

A Simple and Effective Method To Eliminate the Self Language Bias in Multilingual Representations

Ziyi Yang^{1*}, Yinfei Yang², Daniel Cer², Eric Darve¹

¹Stanford University

{ziyi.yang, darve}@stanford.edu

²Google Research

{yinfeiy, cer}@google.com

Abstract

Language agnostic and semantic-language information isolation is an emerging research direction for multilingual representations models. We explore this problem from a novel angle of geometric algebra and semantic space. A simple but highly effective method “Language Information Removal (LIR)” factors out language identity information from semantic related components in multilingual representations pre-trained on multi-monolingual data. A post-training and model-agnostic method, LIR only uses simple linear operations, e.g. matrix factorization and orthogonal projection. LIR reveals that for weak-alignment multilingual systems, the principal components of semantic spaces primarily encodes language identity information. We first evaluate the LIR on a cross-lingual question answer retrieval task (LArEQA), which requires the strong alignment for the multilingual embedding space. Experiment shows that LIR is highly effectively on this task, yielding almost 100% relative improvement in MAP for weak-alignment models. We then evaluate the LIR on Amazon Reviews and XEVAL dataset, with the observation that removing language information is able to improve the cross-lingual transfer performance.

1 Introduction

Recently, large-scale language modeling has expanded from English to the multilingual setting (i.a., Devlin et al. (2019); Conneau and Lample (2019); Conneau et al. (2020)). Although these models are trained with language modeling objectives on monolingual data, i.e. without cross-lingual information, these multilingual systems exhibit impressive zero-shot cross-lingual ability (Hu et al., 2020b). These observations raise many questions and provide insight for multilingual representations learning. First, how is the language

identity information and the semantic information expressed in the representation? Understanding their relations and underlying geometric structure is crucial for insights into designing more effective multilingual embedding systems. Second, how can we factor out the language identity information from the semantic components in representations? In many application, e.g. cross-lingual semantic retrieval, we wish to only keep the semantic information. Third, what is the geometric relation between different languages? Efforts have been made to answer these questions, e.g. Artetxe et al. (2020); Chung et al. (2020); Lauscher et al. (2020). Such prior work has addressed the problem at training time. In this work, we systematically explore a post-training method that can be readily applied to existing multilingual models.

One of the first attempts in this research area, Roy et al. (2020), proposed two concepts for language agnostic models: weak alignment v.s. strong alignment. For multilingual systems with weak alignment, for any item in language L_1 , the nearest neighbor in language L_2 is the most semantically “relevant” item. In the case of strong alignment, for any representation, all semantically relevant items are closer than all irrelevant items, regardless of their language. Roy et al. (2020) show sentence representations from the same language tend to cluster in weak-alignment system. Similar phenomena can be observed on other pre-trained multilingual models like mBERT, XLM-R (Conneau et al., 2020) and CMLM (Yang et al., 2020). Roy et al. (2020) provides carefully-designed training strategies for retrieval-like model to mitigate this issue in order to obtain language agnostic multilingual systems.

We systematically explore a simple post-training method we refer to as Language Information Removal (LIR), to effectively facilitate the language agnosticism in multilingual embedding systems. First introduced in Yang et al. (2020) to reduce same language bias for retrieval tasks, the method

* Work done during internship at Google Research.




Figure 1: Language Information Retrieval (LIR) removes language identification information using principle components of the original representation space. This mechanism is validated by LIR’s effects demonstrated in Fig. 2.

uses only linear algebra factorization and post-training operation. LIR can be conveniently applied to any multilingual model. We show LIR yields surprisingly large improvements in several downstream tasks, including LARQA, a cross-lingual QA retrieval dataset (Roy et al., 2020); Amazon Reviews, a zero-shot cross lingual evaluation dataset; XEVAL, a collection of multilingual sentence embedding tasks. Our results suggest that the principal components of a multilingual system with self-language bias primarily encodes language identification information. Implementation for LIR is available at <https://github.com/ziyi-yang/LIR>.

2 Language Information Removal for Self Language Bias Elimination

In this section we describe Language Information Removal (LIR) to address the self language bias in multilingual embeddings (Yang et al., 2020). The first step is to extract the language identity information for each language space. Given a multilingual embedding system E , e.g. multilingual BERT, and a collection of multilingual texts $\{t_L^i\}$, where t_L^i denotes the i th phrase in the collection for the language L . We construct a language matrix $M_L \in \mathbb{R}^{n \times d}$ for language L , where n denotes the number of sentences in language L and d denotes the dimension of the representation. The row i of M_L is the representation of t_L^i computed by E .

Second, we extract language identification components for each language. One observation in multilingual systems is that representations from the same language tend to cluster together (w.r.t representations in other languages), even though these representations have different semantic meanings. This phenomenon is also known as “weak alignment” (Roy et al., 2020). The mathematical explanation for this clustering phenomenon is that representations in the same language have shared vector space components. We propose that these shared components essentially represent the language identification information. Removing these language

components should leave semantic-related information in the representations.

To remove the shared components, or the language identification from the representations, we leverage singular value decomposition (SVD) which identifies the principal directions of a space. We use SVD instead of PCA since SVD is more stable numerically (e.g. for Läuchli matrix). Specifically, the SVD of a language matrix is $M_L = U_L \Sigma_L V_L^T$, where the columns of $V_L \in \mathbb{R}^{d \times d}$ are the right singular vectors of M_L . We take first r columns of V_L as the language identification components, denoted as $c_L \in \mathbb{R}^{d \times r}$. Different values of r are explored in the next experiments section. Language identification components are removed as follows. Given a multilingual representation e_L in language L , we subtract the projection of e_L onto c_L from e_L , i.e.

$$e_L := e_L - c_L \frac{c_L^T e_L}{\|e_L\|_2} \quad (1)$$

3 Experiments

In the following experiments, sentences used for extracting principle components are sampled from Wiki-40B (Guo et al., 2020). We use 10,000 sentences per language. We notice performance initially increases as more sentences are used but then is almost unchanged after $n > 10,000$. We tried different samplings of $\{t_L^i\}$ and text resources other than Wiki-40B, e.g., Tatoeba (Artetxe and Schwenk, 2019). The minimal differences in performance suggest language components are stable over different domains.

3.1 Cross-lingual Answer Retrieval

We first examine LIR on LARQA, a cross-lingual answer retrieval dataset containing 11 languages (Roy et al., 2020). LARQA consists of two retrieval sub-datasets: XQuAD-R and MLQA-R. XQuAD-R is built by translating 240 paragraphs in the SQuAD v1.1 dev set into 10 languages and converting them to retrieval tasks following

the procedure from ReQA (Ahmad et al., 2019). Similarly, MLQA-R is constructed by converting MLQA (Lewis et al., 2020) to QA retrieval. In other words, each question in LAReQA has 11 relevant answers, one in each language. Two retrieval models with self language bias are presented in the LAReQA original paper, i.e. “En-En” and “X-X”. Specifically, the multilingual model “En-En” finetunes mBERT for QA retrieval on the 80,000 English QA pairs from the SQuAD v1.1 train set using a ranking loss. The model “X-X” trains on the translation (into 11 languages) of the SQuAD train set. In one training example, the question and answer are in the same language. Since given a question query, all positive examples are within-language, “En-En” and “X-X” exhibit strong self-language bias and weak-alignment property.

For evaluation, we first compute the language identification components with “En-En” and “X-X” models released by LAReQA. For testing, language identification components are removed from question and answer embeddings following Eq. (1). Results are shown in Table 1 and the evaluation metric is mean average precision (MAP) of retrieval. Detailed results for each language are provided in the appendix (Table 5). Simply applying LIR results in significant improvements, almost 100% relatively for “X-X” model on XQuAD-R. This huge boost reveals the algebraic structure for multilingual representation space: in weak-alignment multilingual system, the principal components primarily encode language information. In LAReQA, each language has one of the relevant answers. The performance improvement itself already indicates less language bias.

	XQuAD-R		MLQA-R	
	En-En	X-X	En-En	X-X
w/o LIR	27.8	23.3	35.7	26.0
$r = 1$	36.7	45.2	37.0	42.4
$r = 2$	36.7	45.6	36.2	41.6
$r = 3$	36.5	45.9	36.3	41.6
$r = 4$	36.4	45.7	36.1	41.4

Table 1: Mean average precision (MAP) of model “En-En” and “X-X” with and without LIR.

To further illustrate the effect of LIR, we plot the 2D PCA projection of questions and candidates in Chinese and English for the XQuAD-R dataset. Without LIR, as plotted on the left of Fig. 2, Chinese and English embeddings are separated while questions and candidates in the same language are clustering together¹. This weak-alignment prop-

¹Note these two subfigures on the left are reproduced by

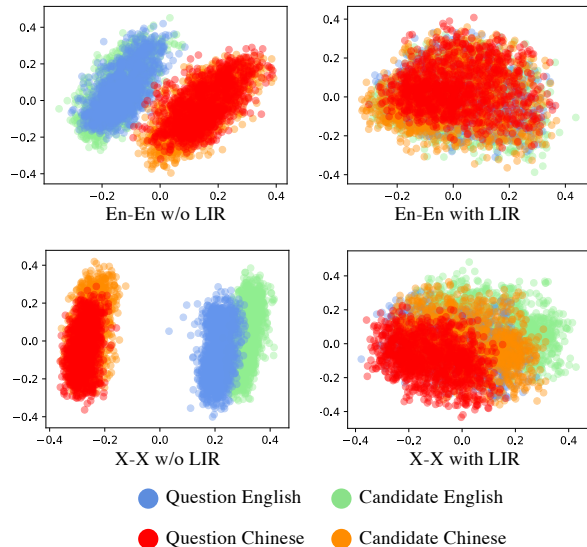


Figure 2: PCA projections of English and Chinese embeddings on the XQuAD-R dataset, with and without LIR. The two subfigures on the left are reproduced by authors to follow Figure 5 in Roy et al. (2020).

erty is especially prominent for model “X-X”. After applying LIR, the separation between the two languages vanishes. Questions and candidates embeddings, no matter which language they are in, group together. Both model “En-En” and “X-X” now exhibit strong cross-lingual alignment.

3.2 Amazon Reviews

We further evaluate LIR on zero-shot transfer learning with Amazon Reviews Dataset (Prettenhofer and Stein, 2010). In this subsection, we use multilingual BERT (Devlin et al., 2019) as the embedding model. Following Chidambaram et al. (2019), the original dataset is converted to a classification benchmark by treating reviews of more than 3 stars as positive and negative otherwise. We split 6000 English reviews in the original training set into 90% for training and 10% for development. A logistic classifier is trained on the English training set and then evaluated on English, French, German and Japanese test sets (each has 6000 examples) using the same trained model, i.e. the evaluation is zero-shot. The weights for mBERT are fixed. The representation of a sentence/phrase is computed as the average pooling of the transformer encoder outputs. LIR is applied in both training and evaluation stage using the corresponding language components.

Results presented in Table 2 show that removing the language components from multilingual representation is beneficial for cross-lingual zero-shot

authors to imitate Figure 5 in Roy et al. (2020) in order to better demonstrate the effects of LIR

transfer learning of mBERT. LIR is expected to leave only semantic-related information in the representation so that the logistic classifier trained on English should be conveniently transferred to other languages. Another interesting observation is that unlike semantic retrieval, the peak performance usually occurs at $r > 1$.

	en	de	fr	jp	Avg.
w/o LIR	80.0	70.4	73.1	71.7	73.8
$r = 1$	80.2	70.9	74.6	72.6	74.6
$r = 2$	80.5	70.9	75.6	73.1	75.0
$r = 3$	80.2	70.8	75.4	71.8	74.5
$r = 5$	80.2	70.8	76.1	72.0	74.8
$r = 8$	80.2	71.1	75.2	70.8	74.3
$r = 10$	80.3	71.0	76.0	71.2	74.6
$r = 12$	80.0	70.9	76.0	71.4	74.6

Table 2: Classification accuracy on Amazon Reviews Dataset.

3.3 XEVAL

We have tested LIR on cross-lingual benchmarks in previous sections. In this section, we apply LIR in XEVAL, a collection of multilingual sentence representation benchmark (Yang et al., 2020). The training set and test set of XEVAL are in the **same** language (i.e. the evaluation is not cross-lingual). Benchmarks on XEVAL include Movie Reviews (Pang and Lee, 2005), binary SST (sentiment analysis, Socher et al. (2013)), MPQA (opinion-polarity, Wiebe et al. (2005)), TREC (question-type, Voorhees and Tice (2000)), CR (product reviews, Hu and Liu (2004)), SUBJ (subjectivity/objectivity, Pang and Lee (2004)) and SICK (both entailment and relatedness (Marelli et al., 2014)). For this evaluation, we use mBERT as the base multilingual encoder. Still the weights of mBERT are fixed during training and only downstream neural structures are trained. The training, cross-validation and evaluation uses SentEval toolkit (Conneau and Kiela, 2018).

Results are presented in Table 3. The metric is the averaging performance across 9 datasets mentioned above. Introducing LIR is beneficial on German, Spanish, French and Chinese. We also notice that for English dataset, removing principal components actually hurts the performance. This observation also echoes with findings in previous English sentence embedding works, e.g. Yang et al. (2019b). We speculate this is because English data are dominant in mBERT training data. Therefore mBERT representations exhibit similar behaviors with monolingual English sentence embeddings.

	en	de	es	fr	zh	Avg.
w/o LIR	80.8	78.1	78.8	79.1	79.3	79.2
$r = 1$	80.4	78.2	79.0	79.1	79.3	79.2
$r = 2$	80.7	78.5	79.4	79.3	79.4	79.5
$r = 5$	80.6	78.0	79.4	78.9	79.3	79.2
$r = 10$	80.2	78.4	79.0	79.0	78.9	79.1

Table 3: Results of applying LIR to XEVAL dataset. The metric is the average of 9 downstream tasks.

3.4 Application to Models without Self-Language Bias

In previous sections, we have shown the great effectiveness of LIR on weak-alignment systems. As an additional analysis, we examine LIR on multilingual models without self language bias, i.e. models “**X-X-mono**” and “**X-Y**” introduced in the original LARQA paper. Model “**X-X-mono**” is modified from “**X-X**” by ensuring that each training batch is monolingual so that in-batch negative and positive examples are in the same language. In model “**X-Y**”, questions and answers are allowed to be translated to different languages, which directly encourage the model to regard answers in different languages from the question as correct. With such designs in training, “**X-X-mono**” and “**X-Y**” are shown to be without self-language bias, i.e. semantically relevant representations are closer than all irrelevant items, regardless of their languages.

The evaluation process is similar as in Section 3.1. Results are presented in Table 4. Applying LIR leads to a slight performance decrease for X-X-mono. While the drop in X-Y is notable and we suspect this is because the training process for X-Y avoids, by design, self-language bias. Rather, the principal components of X-Y contain essential semantic-related information for the retrieval task. This result is **not** negative and actually support our argument, since for “strong alignment” multilingual systems, principal components should both contain semantic and language-related information. Then removing principal components will hinder the semantic retrieval. For weak-alignment models, removing just the first component should be adequate for cross-lingual retrieval (table 1). For tasks like classification and sentiment analysis (tables 2 and 3), the optimal number of components to remove seems to vary on different datasets.

4 Related Work & Our Novelty

Different training methods have been proposed to obtain language agnostic representations. LASER (Artetxe and Schwenk, 2019) leverages translation pairs and BiLSTM encoder for multilingual sen-

	XQuAD-R		MLQA-R	
	X-X-mono	X-Y	X-X-mono	X-Y
w/o LIR	50.8	62.6	48.6	48.5
$r = 1$	50.6	59.5	48.8	46.2
$r = 2$	49.8	58.1	48.0	45.5
$r = 3$	49.3	57.1	47.8	44.8
$r = 4$	48.9	56.5	47.4	44.2

Table 4: Mean average precision (MAP) of “X-X-mono” and “X-Y” models without language bias.

tence representation learning. Multilingual USE (Yang et al., 2019a) uses training data such as translated SNLI, mined multilingual QA and translation pairs to learn multilingual sentence encoder. AMBER (Hu et al., 2020a) aligns contextualized representations of multilingual encoders at different granularities. LaBSE (Feng et al., 2020) finetunes a pretrained language model with the bitext retrieval task and mined cross-lingual parallel data to obtain language agnostic sentence representations. In contrast, LIR does not require any parallel data for semantic alignment.

Faruqui and Dyer (2014) propose a canonical correlation analysis (CCA) based method to add multilingual context to monolingual embeddings. The method is post-processing and requires bilingual word translation pairs to determine the projection vectors. In contrast, LIR is post-training and does not require labeled data. Mrkšić et al. (2017) build semantically specialized cross-lingual vector spaces. Like CCA, their methods requires the additional training to adjust the original embeddings using supervised data: cross-lingual synonyms and antonyms. Libovický et al. (2019) propose that the language-specific information of mBERT is the centroid of each language space (the mean of embeddings). Zhao et al. (2021) propose several training techniques to obtain language-agnostic representations, including segmenting orthographic tokens in training data and aligning monolingual spaces by training. In contrast, LIR is post-training and model-agnostic. Critically, this means LIR can be conveniently and easily applied to any trained multilingual systems without further training.

Previous explorations on principal components of the semantic space for sentence embeddings include Arora et al. (2017) and Yang et al. (2019b), whereby principal component removal is investigated for monolingual models and the evaluation is only conducted on semantic similarity benchmarks. In contrast, our work investigates the multilingual case and the evaluation is more diverse, e.g. cross-lingual transfer learning. Mu and Viswanath

(2018) explore removing top components from English representations. However, it was unclear prior to our work what purpose is served by removing principal components within multilingual and cross-lingual settings. We demonstrate these principal components represent language information for weak-alignment multilingual models.

Compared with Yang et al. (2020), the novelty of this work is two-fold. First, it is unclear in Yang et al. (2020) whether the assumption (i.e. principal components contain language information) holds true for both weak and strong-alignment multilingual models. In this work we clearly show that it is valid for weak-alignment models (Section 3.1). However, for strong-alignment systems, the assumption is not quite true (Table 4). Second, in Yang et al. (2020), the evaluation is only conducted on Tatoeba, a semantic retrieval dataset. While in this work, evaluations are more comprehensive. Besides the cross-lingual retrieval dataset LARQA, our experiments include cross-lingual zero-shot learning (Section 3.2) and monolingual transfer learning (Section 3.3). These extra results establish the effectiveness of LIR beyond the domain of semantic retrieval.

5 Conclusion

In this paper, we investigate the self-language bias in multilingual systems. We explore a simple method “Language Identity Removal (LIR)”. This method identifies and removes the language information in multilingual semantic space by singular value decomposition and orthogonal projection. Although as a simple and linear-algebra-only method, LIR is highly effective in several downstream tasks, including zero-shot transfer learning, sentiment analysis, etc. Especially for cross-lingual retrieval, introducing LIR increases the performance of weak-alignment multilingual systems by almost 100% relatively in MAP.

Acknowledgments

We would like to thank anonymous reviewers for their comments, as well as our teammates from Descartes, Google Brain and other Google teams for their valuable feedback.

References

Amin Ahmad, Noah Constant, Yinfei Yang, and Daniel Cer. 2019. Reqa: An evaluation for end-to-end answer retrieval models. In *Proceedings of the 2nd*

- Workshop on Machine Reading for Question Answering*, pages 137–146.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *5th International Conference on Learning Representations, ICLR 2017*.
- 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.
- Mikel Artetxe and Holger Schwenk. 2019. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Learning cross-lingual sentence representations via a multi-task dual-encoder model](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 250–259, Florence, Italy. Association for Computational Linguistics.
- Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2020. Re-thinking embedding coupling in pre-trained language models. *arXiv preprint arXiv:2010.12821*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard 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.
- Alexis Conneau and Douwe Kiela. 2018. Senteval: An evaluation toolkit for universal sentence representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*, pages 7059–7069.
- 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.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 462–471.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*.
- Mandy Guo, Zihang Dai, Denny Vrandečić, and Rami Al-Rfou. 2020. [Wiki-40b: Multilingual language model dataset](#). In *LREC 2020*.
- Junjie Hu, Melvin Johnson, Orhan Firat, Aditya Siddhant, and Graham Neubig. 2020a. Explicit alignment objectives for multilingual bidirectional encoders. *arXiv preprint arXiv:2010.07972*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020b. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *arXiv preprint arXiv:2003.11080*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. From zero to hero: On the limitations of zero-shot language transfer with multilingual transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499.
- 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*, pages 7315–7330.
- Jindřich Libovický, Rudolf Rosa, and Alexander Fraser. 2019. How language-neutral is multilingual bert? *arXiv preprint arXiv:1911.03310*.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*, pages 216–223.
- Nikola Mrkšić, Ivan Vulić, Diarmuid Ó Séaghdha, Ira Leviant, Roi Reichart, Milica Gašić, Anna Korhonen, and Steve Young. 2017. Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the association for Computational Linguistics*, 5:309–324.
- Jiaqi Mu and Pramod Viswanath. 2018. All-but-the-top: Simple and effective postprocessing for word representations. In *International Conference on Learning Representations*.

- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 271–278.
- Bo Pang and Lillian Lee. 2005. Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 115–124.
- Peter Prettenhofer and Benno Stein. 2010. Cross-language text classification using structural correspondence learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 1118–1127.
- Uma Roy, Noah Constant, Rami Al-Rfou, Aditya Barua, Aaron Phillips, and Yinfei Yang. 2020. [LAREQA: Language-agnostic answer retrieval from a multilingual pool](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5919–5930, Online. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Ellen M Voorhees and Dawn M Tice. 2000. Building a question answering test collection. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 200–207.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39(2-3):165–210.
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernandez Abrego, Steve Yuan, Chris Tar, Yun-Hsuan Sung, et al. 2019a. Multilingual universal sentence encoder for semantic retrieval. *arXiv preprint arXiv:1907.04307*.
- Ziyi Yang, Yinfei Yang, Daniel Cer, Jax Law, and Eric Darve. 2020. Universal sentence representation learning with conditional masked language model. *arXiv preprint arXiv:2012.14388*.
- Ziyi Yang, Chenguang Zhu, and Weizhu Chen. 2019b. Parameter-free sentence embedding via orthogonal basis. 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)*, pages 638–648.
- Wei Zhao, Steffen Eger, Johannes Bjerva, and Isabelle Augenstein. 2021. [Inducing language-agnostic multilingual representations](#). In *Proceedings of *SEM 2021: The Tenth Joint Conference on Lexical and Computational Semantics*, pages 229–240, Online. Association for Computational Linguistics.

A Experimental results for each language of model “X-X” on LAReQA

Here we provide the detailed experiment results of each language on the XQuAD-R dataset. The multilingual encoder is model “X-X”.

	w/o LIR	$r = 1$	$r = 2$	$r = 3$	$r = 4$
ar	20.5	40.5	40.4	40.4	40.0
de	27.5	48.3	49.8	49.7	49.6
el	20.9	43.5	43.9	44.1	44.2
en	27.3	55.1	55.0	55.3	55.3
es	27.6	52.6	52.8	52.7	52.6
hi	18.6	36.5	36.8	37.5	37.3
ru	24.9	48.2	49.6	49.6	49.4
th	16.8	34.7	35.1	34.9	34.6
tr	23.8	45.3	45.4	46.3	46.2
vi	24.8	49.2	48.9	48.9	48.6
zh	24.7	43.8	43.8	45.3	45.2

Table 5: Experimental results for each language of model “X-X” on the XQuAD-R dataset.