT-style templates respectively as introduced in §3.1 and §3.2. In addition, we list the best zero-shot template and the best few-shot template for each of our 18 LLMs in Table 12. Again, as already mentioned in §3.3, our template selection is conducted on a randomly chosen language pair (German, French) where for few-shot templates, the in-context examples are derived from a seed dictionary of size 5K. While we do not have enough computational resources to calculate and do not have enough space to present the performance of each template for each LLM on each XLING BLI direction ($102 \times 18 \times 56 = 102,816$ scores), in the last paragraph of §5.2 we have discussed some preliminary findings only from our template search.

D Full BLI Results

Here we present our full results on both XLING and PanLex-BLI. Table 13, 14, and 16 are our results on all 56 XLING BLI directions in 5K, 1K, and zero-shot (unsupervised) BLI setups respectively. Table 15 and 17 are results for PanLex-BLI lower-resource languages (6 BLI directions) in 1K and zero-shot (unsupervised) BLI setups. Note that an (m)LLM usually cannot support every language, and we use ‘-’ to denote this scenario. Throughout this paper, our expression ‘a language is not supported by an LLM’ means that the language is not used for pretraining the LLM even if the LLM’s tokeniser may still be able to tokenise possibly many input sentences in the language. Table 18 shows the full ablation results for each of our 18 mLLMs.

E Translation Examples

To illustrate how few-shot learning improves BLI, we present some of our BLI results with LLaMA_{13B} in Table 19 comparing five-shot and zero-shot prompting on HR→EN and IT→EN BLI test sets.

Mask-Filling-Style Templates (Zero-Shot Prompting)	
1	The word ' w^{x^1} ' in L^y is: <mask>.
3	The word ' w^{x^1} ' in L^y is: <mask>.
5	The L^x word ' w^{x^1} ' in L^y is: <mask>.
7	The L^x word ' w^{x^1} ' in L^y is: <mask>.
9	The L^x word ' w^{x^1} ' in L^y is: <mask>.
11	The L^x word ' w^{x^1} ' in L^y is: <mask>.
13	' w^{x^1} ' in L^y is: <mask>.
15	' w^{x^1} ' in L^y is: <mask>.
17	What is the translation of the word ' w^{x^1} ' into L^y ? <mask>.
19	What is the translation of the L^x word ' w^{x^1} ' into L^y ? <mask>.
21	The translation of the word ' w^{x^1} ' into L^y is <mask>.
23	The translation of the L^x word ' w^{x^1} ' into L^y is <mask>.
25	How do you say ' w^{x^1} ' in L^y ? <mask>.
27	How do you say the L^x word ' w^{x^1} ' in L^y ? <mask>.
29	Translate the word ' w^{x^1} ' into L^y : <mask>.
31	Translate ' w^{x^1} ' into L^y : <mask>.
33	Translate the L^x word ' w^{x^1} ' into L^y : <mask>.
35	Translate from L^x to L^y : ' w^{x^1} '-> <mask>.
37	Translate from L^x to L^y : ' w^{x^1} '=> <mask>.
2	The word ' w^{x^2} ' in L^y is: <mask>.
4	The word ' w^{x^2} ' in L^y is: <mask>.
6	The L^x word ' w^{x^2} ' in L^y is: <mask>.
8	The L^x word ' w^{x^2} ' in L^y is: <mask>.
10	The L^x word ' w^{x^2} ' in L^y is: <mask>.
12	The L^x word ' w^{x^2} ' in L^y is: <mask>.
14	' w^{x^2} ' in L^y is: <mask>.
16	' w^{x^2} ' in L^y is: <mask>.
18	What is the translation of the word ' w^{x^2} ' into L^y ? <mask>.
20	What is the translation of the L^x word ' w^{x^2} ' into L^y ? <mask>.
22	The translation of the word ' w^{x^2} ' into L^y is <mask>.
24	How do you say ' w^{x^2} ' in L^y ? <mask>.
26	How do you say the L^x word ' w^{x^2} ' in L^y ? <mask>.
28	Translate the word ' w^{x^2} ' into L^y : <mask>.
30	Translate the word ' w^{x^2} ' into L^y : <mask>.
32	Translate the L^x word ' w^{x^2} ' into L^y : <mask>.
34	Translate from L^x to L^y : ' w^{x^2} '-> <mask>.
36	Translate from L^x to L^y : ' w^{x^2} '=> <mask>.
GPT-Style Templates (Zero-Shot Prompting)	
38	The word ' w^{x^1} ' in L^y is:
40	The word ' w^{x^2} ' in L^y is:
42	The L^x word ' w^{x^1} ' in L^y is:
44	The L^x word ' w^{x^1} ' in L^y is:
46	' w^{x^1} ' in L^y is:
47	Translate the word ' w^{x^1} ' into L^y :
49	Translate from L^x to L^y : ' w^{x^1} '->
51	Translate ' w^{x^1} ' into L^y :
53	Translate the L^x word ' w^{x^1} ' into L^y :
55	What is the translation of the word ' w^{x^1} ' into L^y ?
57	The translation of the word ' w^{x^1} ' into L^y is:
59	The translation of the L^x word ' w^{x^1} ' into L^y is:
61	How do you say ' w^{x^1} ' in L^y ?
62	How do you say ' w^{x^1} ' in L^y ?
64	How do you say the L^x word ' w^{x^1} ' in L^y ?
65	How do you say the L^x word ' w^{x^1} ' in L^y ?
66	Q: What is the L^y translation of ' w^{x^1} '? A:

Table 10: Our 66 templates for zero-shot prompting. These include 37 mask-filling-style templates (template IDs: 1 ~ 37) and 29 GPT-style templates (template IDs: 38 ~ 66). In our experiments, the '<mask>' token is '<extra_id_0>' for mT5 and mT0.

Mask-Filling-Style Templates (Few-Shot Prompting)	
67	Translate from L^x to L^y : $w_1^x \rightarrow w_1^y$ $w_2^x \rightarrow w_2^y$ $w^x \rightarrow$ <mask>.
68	Translate from L^x to L^y : $w_1^x \rightarrow w_1^y$, $w_2^x \rightarrow w_2^y$, $w^x \rightarrow$ <mask>.
69	Translate from L^x to L^y : $w_1^x \rightarrow w_1^y$ $w_2^x \rightarrow w_2^y$ $w^x \rightarrow$ <mask>
70	Translate from L^x to L^y : $w_1^x \rightarrow w_1^y$, $w_2^x \rightarrow w_2^y$, $w^x \rightarrow$ <mask>
71	Translate from L^x to L^y : $w_1^x \Rightarrow w_1^y$ $w_2^x \Rightarrow w_2^y$ $w^x \Rightarrow$ <mask>.
72	Translate from L^x to L^y : $w_1^x \Rightarrow w_1^y$, $w_2^x \Rightarrow w_2^y$, $w^x \Rightarrow$ <mask>.
73	Translate from L^x to L^y : $w_1^x \Rightarrow w_1^y$ $w_2^x \Rightarrow w_2^y$ $w^x \Rightarrow$ <mask>
74	Translate from L^x to L^y : $w_1^x \Rightarrow w_1^y$, $w_2^x \Rightarrow w_2^y$, $w^x \Rightarrow$ <mask>
75	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . The word w^x in L^y is <mask>.
76	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is <mask>.
77	The L^x word ' w_1^x ' in L^y is ' w_1^y '. The L^x word ' w_2^x ' in L^y is ' w_2^y '. The L^x word ' w^x ' in L^y is <mask>.
78	The L^x word w_1^x in L^y is w_1^y , The L^x word w_2^x in L^y is w_2^y , The L^x word w^x in L^y is <mask>.
GPT-Style Templates (Few-Shot Prompting)	
79	Translate from L^x to L^y : $w_1^x \rightarrow w_1^y$ $w_2^x \rightarrow w_2^y$ $w^x \rightarrow$
80	Translate from L^x to L^y : $w_1^x \rightarrow w_1^y$, $w_2^x \rightarrow w_2^y$, $w^x \rightarrow$
81	Translate from L^x to L^y : $w_1^x \Rightarrow w_1^y$ $w_2^x \Rightarrow w_2^y$ $w^x \Rightarrow$
82	Translate from L^x to L^y : $w_1^x \Rightarrow w_1^y$, $w_2^x \Rightarrow w_2^y$, $w^x \Rightarrow$
83	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . The word w^x in L^y is
84	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . The word w^x in L^y is:
85	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
86	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is:
87	The word w_1^x in L^y is w_1^y , The word w_2^x in L^y is w_2^y , The word w^x in L^y is
88	The word w_1^x in L^y is w_1^y , The word w_2^x in L^y is w_2^y , The word w^x in L^y is:
89	The L^x word w_1^x in L^y is w_1^y , The L^x word w_2^x in L^y is w_2^y , The L^x word w^x in L^y is
90	The L^x word w_1^x in L^y is w_1^y , The L^x word w_2^x in L^y is w_2^y , The L^x word w^x in L^y is:
91	The word ' w_1^x ' in L^y is w_1^y . The word ' w_2^x ' in L^y is w_2^y . The word ' w^x ' in L^y is
92	The word ' w_1^x ' in L^y is w_1^y . The word ' w_2^x ' in L^y is w_2^y . The word ' w^x ' in L^y is:
93	The L^x word ' w_1^x ' in L^y is w_1^y . The L^x word ' w_2^x ' in L^y is w_2^y . The L^x word ' w^x ' in L^y is
94	The L^x word ' w_1^x ' in L^y is w_1^y . The L^x word ' w_2^x ' in L^y is w_2^y . The L^x word ' w^x ' in L^y is:
95	The word ' w_1^x ' in L^y is w_1^y , The word ' w_2^x ' in L^y is w_2^y , The word ' w^x ' in L^y is
96	The word ' w_1^x ' in L^y is w_1^y , The word ' w_2^x ' in L^y is w_2^y , The word ' w^x ' in L^y is:
97	The L^x word ' w_1^x ' in L^y is w_1^y , The L^x word ' w_2^x ' in L^y is w_2^y , The L^x word ' w^x ' in L^y is
98	The L^x word ' w_1^x ' in L^y is w_1^y , The L^x word ' w_2^x ' in L^y is w_2^y , The L^x word ' w^x ' in L^y is:
99	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . How do you say w^x in L^y ?
100	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . How do you say the L^x word w^x in L^y ?
101	The word ' w_1^x ' in L^y is w_1^y . The word ' w_2^x ' in L^y is w_2^y . How do you say ' w^x ' in L^y ?
102	The L^x word ' w_1^x ' in L^y is w_1^y . The L^x word ' w_2^x ' in L^y is w_2^y . How do you say the L^x word ' w^x ' in L^y ?

Table 11: Our 36 templates for few-shot prompting. For simplicity, we present only two in-context examples in each template. These include 12 mask-filling-style templates (template IDs: 67 ~ 78) and 24 GPT-style templates (template IDs: 79 ~ 102). In our experiments, the '<mask>' token is '<extra_id_0>' for mT5 and mT0.

LLM	Best Template (Zero-Shot)
mT5 _{small}	The word ' w^x ' in L^y is: <mask>.
mT5 _{base}	Translate the word ' w^x ' into L^y : <mask>.
mT5 _{large}	The L^x word ' w^x ' in L^y is: <mask>.
mT5 _{xl}	The L^x word ' w^x ' in L^y is: <mask>.
mT5 _{xxl}	The L^x word ' w^x ' in L^y is: <mask>
mT0 _{small}	Translate from L^x to L^y : $w^x \Rightarrow$ <mask>.
mT0 _{base}	Translate from L^x to L^y : $w^x \Rightarrow$ <mask>.
mT0 _{large}	Q: What is the L^y translation of w^x A:
mT0 _{xl}	How do you say the L^x word ' w^x ' in L^y ?
mT0 _{xxl}	How do you say the L^x word ' w^x ' in L^y ?
XGLM _{564M}	The L^x word w^x in L^y is:
XGLM _{1.7B}	The L^x word w^x in L^y is:
XGLM _{2.9B}	The L^x word w^x in L^y is:
XGLM _{4.5B}	The L^x word w^x in L^y is:
XGLM _{7.5B}	The L^x word w^x in L^y is:
mGPT	Translate the L^x word w^x into L^y :
LLaMA _{7B}	The L^x word w^x in L^y is:
LLaMA _{13B}	Translate from L^x to L^y : $w^x \Rightarrow$
LLM	Best Template (Few-Shot)
mT5 _{small}	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . The word w^x in L^y is <mask>.
mT5 _{base}	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . The word w^x in L^y is <mask>.
mT5 _{large}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is <mask>.
mT5 _{xl}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is <mask>.
mT5 _{xxl}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is <mask>.
mT0 _{small}	The word ' w_1^x ' in L^y is w_1^y . The word ' w_2^x ' in L^y is w_2^y . How do you say ' w^x ' in L^y ?
mT0 _{base}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
mT0 _{large}	The L^x word ' w_1^x ' in L^y is ' w_1^y '. The L^x word ' w_2^x ' in L^y is ' w_2^y '. The L^x word ' w^x ' in L^y is '<mask>'.
mT0 _{xl}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is:
mT0 _{xxl}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is:
XGLM _{564M}	The word w_1^x in L^y is w_1^y . The word w_2^x in L^y is w_2^y . The word w^x in L^y is
XGLM _{1.7B}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
XGLM _{2.9B}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
XGLM _{4.5B}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
XGLM _{7.5B}	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
mGPT	The L^x word w_1^x in L^y is w_1^y . The L^x word w_2^x in L^y is w_2^y . The L^x word w^x in L^y is
LLaMA _{7B}	The L^x word ' w_1^x ' in L^y is w_1^y . The L^x word ' w_2^x ' in L^y is w_2^y . The L^x word ' w^x ' in L^y is
LLaMA _{13B}	The L^x word ' w_1^x ' in L^y is w_1^y . The L^x word ' w_2^x ' in L^y is w_2^y . The L^x word ' w^x ' in L^y is

Table 12: The best template for each LLM respectively for zero-shot and few-shot prompting. For simplicity, we present only two in-context examples in each template in the few-shot setup.

[5K Setup]	VecMAP	CONTRASTIVE BLI	mT5 _{small}	mT5 _{base}	mT5 _{large}	mT5 _{xl}	mT5 _{xxl}	mT0 _{small}	mT0 _{base}	mT0 _{large}	mT0 _{xl}	mT0 _{xxl}	XGLM _{664M}	XGLM _{1.7B}	XGLM _{2.9B}	XGLM _{4.1B}	XGLM _{7.0B}	mGPT	LLaMA _{7B}	LLaMA _{13B}
DE→FI	33.59	44.65	7.82	11.79	25.56	31.72	35.63	0.63	4.49	3.91	15.81	27.18	20.55	31.61	35.89	32.97	39.96	19.67	-	-
FI→DE	38.73	47.03	10.98	17.76	32.79	37.89	41.99	2.52	7.09	6.15	19.81	27.75	16.5	31.48	40.25	40.04	40.15	24.96	-	-
DE→FR	50.44	55.56	14.08	23.53	39.12	47.1	51.07	2.45	10.28	11.95	28.95	36.83	19.72	34.95	43.09	49.5	48.57	37.04	56.13	59.68
FR→DE	47.75	53.29	12.78	21.11	37.35	41.54	44.49	2.17	7.71	7.4	21.68	32.95	14.38	27.42	40.82	45.89	42.89	31.45	51.58	53.75
DE→HR	32.08	42.41	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	39.65	44.08
HR→DE	37.24	48.29	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	43.08
DE→IT	50.55	54.77	12.26	22.27	36.67	43.45	48.15	1.41	6.26	7.2	24.36	35.47	18.41	32.81	43.77	48.51	45.23	39.96	53.63	57.75
IT→DE	47.29	53.8	11.78	19.28	34.88	38.81	43.67	2.58	7.49	7.03	20.31	28.94	14.06	26.1	36.74	41.55	38.4	29.15	47.86	51.52
DE→RU	34.38	46.79	8.61	22.33	35.16	41.84	43.24	0.0	0.37	0.26	17.63	31.82	19.77	34.85	41.99	44.08	44.03	32.19	47.57	49.45
RU→DE	43.32	49.71	10.11	21.9	37.3	41.17	44.21	1.52	4.61	8.38	21.11	29.39	14.56	24.83	40.54	43.84	41.44	31.22	48.35	50.39
DE→TR	27.18	38.86	7.3	12.31	25.2	29.53	36.41	1.46	4.54	4.49	15.18	21.44	13.15	19.25	28.59	30.41	31.92	22.9	-	-
TR→DE	29.93	40.95	8.31	13.26	28.75	32.48	38.18	1.6	6.28	5.38	17.2	24.87	11.08	18.8	26.09	32.11	28.86	24.07	-	-
EN→DE	51.0	57.75	16.8	30.25	45.9	49.35	53.1	3.3	8.55	8.45	25.5	39.35	21.15	38.6	52.5	54.95	50.3	38.05	62.3	63.45
DE→EN	55.24	58.95	19.51	30.26	42.62	49.5	53.94	0.37	9.65	12.0	29.94	38.81	29.11	42.31	47.94	51.64	51.17	40.58	57.28	58.89
EN→FI	37.75	47.15	9.6	18.75	34.25	40.8	44.65	1.3	4.35	3.4	19.15	30.25	25.4	39.9	48.1	41.7	48.65	21.3	-	-
FI→EN	43.51	50.55	10.98	21.97	36.63	43.51	48.82	0.16	4.83	6.36	21.39	34.05	28.64	44.19	50.24	46.14	53.18	27.59	-	-
EN→FR	63.1	67.2	23.95	39.5	55.6	62.45	66.8	5.0	15.5	16.95	38.75	50.95	32.35	55.95	63.8	67.8	66.75	50.25	72.95	76.25
FR→EN	62.75	65.49	27.16	42.42	50.75	57.84	61.41	0.47	14.85	17.95	37.4	47.28	40.87	54.22	57.37	60.99	60.01	50.44	64.82	67.3
EN→HR	34.05	47.2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	48.15	54.7
HR→EN	39.08	49.08	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	48.13
EN→IT	60.4	65.6	21.45	38.3	54.75	61.75	63.85	2.3	5.9	6.35	34.05	48.0	29.7	51.1	60.65	64.55	62.9	52.4	71.4	74.0
IT→EN	62.17	65.27	20.26	35.35	47.03	54.32	59.43	0.31	10.8	14.42	33.49	43.0	34.68	48.17	53.23	57.73	57.11	48.63	63.05	64.44
EN→RU	39.65	50.5	11.85	29.6	45.6	52.8	53.7	0.2	0.35	0.25	22.0	38.5	30.95	46.4	54.0	57.5	55.05	42.05	58.3	62.3
RU→EN	49.35	54.16	12.31	29.39	43.16	51.86	54.58	0.16	4.3	11.05	27.76	36.83	30.12	43.79	52.28	53.95	53.64	40.96	56.1	58.77
EN→TR	32.05	44.75	10.55	20.65	35.45	40.15	45.05	2.65	4.2	3.8	18.6	27.05	16.9	29.35	38.4	40.6	40.85	31.05	-	-
TR→EN	39.24	44.78	9.42	20.77	32.43	39.51	44.36	0.16	3.83	7.61	18.0	30.4	19.81	33.07	39.24	39.83	42.17	29.71	-	-
FI→FR	38.26	45.24	8.57	18.5	35.47	42.35	46.51	1.63	6.2	5.83	24.54	32.05	20.91	37.94	46.66	45.45	48.77	28.69	-	-
FR→FI	34.3	43.2	6.93	12.42	27.32	34.2	37.61	0.83	3.16	2.64	15.99	25.3	20.49	29.07	39.16	35.28	42.52	19.66	-	-
FI→HR	31.58	38.31	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
HR→FI	31.72	39.56	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
FI→IT	37.99	46.3	10.3	17.5	32.16	39.83	44.51	1.52	6.78	5.25	21.97	30.06	21.86	36.57	44.04	44.09	47.08	27.75	-	-
IT→FI	34.32	43.57	7.7	11.73	25.22	32.09	35.09	0.88	4.55	4.29	16.18	23.41	17.26	29.35	36.43	31.47	40.67	18.35	-	-
FI→RU	34.16	40.99	6.36	16.34	30.53	38.05	41.04	0.05	0.53	0.68	15.55	27.01	19.23	36.1	42.3	39.99	44.82	24.38	-	-
RU→FI	33.53	40.91	4.56	13.51	26.72	34.31	35.57	0.37	1.94	2.78	14.14	24.1	16.29	27.45	34.36	34.47	40.81	18.54	-	-
HR→FR	40.24	49.29	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	46.66	50.34
FR→HR	33.21	44.08	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	40.3	46.51
HR→IT	40.24	48.97	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	46.87	50.6
IT→HR	34.32	44.75	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	38.5	44.75
HR→RU	37.98	46.4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	43.29	46.87
RU→HR	39.5	45.47	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	40.07	45.42
IT→FR	65.89	67.86	16.43	31.63	49.87	56.59	59.84	2.53	12.56	13.49	34.83	44.03	23.67	41.76	53.44	58.81	55.71	49.04	63.51	66.77
FR→IT	64.72	67.2	17.64	33.83	51.37	57.68	60.42	1.5	5.33	5.59	32.64	44.85	23.07	41.54	57.37	60.27	59.18	54.58	65.08	68.03
RU→FR	47.51	52.7	11.26	24.2	41.12	48.82	51.44	1.31	6.23	9.17	26.51	35.78	21.74	35.57	46.36	52.38	51.13	39.55	55.32	58.04
FR→RU	38.23	48.06	8.64	23.33	38.18	44.65	47.54	0.0	0.47	0.26	19.61	33.73	25.92	39.89	46.82	50.08	49.77	39.63	51.22	54.16
RU→IT	46.78	51.96	11.79	26.14	41.59	47.77	49.29	0.89	4.45	7.18	24.83	34.36	23.47	32.58	46.88	50.13	50.6	41.28	53.54	57.83
IT→RU	38.71	47.49	8.22	22.02	36.9	43.82	46.41	0.1	0.36	0.26	18.66	30.96	22.58	36.9	43.26	47.39	47.75	35.45	46.61	51.89
TR→FI	28.59	34.77	5.54	8.47	20.18	27.1	29.77	0.43	3.3	2.08	14.11	18.85	13.47	22.79	26.94	25.03	32.8	16.61	-	-
FI→TR	29.8	35.68	7.67	12.19	22.86	31.0	36.21	1.68	5.52	3.84	14.66	20.13	14.61	23.44	31.21	30.22	35.73	20.18	-	-
TR→FR	36.58	43.88	8.95	17.25	31.79	39.99	44.62	1.38	4.95	5.86	22.63	29.5	15.55	25.83	33.76	37.7	38.82	29.55	-	-
FR→TR	31.76	42.06	8.48	14.38	30.42	34.97	40.61	2.28	4.81	3.36	16.81	23.49	13.4	22.76	31.71	33.94	34.71	27.78	-	-
TR→HR	25.99	36.32	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
HR→TR	27.35	37.09	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
TR→IT	34.24	42.17	7.35	15.28	30.46	38.13	42.55	1.22	5.32	4.58	19.97	28.12	14.64	24.17	35.46	36.85	37.01	31.63	-	-
IT→TR	30.7	40.62	8.11	14.21	26.25	31.94	37.26	2.12	5.37	3.36	16.12	22.02	12.61	21.71	31.01	32.09	34.63	25.06	-	-
TR→RU	26.2	36.16	4.47	13.53	29.39	35.57	38.02	0.0	0.27	0.43	13.63	25.03	13.26	27.32	34.45	34.93	36.47	27.21	-	-
RU→TR	27.08	35.78	5.97	14.35	27.08	31.69	36.09	1.47	3.25	4.24	14.51	19.91	11.58	19.75	30.23	32.63	35.73	24.62	-	-

Table 13: Full 5-shot BLI results (P@1×100%) on 56 XLING BLI directions with 5K seed translation pairs. Off-the-shelf LLMs are used without any fine-tuning. ‘-’: a language in the pair is not supported by the LLM.

Ground-Truth Translation Pair	Method	In-Context Examples	Prediction
brisati (HR) → erase (EN)	0-Shot	-	break
	5-Shot	izbrisati→delete, kopirati→copy, ispraviti→amend, dopuniti→supplement, pisati→write	erase
jestiv (HR) → edible (EN)	0-Shot	-	joyful
	5-Shot	jesti→eat, češnjak→garlic, povrće→vegetable, grašak→peas, biljka→plant	edible
slap (HR) → waterfall (EN)	0-Shot	-	slap
	5-Shot	slapovi→falls, kanjon→canyon, jezero→lake, potok→creek, potoci→streams	waterfall
valcer (HR) → waltz (EN)	0-Shot	-	valve
	5-Shot	orkestar→orchestra, koncert→concert, balet→ballet, simfonija→symphony, klavir→piano	waltz
šaran (HR) → carp (EN)	0-Shot	-	shark
	5-Shot	pastrva→trout, riba→fish, ribolovac→fisher, gavran→raven, bizaran→bizarre	carp
sove (HR) → owls (EN)	0-Shot	-	slept
	5-Shot	sova→owl, životinje→animals, ptica→bird, mačke→cats, životinja→animal	owls
gušter (HR) → lizard (EN)	0-Shot	-	gusher
	5-Shot	štakor→rat, zmija→snake, zvijer→beast, kornjača→turtle, majmun→monkey	lizard
sezam (HR) → sesame (EN)	0-Shot	-	cumin
	5-Shot	grašak→peas, rajčica→tomato, povrće→vegetable, rajčice→tomatoes, smokva→fig	sesame
odgođena (HR) → postponed (EN)	0-Shot	-	postponed
	5-Shot	odgoditi→delay, bila→been, trebala→supposed, najave→announcements, provedena→conducted	delayed
tajnost (HR) → secrecy (EN)	0-Shot	-	secrecy
	5-Shot	krajnost→extreme, odanost→loyalty, odlučnost→resolve, tajna→secret, točnost→accuracy	discretion
prognoza (HR) → prognosis (EN)	0-Shot	-	prognosis
	5-Shot	prognoze→forecasts, dijagnoza→diagnosis, terapija→therapy, liječenje→treatment, astma→asthma	forecasts
otapanje (HR) → dissolution (EN)	0-Shot	-	dissolution
	5-Shot	doziranje→dosage, kiseline→acids, filtriranje→filtering, dobivanje→obtaining, curenje→leak	melting
sopportare (IT) → endure (EN)	0-Shot	-	support
	5-Shot	alleviare→ease, portare→bring, superare→exceed, odiare→hate, minimizzare→minimize	endure
snello (IT) → slender (EN)	0-Shot	-	quick
	5-Shot	flessibile→flexible, coda→tail, sottile→subtle, leggermente→slightly, aspetto→appearance	slender
hindi (IT) → hindi (EN)	0-Shot	-	english
	5-Shot	india→india, indiano→indian, thailandese→thai, lingua→tongue, arabo→arabic	hindi
velo (IT) → veil (EN)	0-Shot	-	bicycle
	5-Shot	cappello→hat, mantellina→cape, viso→face, cuscino→pillow, vestito→dress	veil
prugna (IT) → plum (EN)	0-Shot	-	prickly
	5-Shot	ciliegia→cherry, vaniglia→vanilla, zenzero→ginger, marmellata→jam, aglio→garlic	plum
propano (IT) → propane (EN)	0-Shot	-	methane
	5-Shot	idrogeno→hydrogen, acido→acid, carbonio→carbon, combustibili→fuels, polimero→polymer	propane
pomello (IT) → knob (EN)	0-Shot	-	apple
	5-Shot	maniglia→handle, regolabile→adjustable, maniglie→handles, coperchio→lid, cinghia→strap	knob
altopiano (IT) → plateau (EN)	0-Shot	-	high
	5-Shot	montagne→mountains, pianure→plains, vallata→dale, pianura→plain, valle→valley	plateau
eventuale (IT) → eventual (EN)	0-Shot	-	eventual
	5-Shot	necessario→necessary, valutare→gauge, possibile→possible, necessaria→needed, richiedere→require	possible
scopre (IT) → discovers (EN)	0-Shot	-	discovers
	5-Shot	scoprire→discover, crede→believes, chiede→asking, spiega→explains, ragazza→girl	discover
rane (IT) → frogs (EN)	0-Shot	-	frogs
	5-Shot	topi→mice, animali→animals, rana→frog, mosche→flies, creature→creatures	snakes
sorprese (IT) → surprises (EN)	0-Shot	-	surprises
	5-Shot	sorpresa→surprise, inaspettato→unexpected, sorprendente→surprising, aspettandosi→expecting, aspettative→expectations	surprised

Table 19: Translation examples on HR→EN and IT→EN. We include here ground truth translation pairs and show the predictions derived from zero-shot and five-shot prompting with LLaMA_{13B}.