MinLegal at VLSP2025-LegalSLM: A Two-Stage LoRA-to-Full Fine-tuning Approach for Vietnamese Legal Small Language Models

Ngo Dinh Luan, Kiet Nguyen Van
University of Information Technology, Ho Chi Minh city, Vietnam
Vietnam National University Ho Chi Minh city, Vietnam
luannd.19@grad.uit.edu.vn
kietnv@uit.edu.vn

Abstract

Developing effective small-to-medium sized language models (≤ 4B parameters) for Vietnamese legal text processing is challenging due to limited data and computational resources. To address this, we apply a two-stage training methodology that combines task-specific LoRA adaptation with subsequent full fine-tuning. We first augmented a seed dataset of 440 public samples into a high-quality corpus of over 3,400 examples across three legal tasks using a sampling-based data generation technique. Our approach begins with stage-one, where we fine-tune task-specific LoRA adapters and merge them into a single model. This is followed by a stage-two full parameter fine-tuning on the combined dataset, designed to maximize the learning efficiency of parameter-constrained models. Experimental results demonstrate that our methodology enables small models to achieve superior performance over conventional fine-tuning approaches. Our system secured second place in the VLSP 2025-LegalSLM challenge with an average score of 0.7947.

1 Introduction

Legal text processing has emerged as one of the most challenging applications in Natural Language Processing, demanding sophisticated understanding of intricate legal concepts, procedural knowledge, and domain-specific reasoning capabilities (Muresan et al., 2022) (Ariai et al., 2024). While large language models have demonstrated remarkable performance on legal tasks in well-resourced languages such as English, developing effective solutions for Vietnamese legal text processing remains significantly constrained by both linguistic complexity and computational limitations. Vietnamese legal documents present unique challenges including word segmentation ambiguities, complex sentence structures, and intricate logical relationships between legal concepts (Son et al., 2024) (Le et al., 2025). The VLSP2025 Challenge on Vietnamese Legal Small Language Models addresses these needs by focusing on models with ≤4B parameters across three fundamental evaluation tasks: legal citation usefulness classification, multiple-choice legal question answering, and free-text legal reasoning through syllogistic arguments. Figure 1 presents examples of these three legal reasoning tasks that demonstrate the complexity and diversity of the evaluation framework.

The development of Vietnamese legal NLP systems faces several critical challenges. First, the scarcity of high-quality annotated Vietnamese legal datasets significantly limits model development, with available resources remaining substantially smaller compared to English counterparts (**?**). Second, computational constraints in Vietnamese organizations necessitate efficient models that can operate within practical resource limitations while maintaining competitive performance. Third, existing parameter-efficient fine-tuning techniques such as Low-Rank Adaptation (LoRA) have not fully exploited the potential of combining multiple training strategies to maximize small model capabilities in specialized domains (Hu et al., 2022) (Dettmers et al., 2023). Current approaches apply uniform training methodologies across different task types, failing to account for the distinct characteristics and requirements of various legal reasoning scenarios.

This paper presents a two-stage training method that combines task-specific LoRA fine-tuning with full parameter fine-tuning for Vietnamese legal language models. We first train separate LoRA adapters for each task, then merge them and perform additional full fine-tuning on the combined dataset. To address the limited training data, we generate synthetic data using Gemini 2.5 Flash API, expanding from approximately 150 public test samples to create over 3400 through topic-guided

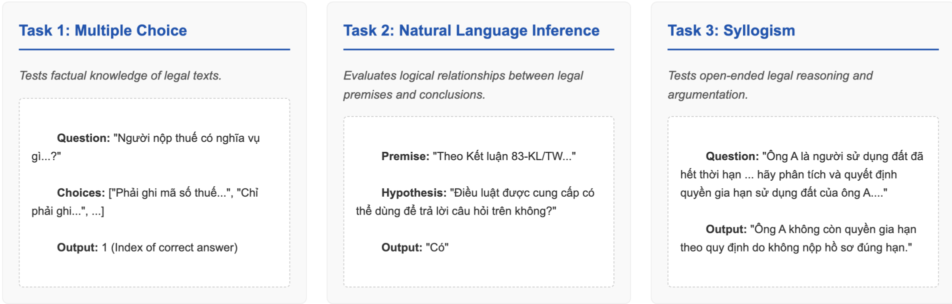

Figure 1: An overview of the three legal reasoning tasks in the VLSP2025 LegalSLM dataset.

search and manual filtering.

The main contributions of this work include:

- **Two-stage training methodology**: We combine task-specific LoRA fine-tuning followed by full parameter fine-tuning, achieving better performance than conventional training approaches.
- **Synthetic data generation**: We develop a systematic approach to generate additional training data from limited public test samples, addressing data scarcity in Vietnamese legal NLP.

The organization of the paper is as follows: In Section 2, reviews related work in parameter-efficient fine-tuning and data augmentation techniques. Section 3, we explain our two-stage training methodology and synthetic data generation approach. Section 4 presents our experimental setup and results analysis. Section 5 is the conclusion and future work.

2 Related Works

2.1 Parameter-Efficient Fine-Tuning

Parameter-efficient fine-tuning has emerged as a crucial technique for adapting large language models to specialized domains while maintaining computational efficiency. Low-Rank Adaptation (LoRA) (Hu et al., 2022) introduced the foundational approach of decomposing weight updates into low-rank matrices ($\Delta W = BA$), enabling efficient adaptation with minimal trainable parameters. Recent advances include QLoRA (Dettmers et al., 2023) which enables fine-tuning of large models through 4-bit quantization, DoRA (Liu et al., 2024) which decomposes weights into magnitude and directional components achieving +3.7 points improvement on reasoning tasks, and AdaLoRA (Zhang et al., 2023) which uses dynamic parameter allocation based on layer importance. Multi-stage training approaches have shown promise for domain specialization, with recent work demonstrating that parameter-efficient methods can achieve performance matching full fine-tuning while using only 0.1% of model parameters in legal domain adaptation (Li et al., 2022).

2.2 Vietnamese Legal NLP

Vietnamese legal text processing has gained significant attention with several foundational datasets addressing the unique challenges of Vietnamese legal language. The VLQA dataset (Nguyen et al., 2025) represents the largest expert-annotated Vietnamese legal QA resource with 3,129 questions paired with 59,636 legal articles across 27 domains. Vietnamese Legal Document Retrieval (Pham Tien et al., 2024) introduced approaches leveraging LLMs to generate over 500,000 synthetic queries for Vietnamese legal passages, while ViBidLQA (Ha et al., 2024) focused on bidding law through LLM synthesis. Technical approaches have converged on hybrid methodologies combining BGE-M3 dense retrieval, BM25 and semantic search, addressing challenges of Vietnamese legal text complexity (Son et al., 2024). Global advances include LegalBench (Guha et al., 2023) with 162 legal reasoning tasks and SaulLM-7B (Colombo et al., 2024), the first 7B parameter model designed for legal comprehension.

2.3 Data Augmentation for Legal NLP

Data augmentation for legal NLP has evolved beyond simple paraphrasing to sophisticated domain-aware approaches. DALE (Ghosh et al., 2023) represents the flagship legal augmentation framework, addressing unique challenges of legal language through encoder-decoder models with selective masking, achieving 1%-50% absolute improvements across 13 datasets. Large language model-based generation has become dominant, with Self-Instruct (Wang et al., 2022) demonstrating generation of 100K examples from 175 seed examples, while recent research emphasizes that high-quality, smaller datasets outperform large unvalidated datasets. Sampling-based generation (Kudalkar et al., 2024), back-translation methods (Köksal et al., 2023), and transformation-based approaches (Gandhi et al., 2024) provide different strategies for synthetic data creation.

The convergence of parameter-efficient fine-tuning methods with sophisticated data augmentation techniques provides the foundation for advanced training approaches in Vietnamese legal NLP. The combination of task-specific adaptation through LoRA-based methods and synthetic data generation addresses both the computational constraints and data scarcity challenges characteristic of Vietnamese legal AI development.

3 Methodology

In this section, we describe our approach to address the Vietnamese legal reasoning tasks, including the following components: 1) Synthetic Data Generation using Sampling-based Approach; 2) Two-Stage Training Architecture; 3) Model Configuration and Training Setup.

The primary challenge in this competition stems from the limited data availability for each task, with only approximately 150 samples per task in the public test set. To address this constraint, we developed a systematic approach combining synthetic data generation with an efficient two-stage training methodology designed to maximize learning efficiency of parameter-constrained models.

3.1 Synthetic Data Generation

To overcome the data scarcity challenge, we employed a sampling-based generation approach leveraging the Gemini 2.5 Flash API with Grounding with Google Search[1] functionality. Our methodology follows these systematic steps:

Few-shot Example Selection: For each generated sample, we randomly selected 5 samples from the 150 available public test samples for the corresponding task. This random sampling ensures diverse coverage of legal domains and question patterns while maintaining reproducibility through code-based selection rather than manual curation.

Legal Domain Identification: Using the randomly selected few-shot examples, the model proposes specific legal domains relevant to Vietnamese law. For instance, the multiple-choice task might suggest "traffic regulations for individuals under 18 years" as a focus area based on the patterns observed in the few-shot examples.

Grounding-based Content Retrieval: We utilized "Grounding with Google Search" to automatically search for legal articles, regulations, and case studies related to the proposed legal domain. This approach ensures that generated content is grounded in actual Vietnamese legal documents and current legal practices.

Task-specific Data Generation: Based on the retrieved legal context, the model generates training data in the required format:

For Multiple Choice Questions: Generate questions with four answer choices based on legal articles, including correct answers.

For Natural Language Inference: Create legal document excerpts paired with specific legal questions, where the task asks "Can the provided legal provision be used to answer the above question?" requiring yes/no classification.

For Syllogism Questions: Generate complex legal scenarios that require multi-step reasoning and logical analysis to reach well-justified conclusions.

Quality Control: All generated samples underwent manual filtering and validation to ensure legal accuracy, linguistic quality, and alignment with Vietnamese legal standards. This step was crucial for maintaining dataset integrity and preventing propagation of legal inaccuracies.

3.2 Two-Stage Training Architecture

Our experimental analysis revealed that a two-stage training approach significantly outperforms conventional fine-tuning methods. The methodology addresses the challenge of uneven data distribution across tasks while maximizing parameter efficiency.

Stage One - Task-Specific LoRA Training: We train separate LoRA adapters for each of the three legal reasoning tasks using their respective synthetic datasets:

Independent Training: Each LoRA adapter is trained separately on its corresponding task dataset (Multiple Choice, NLI, or Syllogism).

Optimal Checkpoint Selection: Due to uneven data distribution across tasks, each LoRA adapter reaches optimal performance at different training steps. We select the best checkpoint for each

[1] https://ai.google.dev/gemini-api/docs/google-search

adapter based on validation performance on the development set.

LoRA Configuration: Rank=16, Alpha=32, Dropout=0.1, targeting approximately 0.1% of total model parameters for efficient adaptation.

Adapter Merging: The three specialized LoRA adapters are merged using linear interpolation to create a unified model incorporating all task-specific knowledge.

Stage Two - Full Parameter Fine-tuning: Starting from the merged LoRA model, we perform full parameter fine-tuning on the combined dataset:

Combined Dataset Training: Full fine-tuning on all three task datasets simultaneously for 2 additional epochs.

Performance Optimization: This stage achieves lower training loss and higher evaluation scores compared to alternative approaches such as direct full fine-tuning or LoRA-only training.

3.3 Experimental Setup

LoRA Configuration: We applied LoRA to all linear modules in the transformer architecture with rank=16, alpha=32, and dropout=0.1. This configuration provides an optimal balance between parameter efficiency and adaptation capability for legal domain specialization.

Training Hyperparameters: The learning rate was set to 2e-4 for LoRA training in Stage 1 and 5e-5 for full parameter fine-tuning in Stage 2. We used a context length of 4096 tokens to accommodate the longer legal documents and complex reasoning scenarios present in the dataset.

Training Procedure: In **Stage 1 (LoRA fine-tuning)**, we trained the model with LoRA adapters and selected the best-performing checkpoint based on performance on the public test set. In **Stage 2 (full fine-tuning)**, we trained all model parameters for **2 epochs**, and the training data included both the original training set and the public test set to enhance robustness and generalization in legal reasoning.

Base Model Selection: We experimented with multiple 4B parameter models, ultimately selecting Qwen3-4B-Base due to its superior performance on Vietnamese text processing tasks and strong foundational capabilities for legal reasoning.

Evaluation Framework: We evaluated each checkpoint on the public test set using different methodologies for each task type:

For Multiple Choice and NLI tasks: We utilized the LM-Evaluation-Harness framework[2] which provides standardized evaluation protocols with few-shot prompting and exact match accuracy scoring.

For Syllogism Questions: We employed Qwen3-32B-AWQ as an LLM-as-a-Judge evaluator to assess the quality of generated reasoning against ground truth answers. This approach allows for nuanced evaluation of complex legal reasoning that goes beyond simple text matching.

Hardware and Infrastructure: All experiments were conducted on 2xA30 GPUs with 48GB memory total, using mixed precision bf16 training and Deepspeed ZeRO-3 offload for memory optimization while maintaining model performance.

4 Results and Performance Analysis

4.1 Dataset Overview

Our synthetic data generation approach successfully addressed the data scarcity challenge inherent in the VLSP 2025 LegalSLM dataset. Starting from approximately 150 samples per task in the public test set, we expanded the training corpus significantly across all three legal reasoning tasks.[3]

	Original	**Generated**
Multichoice Questions	146	803
NLI Questions	150	745
Syllogism Questions	144	1989
Total Samples	440	3537

Table 1: Dataset expansion through synthetic data generation using Gemini 2.5 Flash API with Grounding search

After manual review, the synthetic data generation process achieved substantial expansion ratios: 5.5x for Multiple Choice Questions, 5.0x for NLI, and 13.2x for Syllogism Questions. The higher generation ratio for Syllogism tasks reflects both the complexity of legal reasoning scenarios and the model's capability to generate diverse legal case studies through grounded search.

4.2 Training Methodology Comparison

Our experimental results demonstrate the effectiveness of the two-stage training approach compared

[2]https://github.com/EleutherAI/lm-evaluation-harness

[3]The official public test data used in this work is available at: `https://huggingface.co/datasets/luanngo/Vietnamese-Legal-Chat-Dataset`. All data are formatted following the ShareGPT conversation format to ensure compatibility with instruction-tuned LLMs.

Training Method	MC (%)	NLI (%)	Syllogism (%)	Average (%)
Two-stage Training (Ours)	**95.89**	**97.33**	62.50	**85.24**
Full Fine-tuning	94.52	96.00	**63.19**	84.57
LoRA Training Only	88.35	93.33	56.94	79.54
QLoRA Training Only	89.04	92.66	55.56	79.08

Table 2: Performance comparison across different training methodologies on VLSP 2025 LegalSLM public test set

Base Model	MC (%)	NLI (%)	Syllogism (%)	Average (%)
Qwen3-4B-Base	95.89	**97.33**	**62.50**	**85.24**
Qwen3-4B-Legal-Pretrain	**97.26**	96.00	61.11	84.79
Qwen3-4B-Instruct-2507	92.47	92.00	58.33	80.93
Qwen3-4B	89.04	92.67	55.56	79.09
Gemma-3-4B-IT	90.41	89.33	54.17	77.97
Qwen3-4B-Thinking-2507	85.62	90.67	51.39	75.89
Hunyuan-4B-Pretrain	80.82	86.00	47.22	71.35
Hunyuan-4B-Instruct	76.71	82.67	54.17	71.18
Gemma-3-4B-PT	78.77	84.00	45.83	69.53

Table 3: Performance comparison across different 4B parameter base models using two-stage training methodology

to conventional fine-tuning methods across all three legal reasoning tasks.

The results reveal several key insights about training methodologies for Vietnamese legal NLP. Our two-stage approach achieves the highest average performance at 85.24%, outperforming full fine-tuning by 0.67 percentage points and significantly surpassing parameter-efficient methods by over 5 percentage points. The approach proves particularly effective for Multiple Choice and NLI tasks, achieving over 95% accuracy on both, while maintaining competitive performance on the more challenging Syllogism reasoning task.

4.3 Base Model Analysis

We conducted comprehensive experiments across multiple 4B parameter base models to identify optimal foundations for Vietnamese legal reasoning tasks.

The base model analysis reveals significant performance variations across different foundation models. Qwen3-4B-Base achieves optimal overall performance, followed closely by the legal-pretrained variant. Interestingly, the legal-pretrained model shows superior performance on Multiple Choice questions but slightly lower performance on NLI and Syllogism tasks, suggesting that specialized legal pretraining benefits factual legal knowledge but may not universally improve complex reasoning capabilities.

The Qwen model family consistently outperforms alternatives, with Gemma and Hunyuan models showing progressively lower performance. This pattern indicates that model architecture and pretraining data quality significantly impact Vietnamese legal reasoning capabilities, with Qwen's multilingual pretraining providing advantages for Vietnamese legal text processing.

4.4 Task-Specific Performance Analysis

The performance distribution across tasks reveals fundamental differences in Vietnamese legal reasoning complexity. Multiple Choice and NLI tasks achieve consistently high performance (>95% for top methods), while Syllogism reasoning presents substantial challenges across all approaches. This disparity reflects the transition from factual legal knowledge assessment to complex legal reasoning and argumentation skills.

The two-stage training approach shows particular strength in bridging this gap, maintaining competitive performance on Syllogism tasks while achieving near-perfect accuracy on classification tasks. This balance demonstrates the methodology's effectiveness in handling the diverse requirements of Vietnamese legal AI applications.

The synthetic data generation approach proved crucial for enabling effective training despite original data limitations. Quality control measures ensured legal accuracy while providing sufficient diversity for robust model development. This approach enabled our team to achieve second place in the VLSP 2025 LegalSLM challenge with an average score of 0.7947 on the private test set.

5 Conclusion and Future Work

In this paper, we presented a two-stage training methodology for Vietnamese legal small language models that effectively addresses the challenges of data scarcity and computational constraints. Our approach combines sampling-based synthetic data generation using Gemini 2.5 Flash API with Grounding search and a novel two-stage training strategy that sequentially applies task-specific LoRA adaptation followed by full parameter fine-tuning. We demonstrated that this methodology significantly outperforms conventional fine-tuning approaches, achieving 85.24% average accuracy across three legal reasoning tasks and securing second place in the VLSP 2025 LegalSLM challenge with a final score of 0.7947.

In future work, we plan to explore more sophisticated synthetic data generation techniques, including multi-agent approaches for generating more diverse legal scenarios and incorporating domain expert feedback loops for improved quality control. Additionally, we will investigate the application of our two-stage methodology to larger parameter models when computational resources permit, and extend our approach to other Vietnamese legal NLP tasks such as legal document summarization and contract analysis. We also aim to develop more robust evaluation frameworks specifically designed for Vietnamese legal reasoning to better assess model capabilities in real-world legal applications.

References

Farid Ariai, Joel Mackenzie, and Gianluca Demartini. 2024. Natural language processing for the legal domain: A survey of tasks, datasets, models, and challenges. *arXiv preprint arXiv:2410.21306*.

Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, Andre FT Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, et al. 2024. Saullm-7b: A pioneering large language model for law. *arXiv preprint arXiv:2403.03883*.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115.

Saumya Gandhi, Ritu Gala, Vijay Viswanathan, Tongshuang Wu, and Graham Neubig. 2024. Better synthetic data by retrieving and transforming existing datasets. *arXiv preprint arXiv:2404.14361*.

Sreyan Ghosh, Chandra Kiran Evuru, Sonal Kumar, S Ramaneswaran, S Sakshi, Utkarsh Tyagi, and Dinesh Manocha. 2023. Dale: Generative data augmentation for low-resource legal nlp. *arXiv preprint arXiv:2310.15799*.

Neel Guha, Julian Nyarko, Daniel Ho, Christopher Ré, Adam Chilton, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, et al. 2023. Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models. *Advances in neural information processing systems*, 36:44123–44279.

Nguyen Thu Ha, Truong-Phuc Nguyen, Khang T Trung, Huu-Loi Le, Le Thi Viet Huong, Chi Thanh Nguyen, and Minh-Tien Nguyen. 2024. Vietnamese legal question answering: An experimental study. In *2024 16th International Conference on Knowledge and System Engineering (KSE)*, pages 440–446. IEEE.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.

Abdullatif Köksal, Timo Schick, Anna Korhonen, and Hinrich Schütze. 2023. Longform: Effective instruction tuning with reverse instructions. *arXiv preprint arXiv:2304.08460*.

Vidisha Kudalkar, Navid Hashemi, Shilpa Mukhopadhyay, Swapnil Mallick, Christof Budnik, Parinitha Nagaraja, and Jyotirmoy V Deshmukh. 2024. Sampling-based and gradient-based efficient scenario generation. In *International Conference on Runtime Verification*, pages 70–88. Springer.

Van-Hoang Le, Duc-Vu Nguyen, Kiet Van Nguyen, and Ngan Luu-Thuy Nguyen. 2025. Optimizing legal document retrieval in vietnamese with semi-hard negative mining. *arXiv preprint arXiv:2507.14619*.

Jonathan Li, Rohan Bhambhoria, and Xiaodan Zhu. 2022. Parameter-efficient legal domain adaptation. *arXiv preprint arXiv:2210.13712*.

Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. 2024. Dora: Weight-decomposed low-rank adaptation, 2024. *URL https://arxiv. org/abs/2402.09353*.

Smaranda Muresan, Preslav Nakov, and Aline Villavicencio. 2022. Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: Long papers). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Tan-Minh Nguyen, Hoang-Trung Nguyen, Trong-Khoi Dao, Xuan-Hieu Phan, Ha-Thanh Nguyen, and Thi-Hai-Yen Vuong. 2025. Vlqa: The first comprehensive, large, and high-quality vietnamese dataset for legal question answering. *arXiv preprint arXiv:2507.19995*.

Son Pham Tien, Hieu Nguyen Doan, An Nguyen Dai, and Sang Dinh Viet. 2024. Improving vietnamese legal document retrieval using synthetic data. *arXiv e-prints*, pages arXiv–2412.

Pham Tien Son, Nguyen Doan Hieu, Nguyen Dai An, and Dinh Viet Sang. 2024. Improving vietnamese legal document retrieval using synthetic data. In *International Symposium on Information and Communication Technology*, pages 378–393. Springer.

Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*.

Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. 2023. Adalora: Adaptive budget allocation for parameter-efficient finetuning. *arXiv preprint arXiv:2303.10512*.