

Multilingual LAMA: Investigating Knowledge in Multilingual Pretrained Language Models

Nora Kassner*, Philipp Dufter*, Hinrich Schütze

Center for Information and Language Processing (CIS), LMU Munich, Germany

{kassner, philipp}@cis.lmu.de

Abstract

Recently, it has been found that monolingual English language models can be used as knowledge bases. Instead of structural knowledge base queries, masked sentences such as “Paris is the capital of [MASK]” are used as probes. We translate the established benchmarks TReX and GoogleRE into 53 languages. Working with mBERT, we investigate three questions. (i) Can mBERT be used as a multilingual knowledge base? Most prior work only considers English. Extending research to multiple languages is important for diversity and accessibility. (ii) Is mBERT’s performance as knowledge base language-independent or does it vary from language to language? (iii) A multilingual model is trained on more text, e.g., mBERT is trained on 104 Wikipedias. Can mBERT leverage this for better performance? We find that using mBERT as a knowledge base yields varying performance across languages and pooling predictions across languages improves performance. Conversely, mBERT exhibits a language bias; e.g., when queried in Italian, it tends to predict Italy as the country of origin.

1 Introduction

Pretrained language models (LMs) (Peters et al., 2018; Howard and Ruder, 2018; Devlin et al., 2019) can be finetuned to a variety of natural language processing (NLP) tasks and generally yield high performance. Increasingly, these models and their generative variants are used to solve tasks by simple text generation, without any finetuning (Brown et al., 2020). This motivated research on how much knowledge is contained in LMs: Petroni et al. (2019) used models pretrained with masked language to answer fill-in-the-blank templates such as “Paris is the capital of [MASK].”

* Equal contribution - random order.

Query	Two most frequent predictions
en X was created in MASK.	[Japan (170), Italy (56), ...]
de X wurde in MASK erstellt.	[Deutschland (217), Japan (70), ...]
it X è stato creato in MASK.	[Italia (167), Giappone (92), ...]
nl X is gemaakt in MASK.	[Nederland (172), Italië (50), ...]
en X has the position of MASK.	[bishop (468), God (68), ...]
de X hat die Position MASK.	[WW (261), Ratsherr (108), ...]
it X ha la posizione di MASK.	[pastore (289), papa (138), ...]
nl X heeft de positie van MASK.	[burgemeester (400), bisschop (276) , ...]

Table 1: Language bias when querying (TyQ) mBERT. Top: For an Italian cloze question, Italy is favored as country of origin. Bottom: There is no overlap between the top-ranked predictions, demonstrating the influence of language – even though the facts are the same: the same set of triples is evaluated across languages. Table 3 shows that pooling predictions across languages addresses bias and improves performance. WW = “Wirtschaftswissenschaftler”.

This research so far has been exclusively on English. In this paper, we focus on using *multilingual* pretrained LMs as knowledge bases. Working with mBERT, we investigate three questions. (i) Can mBERT be used as a multilingual knowledge base? Most prior work only considers English. Extending research to multiple languages is important for diversity and accessibility. (ii) Is mBERT’s performance as knowledge base language-independent or does it vary from language to language? To answer these questions, we translate English datasets and analyze mBERT for 53 languages. (iii) A multilingual model is trained on more text, e.g., BERT’s training data contains the English Wikipedia, but mBERT is trained on 104 Wikipedias. Can mBERT leverage this fact? Indeed, we show that pooling across languages helps performance.

In summary our contributions are: **i)** We automatically create a multilingual version of TReX and GoogleRE covering 53 languages. **ii)** We use an alternative to fill-in-the-blank querying – ranking entities of the type required by the template (e.g., cities) – and show that it is a better tool

to investigate knowledge captured by pretrained LMs. **iii)** We show that mBERT answers queries across languages with varying performance: it works reasonably for 21 and worse for 32 languages. **iv)** We give evidence that the query language affects results: a query formulated in Italian is more likely to produce Italian entities (see Table 1). **v)** Pooling predictions across languages improves performance by large margins and even outperforms monolingual English BERT. Code and data are available online (<https://github.com/norakassner/mlama>).

2 Data

2.1 LAMA

We follow the LAMA setup introduced by Petroni et al. (2019). More specifically, we use data from TReX (Elsahar et al., 2018) and GoogleRE.¹ Both consist of triples of the form (object, relation, subject). The underlying idea of LAMA is to query knowledge from pretrained LMs using templates without any finetuning: the triple (Paris, capital-of, France) is queried with the template “Paris is the capital of [MASK].” In LAMA, TReX has 34,039 triples across 41 relations, GoogleRE 5528 triples and 3 relations. Templates for each relation have been manually created by Petroni et al. (2019). We call all triples from TReX and GoogleRE together *LAMA*.

LAMA has been found to contain many “easy-to-guess” triples; e.g., it is easy to guess that a person with an Italian sounding name is born in Italy. *LAMA-UHN* is a subset of triples that are hard to guess introduced by Poerner et al. (2020).

2.2 Translation

We translate both entities and templates. We use Google Translate to translate templates in the form “[X] is the capital of [Y]”. After translation, all templates were checked for validity (i.e., whether they contain “[X]”, “[Y]” exactly once) and corrected if necessary. In addition, German, Hindi and Japanese templates were checked by native speakers to assess translation quality (see Table 2). To translate the entity names, we used Wikidata and Google knowledge graphs.

mBERT covers 104 languages. Google Translate covers 77 of these. Wikidata and Google Knowledge Graph do not provide entity translations for all

¹code.google.com/archive/p/relation-extraction-corpus/




Figure 1: x-axis is the number of translated triples, y-axis the number of languages. There are 39,567 triples in the original LAMA (TReX and GoogleRE).

languages and not all entities are contained in the knowledge graphs. For English we can find a total of 37,498 triples which we use from now on. On average, 34% of triples could be translated (macro average over languages). We only consider languages with a coverage above 20%, resulting in the final number of languages we include in our study: 53. The macro average of translated triples in these 53 languages is 43%. Figure 1 gives statistics. We call the translated dataset *mLAMA*.

3 Experiments

3.1 Model

We work with mBERT (Devlin et al., 2019), a model pretrained on the 104 largest Wikipedias. We denote mBERT queried in language x as mBERT[x]. As comparison we use the English BERT-Base model and refer to it as BERT. In initial experiments with XLM-R (Conneau et al., 2020) we observed worse performance, similar to Jiang et al. (2020a). Thus, for simplicity we only report results on mBERT.

3.2 Typed and Untyped Querying

Petroni et al. (2019) use templates like “Paris is the capital of [MASK]” and give $\arg \max_{w \in V} p(w|t)$ as answer where V is the vocabulary of the LM and $p(w|t)$ is the (log-)probability that word w gets predicted in the template t . Thus the object of a triple must be contained in the vocabulary of the language model. This has two drawbacks: it reduces the number of triples that can be considered drastically and hinders performance comparisons across LMs with different vocabularies. We refer to this procedure as *UnTyQ*.

We propose to use typed querying, *TyQ*: for each relation a candidate set \mathcal{C} is created and the prediction becomes $\arg \max_{c \in \mathcal{C}} p(c|t)$. For templates like “[X] was born in [MASK]”, we know which entity type to expect, in this case cities. We observed that (English-only) BERT-base predicts city

names for MASK whereas mBERT predicts years for the same template. TyQ prevents this.

We choose as \mathcal{C} the set of objects across all triples for a single relation. The candidate set could also be obtained from an entity typing system (e.g., (Yaghoobzadeh and Schütze, 2016)), but this is beyond the scope of this paper. Variants of TyQ have been used before (Xiong et al., 2020).

3.3 Singletoken vs. Multitoken Objects

Assuming that objects are in the vocabulary (Petroni et al., 2019) is a restrictive assumption, even more in the multilingual case as e.g., “Hamburg” is in the mBERT vocabulary, but French “Hambourg” is tokenized to [“Ham”, “##bourg”]. We consider multitoken objects by including multiple [MASK] tokens in the templates. For both TyQ and UnTyQ we compute the score that a multitoken object is predicted by taking the average of the log probabilities for its individual tokens.

Given a template t (e.g., “[X] was born in [Y].”) let t_1 be the template with one mask token, (i.e., “[X] was born in [MASK].”) and t_k be the template with k mask tokens (i.e., “[X] was born in [MASK] [MASK] ... [MASK].”). We denote the log probability that the token $w \in V$ is predicted at i th mask token as $p(m_i = w|t_k)$, where V is the vocabulary of the LM. To compute $p(e|t)$ for an entity e that is tokenized into l tokens $\epsilon_1, \epsilon_2, \dots, \epsilon_l$ we simply average the log probabilities across tokens:

$$p(e|t) = \frac{1}{l} \sum_{i=1}^l p(m_i = \epsilon_i|t_l).$$

If k is the maximum number of tokens of any entity $e \in \mathcal{C}$ gets split into, we consider all templates t_1, \dots, t_k , with \mathcal{C} being the candidate set. The prediction is then the word with the highest average log probability across all templates t_1, \dots, t_k .

Note that for UnTyQ the space of possible predictions is $V \times V \times \dots \times V$ whereas for TyQ it is the candidate set \mathcal{C} .

3.4 Evaluation

We compute precision at one for each relation, i.e., $1/|T| \sum_{t \in T} \mathbb{1}\{\hat{t}_{object} = t_{object}\}$ where T is the set of all triples and \hat{t}_{object} is the object predicted by TyQ or UnTyQ. Note that T is different for each language. Our final measure (p1) is then the precision at one averaged over relations (i.e., macro average). Results for multiple languages are the macro average p1 across languages.




Figure 2: Distribution of p1 scores for 53 languages in UnTyQ vs. TyQ. Left: singletoken (object = 1 token). Right: multitoken (object > 1 token).

4 Results and Discussion

We first investigate TyQ and UnTyQ and find that TyQ is better suited for investigating knowledge in LMs. After exploring the translation quality, we use TyQ on mLAMA and observe rather stable performance for 21 and poor performance for 32 languages. When investigating the languages more closely, we find that prediction results highly depend on the language. Finally, we validate our initial hypothesis that mBERT can leverage its multilinguality by pooling predictions: pooling indeed performs better.

4.1 UnTyQ vs. TyQ

Figure 2 shows the distribution of p1 scores for single and multitoken objects. As expected, TyQ works better, both for single and multitoken objects. With UnTyQ, performance not only depends on the model’s knowledge, but on at least three extraneous factors: **(i)** Does the model understand the type constraints of the template (e.g., in “X is the capital of Y”, Y must be a country)? **(ii)** How “fluent” a substitution is an object under linguistic constraints (e.g., morphology) that can be viewed as orthogonal to knowledge? Many English templates cannot be translated into a single template in many languages, e.g., “in X” (with X a country) has different translations in French: “à Chypre”, “au Mexique”, “en Inde”. But the LAMA setup requires a single template. By enforcing the type, we reduce the number of errors that are due to surface fluency. **(iii)** The inadequacy of the original LAMA setup for multitoken answers. Figure 2 (right) shows that the original UnTyQ struggles with multitokens (mean p1 .03 vs. .17 for TyQ).

Overall, TyQ allows us to focus the evaluation on the core question: what knowledge is contained in LMs? From now on, we report numbers in the TyQ setting.

Manual template tuning or automatic template

	machine translated	manually corrected	manually paraphrased
de	18.1	19.4 (6)	20.9 (18)
hi	5.4	6.2 (14)	6.2 (1)
ja	0.4	0.4 (14)	0.7 (5)

Table 2: Effect of manual template modification on UnTyQ. Shown is p1, number of templates modified (in brackets). Templates are modified to correct mistakes from machine translation and paraphrased to achieve the correct object type. Manual template correction has a small effect on UnTyQ.

mining (Jiang et al., 2020b) has been investigated in the literature to approach the typing problem. We had native speakers check templates for German, Hindi and Japanese, correct mistakes in the automatic translation and paraphrase the template to obtain predictions with the correct type. Table 2 shows that corrections do not yield strong improvements. We conclude that template modifications are not an effective solution for the typing problem.

4.2 Translation Quality

Contemporaneous work by Jiang et al. (2020a) provides manual translations of LAMA templates for 23 languages respecting grammatical gender and inflection constraints. We evaluate our machine translated templates by comparing performance on a common subset of 14 languages using TyQ querying on the TReX subset. Surprisingly, we find a performance difference of 1 percentage points (0.23 vs. 0.24, p1 averaged over languages) in favor of the machine translated templates. This indicates that the machine translated templates in combination with TyQ exhibit comparable performance but come with the benefit of larger language coverage (53 vs. 23 languages).

4.3 Multilingual Performance

In mLAMA, not all triples are available in all languages. Thus absolute numbers are not comparable across languages and we adopt a relative performance comparison: we report p1 of a model-language combination divided by p1 of mBERT’s performance in English (mBERT[en]) on the exact same set of triples and call this *rel-p1*. A rel-p1 score of 0.5 for mBERT[fi] means that p1 of mBERT on Finnish is half of mBERT[en]’s performance on the same triples. rel-p1 of English BERT is usually greater than 1 as monolingual BERT tends to outperform mBERT[en].

Figure 3 shows that mBERT performs reasonably well for 21 languages, but for 32 languages

	LAMA	LAMA-UHN
BERT	38.5	29.0
mBERT[en]	35.0	25.7
mBERT[pooled]	41.1	32.1

Table 3: p1 for BERT, mBERT queried in English, mBERT pooled on LAMA and LAMA-UHN.

rel-p1 is less than 0.6 (i.e., their p1 is 60% of English’s p1). We conclude that mBERT does not exhibit a stable performance across languages. The variable performance (from 20% to almost 100% rel-p1) indicates that mBERT has no common representation for, say, “Paris” across languages, i.e., mBERT representations are language-dependent.

4.4 Bias

If mBERT captured knowledge independent of language, we should get similar answers across languages for the same relation. However, Table 1 shows that mBERT exhibits language-specific biases; e.g., when queried in Italian, it tends to predict Italy as the country of origin. This effect occurs for several relations: Table 4 in the supplementary presents data for ten relations and four languages.

4.5 Pooling

We investigate pooling of predictions across languages by picking the object predicted by the majority of languages. Table 3 shows that pooled mBERT outperforms mBERT[en] by 6 percentage points on LAMA, presumably in part because the language-specific bias is eliminated. mBERT[pooled] even outperforms BERT by 3 percentage points on LAMA-UHN. This indicates that mBERT can leverage the fact that it is trained on 104 Wikipedias vs. just one and even outperforms the much stronger model BERT.

5 Related Work

Petroni et al. (2019) first asked the question: can pretrained LMs function as knowledge bases? Subsequent analyses focused on different aspects, such as negation (Kassner and Schütze, 2020), easy to guess names (Poerner et al., 2020), integrating adapters (Wang et al., 2020) or finding alternatives to a “fill-in-the-blank” approach with single-token answers (Bouraoui et al., 2020; Heinzerling and Inui, 2020; Jiang et al., 2020b). Other work combines pretrained LM with information retrieval (Guu et al., 2020; Lewis et al., 2020a; Izacard and Grave, 2020; Kassner and Schütze, 2020; Petroni




Figure 3: p_1 of BERT (red) vs mBERT[x] (blue) divided by p_1 of mBERT[en] on the same set of triples in each language x . mBERT captures less factual knowledge than monolingual English BERT. While performance is reasonable for 21 languages, it is below 60% for 32 languages. Dashed line is $\text{rel-}p_1$ of mBERT[en] (by definition equal to 1.0). Performance of BERT varies slightly as the set of triples is different for each language. Note that the Wikipedia of Cebuano (ceb) consists mostly of machine translated articles.

et al., 2020). None of this work addresses languages other than English.

Multilingual models like mBERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) perform well for zero-shot crosslingual transfer (Hu et al., 2020). However, we are not aware of any prior work that analyzed to what degree pretrained multilingual models can be used as knowledge bases. There are many multilingual question answering datasets such as XQuAD (Artetxe et al., 2020), TiDy (Clark et al., 2020), MKQA (Longpre et al., 2020) and MLQA (Lewis et al., 2020b). Usually, multilingual models are finetuned to solve such tasks. Our goal is not to improve question answering or create an alternative multilingual question answering dataset, but instead to investigate which knowledge is contained in pretrained multilingual LMs without any kind of supervised finetuning.

There is a range of alternative multilingual knowledge bases that could be used for evaluation. Those include ConceptNet (Speer et al., 2017) or BabelNet (Navigli and Ponzetto, 2010). We decided to provide a translated versions of TReX and GoogleRE for the sake of comparability across languages. By translating manually created templates and entities we can ensure comparability across languages. This is not possible for crowd-sourced databases like ConceptNet.

In contemporaneous work, Jiang et al. (2020a) create and investigate a multilingual version of LAMA. They provide human template translations for 23 languages, propose several methods for multitoken decoding and code-switching, and experiment with a number of PLMs. In contrast to their work, we investigate typed querying, focus on comparability and pooling across languages, and explore language biases.

6 Conclusion

We presented mLAMA, a dataset to investigate knowledge in language models (LMs) in a multilingual setting covering 53 languages. While our results suggest that correct entities can be retrieved for many languages, there is a clear performance gap between English and, e.g., Japanese and Thai. This suggests that mBERT is not storing entity knowledge in a language-independent way. Experiments investigating language bias confirm this finding. We hope that this paper and the dataset we publish will stimulate research on investigating knowledge in LMs *multilingually* rather than just in English.

Acknowledgements

This work was supported by the European Research Council (# 740516) and the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A. The authors of this work take full responsibility for its content. The second author was supported by the Bavarian research institute for digital transformation (bidt) through their fellowship program. We thank Yannick Couzinié and Karan Tiwana for correcting the Japanese and Hindi templates. We thank the anonymous reviewers for valuable comments.

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.
- Zied Bouraoui, José Camacho-Collados, and Steven Schockaert. 2020. [Inducing relational knowledge from BERT](#). In *The Thirty-Fourth AAAI Conference*

- on Artificial Intelligence, AAI 2020, New York, NY, USA, February 7-12, 2020, pages 7456–7463. AAAI Press.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Jonathan H. Clark, Jennimaria Palomaki, Vitaly Nikolaev, Eunsol Choi, Dan Garrette, Michael Collins, and Tom Kwiatkowski. 2020. [TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- 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.
- 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.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. 2018. [T-REx: A large scale alignment of natural language with knowledge base triples](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: retrieval-augmented language model pre-training](#). *Computing Research Repository*, arXiv:2002.08909.
- Benjamin Heinzerling and Kentaro Inui. 2020. [Language models as knowledge bases: On entity representations, storage capacity, and paraphrased queries](#). *Computing Research Repository*, arXiv:2008.09036.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Gautier Izacard and E. Grave. 2020. [Leveraging passage retrieval with generative models for open domain question answering](#). *ArXiv*, abs/2007.01282.
- Zhengbao Jiang, Antonios Anastasopoulos, Jun Araki, Haibo Ding, and Graham Neubig. 2020a. [X-FACTR: Multilingual factual knowledge retrieval from pretrained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5943–5959, Online. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020b. [How can we know what language models know](#). *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Nora Kassner and Hinrich Schütze. 2020. [BERT-kNN: Adding a kNN search component to pretrained language models for better QA](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 3424–3430. Association for Computational Linguistics.
- Nora Kassner and Hinrich Schütze. 2020. [Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020a. [Pre-training via paraphrasing](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020b. [MLQA: Evaluating cross-lingual extractive question answering](#). In

- Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.
- Shayne Longpre, Yi Lu, and Joachim Daiber. 2020. [MKQA: A linguistically diverse benchmark for multilingual open domain question answering](#). *CoRR*, abs/2007.15207.
- Roberto Navigli and Simone Paolo Ponzetto. 2010. [Babelnet: Building a very large multilingual semantic network](#). In *ACL 2010, Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, July 11-16, 2010, Uppsala, Sweden*, pages 216–225. The Association for Computer Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. [How context affects language models’ factual predictions](#). In *Automated Knowledge Base Construction*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. 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, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2463–2473. Association for Computational Linguistics.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2020. [E-BERT: Efficient-yet-effective entity embeddings for BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 803–818, Online. Association for Computational Linguistics.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. [Conceptnet 5.5: An open multilingual graph of general knowledge](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. 2020. [K-adapter: Infusing knowledge into pre-trained models with adapters](#). *Computing Research Repository*, arXiv:2002.01808.
- Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. 2020. [Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yadollah Yaghoobzadeh and Hinrich Schütze. 2016. [Intrinsic subspace evaluation of word embedding representations](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 236–246, Berlin, Germany. Association for Computational Linguistics.

A Language Bias

Table 4 shows the language bias for 10 relations. For each relation we aggregated the predictions across all triples and show the most common two predicted entities together with its count (in brackets). The querying language clearly affects results. The effect is drastic for relations that ask for a country (e.g., P495 or P1001). P39 yields very different results without exhibiting a clear pattern. Other relations such as P463 or P178 are rather stable.

B Data Samples

Table 4 and Table 5 show randomly sampled entries from the data.

C Pretraining Data

We investigate whether performance across languages is correlated with the amount of pretraining data for each language. To this end we investigate the number of articles per language as of January 2021² and p1 for TyQ in Figure 6. We do not have access to the original pretraining data of mBERT. Thus, the number of articles we consider in the analysis might be different to the actual data used to train mBERT.

	Microsoft Word word ontwikkel deur Microsoft.
af	Aix-en-Provence is die hoofstad van Provence. Die hoofstad van Alpes-Maritimes is Nice.
ar	مات جون ديورانت بريفال في باريس يتم إنتاج كيندل فاير بواسطة أمازون. اللغة الرسمية لـ كوسوفو هي الصربية.
az	Madrid və Sofiya əkilz şəhərlərdir. İbn an-Nədim islam dinə bağlıdır. Baş qərqargah Belavia Minsk -dədir.
be	Афіцыйнай мовай Фінляндская праваслаўная царква з'яўляецца фінская мова. Сталіцай Правінцыя Цзянсу з'яўляецца Нанкін. Артур Гардэн нарадзіўся ў Манчэстэр. Нік Філіп е роден в Лондон.
bg	Наситена мазнина се состои от въглерод. Официалният език на Организация на обединените нации е руски език. মন্টনা আইডাহো এর সাথে সীমানা ভাগ করে দেয়।
bn	দৌরিকনিয়া মালি এর সাথে সীমানা ভাগ করে দেয়। খার্কুম এবং কামরো হ'ল দুটি দেশ।
ca	Rich Gannon juga en la posició quarterback. La llengua oficial de Nastola és finès. Nova York i Jerusalem són ciutats bessones. Ang Mario Aldo Montano usa ka lungsuranon sa Italya.
ceb	Ang opisyal nga sinultian sa Belhika mao ang Inolandes. Ang Praga ug Berlin mga kaluha nga lungsod. Mahárádža je právní termin v Indie.
cs	Felix Magath hraje v poloze záložník. Philipp Eduard Anton von Lenard pracuje v oblasti fyzika. Mae Maes Awyr Rhyngwladol Des Moines wedi'i leoli yn Iowa.
cy	Iaith swyddogol Aragón yw Sbaeneg. Enwir Flins-sur-Seine ar ôl Afon Seine. Dylan Taite blev født i Liverpool.
da	Johann Andreas Schmeller døde i München. Hovedstaden i Yolo County er Davis. The Bill wurde in Englisch geschrieben.
de	Buenaventura Sitjar starb in Kalifornien. Mai Jones funktioniert für British Broadcasting Corporation. Ρεύμα είναι το είδος του ηλεκτρισμού.
el	Το Μπρέονα Σαράγεβο βρίσκεται στο Σαράγεβο. Το Μεταδιθεωδές καλο αποτελείται από θεο. Vienna bread is a subclass of bread .
en	Stevie Wonder is represented by music label Motown . Kinji Fukasaku is Japan citizen . Luxemburgo comparte frontera con Francia.
es	Charles Schreiber nació en Colchester. Polonia comparte frontera con Alemania. Kohtumõistjate raamat on osa Piibel -st.
et	Israël hoiab diplomaatilisi suhteid Jordaania -ga. Sambia jagab piiri Angola -ga. Ameriketako Estatu Batuak -ek Albania -ekin harreman diplomatikoak mantentzen ditu.
eu	Tajikistango presidente legezko terminoa da Tajikistan -n. Carla Bruni Paris -n lan egiteko erabiltzen zen.
fa	اسٹائن آنتورپ بہ آنتورپ نامگذاری شده است. سوزوکی کیزاشی توسط سوزوکی تولید می شود. رئیسجمهور زیر کلاس سولستمدار است.
	Pääoman Genovan maakunta pääoma on Genova.
fi	Edward Joseph Kelly syntyi Chicago. The Home Depot perustettiin Atlanta. Mel Charles est pays de Galles citizen.
fr	Audi A5 est produit par Audi. Honda XR est produit par Honda. Fuair Walter Gay bás i Páras.
ga	Tá Ollscoil Concordia suite i Montréal. Is cúpla cathair iad Vín agus An Bhrataisláiv. Hannover e Bristol son cidades xemelgas.
gl	O idioma oficial de Mercia é Lingua latina. Afloramento de algas está composto por Alga.
he	אוביבסיטט ווינגטון נמצא ב- סיאטל. ויקיפדיה היביתית נכתב ב- צ'כית. תירות מין היא תת-סוג של תירות.
hi	कीनिया युग्राण्ड के साथ राजनयिक संबंध बनाए रखता है। कावेरी नदी एशिया में स्थित है। फिलीपीन्स का ध्वज फिलीपीन्स में एक कानूनी शब्द है।
hr	Istanbul i Bukurešt su gradovi blizanci. Gro Harlem Brundtland se nekada radila u Oslo. Izvorni jezik Čelava pjevačica je francuski jezik.
hu	Gilles Grimandi Gap -ben született. Edo-kor elnevezése Edo. Joseph-Marie, comte Portalis anyanyelve francia.
hy	Իրան գտնվում է Ասիա -ում: Բարսելոնա և Դուբայն գոյություն ունեցող երկրներ են: Ադրբեջան -ի մայրաքաղաքն է Bաքու:

Figure 4: Three randomly sampled data entries from mLAMA per language. Due to the automatic generation of the dataset not all of them are fully correct.

²https://meta.wikimedia.org/wiki/List_of_Wikipedias

	en	de	nl	it
P495: "[X] was created in [Y]"	Japan (170), Italy (56)	Deutschland (217), Japan (70)	Nederland (172), Italië (50)	Italia (167), Giappone (92)
P101: "[X] works in the field of [Y]"	art (205), science (135)	Kunst (384), Film (64)	psychologie (263), kunst (120)	fisiologia (168), caccia (135)
P106: "[X] is [Y] by profession"	politician (423), composer (80)	Politiker (323), Journalist (128)	politicus (339), acteur (247)	giornalista (420), giurista (257)
P1001: "[X] is a legal term in [Y]"	India (12), Germany (11)	Deutschland (36), Russland (9)	Nederland (22), België (12)	Italia (31), Germania (16)
P39: "[X] has the position of [Y]"	bishop (468), God (68)	WW (261), Ratsherr (108)	burgemeester (400), bisschop (276)	pastore (289), papa (138)
P527 "[X] consists of [Y]"	sodium (125), carbon (88)	Wasserstof (398), C (49)	vet (216), aluminium (130)	calcio (165), atomo (96)
P1303 "[X] plays [Y]"	guitar (431), piano (165)	Gitarre (312), Klavier (204)	piano (581), harp (42)	arpa (188), pianoforte (139)
P178 "[X] is developed by [Y]"	Microsoft (177), IBM (55)	Microsoft (153), Apple (99)	Microsoft (200), Nintendo (69)	Microsoft (217), Apple (49)
P264 "[X] is represented by music label [Y]"	EMI (267), Swan (32)	EMI (202), Paramount Records (59)	EMI (225), Swan (50)	EMI (217), Swan (99)
P463 "[X] is a member of [Y]"	FIFA (126), NATO (33)	FIFA (118), NATO (38)	FIFA (157), WWE (16)	FIFA (121), NATO (36)

Table 4: Most frequent object predictions (TyQ) in different languages. Some relations exhibit language specific biases. WW = "Wirtschaftswissenschaftler".

id	Landskrona BoIS terletak di Swedia.
id	Roy Hargrove memainkan Trompet.
it	Bahasa resmi Ossetia Selatan adalah Rusia.
it	Federazione calcistica di Vanuatu è un membro di FIFA.
it	Mikhail Gromov funziona nel campo di geometria.
it	letteratura fantasy è una sottoclasse di fantasy.
ja	ポルドーとカサブランカは双子の都市です。
ja	アーロン・マレーはクォーターバック位置で再生されます。
ja	エンツォ・フェラーリはイタリア市民です。
ka	საქართველოს F მუდგობა ნახშირბადი უნდა.
ka	ანრი რუსთა ღარდაღიდან დაარსდა -ში.
ka	კატალიონის მდებარეობს ესპანეთში -ში.
ko	바하칼리포르니아주는 캘리포니아주와 (과) 국경을 공유합니다.
ko	레오 10세의 위치는 교황입니다.
ko	캐나다는 이탈리아와의 외교 관계를 유지합니다.
la	Prophetia Michaeae est pars Biblia.
la	Carentonium Cum shares terminus Lutetia.
la	Carolus Bildt in communicate ad lingua Suecica.
it	Pensilvanija dalljasi riba su Meriandans.
it	Viskonsinas sostine yra Madisonas.
lv	Park Chan-wook naudojamas bendrauti Korėjiečių kalba.
lv	Altaja Republika oficiālā valoda ir krievu valoda.
lv	Francs Kafka dzimtā valoda ir vācu valoda.
ms	Audi Allroad Quattro ražo Audi.
ms	Tour de France dinamakan Perancis.
ms	Goa berkongsi sempadan dengan Maharashtra.
ms	Modal Amerika British ialah London.
nl	allmennaksjeselskap is een juridische term in Noorwegen.
nl	San Marinese voetbalbond is lid van FIFA.
nl	De oorspronkelijke taal van Nouvelle Star is Frans.
nl	Armand Marrast pracował w Paryż.
pl	Irlandia nosi imię Irlandia.
pl	Oficjalnym językiem Kanton Jura jest język francuski.
pl	Denver e Nairóbi są ciadades gêmeas.
pt	Arábia Saudita mantém relações diplomáticas com México.
pt	Nyepi está localizado em Bali.
ro	Districtul Darnah este localizat în Libia.
ro	Sediul central al Toyota este în Toyota.
ro	Roy Orbison folosit pentru a comunica în limba engleză.
ru	Венацый Фортунат имее позицию епископ.
ru	Renault 21 производится Renault.
ru	Личчотто, Г и играет гитара.
sk	Austrália udržiaa diplomatické vzťahy s Nórsko.
sk	Alanin pozostáva z dusík.
sk	Optický ďalekohľad je podtriada teleskop.
sk	Fergus Morton se je rodil v Glasgow.
sl	Nemčija je član NATO.
sl	Cenk Renda se je rodil v Turčija.
sl	Nic Chagalli interpreton Trance muziké.
sq	Marie Lijedahl është Suedia qytetar.
sq	Pallati i Fontainebleau është në pronësi të Francë.
sr	Нортхемптоншир дели границу са Бакинџемшир.
sr	периаписис је део орбита.
sr	Патрик Дотерс је рођен у Беркли.
sv	Det officiella språket för Savukoski är finska.
sv	The Upsetters spelar reggae musik.
sv	Arakidonsyra består av kol.
ta	Չեչենիստանը Դաշնության Երկրամասն է.
ta	தையால் கந்தகம் இக் கொண்டுள்ளது.
ta	எக்ககோட் ஆப்பிள் நிறுவனம் ஆல் உருவாக்கப்பட்டது.
th	ປະເທດສະຫຍາຍ ຄູ່ໃນ ຫນົບລາວ
th	ภาษาธรรมชาติของ นิวเดลีคือ ภาษาฮินดี
th	ကျွန်ုပ်တို့သည် ပျံ့နှံ့မှုများကို ပြောဆိုပါသည်
tr	Markus Feldmann Bern 'da çalışırdı.
tr	Graduate Institute of International and Development Studies şirketinin genel merkezi Cenevre dedir.
tr	Afganistan Demokratik Cumhuriyeti 'un başkenti Kâbil' dir.
uk	Столиця Сирія - Дамаск.
uk	Комплекс Наполеона названий на честь Наполеон I Бонапарт.
uk	Нижня Канада ділиться межею з Вермонт.
ur	مملکت سرینا کا دار الحکومت بلراد ہے۔
ur	میکسیکو قومی فٹ بال ٹیم فیفا کا ممبر ہے۔
ur	یورپی اتحاد بیلاروس کے ساتھ سفارتی تعلقات برقرار رکھے ہوئے ہے۔
vi	Ngôn ngữ chính thức của Lampung là tiếng Indonesia.
vi	Ngôn ngữ chính thức của Vittasaari là tiếng Phần Lan.
vi	Ả Rập Saudi duy trì quan hệ ngoại giao với Yemen.
zh	蒙马特圣伯多禄堂以西門彼得命名。
zh	Sun Media集團的总部位于多伦多中。
zh	多米尼克·杜卡曾经在布拉格中工作。

Figure 5: Data samples continued.




Figure 6: Scatter plot of p1 TyQ and number of articles in the corresponding Wikipedia. There is no clear trend visible.