

mLUKE: The Power of Entity Representations in Multilingual Pretrained Language Models

Ryokan Ri^{1,2*}

ryo0123@ousia.jp

Ikuya Yamada^{1,3}

ikuya@ousia.jp

Yoshimasa Tsuruoka²

tsuruoka@logos.t.u-tokyo.ac.jp

¹Studio Ousia, Tokyo, Japan

²The University of Tokyo, Tokyo, Japan

³RIKEN AIP, Tokyo, Japan

Abstract

Recent studies have shown that multilingual pretrained language models can be effectively improved with cross-lingual alignment information from Wikipedia entities. However, existing methods only exploit entity information in pretraining and do not explicitly use entities in downstream tasks. In this study, we explore the effectiveness of leveraging entity representations for downstream cross-lingual tasks. We train a multilingual language model with 24 languages with entity representations and show the model consistently outperforms word-based pretrained models in various cross-lingual transfer tasks. We also analyze the model and the key insight is that incorporating entity representations into the input allows us to extract more language-agnostic features. We also evaluate the model with a multilingual *cloze prompt* task with the mLAMA dataset. We show that entity-based prompt elicits correct factual knowledge more likely than using only word representations. Our source code and pretrained models are available at <https://github.com/studio-ousia/luke>.

1 Introduction

Pretrained language models have become crucial for achieving state-of-the-art performance in modern natural language processing. In particular, multilingual language models (Conneau and Lample, 2019; Conneau et al., 2020a; Doddapaneni et al., 2021) have attracted considerable attention particularly due to their utility in cross-lingual transfer.

In zero-shot cross-lingual transfer, a pretrained encoder is fine-tuned in a single resource-rich language (typically English), and then evaluated on other languages never seen during fine-tuning. A key to solving cross-lingual transfer tasks is to obtain representations that generalize well across languages. Several studies aim to improve multilingual models with cross-lingual supervision such as

bilingual word dictionaries (Conneau et al., 2020b) or parallel sentences (Conneau and Lample, 2019).

Another source of such information is the cross-lingual mappings of Wikipedia entities (articles). Wikipedia entities are aligned across languages via inter-language links and the text contains numerous entity annotations (hyperlinks). With these data, models can learn cross-lingual correspondence such as the words *Tokyo* (English) and *東京* (Japanese) refers to the same entity. Wikipedia entity annotations have been shown to provide rich cross-lingual alignment information to improve multilingual language models (Calixto et al., 2021; Jiang et al., 2022). However, previous studies only incorporate entity information through an auxiliary loss function during pretraining, and the models do not explicitly have entity representations used for downstream tasks.

In this study, we investigate the effectiveness of entity representations in multilingual language models. Entity representations are known to enhance language models in mono-lingual settings (Zhang et al., 2019; Peters et al., 2019; Wang et al., 2021; Xiong et al., 2020; Yamada et al., 2020) presumably by introducing real-world knowledge. We show that using entity representations facilitates cross-lingual transfer by providing language-independent features. To this end, we present a multilingual extension of LUKE (Yamada et al., 2020). The model is trained with the multilingual masked language modeling (MLM) task as well as the masked entity prediction (MEP) task with Wikipedia entity embeddings.

We investigate two ways of using the entity representations in cross-lingual transfer tasks: (1) perform entity linking for the input text, and append the detected entity tokens to the input sequence. The entity tokens are expected to provide language-independent features to the model. We evaluate this approach with cross-lingual question answering (QA) datasets: XQuAD (Artetxe et al., 2020)

* Work done as an intern at Studio Ousia.

and MLQA (Lewis et al., 2020); (2) use the entity [MASK] token from the MEP task as a language-independent feature extractor. In the MEP task, word tokens in a mention span are associated with an entity [MASK] token, the contextualized representation of which is used to train the model to predict its original identity. Here, we apply similar input formulations to tasks involving mention-span classification, relation extraction (RE) and named entity recognition (NER): the attribute of a mention or a pair of mentions is predicted using their contextualized entity [MASK] feature. We evaluate this approach with the RELX (Köksal and Özgür, 2020) and CoNLL NER (Tjong Kim Sang, 2002; Tjong Kim Sang and De Meulder, 2003) datasets.

The experimental results show that these entity-based approaches consistently outperform word-based baselines. Our analysis reveals that entity representations provide more language-agnostic features to solve the downstream tasks.

We also explore solving a multilingual zero-shot *cloze prompt* task (Liu et al., 2021) with the entity [MASK] token. Recent studies have shown that we can address various downstream tasks by querying a language model for blanks in prompts (Petroni et al., 2019; Cui et al., 2021). Typically, the answer tokens are predicted from the model’s word-piece vocabulary but here we incorporate the prediction from the entity vocabulary queried by the entity [MASK] token. We evaluate our approach with the mLAMA dataset (Kassner et al., 2021) in various languages and show that using the entity [MASK] token reduces language bias and elicits correct factual knowledge more likely than using only the word [MASK] token.

2 Multilingual Language Models with Entity Representations

2.1 Model: multilingual LUKE

To evaluate the effectiveness of entity representations for cross-lingual downstream tasks, we introduce a new multilingual language model based on a bidirectional transformer encoder: Multilingual LUKE (mLUKE), a multilingual extension of LUKE (Yamada et al., 2020). The model is trained with the masked language modeling (MLM) task (Vaswani et al., 2017) as well as the masked entity prediction (MEP) task. In MEP, some of the input entity tokens are randomly masked with the special entity [MASK] token, and the model is trained to predict the original entities. Note that the entity

[MASK] token is different from the word [MASK] token for MLM.

The model takes as input a tokenized text (w_1, w_2, \dots, w_m) and the entities appearing in the text (e_1, e_2, \dots, e_n) , and compute the contextualized representation for each token $(\mathbf{h}_{w_1}, \mathbf{h}_{w_2}, \dots, \mathbf{h}_{w_m}$ and $\mathbf{h}_{e_1}, \mathbf{h}_{e_2}, \dots, \mathbf{h}_{e_n}$). The word and entity tokens equally undergo self-attention computation (*i.e.*, no entity-aware self-attention in Yamada et al. (2020)) after embedding layers.

The word and entity embeddings are computed as the summation of the following three embeddings: token embeddings, type embeddings, and position embeddings (Devlin et al., 2019). The entity tokens are associated with the word tokens through position embeddings: the position of an entity token is defined as the positions of its corresponding word tokens, and the entity position embeddings are summed over the positions.

Model Configuration. The model configurations of mLUKE follow the *base* and *large* configurations of XLM-RoBERTa (Conneau et al., 2020a), a variant of BERT (Devlin et al., 2019) trained with CommonCrawl data from 100 languages. Before pretraining, the parameters in common (*e.g.*, the weights of the transformer encoder and the word embeddings) are initialized using the checkpoint from the Transformers library.¹

The size of the entity embeddings is set to 256 and they are projected to the size of the word embeddings before being fed into the encoder.

2.2 Training Corpus: Wikipedia

We use Wikipedia dumps in 24 languages (Appendix A) as the training data. These languages are selected to cover reasonable numbers of languages that appear in downstream cross-lingual datasets. We generate input sequences by splitting the content of each page into sequences of sentences comprising ≤ 512 words with their entity annotations (*i.e.*, hyperlinks). During training, data are sampled from each language with n_i items with the following multinomial distribution:

$$p_i = \frac{n_i^\alpha}{\sum_{k=1}^N n_k^\alpha}, \quad (1)$$

where α is a smoothing parameter and set to 0.7 following multilingual BERT.²

¹<https://huggingface.co/transformers/>

²<https://github.com/google-research/bert/blob/master/multilingual.md>

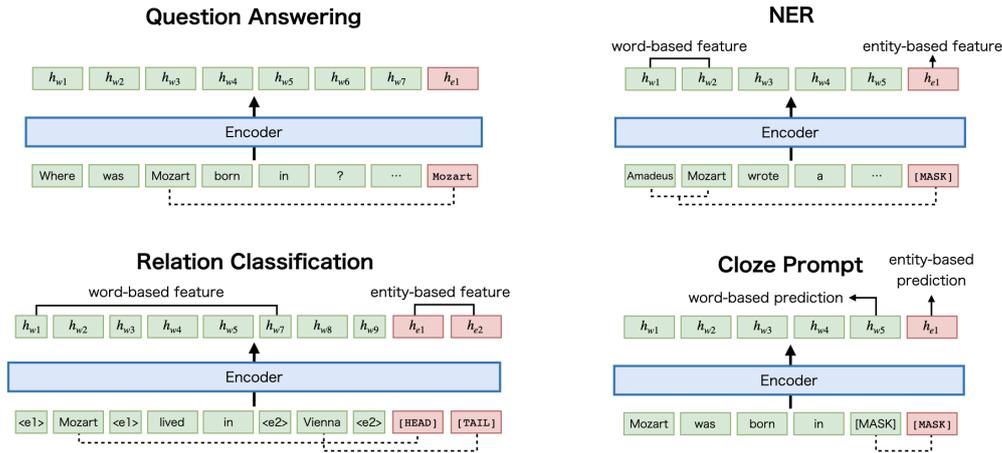


Figure 1: How to use entity representations in downstream tasks. The input entity embeddings are associated with their mentions (indicated by dotted lines) via positional embeddings.

Entity Vocabulary. Entities used in mLUKE are defined as Wikipedia articles. The articles from different languages are aligned through inter-language links³ and the aligned articles are treated as a single entity. We include in the vocabulary the most frequent 1.2M entities in terms of the number of hyperlinks that appear across at least three languages to facilitate cross-lingual learning.

Optimization. We optimize the models with a batch size of 2048 for 1M steps in total using AdamW (Loshchilov and Hutter, 2019) with warmup and linear decay of the learning rate. To stabilize training, we perform pretraining in two stages: (1) in the first 500K steps, we update only those parameters that are randomly initialized (e.g., entity embeddings); (2) we update all parameters in the remaining 500K steps. The learning rate scheduler is reset at each training stage. For further details on hyperparameters, see Appendix A.

2.3 Baseline Models

We compare the primary model that we investigate, multilingual LUKE used with entity representations (**mLUKE-E**), against several baselines pretrained models and an ablation model based on word representations:

mBERT (Devlin et al., 2019) is one of the earliest multilingual language models. We provide these results as a reference.

XLM-R (Conneau et al., 2020a) is the model that mLUKE is built on. This result indicates how our additional pretraining step and entity representa-

tion impact the performance. Since earlier studies (Liu et al., 2019; Lan et al., 2020) indicated longer pretraining would simply improve performance, we train another model based on XLM-R_{base} with extra MLM pretraining following the same configuration of mLUKE.

mLUKE-W is an ablation model of mLUKE-E. This model discards the entity embeddings learned during pretraining and only takes word tokens as input as with the other baseline models. The results from this model indicate the effect of MEP only as an auxiliary task in pretraining, and the comparison with this model will highlight the effect of using entity representations for downstream tasks in mLUKE-E.

The above models are fine-tuned with the same hyperparameter search space and computational budget as described in Appendix B.

We also present the results of **XLM-K** (Jiang et al., 2022) for ease of reference. XLM-K is based on XLM-R_{base} and trained with entity information from Wikipedia but does not use entity representations in downstream tasks. Notice that their results are not strictly comparable to ours, because the pretraining and fine-tuning settings are different.

3 Adding Entities as Language-Agnostic Features in QA

We evaluate the approach of adding entity embeddings to the input of mLUKE-E with cross-lingual extractive QA tasks. The task is, given a question and a context passage, to extract the answer span from the context. The entity embeddings provide language-agnostic features and thus should facilitate cross-lingual transfer learning.

³https://en.wikipedia.org/wiki/Help:Interlanguage_links. We build an inter-language database from the wikidatawiki dump from November 30, 2020.

XQuAD	en	es	de	el	ru	tr	ar	vi	th	zh	hi	avg.
mBERT	84.5	76.1	73.1	59.0	70.2	53.2	62.1	68.5	40.7	58.3	57.0	63.9
XLM-R _{base}	84.0	76.5	76.4	73.9	74.4	67.8	68.1	74.2	66.8	61.5	68.7	72.0
+ extra training	86.1	76.9	76.5	73.7	74.7	66.3	68.2	74.5	67.7	64.7	66.6	72.4
mLUKE-W _{base}	85.7	78.0	77.4	74.7	75.7	68.3	71.7	75.9	67.1	65.1	69.9	73.6
mLUKE-E _{base}	86.3	78.9	78.9	73.9	76.0	68.8	71.4	76.4	67.5	65.9	72.2	74.2
XLM-R _{large}	88.5	82.4	82.0	81.4	81.2	75.5	75.9	80.7	72.3	67.6	77.2	78.6
mLUKE-W _{large}	89.0	83.1	82.4	81.3	81.3	75.3	77.9	81.2	75.1	71.5	77.3	79.6
mLUKE-E _{large}	88.6	83.0	81.7	81.4	80.8	75.8	77.7	81.9	75.4	71.9	77.5	79.6

MLQA	en	es	de	ar	hi	vi	zh	avg.	G-XLT avg.
mBERT	79.1	65.9	58.6	48.6	44.8	58.5	58.1	59.1	40.9
XLM-R _{base}	79.7	67.7	62.2	55.8	59.9	65.3	62.5	64.7	33.4
+ extra training	81.3	69.8	65.0	54.8	59.3	65.6	64.2	65.7	50.2
mLUKE-W _{base}	81.3	69.7	65.4	60.4	63.2	68.3	66.1	67.8	54.0
mLUKE-E _{base}	80.8	70.0	65.5	60.8	63.7	68.4	66.2	67.9	55.6
XLM-K (Jiang et al., 2022)	80.8	69.2	63.8	60.0	65.3	70.1	63.8	67.7	-
XLM-R _{large}	83.9	74.7	69.9	64.9	69.9	73.3	70.3	72.4	65.3
mLUKE-W _{large}	84.0	74.3	70.3	66.2	70.2	74.2	69.7	72.7	67.4
mLUKE-E _{large}	84.1	74.5	70.5	66.2	71.4	74.3	70.5	73.1	67.7

Table 1: F1 scores on the XQuAD and MLQA dataset in the cross-lingual transfer settings. The scores without reference are from the best model tuned with the English development data.

3.1 Main Experiments

Datasets. We fine-tune the pretrained models with the SQuAD 1.1 dataset (Rajpurkar et al., 2016), and evaluate them with the two multilingual datasets: XQuAD (Artetxe et al., 2020) and MLQA (Lewis et al., 2020). XQuAD is created by translating a subset of the SQuAD development set while the source of MLQA is natural text in Wikipedia. Besides multiple monolingual evaluation data splits, MLQA also offers data to evaluate generalized cross-lingual transfer (G-XLT), where the question and context texts are in different languages.

Models. All QA models used in this experiment follow Devlin et al. (2019). The model takes the question and context word tokens as input and predicts a score for each span of the context word tokens. The span with the highest score is predicted as the answer to the question.

mLUKE-E takes entity tokens as additional features in the input (Figure 1) to enrich word representations. The entities are automatically detected using a heuristic string matching based on the original Wikipedia article from which the dataset instance is created. See Appendix C for more details.

Results. Table 1 summarizes the model’s F1 scores for each language. First, we discuss the *base* models. On the effectiveness of entity representations, mLUKE-E_{base} performs better than its word-based counterpart mLUKE-W_{base} (0.6 average points improvement in the XQuAD average score, 0.1 points in MLQA) and XLM-K (0.2 points improvement

in MLQA), which indicates the input entity tokens provide useful features to facilitate cross-lingual transfer. The usefulness of entities is demonstrated especially in the MLQA’s G-XLT setting (full results available in Appendix F); mLUKE-E_{base} exhibits a substantial 1.6 point improvement in the G-XLT average score over mLUKE-W_{base}. This suggests that entity representations are beneficial in a challenging situation where the model needs to capture language-agnostic semantics from text segments in different languages.

We also observe that XLM-R_{base} benefits from extra training (0.4 points improvement in the average score on XQuAD and 2.1 points in MLQA). The mLUKE-W_{base} model further improves the average score from XLM-R_{base} with extra training (1.2 points improvement in XQuAD and 2.1 points in MLQA), showing the effectiveness of the MEP task for cross-lingual QA.

By comparing *large* models, we still observe substantial improvements from XLM-R_{large} to the mLUKE models. Also we can see that mLUKE-E_{large} overall provides better results than mLUKE-W_{large} (0.4 and 0.3 points improvements in the MLQA average and G-XLT scores; comparable scores in XQuAD), confirming the effectiveness of entity representations.

3.2 Analysis

How do the entity representations help the model in cross-lingual transfer? In the mLUKE-E model,

the input entity tokens annotate mention spans on which the model performs prediction. We hypothesize that this allows the encoder to inject language-agnostic entity knowledge into span representations, which help better align representations across languages. To support this hypothesis, we compare the degree of alignment between span representations before and after adding entity embeddings in the input, *i.e.*, mLUKE-W and mLUKE-E.

Task. We quantify the degree of alignment as performance on the contextualized word retrieval (CWR) task (Cao et al., 2020). The task is, given a word within a sentence in the query language, to find the word with the same meaning in the context from a candidate pool in the target language.

Dataset. We use the MLQA dev set (Lewis et al., 2020). As MLQA is constructed from parallel sentences mined from Wikipedia, some sentences and answer spans are aligned and thus the dataset can be easily adapted for the CWR task. As the query and target word, we use the answer span⁴ annotated in the dataset, which is also parallel across the languages. We use the English dataset as the query language and other languages as the target. We discard query instances that do not have their parallel data in the target language. The candidate pool is all answer spans in the target language data.

Models. We evaluate the mLUKE-W_{base} and mLUKE-E_{base} models without fine-tuning. The retrieval is performed by ranking the cosine similarity of contextualized span representations, which is computed by mean-pooling the output word vectors in the span.

Results. Table 2 shows the retrieval performance in terms of the mean reciprocal rank score. We observe that the scores of mLUKE-E_{base} are higher than mLUKE-W_{base} across all the languages. This demonstrates that adding entities improves the degree of alignment of span representations, which may explain the improvement of mLUKE-E in the cross-lingual QA task.

	ar	de	es	hi	vi	zh	avg.
mLUKE-W _{base}	55.6	66.1	68.4	60.4	69.7	56.1	62.7
mLUKE-E _{base}	56.9	68.1	70.4	61.5	71.2	60.0	64.7

Table 2: The mean reciprocal rank score of the CWR task with the MLQA dev set.

⁴Answer spans are not necessarily a word, but here we generalize the task as span retrieval for our purpose.

4 The Entity MASK Token as Feature Extractor in RE and NER

In this section, we evaluate the approach of using the entity [MASK] token to extract features from mLUKE-E for two entity-related tasks: relation extraction and named entity recognition.

We formulate both tasks as the classification of mention spans. The baseline models extract the feature of spans as the contextualized representations of word tokens, while mLUKE-E extracts the feature as the contextualized representations of the special language-independent entity tokens associated with the mentions (Figure 1). We demonstrate that this approach consistently improves the performance in cross-lingual transfer.

4.1 Relation Extraction

Relation Extraction (RE) is a task to determine the correct relation between the two (head and tail) entities in a sentence. Adding entity type features have been shown to be effective to cross-lingual transfer in RE (Subburathinam et al., 2019; Ahmad et al., 2021), but here we investigate an approach that does not require predefined entity types but utilize special entity embeddings learned in pretraining.

Datasets. We fine-tune the models with the English KBP-37 dataset (Zhang and Wang, 2015) and evaluate the models with the RELX dataset (Köksal and Özgür, 2020), which is created by translating a subset of 502 sentences from KBP-37’s test set into four different languages. Following Köksal and Özgür (2020), we report the macro average of F1 scores of the 18 relations.

Models. In the input text, the head and tail entities are surrounded with special markers (<ent>, <ent2>). The baseline models extract the feature vectors for the entities as the contextualized vector of the first marker followed by their mentions. The two entity features are concatenated and fed into a linear classifier to predict their relation.

For mLUKE-E, we introduce two special entities, [HEAD] and [TAIL], to represent the head and tail entities (Yamada et al., 2020). Their embeddings are initialized with the entity [MASK] embedding. They are added to the input sequence being associated with the entity mentions in the input, and their contextualized representations are extracted as the feature vectors. As with the word-based models, the features are concatenated and input to a linear classifier.

	RE						NER				
	en	de	es	fr	tr	avg.	en	de	nl	es	avg.
mBERT	65.0	57.3	61.6	58.9	56.2	59.8	89.7	70.0	75.2	77.1	78.0
XLM-R _{base}	66.5	60.8	62.9	60.9	57.7	61.7	91.5	74.3	80.7	79.8	81.6
+ extra training	67.0	61.3	62.9	64.3	61.9	63.5	91.8	75.7	80.3	79.8	81.9
mLUKE-W _{base}	68.7	64.3	65.8	62.1	65.0	65.2	91.6	75.1	80.2	79.2	81.5
mLUKE-E _{base}	69.3	64.5	65.2	64.7	68.7	66.5	93.6	77.2	81.8	77.7	82.6
XLM-K (Jiang et al., 2022)	-	-	-	-	-	-	90.7	73.3	80.0	76.6	80.1
XLM-R _{large}	68.0	65.3	65.0	63.3	64.1	65.1	92.5	75.1	82.9	80.5	82.8
mLUKE-W _{large}	66.2	65.3	68.1	66.5	64.7	66.2	92.3	76.5	82.6	80.7	83.0
mLUKE-E _{large}	68.1	65.8	67.8	66.4	64.4	66.5	94.0	78.3	83.5	81.4	84.3

Table 3: F1 scores on relation extraction (RE) and named entity recognition (NER).

4.2 Named Entity Recognition

Named Entity Recognition (NER) is the task to detect entities in a sentence and classify their type. We use the CoNLL-2003 English dataset (Tjong Kim Sang and De Meulder, 2003) as the training data, and evaluate the models with the CoNLL-2003 German dataset and the CoNLL-2002 Spanish and Dutch dataset (Tjong Kim Sang, 2002).

Models. We adopt the model of Sohrab and Miwa (2018) as the baseline model, which enumerates all possible spans in a sentence and classifies them into the target entity types or *non-entity* type. In this experiment, we enumerate spans with at most 16 tokens. For the baseline models, the span features are computed as the concatenation of the word representations of the first and last tokens. The span features are fed into a linear classifier to predict their entity type.

The input of mLUKE-E contains the entity [MASK] tokens associated with all possible spans. The span features are computed as the contextualized representations of the entity [MASK] tokens. The features are input to a linear classifier as with the word-based models.

4.3 Main Results

The results are shown in Table 3. The mLUKE-E models outperform their word-based counterparts mLUKE-W in the average score in all the comparable settings (the *base* and *large* settings; the RE and NER tasks), which shows entity-based features are useful in cross-lingual tasks. We also observe that XLM-R_{base} benefits from extra training (1.8 average points improvement in RE and 0.3 points in NER), but mLUKE-E still outperforms the results.

4.4 Analysis

The performance gain of mLUKE-E over mLUKE-W can be partly explained as the entity [MASK]

	de	es	fr	tr
mLUKE-W _{base}	0.71	0.74	0.74	0.84
mLUKE-E _{base}	0.25	0.28	0.24	0.36

Table 4: The modularity of word and entity features computed with the same mLUKE model. The data are from pairs of English and the other languages in the RELX dataset.

token extracts better features for predicting entity attributes because it resembles how mLUKE is pre-trained with the MEP task. We hypothesize that there exists another factor for the improvement in cross-lingual performance: language neutrality of representations.

The entity [MASK] token is shared across languages and their contextualized representations may be less affected by the difference of input languages, resulting in features that generalize well for cross-lingual transfer. To find out if the entity-based features are actually more language-independent than word-based features, we evaluate the *modularity* (Fujinuma et al., 2019) of the features extracted for the RELX dataset.

Modularity is computed for the k -nearest neighbor graph of embeddings and measures the degree to which embeddings tend to form clusters within the same language. We refer readers to Fujinuma et al. (2019) for how to compute the metric. Note that the maximum value of modularity is 1, and 0 means the embeddings are completely randomly distributed regardless of language.

We compare the modularity of the word features from mLUKE-W_{base} and entity features from mLUKE-E_{base} before fine-tuning. Note that the features here are concatenated vectors of head and tail features. Table 4 shows that the modularity of mLUKE-E_{base} is much lower than mLUKE-W_{base},

	ar	en	fi	fr	id	ja	ru	vi	zh	avg.
mBERT	17.1	36.8	24.0	24.3	42.9	14.3	19.5	39.4	26.2	27.2
XLM-R _{base}	14.2	27.2	16.2	14.9	28.2	11.9	11.7	25.1	17.6	18.5
+ extra training	21.2	35.0	23.0	22.2	46.8	19.6	17.5	34.4	30.7	27.8
mLUKE-W _{base}	22.3	31.3	18.4	19.6	46.7	18.4	16.7	31.9	29.3	26.1
mLUKE-E _{base} ([Y])	27.8	37.5	30.4	28.4	44.2	28.9	25.8	42.1	33.4	33.2
mLUKE-E _{base} ([X] & [Y])	42.4	47.5	44.2	35.9	56.2	40.3	35.5	55.2	46.7	44.9

Table 5: The top-1 accuracies from 9 languages from the mLAMA dataset.

demonstrating that entity-based features are more language-neutral. However, with entity-based features, the modularities are still greater than zero. In particular, the modularity computed with Turkish, which is the most distant language from English here, is significantly higher than the others, indicating that the contextualized entity-based features are still somewhat language-dependent.

5 Cloze Prompt Task with Entity Representations

In this section, we show that using the entity representations is effective in a cloze prompt task (Liu et al., 2021) with the mLAMA dataset (Kassner et al., 2021). The task is, given a cloze template such as “[X] was born in [Y]” with [X] filled with an entity (e.g., *Mozart*), to predict a correct entity in [Y] (e.g., *Austria*). We adopt the typed querying setting (Kassner et al., 2021), where a template has a set of candidate answer entities and the prediction becomes the one with the highest score assigned by the language model.

Model. As in Kassner et al. (2021), the word-based baseline models compute the candidate score as the log-probability from the MLM classifier. When a candidate entity in [Y] is tokenized into multiple tokens, the same number of the word [MASK] tokens are placed in the input sequence, and the score is computed by taking the average of the log-probabilities for its individual tokens.

On the other hand, mLUKE-E computes the log-probability of the candidate entity in [Y] with the entity [MASK] token. Each candidate entity is associated with an entity in mLUKE’s entity vocabulary via string matching. The input sequence has the entity [MASK] token associated with the word [MASK] tokens in [Y], and the candidate score is computed as the log-probability from the MEP classifier. We also try additionally appending the entity token of [X] to the input sequence if the entity is found in the vocabulary.

To accurately measure the difference between

word-based and entity-based prediction, we restrict the candidate entities to the ones found in the entity vocabulary and exclude the questions if their answers are not included in the candidates (results with full candidates and questions in the dataset are in Appendix G).

Results. We experiment in total with 16 languages which are available both in the mLAMA dataset and the mLUKE’s entity vocabulary. Here we only present the top-1 accuracy results from 9 languages on Table 5, as we can make similar observations with the other languages.

We observe that XLM-R_{base} performs notably worse than mBERT as mentioned in Kassner et al. (2021). However, with extra training with the Wikipedia corpus, XLM-R_{base} shows a significant 9.3 points improvement in the average score and outperforms mBERT (27.8 vs. 27.2). We conjecture that this shows the importance of the training corpus for this task. The original XLM-R is only trained with the CommonCrawl corpus (Conneau et al., 2020a), text scraped from a wide variety of web pages, while mBERT and XLM-R + training are trained on Wikipedia. The performance gaps indicate that Wikipedia is particularly useful for the model to learn factual knowledge.

The mLUKE-W_{base} model lags behind XLM-R_{base} + extra training by 1.7 average points but we can see 5.4 points improvement from XLM-R_{base} + extra training to mLUKE-E_{base} ([Y]), indicating entity representations are more suitable to elicit correct factual knowledge from mLUKE than word representations. Adding the entity corresponding to [X] to the input (mLUKE-E_{base} ([X] & [Y])) further pushes the performance by 11.7 points to 44.9 %, which further demonstrates the effectiveness of entity representations.

Analysis of Language Bias. Kassner et al. (2021) notes that the prediction of mBERT is biased by the input language. For example, when queried in Italian (e.g., “[X] e stato creato in [MASK].”), the model tends to predict entities that often appear in Italian text (e.g., *Italy*) for any question to answer

	en	ja	fr
mBERT	The Bahamas, 41% (355/870)	Japan, 82% (361/439)	Pays-Bas, 71% (632/895)
XLM-R _{base}	London, 78% (664/850)	Japan, 99% (437/440)	Allemagne, 96% (877/916)
+ extra training	Australia, 27% (247/899)	Japan, 99% (437/442)	Allemagne, 93% (854/917)
mLUKE-W _{base}	Germany, 22% (198/895)	Japan, 97% (428/442)	Allemagne, 99% (906/918)
mLUKE-E _{base} ([Y])	London, 37% (310/846)	Japan, 56% (241/430)	Suède, 40% (362/908)
mLUKE-E _{base} ([X] & [Y])	London, 27% (213/797)	Japan, 44% (176/401)	Suède, 30% (266/895)

Table 6: The top incorrect predictions in three languages for the template “[X] was founded in [Y].” for each model. The predictions in the original language are translated into English.

location. We expect that using entity representations would reduce language bias because entities are shared among languages and less affected by the frequency in the language of questions.

We qualitatively assess the degree of language bias in the models looking at their incorrect predictions. We show the top incorrect prediction for the template “[X] was founded in [Y].” for each model in Table 6, together with *the top-1 incorrect ratio*, that is, the ratio of the number of the most common incorrect prediction to the total false predictions, which indicates how much the false predictions are dominated by few frequent entities.

The examples show that the different models exhibit bias towards different entities as in English and French, although in Japanese the model consistently tends to predict *Japan*. Looking at the degree of language bias, mLUKE-E_{base} ([X] & [Y]) exhibits lower top-1 incorrect ratios overall (27% in fr, 44% in ja, and 30% in en), which indicates using entity representations reduces language bias. However, lower language bias does not necessarily mean better performance: in French (fr), mLUKE-E_{base} ([X] & [Y]) gives a lower top-1 incorrect ratio than mBERT (30% vs. 71%) but their numbers of total false predictions are the same (895). Language bias is only one of several factors in the performance bottleneck.

6 Related Work

6.1 Multilingual Pretrained Language Models

Multilingual pretrained language models have recently seen a surge of interest due to their effectiveness in cross-lingual transfer learning (Conneau and Lample, 2019; Liu et al., 2020). A straightforward way to train such models is multilingual masked language modeling (mMLM) (Devlin et al., 2019; Conneau et al., 2020a), i.e., training a single model with a collection of monolingual corpora in multiple languages. Although models trained

with mMLM exhibit a strong cross-lingual ability without any cross-lingual supervision (K et al., 2020; Conneau et al., 2020b), several studies aim to develop better multilingual models with explicit cross-lingual supervision such as bilingual word dictionaries (Conneau et al., 2020b) or parallel sentences (Conneau and Lample, 2019). In this study, we build a multilingual pretrained language model on the basis of XLM-RoBERTa (Conneau et al., 2020a), trained with mMLM as well as the masked entity prediction (MEP) (Yamada et al., 2020) with entity representations.

6.2 Pretrained Language Models with Entity Knowledge

Language models trained with a large corpus contain knowledge about real-world entities, which is useful for entity-related downstream tasks such as relation classification, named entity recognition, and question answering. Previous studies have shown that we can improve language models for such tasks by incorporating entity information into the model (Zhang et al., 2019; Peters et al., 2019; Wang et al., 2021; Xiong et al., 2020; Févry et al., 2020; Yamada et al., 2020).

When incorporated into multilingual language models, entity information can bring another benefit: entities may serve as anchors for the model to align representations across languages. Multilingual knowledge bases such as Wikipedia often offer mappings between different surface forms across languages for the same entity. Calixto et al. (2021) fine-tuned the top two layers of multilingual BERT by predicting language-agnostic entity ID from hyperlinks in Wikipedia articles. As our concurrent work, Jiang et al. (2022) trained a model based on XLM-RoBERTa with an entity prediction task along with an object entailment prediction task. While the previous studies focus on improving cross-lingual language representations by pretraining with entity information, our work investigates a multilingual model not only pretrained

with entities but also explicitly having entity representations and how to extract better features from such model.

7 Conclusion

We investigated the effectiveness of entity representations in multilingual language models. Our pretrained model, mLUKE, not only exhibits strong empirical results with the word inputs (mLUKE-W) but also shows even better performance with the entity representations (mLUKE-E) in cross-lingual transfer tasks. We also show that a cloze-prompt-style fact completion task can effectively be solved with the query and answer space in the entity vocabulary. Our results suggest a promising direction to pursue further on how to leverage entity representations in multilingual tasks. Also, in the current model, entities are represented as individual vectors, which may incur a large memory footprint in practice. One can investigate an efficient way of having entity representations.

References

- Wasi Ahmad, Nanyun Peng, and Kai-Wei Chang. 2021. [Gate: Graph attention transformer encoder for cross-lingual relation and event extraction](#). In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- 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*.
- Iacer Calixto, Alessandro Raganato, and Tommaso Pasini. 2021. [Wikipedia Entities as Rendezvous across Languages: Grounding Multilingual Language Models by Predicting Wikipedia Hyperlinks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. [Multilingual Alignment of Contextual Word Representations](#). In *International Conference on Learning Representations*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [Unsupervised Cross-lingual Representation Learning at Scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual Language Model Pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Emerging Cross-lingual Structure in Pretrained Language Models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Leyang Cui, Yu Wu, Jian Liu, Sen Yang, and Yue Zhang. 2021. [Template-Based Named Entity Recognition Using BART](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.
- 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*.
- Sumanth Doddapaneni, Gowtham Ramesh, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh M. Khapra. 2021. [A Primer on Pretrained Multilingual Language Models](#). *ArXiv*, abs/2107.00676.
- Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. 2020. [Entities as Experts: Sparse Memory Access with Entity Supervision](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Yoshinari Fujinuma, Jordan Boyd-Graber, and Michael J. Paul. 2019. [A Resource-Free Evaluation Metric for Cross-Lingual Word Embeddings Based on Graph Modularity](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Xiaozhe Jiang, Yaobo Liang, Weizhu Chen, and Nan Duan. 2022. [XLM-K: Improving Cross-Lingual Language Model Pre-Training with Multilingual Knowledge](#). In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-Lingual Ability of Multilingual BERT: An Empirical Study](#). In *International Conference on Learning Representations*.
- Nora Kassner, Philipp Dufter, and Hinrich Schütze. 2021. [Multilingual LAMA: Investigating Knowledge in Multilingual Pretrained Language Models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*.
- Abdullatif Köksal and Arzucan Özgür. 2020. [The RELX Dataset and Matching the Multilingual Blanks for Cross-Lingual Relation Classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A Lite BERT for Self-supervised Learning of Language Representations](#). In *International Conference on Learning Representations*.

- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. [MLQA: Evaluating Cross-lingual Extractive Question Answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. [Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing](#). *ArXiv*, abs/2107.13586.
- Y. Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *ArXiv*, abs/1907.11692.
- 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 *International Conference on Learning Representations*.
- Matthew E. Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. 2019. [Knowledge Enhanced Contextual Word Representations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language Models as Knowledge Bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Mohammad Golam Sohrab and Makoto Miwa. 2018. [Deep Exhaustive Model for Nested Named Entity Recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Ananya Subburathinam, Di Lu, Heng Ji, Jonathan May, Shih-Fu Chang, Avirup Sil, and Clare Voss. 2019. [Cross-lingual structure transfer for relation and event extraction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Erik F. Tjong Kim Sang. 2002. [Introduction to the CoNLL-2002 Shared Task: Language-Independent Named Entity Recognition](#). In *COLING-02: The 6th Conference on Natural Language Learning 2002*.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*.
- 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.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhiyuan Liu, Juan-Zi Li, and J. Tang. 2021. [KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation](#). *Transactions of the Association for Computational Linguistics*, 9:176–194.
- Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. 2020. [Pretrained Encyclopedia: Weakly Supervised Knowledge-Pretrained Language Model](#). In *International Conference on Learning Representations*.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. [LUKE: Deep Contextualized Entity Representations with Entity-aware Self-attention](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Dongxu Zhang and Dong Wang. 2015. [Relation Classification via Recurrent Neural Network](#). *ArXiv*, abs/1508.01006.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [ERNIE: Enhanced Language Representation with Informative Entities](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

Appendix for “mLUKE: The Power of Entity Representations in Multilingual Pretrained Language Models”

A Details of Pretraining

Dataset. We download the Wikipedia dumps from December 1st, 2020. We show the 24 languages included in the dataset on Table 7, along with the data size and the number of entities in the vocabulary.

Language Code	Size	# entities in vocab	Language Code	Size	# entities in vocab
ar	851M	427,460	ko	537M	378,399
bn	117M	62,595	nl	1.1G	483,277
de	3.5G	540,347	pl	1.3G	489,109
el	315M	135,277	pt	1.0G	537,028
en	6.9G	613,718	ru	2.5G	529,171
es	2.1G	587,525	sv	1.1G	390,313
fi	480M	300,333	sw	27M	30,129
fr	3.1G	630,355	te	66M	14,368
hi	90M	54,038	th	153M	100,231
id	327M	217,758	tr	326M	297,280
it	1.9G	590,147	vi	516M	263,424
ja	2.3G	369,470	zh	955M	332,970
			Total	31.4G	8,374,722

Table 7: Training Data Statistics: the size of training data, and the number of entities found in the 1.2M entity vocabulary.

Optimization. We optimize the mLUKE models for 1M steps in total using AdamW (Loshchilov and Hutter, 2019) with learning rate warmup and linear decay of the learning rate. The pretraining consists of two stages: (1) in the first 500K steps, we update only those parameters that are randomly initialized (e.g., entity embeddings); (2) we update all parameters in the remaining 500K steps. The learning rate scheduler is reset at each training stage. The detailed hyper-parameters are shown in Table 8.

Maximum word length	512	Mask probability for entities	15%
Batch size	2048	The size of word token embeddings	768
Peak learning rate	1e-4	The size of entity token embeddings	256
Peak learning rate (first 500K steps)	5e-4	Dropout	0.1
Learning rate decay	linear	Weight decay	0.01
Warmup steps	2500	Adam β_1	0.9
Mask probability for words	15%	Adam β_2	0.999
Random-word probability for words	10%	Adam ϵ	1e-6
Unmasked probability for words	10%	Gradient clipping	none

Table 8: Hyper-parameters used to pretrain mLUKE.

Computing Infrastructure. We run the pretraining on NVIDIA’s PyTorch Docker container 19.02 hosted on a server with two Intel Xeon Platinum 8168 CPUs and 16 NVIDIA Tesla V100 GPUs. The training takes approximately 2 months.

B Details of Downstream Experiments

Hyperparameter Search. For each downstream task, we perform hyperparameter searching for all the models with the same computational budget to ensure a fair comparison. For each task, we use the final evaluation metric on the validation split of the training English corpus as the validation score. The models are optimized with the AdamW optimizer (Loshchilov and Hutter, 2019) with the weight decay term set to 0.01 and a linear warmup scheduler. The learning rate is linearly increased to a specified value in the first 6 % of training steps, and then gradually decreased to zero towards the end. Table 9 summarizes the task-specific hyperparameter search spaces.

	QA (SQuAD)	Relation Classification (KBP37)	NER (CoNLL 2003)
Learning rate	2e-5	2e-5	2e-5
Batch size	{16, 32}	{4, 8, 16}	{4, 8, 16}
Epochs	2	5	5
# of random seeds	3	3	3
Validation metric	F1	F1	F1

Table 9: The hyperparameters search spaces and other details of downstream experiments.

Computing Infrastructure. We run the fine-tuning on a server with a Intel(R) Core(TM) i7-6950X CPU and 4 NVIDIA GeForce RTX 3090 GPUs.

C Detecting Entities in the QA datasets

For each question–passage pair in the QA datasets, we first create a mapping from the entity mention strings (e.g., “U.S.”) to their referent Wikipedia entities (e.g., United States) using the entity hyperlinks on the source Wikipedia page of the passage. We then perform simple string matching to extract all entity names in the question and the passage and treat all matched entity names as entity annotations for their referent entities. We ignore an entity name if the name refers to multiple entities on the page. Further, to reduce noise, we also exclude an entity name if its link probability, the probability that the name appears as a hyperlink in Wikipedia, is lower than 1%.

The XQuAD datasets are created by translating English Wikipedia articles into target languages. For each translated article, we create the mention-entity mapping from the source English article by the following procedure: for all the entities found in the source article, we find the corresponding entity in the target language through inter-language links, and then collect its possible mention strings (*i.e.*, hyperlinks to the entity) from a Wikipedia dump of the target language; the entity and the collected mention strings form the mention-entity mapping for the translated article.

D The Model Size

	# of layers	hidden size	# of heads	vocabulary size	# of parameters
mBERT	12	768	12	120K	177M
XLM-R _{base}	12	768	8	250K	278M
mLUKE-E _{base}	12	768	8	250K	585M
XLM-R _{large}	24	1024	16	250K	559M
mLUKE-E _{large}	24	1024	16	250K	867M

Table 10: The model sizes of the pretrained models.

E Ablation Study of Entity Embeddings

In Section 3 and 4, we have shown that using entity representations in mLUKE improves the cross-lingual transfer performance in QA, RE, and NER. Here we conduct an additional ablation study to investigate whether the learned entity embeddings are crucial to the success of our approach. We train an ablated model of mLUKE-E whose entity embeddings are re-initialized randomly before fine-tuning (-ablation). Table 11 and Table 12 show that the ablated model performs significantly worse than the full model (mLUKE-E), indicating that using pretrained entity embeddings is crucial rather than applying our approach during fine-tuning in an ad-hoc manner without entity-aware pretraining.

XQuAD	en	es	de	el	ru	tr	ar	vi	th	zh	hi	avg.
mLUKE-E	86.3	78.9	78.9	73.9	76.0	68.8	71.4	76.4	67.5	65.9	72.2	74.2
- ablation	84.3	76.8	76.4	71.9	74.3	67.4	70.2	75.3	67.1	64.4	68.4	72.4

MLQA	en	es	de	ar	hi	vi	zh	avg.	G-XLT avg.
mLUKE-E _{base}	80.8	70.0	65.5	60.8	63.7	68.4	66.2	67.9	55.6
- ablation	80.3	69.4	64.5	59.1	59.2	66.5	63.6	66.1	50.7

Table 11: F1 scores on the XQuAD and MLQA datasets in the cross-lingual transfer settings.

	RE						NER				
	en	de	es	fr	tr	avg.	en	de	du	es	avg.
mLUKE-E _{base}	69.3	64.5	65.2	64.7	68.7	66.5	93.6	77.2	81.8	77.7	82.6
- ablation	62.5	59.3	60.7	61.0	60.5	50.8	93.0	76.3	80.8	76.1	81.6

Table 12: F1 scores on relation extraction (RE) and named entity recognition (NER).

F Full Results of MLQA

c/q	en	es	de	ar	hi	vi	zh
en	79.1	65.4	63.4	37.9	29.7	47.1	43.2
es	67.7	65.9	58.2	38.2	24.4	43.6	39.5
de	61.7	55.9	58.6	32.3	29.7	38.4	36.8
ar	49.9	43.2	44.6	48.6	23.4	29.4	27.1
hi	47.0	37.8	39.1	26.2	44.8	28.0	23.0
vi	59.9	49.4	48.6	26.7	25.6	58.5	40.7
zh	55.3	44.2	45.3	28.3	22.7	38.7	58.1

Table 13: MLQA full results of mBERT

c/q	en	es	de	ar	hi	vi	zh
en	79.6	52.3	59.6	30.8	43.2	40.0	36.0
es	67.0	67.7	52.0	25.2	31.8	32.9	31.5
de	59.5	41.7	62.1	22.2	27.8	29.2	29.5
ar	49.6	23.2	30.9	55.8	10.6	11.6	10.3
hi	58.5	34.6	42.3	17.8	59.8	22.4	23.0
vi	61.1	28.1	39.5	17.0	27.5	65.2	26.5
zh	55.2	22.7	28.1	9.26	21.1	17.5	62.4

Table 14: MLQA full results of XLM-R_{base}

c/q	en	es	de	ar	hi	vi	zh
en	81.3	71.2	70.1	40.6	52.3	54.8	48.2
es	70.6	69.8	66.2	43.3	47.9	52.8	49.0
de	64.4	60.4	64.9	36.8	42.3	44.3	42.9
ar	59.3	52.3	52.2	54.8	30.3	37.1	31.5
hi	65.0	56.5	56.8	33.8	59.3	43.0	39.9
vi	67.0	57.1	58.2	31.7	43.8	65.5	44.0
zh	62.4	53.7	54.2	33.3	40.2	44.8	64.2

Table 15: MLQA full results of XLM-R_{base} + training

c/q	en	es	de	ar	hi	vi	zh
en	81.2	69.5	69.1	53.6	60.8	60.4	58.4
es	70.3	69.6	65.5	52.1	52.9	56.1	56.4
de	64.7	59.8	65.3	45.4	48.9	49.9	49.3
ar	60.4	52.3	54.3	60.3	34.0	43.4	41.3
hi	65.5	56.9	58.3	35.4	63.1	49.0	44.6
vi	66.8	54.4	57.1	39.7	49.3	68.3	52.4
zh	63.2	55.1	56.6	39.8	43.3	49.6	66.1

Table 16: MLQA full results of mLUKE-W_{base}

c/q	en	es	de	ar	hi	vi	zh
en	80.8	71.3	69.9	55.9	61.9	62.8	62.1
es	70.6	69.9	66.4	52.6	53.7	57.6	58.0
de	65.2	61.2	65.4	47.2	49.3	51.8	51.7
ar	61.1	54.6	56.9	60.7	39.5	47.0	44.8
hi	65.1	58.4	59.2	38.3	63.7	50.5	46.2
vi	66.7	56.5	59.5	44.3	51.1	68.4	54.2
zh	62.7	56.3	56.2	41.1	44.3	51.7	66.2

Table 17: MLQA full results of mLUKE-E_{base}

c/q	en	es	de	ar	hi	vi	zh
en	83.9	79.6	79.0	62.0	70.6	70.5	69.5
es	75.2	74.7	73.0	60.3	63.4	66.6	65.9
de	69.4	69.0	69.9	58.9	59.7	62.0	60.6
ar	67.0	63.6	66.2	64.9	54.5	58.9	57.7
hi	72.1	67.3	67.2	56.1	69.9	61.0	62.1
vi	73.5	69.6	70.7	57.1	63.0	73.3	64.5
zh	69.1	64.0	65.7	53.4	58.2	62.7	70.3

Table 18: MLQA full results of XLM-R_{large}

c/q	en	es	de	ar	hi	vi	zh
en	84.0	80.1	79.9	71.5	74.2	72.8	72.8
es	74.6	74.3	74.6	65.5	64.3	66.0	66.0
de	70.1	69.5	70.3	63.9	60.8	61.7	62.6
ar	67.9	65.0	67.9	66.2	58.6	60.2	58.7
hi	72.9	69.7	70.3	60.8	70.2	63.1	62.6
vi	73.9	69.5	72.2	65.5	64.9	74.2	67.3
zh	69.6	66.5	68.5	61.5	58.3	64.5	69.7

Table 19: MLQA full results of mLUKE-W_{large}

c/q	en	es	de	ar	hi	vi	zh
en	84.1	80.5	80.2	70.0	75.0	75.0	73.5
es	75.2	74.5	74.8	62.4	65.3	67.6	66.5
de	71.1	70.2	70.5	62.2	61.0	63.5	62.3
ar	68.4	65.6	68.4	66.2	57.7	62.3	58.0
hi	72.9	70.9	71.6	59.1	71.4	65.6	62.1
vi	74.7	71.0	73.1	61.7	64.7	74.3	66.8
zh	70.1	66.1	68.8	59.2	60.9	66.3	70.5

Table 20: MLQA full results of mLUKE-E_{large}

G Full Results of mLAMA

Table 5 shows the results from the setting where the entity candidates not in the mLUKE’s entity vocabulary are excluded. Here we provide in Table 21 the results with the full candidate set provided in the dataset for ease of comparison with other literature. When the candidate entity is not found in the mLUKE’s entity vocabulary, the log-probability from the word [MASK] tokens are used instead.

	ar	bn	de	el	en	es	fi	fr	
mBERT	15.1	12.7	28.6	19.4	34.8	30.2	19.2	27.1	
XLM-R _{base}	14.9	7.5	18.4	12.7	24.2	18.5	14.5	16.1	
+ extra training	20.7	14.0	29.3	18.2	31.6	26.4	19.2	25.0	
mLUKE-W _{base}	21.3	12.9	25.7	17.5	27.1	23.3	15.9	23.0	
mLUKE-E _{base} ([Y])	25.6	21.6	32.9	25.2	34.9	28.5	24.7	27.7	
mLUKE-E _{base} ([X] & [Y])	37.3	32.3	43.7	34.4	43.2	36.4	35.3	34.2	
	id	ja	ko	pl	pt	ru	vi	zh	avg.
mBERT	37.4	14.2	17.8	21.9	32.0	17.4	36.5	24.2	24.3
XLM-R _{base}	24.6	11.4	10.9	16.6	22.2	12.6	23.0	15.5	16.5
+ extra training	38.2	19.1	21.4	20.5	29.6	20.6	33.8	28.1	24.7
mLUKE-W _{base}	36.6	18.0	17.9	20.2	29.4	19.6	31.0	26.9	22.9
mLUKE-E _{base} ([Y])	35.3	27.2	26.3	25.7	34.7	23.8	39.1	29.5	28.9
mLUKE-E _{base} ([X] & [Y])	47.6	37.7	41.6	37.7	44.8	31.4	50.1	41.6	39.3

Table 21: The average of Top-1 accuracies from 16 languages from the mLAMA dataset.