

Adapting Multilingual LLMs to Low-Resource Languages with Knowledge Graphs via Adapters

Daniil Gurgurov^{1,2} Mareike Hartmann² Simon Ostermann¹

¹German Research Center for Artificial Intelligence (DFKI)

²Department of Language Science and Technology, Saarland University

daniil.gurgurov@dfki.de, mareikeh@coli.uni-saarland.de, simon.ostermann@dfki.de

Abstract

This paper explores the integration of graph knowledge from linguistic ontologies into multilingual Large Language Models (LLMs) using adapters to improve performance for low-resource languages (LRLs) in sentiment analysis (SA) and named entity recognition (NER). Building upon successful parameter-efficient fine-tuning techniques, such as K-ADAPTER (Wang et al., 2021) and MAD-X (Pfeiffer et al., 2020), we propose a similar approach for incorporating knowledge from multilingual graphs, connecting concepts in various languages with each other through linguistic relationships, into multilingual LLMs for LRLs. Specifically, we focus on eight LRLs —Maltese, Bulgarian, Indonesian, Nepali, Javanese, Uyghur, Tibetan, and Sinhala — and employ language-specific adapters fine-tuned on data extracted from the language-specific section of ConceptNet, aiming to enable knowledge transfer across the languages covered by the knowledge graph. We compare various fine-tuning objectives, including standard Masked Language Modeling (MLM), MLM with full-word masking, and MLM with targeted masking, to analyze their effectiveness in learning and integrating the extracted graph data. Through empirical evaluation on language-specific tasks, we assess how structured graph knowledge affects the performance of multilingual LLMs for LRLs in SA and NER, providing insights into the potential benefits of adapting language models for low-resource scenarios.

1 Introduction

In recent years, the advancement of multilingual Large Language Models (LLMs) (Devlin et al., 2019; Conneau et al., 2020; Xue et al., 2021) has revolutionized the field of natural language processing (NLP), enabling impressive performance across various languages. However, these models often struggle with low-resource languages (LRLs), where limited data availability affects

their effectiveness (Wu and Dredze, 2020). To address this limitation, researchers have explored integrating external knowledge sources into multilingual LLMs to enhance their performance in both high-resource and low-resource contexts (Wang et al., 2021; Lauscher et al., 2020; Pfeiffer et al., 2020) via adapters (Houlsby et al., 2019) and full fine-tuning.

Adapters, introduced by Houlsby et al. (2019), are small modules inserted between the layers of a model and trained while the model is kept frozen. Previous work has used such Adapters to integrate external knowledge into LLMs. For instance, Wang et al. (2021) demonstrated improvements in relation classification, entity typing, and question answering tasks by integrating graph knowledge from Wikidata (Vrandečić and Krötzsch, 2014) into RoBERTa (Liu et al., 2019) using adapters. Similarly, Lauscher et al. (2020) enhanced BERT (Devlin et al., 2019) with graph knowledge from ConceptNet (Speer et al., 2017), achieving significant performance gains on tasks requiring common-sense knowledge. However, these efforts primarily focused on the English language. In contrast, Pfeiffer et al. (2020) addressed low-resource languages by integrating textual knowledge from Wikipedia into XLM-R (Conneau et al., 2020) via language adapters. Their approach demonstrated improvements over the baseline model for named entity recognition (NER) task.

Motivated by recent advancements in the integration of graph knowledge into language models, particularly for English, this paper investigates the incorporation of graph knowledge from linguistic ontologies, specifically ConceptNet, into multilingual LLMs particularly for LRLs. Injecting such data might be beneficial due to the scarcity of training data for these languages and the additional semantic and multilingual information provided by knowledge graphs (Miller, 1995; Speer

et al., 2017). Our focus is on a subset of LRLs, aiming to extend the success observed in graph knowledge integration to linguistically diverse and resource-scarce contexts. We work with Maltese, Bulgarian, Indonesian, Nepali, Javanese, Uyghur, Tibetan, and Sinhala, identified as low-resource according to Joshi et al. (2020). Our primary objective is to evaluate whether injecting multilingual graph knowledge, connecting various languages through linguistic relationships, into pre-trained multilingual LLMs through adapters improves performance for LRLs. We train language-specific adapters on ConceptNet data using different objective functions, including standard Masked Language Modeling (MLM) (Devlin et al., 2019), MLM with full-word masking (Cui et al., 2021), and MLM with targeted masking, and evaluate the downstream performance of the adapted model on sentiment analysis (SA) and NER tasks.

Our work extends existing advancements by proposing an approach that utilizes adapters to integrate graph knowledge specifically for LRLs, following a modular design similar to the one introduced by Pfeiffer et al. (2020). Our contributions include:

- **Low-Resource Languages Focus:** Unlike prior works on graph knowledge integration (Lauscher et al., 2020; Wang et al., 2021), our research concentrates explicitly on improving multilingual LLMs through the external graph knowledge injection for low-resource scenarios.
- **Exploiting Various Knowledge Sources and Types:** We investigate the integration of language adapters based on Wikipedia and ConceptNet, both individually and in combination. This expands the approach of Pfeiffer et al. (2020), which solely utilized Wikipedia data, enabling a comprehensive assessment of different knowledge sources’ impact on model performance. We assume that language models can benefit from the multilingual connections in ConceptNet.
- **Single-Language Training Approach:** In contrast to the multilingual transfer learning approach used by Pfeiffer et al. (2020), which primarily focuses on cross-lingual adaptation, our methodology involves training language and task adapters using data in the same LRL. This training strategy aims to maximize

model performance and adaptability to the specific linguistic characteristics of each target language.




Figure 1: Proposed method. One of the Wiki or ConceptNet language adapters is used during inference. The outputs then go to a task adapter, which is followed by a classification head. If fusion is specified, the fusion mechanism is activated.

Our study provides insights into the potential benefits of adapting multilingual LLMs for low resource scenarios, contributing to the ongoing exploration of multilingual language model adaptation and graph knowledge integration.

The complete code for our experiments is publicly available on GitHub¹.

2 Related Work

Our work extends recent advancements in integrating external graph knowledge into pre-trained LLMs (Lauscher et al., 2020; Wang et al., 2021) and adapting pre-trained multilingual LLMs to specific languages (Pfeiffer et al., 2020). We use the methodologies proposed in these studies as a foundation for enhancing the performance of multilingual LLMs on downstream tasks, particularly for LRLs. To this end, we provide the overview of these methods as well as other strategies of adapting multilingual LLMs to LRLs (Artetxe et al., 2020; Muller et al., 2021; Vernikos and Popescu-Belis, 2021; Pfeiffer et al., 2022).

2.1 Adapter-based Knowledge Integration

In our approach, we draw inspiration from recently proposed adapter-based knowledge integration techniques, particularly the concept of K-Adapters (Wang et al., 2021) and the work by

¹<https://github.com/d-gurgurov/Injecting-Commonsense-Knowledge-into-LLMs>

Lauscher et al. (2020). K-Adapters introduced a novel approach for injecting knowledge into pre-trained models like RoBERTa (Liu et al., 2019) without modifying their original parameters. This method utilizes two types of adapters dedicated to factual and linguistic knowledge, demonstrating improvements in tasks such as entity typing and question answering. Factual adapters are trained with a relation classification objective using data aligned from Wikipedia text to Wikidata triplets (Vrandečić and Krötzsch, 2014), while linguistic adapters are trained with a dependency relation prediction objective using linguistic information obtained from available dependency parsers.

Similarly, Lauscher et al. (2020) explored injecting external knowledge, specifically from ConceptNet (Speer et al., 2017) and the Open Mind Common Sense (OMCS) corpus (Singh et al., 2002), into language models. They introduced two boosted models: CN-ADAPT and OM-ADAPT. CN-ADAPT involves creating a synthetic corpus through random traversal of the ConceptNet graph, with adapter parameters learned through Masked Language Modeling (MLM) training on this synthetic corpus. In OM-ADAPT, adapter parameters are learned directly through MLM training on the OMCS corpus. Both models employ a parameter-efficient adapter-based architecture (Houlsby et al., 2019), injecting bottleneck adapters into BERT’s (Devlin et al., 2019) transformer layers.

An approach similar to ours is presented in Hou et al. (2022), who however make stronger assumptions on the training data (such as the alignment of entities in the data with a knowledge graph) and rather train models to represent entities explicitly.

2.2 Language Adapters

In our exploration of injecting graph knowledge for LRL scenarios, we follow the MAD-X architecture presented by Pfeiffer et al. (2020). MAD-X offers an efficient approach to adapt pre-trained language models to LRLs by utilizing a modular structure consisting of language adapters, task-specific adapters, and invertible adapters.

The MAD-X framework utilizes language adapters as a fundamental component to adapt the model to specific languages. These adapters are trained on language-specific Wikipedia data and stacked onto the pre-trained model, allowing the model to capture language-specific nuances and patterns effectively. Following the enhanced bottleneck architecture (Pfeiffer et al., 2021),

the language adapter involves down- and up-projections with ReLU activation and is trained via MLM on unlabeled data.

During downstream task training, such as named entity recognition (NER), the fixed language adapter corresponding to the source language is used, ensuring adaptability to different languages without changing the underlying multilingual model. The embeddings are passed through the fixed language adapter before entering the task adapter, facilitating efficient adaptation to diverse linguistic contexts and specific task requirements.

Additionally, MAD-X introduces invertible adapters to mitigate the mismatch between multilingual and target language vocabulary. These invertible adapters are stacked on top of the embedding layer, with their respective inverses preceding the output embedding layer.

Task-specific adapters are then stacked on top of the language and invertible adapters to capture task-specific knowledge and specialize a language model in a certain task.

These insights from the MAD-X framework serve as a valuable reference for our research on injecting structured graph knowledge into multilingual LLMs for low-resource cases.

2.3 Adapting LLMs to Low-Resource Languages

Various other strategies have been proposed to address the challenges of adapting multilingual LLMs to LRLs, particularly for languages with limited pre-training data. Pfeiffer et al. (2022) propose X-MOD, a modular multilingual architecture that integrates shared and language-specific parameters to overcome the curse of multilinguality (Conneau et al., 2020), allowing efficient handling of linguistic diversity and supporting the extension to new languages with minimal performance impact on pre-trained languages. Additionally, Artetxe et al. (2020) train a new embedding layer with a corresponding target-language tokenizer to extend monolingual models to new languages, aiding language extension while maintaining model stability. Moreover, approaches based on transliteration and subword mappings have been proposed to incorporate additional languages into multilingual models, contributing to the expansion of multilingual capabilities of LLMs (Muller et al., 2021; Vernikos and Popescu-Belis, 2021). Hangya et al. (2022) present a bootstrapping-based approach for

ConceptNet Relationship	Natural Language Predicate
Antonym	is the opposite of
DerivedFrom	is derived from
EtymologicallyDerivedFrom	is etymologically derived from
EtymologicallyRelatedTo	is etymologically related to
FormOf	is a form of
HasContext	has context of
IsA	is a type of
RelatedTo	is related to
SimilarTo	is similar to
Synonym	is a synonym of
SymbolOf	is a symbol of
DistinctFrom	is distinct from

Table 1: Predefined mapping from ConceptNet relations to natural language predicates used for training ConceptNet-based Language Adapters.

enhancing low-resource languages in multilingual LLMs, which relies on unsupervised word translation pairs from monolingual corpora.

3 Injecting External Knowledge into LLMs for LRLs

This section describes our approach for enhancing multilingual LLMs for LRLs by injecting external knowledge. We discuss the use of language adapters trained on ConceptNet and Wikipedia data, explore Adapter Fusion (Pfeiffer et al., 2021) for combining knowledge sources, and describe task adapters for fine-tuning multilingual LLMs for specific tasks. Our proposed method is illustrated in Figure 1.

3.1 Language Adapters

We use language adapters for integrating external knowledge into multilingual LLMs and adapting these models to a specific language. In our study, two types of language adapters are employed: those trained on ConceptNet data and those trained on Wikipedia.

3.1.1 ConceptNet Data Preparation

ConceptNet-based language adapters are trained on knowledge extracted from ConceptNet, providing a rich source of linguistic relationships and semantic information across various languages. Data preparation involves retrieving and formatting data from ConceptNet and converting it into natural text². The number of triples extracted for the cho-

²For the extraction process, we utilized a dedicated self-built module for fetching data from CN, built on top of the CN API (<https://github.com/commonsense/conceptnet5/wiki/API>) for an easier extraction of per-language data. Code available on <https://github.com/d-gurgurov/Conceptnet-Embeddings>

sen languages are given in Table 2. These triples were converted into natural language using a predefined mapping from ConceptNet relationships to natural language predicates. This mapping allows for a straightforward method for injecting the graph knowledge through MLM-like objectives. The relationship mapping includes all possible connections from the ontology for the selected languages and is as specified in Table 1. An example of constructing a natural language sentence from an extracted triple is as follows: the triple (*kiel, RelatedTo, eat*) is converted into the sentence "*kiel is related to eat*". In this context, "kiel" is the Maltese word for "eat." The natural language predicates are always kept in English, resulting in the generated text for a triple being multilingual.

3.1.2 ConceptNet-based Language Adapters

The ConceptNet language adapters are sequential bottleneck adapters (Pfeiffer et al., 2021), similar to the ones used in MAD-X, with modifications to exclude invertible adapter layers for simplicity. Further, different objective functions were used for various downstream tasks at hand. For Sentiment Analysis (SA), we used the standard MLM objective, whereas for Named Entity Recognition (NER) another self-designed objective function, targeted Masked Language Modeling (TLM), was used for training the language adapters on the graph knowledge. The latter objective implies predicting the masked tokens not included in a natural language predicates specified in Table 1. Following the earlier provided example with the sentence "*kiel is related to eat*", only either the word "*kiel*" or "*eat*" would be masked.

3.1.3 Wikipedia-based Language Adapters

In contrast, the language adapters trained on Wikipedia data utilize the Wikimedia dataset³ for selected LRLs. This dataset provides a diverse and extensive collection of textual information scraped from Wikipedia for each language of interest. The number of articles available for each language is as in Table 2. The adapter architecture for Wikipedia language adapters is the same as the ConceptNet language adapters and uses the standard MLM objective function⁴.

3.2 Fusion of Language Adapters

In our search of enhancing multilingual LLMs for LRLs, we extend our investigation to the fusion of knowledge sources through Adapter Fusion (Pfeiffer et al., 2021). The Adapter Fusion mechanism facilitates the integration of knowledge extracted from ConceptNet and Wikipedia-based language adapters, providing a non-destructive method to combine multiple pre-trained adapters for new downstream tasks.

An adapter fusion block introduces a set of parameters that dynamically combines adapters and the shared pre-trained model at each layer of the transformer. The fusion layer incorporates Key, Value, and Query matrices at each layer to learn contextual activation of each adapter. This dynamic combination is achieved through a contextual activation mechanism similar to attention mechanisms (Vaswani et al., 2017).

We activate the fusion layer with two language adapters for each language - the Wikipedia adapter and the ConceptNet adapter. This fusion layer is introduced to allow the model to learn the optimal way to dynamically compose the knowledge from different sources. The learnable weights (Query, Key, and Value) within adapter fusion should enable the model to identify and activate the most relevant information from each adapter based on the context of the task.

3.3 Task Adapters

To fine-tune multilingual LLMs for specific downstream tasks such as sentiment analysis (SA) and named entity recognition (NER), we employ task adapters stacked on top of language adapters.

³<https://huggingface.co/datasets/wikimedia/wikipedia>

⁴All extracted ConceptNet data used for the language adapters, along with the language adapters themselves, can be found on HuggingFace (<https://huggingface.co/DGurgurov>).

Language	ISO	CN	Wiki	mBERT?
Bulgarian	bg	58060	297516	✓
Indonesian	ms	44190	689034	✓
Nepali	ne	7497	33040	✓
Javanese	jv	5082	73311	✓
Maltese	mt	8578	6310	
Uyghur	ug	3225	5979	
Tibetan	bo	9532	7090	
Sinhala	si	3350	20454	

Table 2: Number of ConceptNet triples and Wikipedia articles per language. The last column indicates if the respective language was included in the mBERT pre-training data.

The architecture of the task adapters follows the established design, featuring a stack of task-specific adapters on the top layers of a multilingual LLM and language adapter. For our study, we specifically focus on SA and NER, aiming to evaluate the impact of external knowledge injection on sentiment classification and entity recognition in LRLs, as these tasks are the most accessible in terms of labeled data availability for the selected languages.

To maintain the knowledge of the pre-trained model and language adapters during task adaptation, we adopt a weight freezing strategy, as in MAD-X. This involves preventing further fine-tuning of the weights in the pre-trained model and language adapters when training the task adapters. By doing so, we ensure that the foundational knowledge captured by the language adapters, whether sourced from Wikipedia, ConceptNet, or their fusion, remains unchanged.

The task adapters are stacked on top of the language adapters, like in the original MAD-X architecture. This stacking configuration facilitates the flow of information from the base model and the external knowledge sources through the language adapters to the task-specific adapters.

4 Experiments

In this section, we detail the experiments and data used for conducting the study.

4.1 Languages

The focus languages, classified as low-resource according to Joshi et al. (2020), are Maltese, Bulgarian, Indonesian, Nepali, Javanese, Uyghur, Tibetan, and Sinhala, as presented in Table 2. These languages serve as a subset of underrepresented languages for injecting external knowl-

Model/Language	bg	ms	ne	jv	mt	ug	bo	si
Sentiment Analysis (SA)								
mBERT	0.860	0.888	0.565	0.728	0.557	0.696	0.687	0.646
mBERT+TA	0.885	0.917	0.565	0.761	0.598	0.734	0.801	0.661
mBERT+Wiki+TA	0.893	0.919	0.584	0.746	0.702	0.706	0.816	0.663
mBERT+CN+TA	0.893	0.915	0.636	0.751	0.658	0.699	0.803	0.653
mBERT+F(CN&Wiki)+TA	0.882	0.906	0.627	0.750	0.662	0.784	0.804	0.689
Named Entity Recognition (NER)								
mBERT	0.919	0.934	0.694	0.575	0.595	0.402	0.520	0.197
mBERT+TA	0.917	0.934	0.644	0.564	0.601	0.383	0.575	0.172
mBERT+Wiki+TA	0.915	0.932	0.610	0.543	0.603	0.411	0.576	0.172
mBERT+CN+TA	0.915	0.928	0.649	0.571	0.576	0.403	0.544	0.244
mBERT+F(CN&Wiki)+TA	0.888	0.889	0.713	0.503	0.563	0.401	0.540	0.165

Table 3: Experimental Results. All numbers are averaged over 3 independent runs. The scores in **bold** are the ones that outperform both baseline models.

edge through ConceptNet and Wikipedia-based language adapters. The choice is bounded to the ConceptNet and downstream tasks data available for the languages. While Bulgarian, Indonesian, Nepali and Javanese data were used for pre-training mBERT (Devlin et al., 2019), which is the multilingual LLM we will use for our experiments, Maltese, Tibetan, Uyghur and Sinhala were not included in the pre-training dataset.

4.2 Tasks

Two tasks considered for empirical evaluation are Sentiment Analysis (SA) and Named Entity Recognition (NER). Datasets for SA for all the languages are acquired from different sources (Martínez-García et al., 2021; Purwarianti and Crisdayanti, 2019; Cortis and Davis, 2019; Dingli and Sant, 2016; Singh et al., 2020; Wongso et al., 2021; Li et al., 2022; Zhu et al., 2023; Ranathunga and Liyanage, 2021) and available in our HuggingFace repositories⁵. For NER, the datasets are obtained from the WikiANN project (Pan et al., 2017). All datasets for both SA and NER contain various amounts of data, depending on language and task, and described in more detail in Appendix A. We use the vanilla F1 score (Sokolova et al., 2006) for SA performance monitoring and the "se-geval" F1 score (Nakayama, 2018) for NER.

While these datasets do not allow for a full assessment of the impact of injected graph knowledge on LRLs due to a lack of labeled data for other tasks, they serve as a good starting point for mea-

suring the effects of multilingual graph knowledge integration on LRLs.

4.3 Baselines

In establishing baseline models for our study, we fine-tune the widely used mBERT (bert-base-multilingual-cased) (Devlin et al., 2019) from the transformer library (Wolf et al., 2020). The focus is on two baseline scenarios for each task—SA and NER.

The first baseline involves fine-tuning mBERT directly on the respective datasets for SA and NER. We employ common hyperparameters for training, including a learning rate of $\{1e-4, 2e-4\}$, a batch size of 64, $\{50, 100\}$ epochs with best-model-saving, and a dropout rate of $\{0.5, 0.2\}$, respectively for each task. The choice of hyperparameters is aligned with standard practices in transformer-based model training and our own experiments on the given datasets.

For the second baseline (mBERT+TA), we introduce a single task adapter on top of mBERT and fine-tune it on the SA and NER datasets, while keeping the parameters of mBERT frozen, using the same hyperparameters as used for the first baseline configuration. This single adapter architecture allows us to explore the effectiveness of a more compact adaptation strategy compared to the traditional fine-tuning approach.

For all models, the best checkpoint for evaluation is selected based on validation loss performance.

⁵<https://huggingface.co/DGurgurov>

4.4 Language Adapter Training

In this section, we provide technical details of the training process for language adapters, focusing on both ConceptNet (*CN*) and Wikipedia (*Wiki*) variants. We use sequential bottleneck architecture without invertible layers for training language adapters.

We maintain identical hyperparameters during the training of both Wikipedia and ConceptNet language adapters: reduction factor 16, learning rate $5e - 05$, train and eval batch size 16, training steps 50,000 for CN and 100,000 for Wiki. These hyperparameters ensure a stable and uniform training environment for both variants. The training process is monitored through loss and accuracy metrics on the validation sets.

4.5 Task Adapters Training

In this setting, we stack task-specific adapters on top of language adapters and train the task adapters while keeping the language adapters frozen. After empirical investigations, we employ similar hyperparameters to those used for training the baselines: a learning rate of $\{1e - 5, 1e - 4\}$, a batch size of 64, $\{50, 100\}$ epochs with best-model saving, and a dropout of $\{0.5, 0.2\}$ for SA and NER, respectively. The experiment involves stacking the task adapters on top of either Wikipedia-based (*Wiki+TA*) or ConceptNet-based (*CN+TA*) language adapters, as well as on top of the fusion of both ($F(CN\&Wiki)+TA$) (Pfeiffer et al., 2021).

4.6 Objective Functions

In this section, we compare different objective functions used for training language adapters on graph knowledge and examine their influence on the performance on the downstream tasks at hand. We experiment with three objectives for language modeling - standard token Masked Language Modeling (MLM) (Devlin et al., 2019), full-word Masked Language Modeling (FLM) (Cui et al., 2021), and targeted Masked Language Modeling (TLM). MLM and FLM were implemented as provided by the Transformers library, and TLM was self-designed. MLM masks individual tokens with a 15% probability, FLM performs the same but masks full words, and TLM masks targeted words which are not part of the natural language predicates list extracted from ConceptNet with a 50% probability. The implied goal of TLM is to create the connections between the words of a LRL

to the words of other languages, which might come in bigger quantities, within the parameters of a model. The downstream results for all the objectives are as in Table 4. Upon the inspection of the results, different objective functions were chosen for SA and NER, according to the outcomes of the experiments. MLM was utilized as an objective for the language adapters used for SA, and TLM was chosen as an objective for the language adapters used for NER.

5 Results and Discussion

The experimental results, summarized in Table 3, demonstrate the results of different model configurations in improving SA and NER tasks across LRLs. This section discusses the performance of each model configuration and provide insights into the impact of external knowledge through language and task adapters on enhancing multilingual LLMs for LRLs. All scores are an average over three independent runs.

5.1 Sentiment Analysis

In SA, the performance of various model configurations on different languages reveals interesting insights. First, considering the baseline performance of fully fine-tuned mBERT across all languages, we observe moderate to high F1-scores. However, when single task-specific adapters are added to mBERT, we notice consistent improvements across all languages, indicating the effectiveness of using parameter-efficient fine-tuning techniques for adapting the model for specific languages, especially in low-resource scenarios. This confirms the findings by Li and Liang (2021), He et al. (2021), and Jukić and Snajder (2023).

When incorporating language adapters trained on CN and Wiki, relatively good results are observed. For nearly all languages, using language adapters trained on CN and Wiki leads to performance gains compared to the baselines-mBERT and mBERT with a single task adapter. The CN language adapter boosts the performance for Bulgarian, Nepalese, Maltese, and Tibetan over both baselines. As for the Wiki language adapters, they improve the scores for Bulgarian, Indonesian, Nepalese, Maltese, Tibetan, and Sinhala when compared to both mBERT and mBERT with a single adapter. The fusion of language adapters yields improvements over the baselines for Nepalese, Maltese, Uyghur, Tibetan, and Sinhala.

Sentiment Analysis (SA)									
Configuration	Objective	bg	ms	ne	jv	mt	ug	bo	si
CN+TA	MLM	0.893	0.915	0.636	0.751	0.658	0.699	0.803	0.653
	FLM	0.898	0.916	0.575	0.756	0.639	0.685	0.811	0.650
	TLM	0.893	0.918	0.625	0.749	0.661	0.638	0.797	0.653
<i>F(CN&Wiki)+TA</i>	MLM	0.882	0.906	0.627	0.750	0.662	0.784	0.804	0.689
	FLM	0.877	0.906	0.669	0.745	0.640	0.677	0.806	0.661
	TLM	0.884	0.912	0.598	0.742	0.667	0.712	0.816	0.659

Named Entity Recognition (NER)									
Configuration	Objective	bg	ms	ne	jv	mt	ug	bo	si
CN+TA	MLM	0.915	0.932	0.657	0.603	0.471	0.341	0.559	0.196
	FLM	0.918	0.930	0.626	0.578	0.476	0.398	0.572	0.242
	TLM	0.915	0.928	0.649	0.571	0.576	0.403	0.544	0.244
<i>F(CN&Wiki)+TA</i>	MLM	0.889	0.901	0.670	0.504	0.540	0.373	0.514	0.261
	FLM	0.887	0.900	0.688	0.496	0.580	0.387	0.509	0.250
	TLM	0.888	0.889	0.713	0.503	0.563	0.401	0.540	0.165

Table 4: Comparison of various objective functions used for training ConceptNet based Language Adapters-Token Masked Language Modeling (MLM), Full-Word Masked Language Modeling (FLM), and Targeted Masked Language Modeling (TLM). Maximum score per configuration in **bold**. SA results are based on MLM, and NER results are based on TLM.

5.2 Named Entity Recognition

In NER, the performance trends across different model configurations and languages exhibit similar patterns to SA but with some notable differences. mBERT demonstrates moderate to high F1 scores across languages, indicating its ability to recognize named entities to some extent. However, the addition of single task-specific adapters leads to marginal improvements in only some cases, suggesting that named entity recognition might not benefit significantly from single task-specific adapter fine-tuning compared to SA. The improvements are only observed for Maltese and Tibetan.

When incorporating language adapters, particularly those trained on CN and Wiki, we observe mixed results. Utilizing CN language adapters leads to slight improvements over the baselines only in the case of Uyghur and Sinhala. Wiki language adapters, on the other hand, give improvements over both baseline models only for Maltese, Uyghur, and Tibetan. The combination of CN and Wiki adapters shows positive impact only on Nepalese.

5.3 Effects of Data Quantity and Language Presence in LLM pre-training Data

The data quantity of external data sources might play a crucial role in the performance of language adapters and their impact on downstream tasks.

Looking at the data quantities provided in Table 2, languages like Maltese, Nepali, Uyghur, Tibetan, and Sinhala have notably fewer CN and Wiki resources compared to languages like Bulgarian and Indonesian. Despite this, language adapters trained on these limited resources still contribute to performance enhancements in SA and NER tasks for these languages compared to the baseline models. This indicates the effectiveness of leveraging even small amounts of external knowledge for adapting LLMs to low-resource languages.

Another interesting observation is the performance improvement in languages like Maltese, Uyghur, Tibetan, and Sinhala, which are not included in the mBERT pre-training data. This emphasizes that the method might be more useful for languages absent in the pre-training corpus as mBERT benefits from this adaptation using task-specific and language adapters, allowing them to effectively learn from external knowledge sources and adapt to new languages.

5.4 Take-aways

The experimental results shed light on the effectiveness of integrating graph knowledge from linguistic ontologies into multilingual LLMs via adapters for LRLs. Across both SA and NER tasks, we observe that single task-specific adapters generally lead to performance improvements, emphasizing

ing the benefits of parameter-efficient fine-tuning for specific tasks (Li and Liang, 2021; He et al., 2021; Jukić and Snajder, 2023).

In turn, the impact of language adapters trained on external knowledge sources such as CN and Wiki varies across languages and tasks. CN-based adapters generally show promise in enhancing SA but not NER. Wiki language adapters are also more beneficial for SA than NER.

The combination of both ConceptNet and Wikipedia adapters through the Adapter Fusion demonstrates competitive performance, in some cases outperforming individual adapters alone, suggesting that leveraging diverse knowledge sources can effectively enhance the capabilities of multilingual LLMs for low resource scenarios.

Our findings underscore the partial effectiveness of our method in leveraging external graph knowledge to enhance SA and NER tasks for individual LRLs. This highlights the need for further research to develop more effective strategies for adapting multilingual LLMs to low-resource contexts using various types of knowledge. Further, the results emphasize that each LRL needs an individual approach when building the dedicated NLP tools, where some languages might benefit from a certain method and the others might not need it.

6 Conclusion

In this study, we investigated the integration of structured graph knowledge into multilingual LLMs for LRLs using language adapters and task-specific adapters. We explored the use of ConceptNet and Wikipedia data for training language adapters, and we examined Adapter Fusion as a method to combine knowledge sources. Additionally, we implemented task adapters for fine-tuning LLMs for specific downstream tasks such as Sentiment Analysis (SA) and Named Entity Recognition (NER).

Our experiments revealed insights into the effectiveness of different model configurations in improving SA and NER tasks performance across LRLs. We observed a positive effect of incorporating external graph and textual knowledge through language adapters for a number of languages, including Bulgarian, Indonesian, Maltese, Nepali, Uyghur, Tibetan, and Sinhala, some of which did not possess extensive data for training both language adapters and task adapters. Fusion of knowledge sources yielded improvements in less cases,

suggesting the need for further refinement in this area.

Overall, our findings underscore the importance of parameter-efficient fine-tuning methods and the potential benefits of leveraging external knowledge for enhancing multilingual LLMs in low resource contexts. However, there are limitations to our approach, including the choice of objective functions and the need for tasks better suited to leverage external knowledge.

Limitations and Future Work

Our approach shows several limitations that should be taken into consideration in future investigations aiming to integrate graph knowledge into multilingual LLMs for enhancing LRL performance. Firstly, the choice of objective function employed for learning graph knowledge plays a critical role in effectively acquiring underlying knowledge. The objectives we explored may not be optimally suited for this purpose, highlighting the need for more tailored approaches to graph knowledge acquisition. Secondly, the tasks we selected for evaluating the effectiveness of knowledge injection may not inherently require the type of knowledge provided by graph sources. Future work should explore tasks that better leverage the acquired knowledge. Thirdly, our study was limited to a subset of LRLs, and expanding the scope to include a broader range of languages would provide a more comprehensive assessment of our approach’s effectiveness. Lastly, larger models should be explored as backbones to build upon.

Acknowledgments

We are thankful to the anonymous reviewers for their insightful comments and suggestions. This research was supported by the EU-funded LT-Bridge project, GA No. 952194; and the EU-funded project DisAI - Improving scientific excellence and creativity in combating disinformation with artificial intelligence and language technologies, GA No. 101079164.

References

Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, 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.
- Keith Cortis and Brian Davis. 2019. [A social opinion gold standard for the Malta government budget 2018](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 364–369, Hong Kong, China. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, and Ziqing Yang. 2021. Pre-training with whole word masking for chinese bert. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3504–3514.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexiei Dingli and Nicole Sant. 2016. Sentiment analysis on maltese using machine learning. In *Proceedings of The Tenth International Conference on Advances in Semantic Processing (SEMAPRO 2016)*, pages 21–25.
- Viktor Hangya, Hossain Shaikh Saadi, and Alexander Fraser. 2022. [Improving low-resource languages in pre-trained multilingual language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11993–12006, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jiawei Low, Lidong Bing, and Luo Si. 2021. [On the effectiveness of adapter-based tuning for pretrained language model adaptation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2208–2222, Online. Association for Computational Linguistics.
- Yifan Hou, Wenxiang Jiao, Meizhen Liu, Carl Allen, Zhaopeng Tu, and Mrinmaya Sachan. 2022. [Adapters for enhanced modeling of multilingual knowledge and text](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3902–3917.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for nlp](#). In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Josip Jukić and Jan Snajder. 2023. [Parameter-efficient language model tuning with active learning in low-resource settings](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5061–5074, Singapore. Association for Computational Linguistics.
- Anne Lauscher, Olga Majewska, Leonardo F. R. Ribeiro, Iryna Gurevych, Nikolai Rozanov, and Goran Glavaš. 2020. [Common sense or world knowledge? investigating adapter-based knowledge injection into pretrained transformers](#). In *Proceedings of Deep Learning Inside Out (DeeLIO): The First Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 43–49, Online. Association for Computational Linguistics.
- Siyu Li, Kui Zhao, Jin Yang, Xinyun Jiang, Zhengji Li, and Zicheng Ma. 2022. [Senti-exlm: Uyghur enhanced sentiment analysis model based on xlm](#). *Electronics Letters*, 58(13):517–519.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Antonio Martínez-García, Toni Badia, and Jeremy Barnes. 2021. [Evaluating morphological typology in zero-shot cross-lingual transfer](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3136–3153.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. [When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- Hiroki Nakayama. 2018. [segeval: A python framework for sequence labeling evaluation](#). Software available from <https://github.com/chakki-works/segeval>.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe. 2022. [Lifting the curse of multilinguality by pre-training modular transformers](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3479–3495, Seattle, United States. Association for Computational Linguistics.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Ayu Purwarianti and Ida Ayu Putu Ari Crisdayanti. 2019. Improving bi-lstm performance for indonesian sentiment analysis using paragraph vector. In *2019 International Conference of Advanced Informatics: Concepts, Theory and Applications (ICAICTA)*, pages 1–5. IEEE.
- Surangika Ranathunga and Isuru Udara Liyanage. 2021. Sentiment analysis of sinhala news comments. *Transactions on Asian and Low-Resource Language Information Processing*, 20(4):1–23.
- Oyesh Mann Singh, Sandesh Timilsina, Bal Krishna Bal, and Anupam Joshi. 2020. [Aspect based abusive sentiment detection in nepali social media texts](#). In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 301–308.
- Push Singh, Thomas Lin, Erik T Mueller, Grace Lim, Travell Perkins, and Wan Li Zhu. 2002. Open mind common sense: Knowledge acquisition from the general public. In *On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE: Confederated International Conferences CoopIS, DOA, and ODBASE 2002 Proceedings*, pages 1223–1237. Springer.
- Marina Sokolova, Nathalie Japkowicz, and Stan Szpakowicz. 2006. Beyond accuracy, f-score and roc: a family of discriminant measures for performance evaluation. In *Australasian joint conference on artificial intelligence*, pages 1015–1021. Springer.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- Giorgos Vernikos and Andrei Popescu-Belis. 2021. [Subword mapping and anchoring across languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2633–2647, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. 2021. [K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1405–1418, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Wilson Wongso, David Samuel Setiawan, and Derwin Suhartono. 2021. Causal and masked language modeling of javanese language using transformer-based architectures. In *2021 International Conference on Advanced Computer Science and Information Systems (ICACSIS)*, pages 1–7. IEEE.

Shijie Wu and Mark Dredze. 2020. *Are all languages created equal in multilingual BERT?* In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 120–130, Online. 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.

Yulei Zhu, Baima Luosai, Liyuan Zhou, Nuo Qun, and Tashi Nyima. 2023. *Research on sentiment analysis of tibetan short text based on dual-channel hybrid neural network*. In *2023 IEEE 4th International Conference on Pattern Recognition and Machine Learning (PRML)*, pages 377–384.

Appendix

A SA and NER Data Details

Table 5 and 6 provide a more detailed description of the datasets used for training task adapters.

Language	ISO code	Source	#pos	#neg	#train	#val	#test
Bulgarian	bg	Martínez-García et al., 2021	6652	1271	5412	838	1673
Indonesian	ms	Purwarianti and Crisdayanti, 2019	7319	4005	7926	1132	2266
Maltese	mt	Cortis and Davis, 2019; Dingli and Sant, 2016	271	580	595	85	171
Nepali	ne	Singh et al., 2020	680	1019	1189	255	255
Javanese	jv	Wongso et al., 2021	12500	12500	17500	5025	2475
Uyghur	ug	Li et al., 2022	2450	353	1962	311	530
Tibetan	bo	Zhu et al., 2023	5006	5000	7004	1501	1501
Sinhala	si	Ranathunga and Liyanage, 2021	2487	2516	3502	750	751

Table 5: Sentiment Analysis Data Details

Language	ISO code	#train	#val	#test
Bulgarian	bg	20000	10000	10000
Indonesian	ms	20000	1000	1000
Maltese	mt	100	100	100
Nepali	ne	100	100	100
Javanese	jv	100	100	100
Uyghur	ug	100	100	100
Tibetan	bo	100	100	100
Sinhala	si	100	100	100

Table 6: Named Entity Recognition Data Details