

# Advancing Vietnamese Information Retrieval with Learning Objective and Benchmark

Phu-Vinh Nguyen<sup>1,2</sup>, Minh-Nam Tran<sup>1,2</sup>, Long Nguyen<sup>1,2\*</sup>, Dien Dinh<sup>1,2</sup>

<sup>1</sup>Faculty of Information Technology, University of Science, Ho Chi Minh City, Vietnam

<sup>2</sup>Vietnam National University, Ho Chi Minh City, Vietnam

{npvinh20, tmnam20}@apcs.fitus.edu.vn, {nhblong, ddien}@fit.hcmus.edu.vn

## Abstract

With the rapid development of natural language processing, many language models have been invented for multiple tasks. One important task is information retrieval (IR), which requires models to retrieve relevant documents. Despite its importance in many real-life applications, especially in retrieval augmented generation (RAG) systems, this task lacks Vietnamese benchmarks. This situation causes difficulty in assessing and comparing many existing Vietnamese embedding language models on the task and slows down the advancement of Vietnamese natural language processing (NLP) research. In this work, we aim to provide the Vietnamese research community with a new benchmark for information retrieval, which mainly focuses on retrieval and reranking tasks. Furthermore, we also present a new objective function based on the InfoNCE loss function, which is used to train our Vietnamese embedding model. Our function aims to be better than the origin in information retrieval tasks. Finally, we analyze the effect of temperature, a hyper-parameter in both objective functions, on the performance of text embedding models.

## 1 Introduction

With the born of transformer architecture (Vaswani et al., 2017) since 2017, many language models such as BERT (Devlin et al., 2019), GPT (Brown et al., 2020), and T5 (Raffel et al., 2020) have been developed and have strong performance in many natural language tasks. Furthermore, the rise of many large language models (LLMs) recently, such as Llama (Touvron et al., 2023), Mixtral (Jiang et al., 2024), Qwen (Bai et al., 2023), and Phi (Gunasekar et al., 2023), has gained strong attention for the research community due to their exceptional performance in text generation. However, LLMs have one disadvantage, they cannot access the custom data and new information to update

their knowledge, which makes them unable to shift their knowledge to fit different applications. Consequently, Retrieval-Augmented Generation (Lewis et al., 2020) systems (or RAG) are invented to handle the problem by utilizing retrieval systems to search for relevant information from the database before feeding those information to LLMs as an extra context. This shows the necessity and importance of embedding language models for retrieval and reranking tasks in the era of LLMs.

Despite the importance of retrieval systems for LLMs, in Vietnam, the number of existing benchmarks for retrieval and reranking tasks are limited, which leads to the difficulty in comparing and assessing the performance of many Vietnamese embedding language models on those two tasks. Despite there are some Vietnamese benchmarks like ViGLUE (Tran et al., 2024a), ViNLI (Huynh et al., 2022), VMNLU<sup>1</sup>, and VSFC (Nguyen et al., 2018a), none of them evaluate performance of language models on retrieval and reranking tasks. This paper attempts to address the need for those benchmarks by introducing a new benchmark, the Vietnamese Context Search (or the VCS) to evaluate the ability of text embedding models to search for relevant Vietnamese documents. This benchmark is constructed using existing Vietnamese datasets with modifications in their structure and tasks. Despite having a simple construction process, this benchmark effectively provides different inspections of Vietnamese text embedding models. The VCS serves as a standard and high-quality benchmark to evaluate and compare different Vietnamese embedding models on retrieval and reranking tasks.

Furthermore, this work also introduces a new training objective to train Vietnamese embedding language models on retrieval and reranking tasks. This training objective aims to yield better performance of embedding language models compared to

\* Corresponding author.

<sup>1</sup><https://github.com/ZaloAI-Jaist/VMLU.git>

the InfoNCE loss function, which is usually used in contrastive learning. The research will experiment with different training objectives with two training methods, including in-batch negative and curated hard-negative to compare the ability of two loss functions. Next, the evaluation of some existing Vietnamese embedding language models on the VCS benchmark is conducted to examine their ability in context search. Lastly, an empirical study is conducted to understand the effect of temperature  $\tau$  in the loss function on the overall performance of embedding models. Different training methods are included in the study to further investigate the impact of temperature on the loss function.

To conclude, this work includes three primary contributions:

- First, introduce a new Vietnamese benchmark, the VCS, to evaluate Vietnamese language models in their ability to search relevant documents. This benchmark evaluates models on two tasks, retrieval and reranking tasks.
- Second, introduce a new training objective function to train text embedding models on retrieval and reranking tasks
- Lastly, we conduct an empirical study to investigate the impact of temperature, a hyperparameter, of the InfoNCE and our loss functions in the performance of embedding language models on reranking and retrieval tasks.

## 2 Related Work

In the era of large language models (LLMs), not only does the development of different generative language models such as Gemma (Team et al., 2024), SeaLLM (Nguyen et al., 2023), and Mamba (Dao and Gu, 2024) gains the attention from the community, but also do embedding language models, especially those support searching text documents like GTE (Li et al., 2023), NV-Embed (Lee et al., 2024), BGE (Luo et al., 2024), or GritLM (Muennighoff et al., 2024), become more important due to their applications in RAG systems, which provide more context and information for LLMs to generate correct answers. Consequently, many works aim to provide a benchmark to evaluate language models on their ability in information retrieval (IR) such as BEIR (Thakur et al., 2021), MTEB (Muennighoff et al., 2023), BRIGHT (Su et al., 2024), and ReQA (Ahmad et al., 2019). Those benchmarks advance the development of

many text embedding language models and the research of natural language processing (NLP) by supporting the research community with resources to compare and evaluate text embedding models.

However, similar and comparable benchmarks for Vietnamese embedding language models are limited. While there are some benchmarks like ViGLUE (Tran et al., 2024a), ViQuAD (Nguyen et al., 2020), ViSFD (Luc Phan et al., 2021), VMLU, VSMEC (Ho et al., 2020), and VSFC (Nguyen et al., 2018b), they mostly focus on question-answering and natural language understanding aspects of language models and completely ignore the ability of language models in retrieval and reranking tasks. That leads to the difficulty in evaluating and comparing Vietnamese text embedding models in their ability of retrieve relevant information. Despite some Vietnamese embedding language models being created, without a standard benchmark on this field, the Vietnamese research community is unable to know the benefits, pros, and cons of those language models, which can lead to misleading when applying them to applications (RAG systems) or research projects.

In the early days of information retrieval, different systems were created to find relevant text information from large databases such as TF-IDF (Sammut and Webb, 2010), BM25 (Amati, 2009), and BM25F (Pérez-Agüera et al., 2010). However, those methods cannot capture the context of documents and use it for the retrieval process. Consequently, different retrieval methods using deep learning models are utilized to encode a piece of text to a vector that can present different or hidden aspects of it. Many pre-trained language models (PLMs), including BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and DeBERTa (He et al., 2020), are employed to construct new embedding models due to their capabilities in understanding natural language. Some existing embedding language models like DPR (Karpukhin et al., 2020) (Dense Passage Retrieval) and ColBERT (Khattab and Zaharia, 2020) utilize a dual encoder model structure with two separate encoders, one for queries and one for documents. Despite being fast at reference time, training two separated models would take a lot of effort and time. Meanwhile, cross-encoder models use one encoder for both queries and documents, which is more effective for the training process. Moreover, different approaches are invented to improve the performance of retrieval systems and utilize as much data as

they can such as in-batch negative, which uses other examples within the same training batch as negative samples, curated hard-negative training, selecting challenging negative samples that are difficult to distinguish from positive samples, and sim-CSE (Gao et al., 2021) pre-training method, which employs different dropout rate to create different embedding vectors of a text as positive samples.

### 3 Methodology

In this section, we explain our method to create tasks of the Vietnamese Context Search benchmark and go into detail about this benchmark. Furthermore, we introduce and explain our proposed training loss function, a modified version of the InfoNCE loss function, and our training method to create a Vietnamese text embedding model.

#### 3.1 Vietnamese Context Search benchmark

Due to the lack of Vietnamese benchmarks to compare and evaluate text embedding models on information retrieval tasks, this section proposed a new Vietnamese benchmark to tackle the problem.

##### 3.1.1 ViMedRetrieve

Given a database with  $n$  documents  $d$ , the mission of a retrieval system, given a query  $q$ , is to retrieve documents most relevant to the query  $q$  from the given database. As the number of documents  $n$  increases, this task will become more challenging and require text embedding models to understand natural language to embed sequences more precisely with much information. This real-life scenario inspires us to create a new and similar benchmark to evaluate Vietnamese text embedding models.

This dataset includes  $n$  different pairs of  $(q, d)$ , where  $q$  is the question and  $d$  is the document containing relevant information to answer the question  $q$ . In this task, the primary mission of an embedding language model, given question  $q'$  as input, is to search for the expected document, which is the document  $d'$  of the same pair with  $q'$ , after  $k$  tries. This is similar to how a retrieval system would work in real-life scenarios if we consider  $q$  as user input and  $d$  as the document the user expects to retrieve. For further experiments on this task, we try different values of  $k$  in  $\{5, 10, 20\}$  and take accuracy when  $k = 5$ , reflecting the ability of the embedding language model to retrieve the correct document instantly, is the primary score.

To construct this task, we re-use the ViMedAQA (Tran et al., 2024b) dataset, a

collection of Vietnamese questions and answers in healthcare, and create a new task based on it. This dataset includes four distinguished topics (drug, body part, medicine, and disease). We collect a set of questions and corresponding contexts from the dataset and use them as pairs of queries and documents for this task. To evaluate embedding models on this dataset, we use accuracy as the main metric, the model needs to search for the best  $k$  documents from the whole dataset for each question, and if the model can find the relevant document within the first  $k$  documents, its answer is considered to be correct and vice versa. This process creates a dataset with over 44 thousand pairs of queries and documents. The test set of this dataset, which includes over two thousand samples, is employed to evaluate embedding systems.

##### 3.1.2 ViRerank

Given a query and a list of relevant and irrelevant reference texts, the target of an embedding model in the reranking task is to embed all reference texts, and then rank them based on the similarity of reference and query. This final ranking result is used to evaluate the performance of text embedding models. In this research, we utilize the mean Average Precision (mAP) metric to assess language model ability on all reranking tasks, including ViRerank.

To construct the dataset for the reranking task, we employ the ViNLI dataset, a Vietnamese benchmark for natural language inference (NLI). The ViNLI includes pairs of text pieces labeled to show their relationship, which is classified into one of four classes (entailment, contradiction, neutral, and other). The ViRerank dataset utilizes one part of each ViNLI text pair as the query and the corresponding text piece as the reference. Furthermore, each query in the ViRerank has multiple references as the ViNLI uses the same sentence for many text pairs. Positive references are chosen from text pieces labeled as entailment with the query, while negative references are taken from different labels.

The final result of this process is a new dataset with 363 samples for the test set and 367 samples for the development set, while the train set includes over 3000 samples. However, in this work, to prevent biased evaluation results toward the training and development set, we only use the test set to evaluate Vietnamese text embedding models.

### 3.1.3 MNLIR and QNLIR

Similar to the ViReRank dataset, we utilized two tasks from the ViGLUE dataset, MNLIR and QNLIR, for the reranking task. The MNLIR task requires models to determine the relationship between a pair of sentences. In contrast, the QNLIR task involves determining if the answer to a given question can be found in a sentence from a passage. We collect duplicated texts for each task and use them as queries just like in the ViReRank task. The entailment sentences (corresponding to the query) are used as positive examples and different labels are negative.

We do not employ this method for other NLI tasks of the ViGLUE dataset due to the insufficient amount of duplicated samples in those tasks. Applying this method to MNLIR and QNLIR creates two new sub-sets for reranking tasks, MNLIR-R, with over 3,000 samples, and QNLIR-R, with over 1,000 samples. Despite being the same reranking task, MNLIR-R evaluates models on their ability of reranking based on context similarity while QNLIR-R assesses models on their answer-searching capability.

## 3.2 Training Vietnamese Embedding Model

In this section, we introduce our training method and training objective to train a new Vietnamese embedding model for retrieval and reranking tasks.

### 3.2.1 Model architecture

Given a text sequence  $x = (x_1, \dots, x_n)$  consisting of  $n$  tokens, the objective is to extract information from this piece of text and map it into  $R^d$ , a  $d$ -dimensional space. This task can be fulfilled using an embedding model  $E$  such that  $e = E(x) \in R^d$  where  $e$  is a presentation vector of  $x$  in  $R^d$ .

We first use a pre-trained BERT (Bidirectional Encoder Representations from Transformers) model to extract contextual information of every token in text  $x$ . The output of this model is as follows:

$$c = LM(x) \in R^{n \times d} \quad (1)$$

Where the output  $c$  of the language model is an embedding matrix of  $n$  tokens in the sequence  $x$ , each token is represented by a  $d$ -dimensional vector.

After that, a mean pooling layer is employed to gather all contextual representations of tokens and obtain the final embedding for the entire text.

$$e = \frac{1}{n} \sum_1^n c_i \quad (2)$$

Where  $c_i$  is the context embedding of  $x_i$ , the  $i$ -th token of the sequence. This results in a  $d$ -dimensional vector  $e$ , a presentation of input  $x$ .

### 3.2.2 Instruction training

In retriever and re-ranking tasks, two different inputs are query and document. To handle them separately, some previous work used two embedding modules, one to encode queries and another to encode retrieved documents. This solution requires more resources during the training process as we need to train two embedding models separately.

Another solution is to apply different prompts for the query and document. By giving a hint from the input, the model can understand how to perform different calculations to compute embedding for queries and documents. This method significantly reduces the resources used to train embedding models while ensuring that the model will be trained on as much data as possible, which enhances the model's ability to comprehend the natural language. Some text embedding models such as gte-Qwen2-7B-instruct utilize this method and can achieve extremely high performance.

In this research, we employ instruction training to train our text embedding models. For input query, we add  $\langle |query| \rangle$  as the prefix before feeding the whole text to the model. Meanwhile, we keep the retrieved documents the same without any modification. The difference between the two types of input lets the text embedding model know when and how to embed input query and document.

### 3.2.3 Training methods

In this work, we experiment with two different training methodologies: in-batch negatives and curated hard-negative training, and see how different training methods could affect the performance of the model on retrieval and reranking tasks.

In-batch negative sampling is a technique to improve the model's ability to differentiate the positive and negative pair of text. Given a batch of text,  $x = (x_1, \dots, x_n)$  and its positive pair of text  $x^+ = (x_1^+, \dots, x_n^+)$ , in-batch negative sampling consider all text pieces in the batch, except for the corresponding positive one, are negative. The task is to maximize the similarity of positive pairs and minimize the similarity of all the remaining negative pairs. This method has been proven to be highly resource-effective in training embedding models as it can train on  $n^2$  pairs of text with a batch of  $n$  pairs of text. However, as negative pairs

are collected randomly during training, a negative text pair can be too obvious or not exactly negative.

Meanwhile, curated hard-negative requires the dataset to be more precise and challenging. Given a dataset item  $(x, x^+, x^-)$ , where  $(x^+, x)$  is positive pair and the negative pair is  $(x^-, x)$ . Similar to in-batch negative, the target of curated hard-negative is to maximize the similarity of the positive pair and minimize those of the negative pair. The advantage of this type of training is that the negative pairs can be more challenging to differentiate, which forces the model to learn about different aspects of a text.

### 3.2.4 Training objectives

Denote  $s(x, x^+)$  and  $s(x, x^-)$  are predicted similarity scores of positive and negative text pairs.  $p^+$  is comprehended as the probability of the positive pair. To train an embedding model on contrastive objectives and distinguish relevant documents from those that are irrelevant, one popular objective is the InfoNCE loss, which can be written as follows:

$$p^+ = \frac{e^{s(x, x^+)/\tau}}{e^{s(x, x^+)/\tau} + \sum e^{s(x, x^-)/\tau}} \quad (3)$$

$$L = -\log(p^+) \quad (4)$$

The primary objective of this loss function is to increase the similarity of  $(x, x^+)$  pairs while decreasing the similarity of negative text pairs. However, when the positive pair has a higher probability than other negative pairs, this training objective might still put much effort into increasing it and decreasing the likelihood of different text pairs, ignoring their relationship. With that theory, we modify the InfoNCE loss function, the idea is to lessen the loss more, which leads to slower learning speed, as  $p^+$  gets larger. This target can be easily fulfilled by multiplying the final InfoNCE loss with  $(1 - p^+)$ , resulting in the loss function in Equation 5:

$$L_{ours} = -\log(p^+)(1 - p^+) \quad (5)$$

The term  $(1 - p^+)$  added in the function is used as an extra weight to the loss function, which gets smaller as the probability of the correct pair is higher, slowing down the learning speed of the model on correct examples. This extra weight prevents the over-learned scenario of the original loss function by reducing the gradient from the loss value to the model’s weights on those samples.

### 3.2.5 Training datasets

We create two different training datasets for two different training methods, training with in-batch negative and training with curated hard-negative. Despite having different structures, those datasets share the same data-collecting method. We first collect data from three primary resources, the Vietnamese NewsSapo dataset (Duc et al., 2024), the Binhvq News Corpus<sup>2</sup>, and the Vietnamese version of the QQP triplet (NghiemAbe, 2024). The dataset is summarized in Table 1.

Dataset	Number of samples
BKAINewsCorpus	1.5M
Vietnamese QQP triplet	101K
Binhvq News Corpus	1M

Table 1: Dataset summarization for the training set before filtering samples based on text length

Next, we filter this dataset based on text length. Then, to prepare a dataset for in-batch negative training, we remove all negative examples from each data sample, leaving only a text pair of anchor and positive text. Meanwhile, for curated hard-negative training, we keep the original negative examples while adding negative samples to those text groups that do not have any. We do this by randomly selecting a text piece in the dataset that does not belong to the original group. Although this method may not provide a difficult and high-quality training curated hard-negative dataset, the text embedding models can still learn the relationship between the positive and negative text pairs.

## 4 Experiments

In this section, we compare our modified loss function with the InfoNCE loss function with different training methods. Furthermore, we also evaluate the performance of our embedding language models from the previous step and compare them with some existing Vietnamese embedding models on retrieval and reranking tasks of our benchmark. Lastly, we investigate the effect of temperature  $\tau$  on the performance of embedding models as they are trained with ours and the InfoNCE loss function.

<sup>2</sup><https://github.com/binhvq/news-corpus>

#### 4.1 Comparison of training objectives and training method

In this experiment, we fine-tune the pre-trained BERT-based embedding model <sup>3</sup> on the Vietnamese dataset. Despite this model being pre-trained on English datasets, its performance on our Vietnamese benchmark is reasonably high. Furthermore, its small size can provide an empirical study and comparison of different training methods without requiring much computational resources.

We train text embedding models using two loss functions, ours and the InfoNCE loss function. Furthermore, two training methods, including in-batch negative and curated hard-negative training, are employed in this experiment. Finally, we evaluate those models on our benchmark. The result of this experiment is summarized in Table 2 and Table 3.

		ViNLI	MNLI-R	QNLI-R
baseline		62.42	78.92	<b>87.06</b>
InfoNCE	IB	62.07	75.61	85.26
	HN	66.27	83.86	85.56
ours	IB	63.24	77.15	86.22
	HN	<b>67.86</b>	<b>84.51</b>	86.04

Table 2: Experiment results on reranking tasks using the mAP score. **IB** denotes the in-batch negative training method, and **HN** refers to curated hard-negative training. Results are presented as percentages.

From the experiment results of Table 2, training with hard-negative examples results in better performance compared to the in-batch negative training method for all reranking tasks. Furthermore, in some reranking tasks, training with the in-batch negative method might degrade the performance of the model on this task. Next, our training objective reproduces better performance in all reranking tasks and all methods compared to the InfoNCE loss function despite there is still a degradation in task QNLI-R as we compare with the baseline model. Lastly, the baseline model, despite only being trained on the English datasets, has a relatively high performance. As MNLI and QNLI tasks in the ViGLUE dataset are translations from the GLUE benchmark, some English structural patterns, and similar terminology may be retained in the translated versions, which could explain why the baseline model performs well on these tasks despite having limited knowledge of Vietnamese.

<sup>3</sup><https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

		ViMedRetrieve		
		k@5	k@10	k@20
baseline		0.20	0.37	0.53
InfoNCE	IB	0.25	0.32	0.36
	HN	0.24	0.27	0.29
ours	IB	0.26	<b>0.46</b>	<b>0.59</b>
	HN	<b>0.30</b>	0.44	0.50

Table 3: Experiment results on retrieval tasks with varying numbers of retrieved items. **IB** is in-batch negatives, and **HN** refers to curated hard negatives. Results are presented as percentages based on the accuracy metric.

However, the result from Table 3 shows the opposite: for both objective functions, the performance of the in-batch negative method is higher than that of the hard-negative training method. With our training objective applied, the in-batch negative training method can raise a better result with  $k = 10$  and  $k = 20$  while unable to surpass when  $k = 5$ , this shows that the in-batch negative method has better performance if we want to find correct documents with a large number of finding at a time. Moreover, the result of our objective function is still higher than that of the infoNCE loss function with multiple values of  $k$  and different training methods. Furthermore, the low results of other methods on this task depict the difficulty of this task. Lastly, using the infoNCE loss function degrades significantly the performance of the baseline model in the retrieval task, this can be a consequence of low-quality training data in the case of curated hard-negative training. However, in in-batch negative training, the employed loss function plays a crucial role in this reduced performance.

#### 4.2 Comparison of Vietnamese embedding models

Model	Parameters
SimCSE	130M
Bi-encoder	130M
Sbert	130M
ours	20M

Table 4: Number of parameters of Vietnamese embedding language models in the experiment

This experiment will explore the ability of Vietnamese embedding language models to retrieve and rerank tasks by evaluating the VCS benchmark. The models used in this experiment includes

sup-SimCSE-VietNameese-phobert-base<sup>4</sup>, vietnamese-bi-encoder<sup>5</sup>, vietnamese-sbert<sup>6</sup>, and our models. It is worth noticing that the three first models in the list are trained based on phoBERT with 135M parameters and our experimental model has just 20M, this is stated in Table 4. The result of this experiment is reported in Table 5 and Table 6.

	<b>ViRerank</b>	<b>MNLI-R</b>	<b>QNLI-R</b>
SimCSE	<b>69.46</b>	<b>87.74</b>	88.50
Bi-encoder	65.41	82.10	<b>90.30</b>
Sbert	66.9	83.57	<u>88.79</u>
ours	<u>67.86</u>	<u>84.51</u>	86.04

Table 5: Vietnamese embedding models comparison on reranking tasks, measured by mAP metric. **Bold** text expresses the highest score, Underline highlight the second highest score.

From the evaluation result in Table 5, model sup-SimCSE-VietNameese-phobert-base achieves the highest score on the ViRerank and MNLI-R tasks with a score of 69.46 and 87.74 respectively. Our model comes in second place in the same tasks with 67.86 on ViRerank and 84.51 on MNLI-R. For the last reranking task, QNLI-R, model vietnamese-bi-encoder has the highest score with 88.79 while model vietnamese-sbert is in the second place with 88.79.

	<b>ViMedRetrieve</b>		
	k@5	k@10	k@20
SimCSE	0.09	0.11	0.12
Bi-encoder	<u>0.25</u>	<b>0.45</b>	<b>0.73</b>
Sbert	0.18	0.26	0.32
ours	<b>0.30</b>	<u>0.44</u>	<u>0.50</u>

Table 6: Vietnamese embedding models comparison on retrieval task, measure by accuracy. **Bold** text expresses the highest score, Underline highlight the second highest score.

Despite having high performance on reranking tasks, the performance on the retrieval task of sup-SimCSE-VietNameese-phobert-base model is significantly lower compared to other Vietnamese embedding models. Meanwhile, vietnamese-bi-encoder can achieve the highest score when the number of retrieved items  $k$  is set

<sup>4</sup><https://huggingface.co/VoVanPhuc/sup-SimCSE-VietNameese-phobert-base>

<sup>5</sup><https://huggingface.co/bkai-foundation-models/vietnamese-bi-encoder>

<sup>6</sup><https://huggingface.co/keepitreal/vietnamese-sbert>

to 10 or 20, and is the second highest when  $k = 5$ . Our model, on the other hand, gets the highest score as  $k = 5$  and comes in second place as the number of retrieved items  $k$  increases to 10 and 20.

From the results on retrieval and reranking tasks, sup-SimCSE-VietNameese-phobert-base presents a strong ability in reranking tasks, which contain a small number of text. However, in retrieval tasks with a large amount of text, vietnamese-bi-encoder tend to have better performance than different embedding models. Furthermore, our model, with just over 20 million parameters, is on par with three existing Vietnamese embedding language models with larger sizes.

### 4.3 Affect of temperature on performance

This experiment explores the different values of temperate ( $\tau = 0.1, 0.4, 0.7$ ) in the InfoNCE loss function and our loss function. The result of this experiment is visualized in Figure 1.

From Figure 1, the performance of embedding models on reranking tasks decreases as the temperature  $\tau$  increases. This phenomenon happens for both training objectives (InfoNCE loss and our loss function) as well as for both training methods (curated hard-negative and in-batch negative). Furthermore, the performance of models on retrieval tasks significantly decreases when the temperature increases from 0.1 to 0.4. However, when the temperature increases from 0.4 to 0.7, different behaviors are recorded for different combinations of training objectives and training methods. For models trained on curated hard-negative with the InfoNCE loss function and models trained on in-batch negative with our loss function, their performance on retrieval tasks slightly decreased. Meanwhile, the performance of the model trained on the in-batch negative with the InfoNCE loss will increase. Finally, the model trained on curated hard-negative with our loss function remains consistent performance when temperature increases from 0.4 to 0.7.

It is also important to notice that our loss function raises better performance on both retrieval and reranking tasks with different temperatures, except for the retrieval task with in-batch negative training when the InfoNCE loss has better performance with  $\tau = 0.7$ . Furthermore, this experiment shows that the temperature should be low for text embedding models to perform well on retrieval and reranking.




Figure 1: The impact of temperature  $\tau$  on the model’s performance on two tasks: retrieval and reranking, along with two training methods: in-batch negative and curated hard-negative.

## 5 Conclusion

This work constructs the Vietnamese Context Search benchmark to evaluate Vietnamese embedding language models on retrieval and reranking tasks, with three validation datasets (ViMedRetrieve, ViRerank, and ViGLUE-R). Moreover, this work presents a new training objective function, which performs better than the InfoNCE loss function in reranking and retrieval tasks. Lastly, we evaluate the performance of some Vietnamese embedding language models on our benchmark and experiment to study the effect of temperature  $\tau$  on the performance of embedding models with different training methods.

## 6 Limitation and Future works

One limitation of this work is the difficulty of the ViMedRetrieve dataset, which makes the results of many Vietnamese embedding language models extremely low. Moreover, the evaluation score of ViMedRetrieve is conducted based on the accuracy of different numbers of retrieved documents. Despite providing more detail about the model’s performance, this metric poorly summarizes the

model’s overall performance on the retrieval task. Future works aim to add a new metric to evaluate the overall model’s performance on this task.

## 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, Hong Kong, China. Association for Computational Linguistics.
- Giambattista Amati. 2009. *BM25*, pages 257–260. Springer US, Boston, MA.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.



- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Tri Dao and Albert Gu. 2024. Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality. In *International Conference on Machine Learning (ICML)*.
- 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.
- Nguyen Quang Duc, Le Hai Son, Nguyen Duc Nhan, Nguyen Dich Nhat Minh, Le Thanh Huong, and Dinh Viet Sang. 2024. Towards comprehensive vietnamese retrieval-augmented generation and large language models. *arXiv preprint arXiv:2403.01616*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#). *CoRR*, abs/2306.11644.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. DeBERTa: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Vong Anh Ho, Duong Huynh-Cong Nguyen, Danh Hoang Nguyen, Linh Thi-Van Pham, Duc-Vu Nguyen, Kiet Van Nguyen, and Ngan Luu-Thuy Nguyen. 2020. Emotion recognition for vietnamese social media text. In *Computational Linguistics*, pages 319–333, Singapore. Springer Singapore.
- Tin Van Huynh, Kiet Van Nguyen, and Ngan Luu-Thuy Nguyen. 2022. [ViNLI: A Vietnamese corpus for studies on open-domain natural language inference](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3858–3872, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. 2024. [Nv-embed: Improved techniques for training llms as generalist embedding models](#). *ArXiv*, abs/2405.17428.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. [Towards general text embeddings with multi-stage contrastive learning](#). *ArXiv*, abs/2308.03281.
- 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*.
- Luong Luc Phan, Phuc Huynh Pham, Kim Thi-Thanh Nguyen, Sieu Khai Huynh, Tham Thi Nguyen, Luan Thanh Nguyen, Tin Van Huynh, and Kiet Van Nguyen. 2021. Sa2sl: From aspect-based sentiment analysis to social listening system for business intelligence. In *Knowledge Science, Engineering and Management*, pages 647–658, Cham. Springer International Publishing.

- Kun Luo, Zheng Liu, Shitao Xiao, and Kang Liu. 2024. [Bge landmark embedding: A chunking-free embedding method for retrieval augmented long-context large language models](#). *ArXiv*, abs/2402.11573.
- Niklas Muennighoff, Hongjin Su, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, and Douwe Kiela. 2024. [Generative representational instruction tuning](#). *Preprint*, arXiv:2402.09906.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- NghiemAbe. 2024. [Vietnamese qqp triplet](#).
- 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, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Kiet Van Nguyen, Duc-Vu Nguyen, Phu X. V. Nguyen, Tham T. H. Truong, and Ngan Luu-Thuy Nguyen. 2018a. [Uit-vsfc: Vietnamese students’ feedback corpus for sentiment analysis](#). *2018 10th International Conference on Knowledge and Systems Engineering (KSE)*, pages 19–24.
- Kiet Van Nguyen, Vu Duc Nguyen, Phu X. V. Nguyen, Tham T. H. Truong, and Ngan Luu-Thuy Nguyen. 2018b. [Uit-vsfc: Vietnamese students’ feedback corpus for sentiment analysis](#). In *2018 10th International Conference on Knowledge and Systems Engineering (KSE)*, pages 19–24.
- Xuan-Phi Nguyen, Wenxuan Zhang, Xin Li, Mahani Aljunied, Qingyu Tan, Liying Cheng, Guanzheng Chen, Yue Deng, Sen Yang, Chaoqun Liu, Hang Zhang, and Li Bing. 2023. [Seallms - large language models for southeast asia](#). *ArXiv*, abs/2312.00738.
- José R. Pérez-Agüera, Javier Arroyo, Jane Greenberg, Joaquín Pérez-Iglesias, and Víctor Fresno-Fernández. 2010. [Using bm25f for semantic search](#). In *SEM-SEARCH ’10*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Claude Sammut and Geoffrey I. Webb, editors. 2010. [TF-IDF](#), pages 986–987. Springer US, Boston, MA.
- Hongjin Su, Howard Yen, Mengzhou Xia, Weijia Shi, Niklas Muennighoff, Han-yu Wang, Haisu Liu, Quan Shi, Zachary S Siegel, Michael Tang, Ruoxi Sun, Jinsung Yoon, Serkan O Arik, Danqi Chen, and Tao Yu. 2024. [Bright: A realistic and challenging benchmark for reasoning-intensive retrieval](#).
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. [Gemma: Open models based on gemini research and technology](#). *arXiv preprint arXiv:2403.08295*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Minh-Nam Tran, Phu-Vinh Nguyen, Long Nguyen, and Dien Dinh. 2024a. [ViGLUE: A Vietnamese general language understanding benchmark and analysis of Vietnamese language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4174–4189, Mexico City, Mexico. Association for Computational Linguistics.
- Minh-Nam Tran, Phu-Vinh Nguyen, Long Nguyen, and Dien Dinh. 2024b. [ViMedAQA: A Vietnamese medical abstractive question-answering dataset and findings of large language model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 356–364, Bangkok, Thailand. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

## A Appendix

### A.1 Hardware Resources

This research uses the free NVIDIA Tesla P100 PCIe 16 GB 824 provided by Kaggle.

### A.2 Hyperparameters

The hyper-parameters used in the training process are reported in Table 7.

### A.3 Running Time

The running time for the in-batch negative training is 4 hours and 45 minutes while training with the curated hard-negative training method requires 7 hours and 30 minutes.

<b>Hyper-parameter</b>	<b>Value</b>
Batch size	32
Learning rate	5e-5
Max sequence length	224
Epochs	3
Temperature	{0.1, 0.4, 0.7}

Table 7: Hyper-parameters used in in-batch negative and curated hard-negative training

#### **A.4 Datasets and models**

The datasets and models used in this paper are publicly available on Hugging Face<sup>7</sup> and GitHub<sup>8</sup>.

---

<sup>7</sup><https://huggingface.co/ContextSearchLM>

<sup>8</sup><https://github.com/phuvinhnguyen/VietnameseTextSearch>