

ViDeBERTa: A powerful pre-trained language model for Vietnamese

Cong Dao Tran *
FPT Software AI Center
daotc2@fsoft.com.vn

Nhut Huy Pham *
FPT Software AI Center
huypn10@fsoft.com.vn

Anh Nguyen
Microsoft
anhnguyen@microsoft.com

Truong Son Hy †
University of California San Diego
tshy@ucsd.edu

Tu Vu
University of Massachusetts Amherst
tuvu@cs.umass.edu

Abstract

This paper presents ViDeBERTa, a new pre-trained monolingual language model for Vietnamese, with three versions - ViDeBERTa_{small}, ViDeBERTa_{base}, and ViDeBERTa_{large}, which are pre-trained on a large-scale corpus of high-quality and diverse Vietnamese texts using DeBERTa architecture. Although many successful pre-trained language models based on Transformer have been widely proposed for the English language, there are still few pre-trained models for Vietnamese, a low-resource language, that perform good results on downstream tasks, especially Question answering. We fine-tune and evaluate our model on three important natural language downstream tasks, Part-of-speech tagging, Named-entity recognition, and Question answering. The empirical results demonstrate that ViDeBERTa with far fewer parameters surpasses the previous state-of-the-art models on multiple Vietnamese-specific natural language understanding tasks. Notably, ViDeBERTa_{base} with 86M parameters, which is only about 23% of PhoBERT_{large} with 370M parameters, still performs the same or better results than the previous state-of-the-art model. Our ViDeBERTa models are available at: <https://github.com/HySonLab/ViDeBERTa>.

1 Introduction

In recent years, pre-trained language models (PLMs) and Transformer-based architecture models have been essential in the advancement of Natural Language Processing (NLP). Large-scale Transformer-based pre-trained models with the capacity to derive a contextual representation of the languages in the training data include GPT (Radford et al., 2019; Brown et al., 2020), BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2019b), ELECTRA (Clark et al., 2020), T5 (Raffel et al., 2020), and DeBERTa (He

et al., 2020, 2021). Following pre-training, these models performed at the cutting edge on various downstream NLP tasks (Devlin et al., 2019). The development of pre-trained models in other languages, including Vietnamese (PhoBERT (Nguyen and Nguyen, 2020); ViBERT (Tran et al., 2020); ViT5 (Phan et al., 2022)), and Arabic (Antoun et al., 2021), has been spurred on by the success of pre-trained models in English. In order to enhance performance across several languages by learning both general and language-specific representations, multilingual pre-trained models (XLM-R (Conneau et al., 2020), mT5 (Xue et al., 2021), and mBART (Liu et al., 2020) are also being developed.

Most recently, PhoBERT (Nguyen and Nguyen, 2020), the first large pre-trained model for Vietnamese that inherits the RoBERTa (Liu et al., 2019) architecture, has demonstrated the effectiveness of the trained language model compared with current methods modernized in four Vietnamese-specific tasks, including Part of Speech Tagging (POS), Dependency Parsing, Named Entity Recognition (NER), and Natural Language Inference (NLI). Nevertheless, there are still rooms to build an improved pre-trained language model for Vietnamese. Firstly, PhoBERT was pre-trained on a relatively small Vietnamese dataset of 20GB of uncompressed texts, while pre-trained language models can be significantly improved by using more pre-training data (Liu et al., 2019). Secondly, Question answering (QA) is one of the most impactful tasks that has mainly focused on the computational linguistics and artificial intelligence research community within information retrieval and information extraction in recent years. However, there are a few pre-trained models for Vietnamese that produce efficient results in the QA tasks, especially PhoBERT (Nguyen and Nguyen, 2020) and ViT5 (Phan et al., 2022). Last but not least, some previous works point to DeBERTa architecture (He et al., 2020, 2021) using several novel techniques

*: Co-first authors. †: Correspondent author.

that can significantly outperform RoBERTa and improve the efficiency of model pre-training and the performance of downstream tasks in some respects.

Inspired by that, we introduce an improved large-scale pre-trained language model, ViDeBERTa, trained on CC100 Vietnamese monolingual, following the architecture and pre-training methods of DeBERTaV3 (He et al., 2021). We comprehensively evaluate and compare our model with competitive baselines, i.e., the previous SOTA models PhoBERT, ViT5, and the multilingual model XLM-R on three Vietnamese downstream tasks, including POS tagging, NER, and QA. In this work, we focus on two main categories of QA: Machine Reading Comprehension (MRC) and Open-domain Question Answering (ODQA). The experiment results show the performance of our model surpasses all baselines on all tasks. Our main contributions are summarized as follows:

- We present and implement ViDeBERTa with three versions: ViDeBERTa_{small}, ViDeBERTa_{base}, and ViDeBERTa_{large} which are the improved large-scale monolingual language models pre-trained for Vietnamese based on the DeBERTa architecture and pre-training procedure.
- We also conduct extensive experiments to verify the performance of our pre-trained models compared to previous strong models in terms of Vietnamese language modeling. Our empirical results demonstrated the state-of-the-art (SOTA) results on Vietnamese downstream tasks: POS tagging, NER, and QA, thus confirming the effectiveness of our improved pre-trained language model for Vietnamese.
- Our model, ViDeBERTa, which works with *huggingface* and *transformers*, is available to the public. We expect that ViDeBERTa will be an effective pre-trained model for many NLP applications and research in Vietnamese and other low-resource languages.

2 Related work

Pre-trained language models for Vietnamese. PhoBERT (Nguyen and Nguyen, 2020) is the first large-scale PLM for Vietnamese, which has the same architecture as BERT (Devlin et al., 2019) and the same pre-training approach as RoBERTa (Liu et al., 2019) for more robust performance. This

model was trained on a Vietnamese Wikipedia corpus of 20GB word-level texts and produced SOTA results on Vietnamese understanding tasks such as POS, NER, Dependency parsing, and NLI. Following PhoBERT, ViBERT (Tran et al., 2020) and ViELECTRA are public monolingual language models for Vietnamese based on BERT and ELECTRA pre-training techniques (Clark et al., 2020) that are pre-trained on syllable-level Vietnamese textual data. Recent works such as BARTpho (Tran et al., 2021) and ViT5 (Tran et al., 2020) are pre-trained for Vietnamese text summarization.

Fine-tuning tasks. This work utilizes three Vietnamese natural language understanding (NLU) tasks, including POS tagging, NER, and QA, for fine-tuning and evaluating our model’s performance. For POS tagging and NER, PhoBERT still produces better results than ViELECTRA, PhoNLP, and ViT5 (Nguyen and Nguyen, 2020, 2021; Phan et al., 2022). While early QA (Voorhees et al., 1999; Brill et al., 2002; Ferrucci et al., 2010) systems were commonly complex and had many parts, MRC models have evolved and now suggest a simpler two-stage retriever-reader framework (Chen et al., 2017). A context retriever first selects a small subset of passages where some of them contain the answer to the question then a machine reader can carefully review the retrieved contexts and determine the correct answer. The tasks based on QA have gained much attention in recent years in the Vietnamese natural language processing and computational linguistics community. However, to the best of our knowledge, there is only the work (Van Nguyen et al., 2022) that proposes the first Vietnamese retriever-reader QA system employing a transformer-based model (XLM-R) evaluated on the ViQuAD corpus (Nguyen et al., 2020).

3 ViDeBERTa

3.1 Pre-training data

In this work, we use a large corpus CC100 Dataset of 138GB uncompressed texts (Monolingual Datasets from Web Crawl Data) (Conneau et al., 2020) as a pre-training dataset. This corpus includes data for romanized languages and monolingual data for more than 100 languages.

According to Nguyen and Nguyen (2020); Tran et al. (2021), pre-trained language models trained on word-level data can perform better than those trained on syllable-level data for word-level Vietnamese NLP tasks. As a result, we perform word

and sentence segmentation using a Vietnamese toolkit PyVi¹ on the pre-training dataset. After that, we use a pre-trained SentencePiece tokenizer from DeBERTaV3 (He et al., 2021) to segment these sentences with sub-word units, which have a vocabulary of 128K sub-word types.

3.2 Model Architecture

Our model, ViDeBERTa, follows the DeBERTaV3 architecture by He et al. (2021), which is trained using the self-supervised learning objectives of MLM and RTD task and a new weight-sharing Gradient-Disentangled Embedding Sharing (GDES) to enhance the performance of the model. We present three versions of our model, ViDeBERTa_{xsmall}, ViDeBERTa_{base}, and ViDeBERTa_{large} with 22M, 86M, and 304M backbone parameters, respectively.

The details of our model architecture hyper-parameters are listed in Table 1.

Table 1: Statistic of our model hyper-parameters. #layer and #heads denote the numbers of layers and attention heads of ViDeBERTa model versions, respectively.

Model	#layers	#heads	hidden size
ViDeBERTa _{xsmall}	6	12	768
ViDeBERTa _{base}	12	12	768
ViDeBERTa _{large}	24	12	1024

3.3 Optimization

We employ our model based on the DeBERTaV3 implementation from (He et al., 2021). We use Adam (Kingma and Ba, 2015) as the optimizer with weight decay (Loshchilov and Hutter, 2018) and use a global batch size of 8,192 across 32 A100 GPUs (80GB each) and a peak learning rate of 6e-4 for both ViDeBERTa_{xsmall} and ViDeBERTa_{base}, while peak learning rate of 3e-4 was used for ViDeBERTa_{large}. We pre-train ViDeBERTa_{xsmall} and ViDeBERTa_{base} for 500k training iterations and ViDeBERTa_{large} for 250k training iterations.

4 Experiments and Results

4.1 POS tagging and NER

4.1.1 Experimental setup

For POS tagging and NER tasks, we use standard benchmarks of the VLSP POS tagging dataset² and the PhoNER dataset (Truong et al., 2021).

¹<https://pypi.org/project/pyvi/>

²<https://vlsp.org.vn/vlsp2013/eval/ws-pos>

We follow the procedure in Devlin et al. 2019; Nguyen and Nguyen 2020 to fine-tune our pre-trained model for POS tagging and NER tasks. In particular, a linear layer for prediction is appended on top of our model architecture (the last Transformer layer). We then use Adam (Kingma and Ba, 2015) to optimize our model for fine-tuning with a fixed learning rate of 1e-5 and batch size of 16 (He et al., 2021). The final results for each task and each dataset are averaged and reported over five independent runs with different random seeds.

We compare the performance of ViDeBERTa models with the solid baselines, including PhoBERT, XLM-R, and ViT5, for these tasks. Here, XLM-R is a multilingual masked language model pre-trained on 2.5 TB of CommonCrawl dataset of 100 languages, which includes 137GB of Vietnamese texts.

4.1.2 Main results

Model	POS	NER	MRC
	Acc.	F ₁	F ₁
XLM-R _{base}	96.2 [†]	—	82.0 [‡]
XLM-R _{large}	96.3 [†]	93.8*	87.0 [‡]
PhoBERT _{base}	96.7 [†]	94.2*	80.1
PhoBERT _{large}	96.8 [†]	94.5*	83.5
ViT5 _{base1024-length}	—	94.5*	—
ViT5 _{large1024-length}	—	93.8*	—
ViDeBERTa _{xsmall}	96.4	93.6	81.3
ViDeBERTa _{base}	96.8	94.5	85.7
ViDeBERTa _{large}	97.2	95.3	89.9

Table 2: Test results (%) for three tasks POS tagging (POS for short), NER, and MRC on test sets. Note that “Acc.” abbreviates the accuracy. †, *, and ‡ denote scores taken from the PhoBERT paper (Nguyen and Nguyen, 2020), the ViT5 paper (Phan et al., 2022), and the ViQuAD paper (Nguyen et al., 2020), respectively.

Table 2 shows the obtained scores of ViDeBERTa compared to the baselines with the highest reported results. It can be seen clearly that our model produces significantly better results than the baselines and achieves new SOTA performance on both POS tagging and NER tasks.

For POS tagging, ViDeBERTa obtains 0.9% and 0.4% absolute higher accuracy than the large-scale multilingual model XLM-R (Nguyen et al., 2020) and the previous SOTA model PhoBERT (Nguyen and Nguyen, 2020), respectively. Table 2 also shows our ViDeBERTa_{xsmall} obtains 96.4% accuracy that are better than the baseline XLM-R_{large}

and ViDeBERTa_{base} obtains 96.8% that are competitively the same as the PhoBERT_{large}.

For NER, our ViDeBERTa_{large} achieves F₁ score at 95.3% and improves 0.8% absolute higher score than the previous SOTA models ViT5_{base1024-length} and PhoBERT_{large}. Furthermore, ViDeBERTa_{large} and ViDeBERTa_{base} perform 1.5% and 0.7% absolute higher scores than the baseline XLM-R_{large} on the PhoNER corpus.

4.2 Question Answering

4.2.1 Experimental setup

For QA, we evaluate our model on two main tasks: MRC and ODQA. For ODQA, we propose a new framework ViDeBERTa-QA, that uses a BM25 (Robertson et al., 2009) as a retriever and ViDeBERTa as a text reader.

Figure 1 depicts an overview of our ViDeBERTa framework for the Vietnamese Open-domain Question answering task. The statistics of the ViQuAD dataset used for the task, which is introduced by Nguyen et al. (2020), are summarized in Table 3.

Corpus	#article	#passage	#question
Train	138	4,101	18,579
Dev	18	515	2,285
Test	18	493	2,21
Full	174	5,109	23,074

Table 3: Statistics of the ViQuAD dataset for QA. “#article”, “#valid”, and “#test” denote the number of articles, passages, and questions in the ViQuAD, respectively.

We compare ViDeBERTa to the best model XLM-R (Nguyen et al., 2020) and PhoBERT³ for Vietnamese MRC. We also compare our framework, ViDeBERTa-QA, to strong baselines DrQA (Chen et al., 2017), BERTserini (Yang et al., 2019a), and the first Vietnamese ODQA system XLMRQA (Van Nguyen et al., 2022) that uses XLM-R_{large} as a reader. We use the ViQuAD corpus introduced by Nguyen et al. (2020) for assessing these tasks. ViQuAD is a Vietnamese corpus that comprises over 23k triples and each triple includes a question, its answer, and a passage containing the answer.

Similar to POS tagging and NER, we use Adam (Kingma and Ba, 2015) as an optimizer with a learning rate of 2e-5 and a batch size of 16. We report

³We carefully fine-tune PhoBERT for the MRC task following the fine-tuning approach that we use for ViDeBERTa.

the final results as an average over five independent runs with different random seeds.

4.2.2 Main results

Table 2 presents the results obtained by ViDeBERTa and two baselines XLM-R (reported by Nguyen et al. (2020)) and PhoBERT for MRC on ViQuAD corpus. We find that our ViDeBERTa performance outperforms both XLM-R and PhoBERT in terms of F₁ score.

In particular, the previous SOTA model XLM-R_{large} for Vietnamese MRC obtains 87%. Clearly, ViDeBERTa helps boost the XLM-R with about 2.9% absolute improvement, obtaining a new SOTA result at 89.9%. In addition, both versions ViDeBERTa_{base} and ViDeBERTa_{large} also outperform PhoBERT_{base} and PhoBERT_{large} by large margins, respectively. Especially, ViDeBERTa_{xsmall} (22M parameters) produces 1.2% absolute higher score than PhoBERT_{base} (135M parameters) and ViDeBERTa_{base} (86M parameters) produces 2.2% absolute higher score than PhoBERT_{large} (370M parameters) but uses far fewer parameters than PhoBERT.

For ODQA, Table 4 shows the obtained F₁ scores for ViDeBERTa-QA and its baselines on the test set. Obviously, ViDeBERTa-QA achieves better scores than the previous SOTA XLMRQA, BERTserini, and DrQA at the top k passages, selected by retrievers, is 10 and 20. In particular, ViDeBERTa-QA performs 0.85% (at $k = 20$) and 0.4% (at $k = 10$) absolute higher scores than the previous SOTA system. At smaller k ($= 1, 5$), ViDeBERTa performs better BERTserini and DrQA by a large margin; however, XLMRQA does better than ViDeBERTa-QA.

Model	Top k selected passages			
	1	5	10	20
DrQA [*]	37.86	37.86	37.86	37.86
BERTserini [*]	55.55	58.30	57.98	58.09
XLMRQA [*]	61.83	64.99	64.49	64.49
ViDeBERTa _{xsmall}	52.76	56.24	56.93	57.40
ViDeBERTa _{base}	58.55	61.37	61.89	62.43
ViDeBERTa _{large}	61.23	63.57	64.89	65.34

Table 4: Test scores (F₁ in %) for ODQA on ViQuAD corpus with different k values. Note that [*] indicates the results reported following Van Nguyen et al. (2022).

4.3 Discussion

According to the results on both downstream tasks of POS tagging and NER in Table 2, we find that

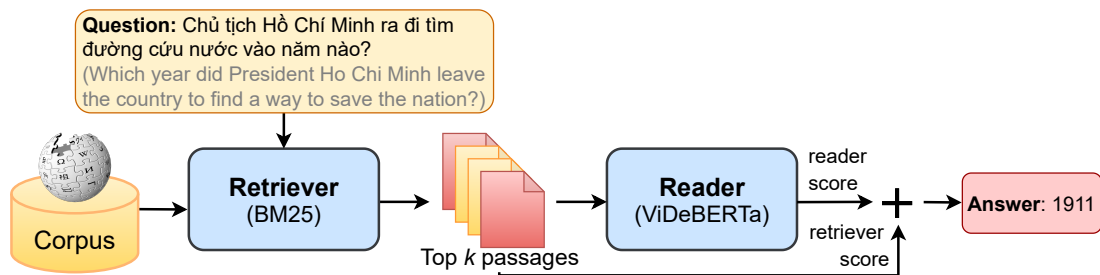


Figure 1: An overview of ViDeBERTa-QA framework for Vietnamese Open-domain Question Answering task.

ViDeBERTa_{small} (86M) with fewer parameters (i.e. only about 15% of XLM-R_{large} 560M and 25% of PhoBERT_{large} 370M) but still performs slightly better than XLM-R_{large} and competitively the same as the previous SOTA PhoBERT_{large}. One possible reason is that our model inherits the robustness of DeBERTaV3 architecture and pre-training techniques, which are demonstrated superior performance by He et al. (2020, 2021). Moreover, using more high-quality pre-training data (138GB) can help ViDeBERTa significantly improve its performance compared to PhoBERT (using 20GB).

For Vietnamese QA, the results on the MRC task show that ViDeBERTa outperforms PhoBERT by a large margin. It is worth noting that PhoBERT set a maximum length of 256 subword tokens for both versions while ViDeBERTa set a larger one of 512. As a result, our models are more scalable than PhoBERT for long contexts. The results obtained by ViDeBERTa-QA on ODQA also suggest that our framework achieves the best performance with large top k passages selected by the retriever (i.e. $k = 10, 20$).

5 Conclusion

In this paper, we have introduced ViDeBERTa, a new pre-trained large-scale monolingual language model for Vietnamese. We demonstrate the effectiveness of our ViDeBERTa by showing that ViDeBERTa with fewer parameters performs better than the recent strong pre-trained language models as XLM-R, PhoBERT, and ViT5, and achieves SOTA performances for three downstream Vietnamese language understanding tasks, including POS tagging, NER, and especially QA. We hope that our public ViDeBERTa model will boost ongoing NLP research and applications for Vietnamese and other low-resource languages.

Limitations

While we have shown that ViDeBERTa can achieve state-of-the-art performance on a variety of NLP tasks for Vietnamese, we believe that more analyses and ablations are required to better understand what facets of ViDeBERTa contributed to its success and what knowledge of Vietnamese that ViDeBERTa captures. We leave these further explorations to future work.

References

- Wissam Antoun, Fady Baly, and Hazem Hajj. 2021. Aragpt2: Pre-trained transformer for arabic language generation. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 196–207.
- Eric Brill, Susan Dumais, and Michele Banko. 2002. An analysis of the askmsr question-answering system. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, pages 257–264.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- 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 Asso-*

- ciation for Computational Linguistics*, pages 8440–8451.
- 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.
- David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, et al. 2010. Building watson: An overview of the deepqa project. *AI magazine*, 31(3):59–79.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- 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.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Dat Quoc Nguyen and Anh-Tuan Nguyen. 2020. Phobert: Pre-trained language models for vietnamese. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1037–1042.
- Kiet Nguyen, Vu Nguyen, Anh Nguyen, and Ngan Nguyen. 2020. A vietnamese dataset for evaluating machine reading comprehension. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2595–2605.
- Linh The Nguyen and Dat Quoc Nguyen. 2021. Phonlp: A joint multi-task learning model for vietnamese part-of-speech tagging, named entity recognition and dependency parsing. *arXiv preprint arXiv:2101.01476*.
- Long Phan, Hieu Tran, Hieu Nguyen, and Trieu H Trinh. 2022. Vit5: Pretrained text-to-text transformer for vietnamese language generation. *arXiv preprint arXiv:2205.06457*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Nguyen Luong Tran, Duong Minh Le, and Dat Quoc Nguyen. 2021. Bartpho: Pre-trained sequence-to-sequence models for vietnamese. *arXiv preprint arXiv:2109.09701*.
- Thi Oanh Tran, Phuong Le Hong, et al. 2020. Improving sequence tagging for vietnamese text using transformer-based neural models. In *Proceedings of the 34th Pacific Asia Conference on Language, Information and Computation*, pages 13–20.
- Thinh Hung Truong, Mai Hoang Dao, and Dat Quoc Nguyen. 2021. Covid-19 named entity recognition for vietnamese. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2146–2153.
- Kiet Van Nguyen, Phong Nguyen-Thuan Do, Nhat Duy Nguyen, Tin Van Huynh, Anh Gia-Tuan Nguyen, and Ngan Luu-Thuy Nguyen. 2022. Xlmrqa: Open-domain question answering on vietnamese wikipedia-based textual knowledge source. *arXiv preprint arXiv:2204.07002*.
- Ellen M Voorhees et al. 1999. The trec-8 question answering track report. In *Trec*, volume 99, pages 77–82.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mt5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019a. End-to-end open-domain question answering with bertserini. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019b. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.

A Background of DeBERTa

DeBERTa enhances BERT with disentangled attention and a more powerful mask decoder. The disentangled attention mechanism is distinct from prior methods in that it uses two distinct vectors to represent each input word: one for the content and one for the location. The words’ attention weights are calculated using disentangled matrices based on both their relative placements and contents. Similar to BERT, DeBERTa has been pre-trained using masked language modeling. The disentangled attention process already accounts for the relative locations and contents of the context words but not for their absolute positions, which are usually crucial for prediction. DeBERTa improves MLM by utilizing a better mask decoder at the MLM decoding layer and absolute position information of the context words.

A.1 Masked Language model

Large-scale Transformer-based PLMs are often pre-trained using a self-supervision aim called Masked Language Model (MLM) (Devlin et al., 2019) to learn contextual word representations in enormous volumes of text. In further detail, we corrupt a given sequence $X = \{x_i\}$ into \tilde{X} by randomly masking 15% of its tokens and train a language model parameterized by θ to reconstruct X by anticipating the masked tokens \tilde{x} conditioned on \tilde{X} :

$$\max_{\theta} \log p_{\theta}(X|\tilde{X}) = \max_{\theta} \sum_{i \in C} \log p_{\theta}(\tilde{x}_i = x_i|\tilde{X}), \quad (1)$$

where C is the sequence’s index set for the masked tokens. The authors of BERT suggest keeping 10% of the masked tokens unchanged, replacing another 10% with tokens chosen at random, and replacing the remaining tokens with the [MASK] token.

A.2 Replaced token detection

Like ELECTRA, which was trained with two transformer encoders in GAN style, DeBERTaV3 (He et al., 2021) improves DeBERTa by using the training loss in the generator is MLM and discriminator is Replaced Token Detection (RTD). The loss function of the generator can be written as follows:

$$L_{MLM} = \mathbb{E} \left(- \sum_{i \in C} \log p_{\theta_G}(\tilde{x}_{i,G} = x_i|\tilde{X}_G) \right), \quad (2)$$

where θ_G and \tilde{X}_G are the parameter and the input of the generator by masking 15% tokens in X , respectively.

The discriminator’s input sequence is constructed by replacing masked tokens with new tokens sampled according to the generator’s output probability:

$$\tilde{x}_{i,D} = \begin{cases} \tilde{x}_i \sim p_{\theta_G}(\tilde{x}_{i,G} = x_i|\tilde{X}_G), & i \in C \\ x_i, & i \notin C \end{cases} \quad (3)$$

The loss function of the discriminator is written as follows:

$$L_{RTD} = \mathbb{E} \left(- \sum_i \log p_{\theta_D}(\mathbb{1}(\tilde{x}_{i,D} = x_i)|\tilde{X}_D) \right), \quad (4)$$

where θ_D is the parameter of the discriminator, $\mathbb{1}(\cdot)$ is the indicator function, and \tilde{X}_D is the input to the discriminator constructed by Equation 4. Then L_{MLM} and L_{RTD} are optimized jointly by the final loss $L = L_{MLM} + \lambda L_{RTD}$, where λ is the weight of the discriminator loss.

Besides using the RTD training loss like ELECTRA (Clark et al., 2020), DeBERTaV3 improves DeBERTa by using a new weight-sharing method called Gradient-Disentangled Embedding Sharing (GDES) (He et al., 2021). The experimental results conducted by He et al. indicate that GDES is an effective weight-sharing method for language model pre-trained with MLM and RTD tasks.