

Enhancing Legal Expertise in Large Language Models through Composite Model Integration: The Development and Evaluation of Law-Neo

Zhihao Liu¹, Yanzhen Zhu¹, Mengyuan Lu¹,

¹Shandong University of Finance and Economics,

Correspondence: 202116440220@mail.sdufe.edu.cn

Abstract

Although large language models (LLMs) like ChatGPT (OpenAI et al., 2024) have demonstrated considerable capabilities in general domains, they often lack proficiency in specialized fields. Enhancing a model’s performance in a specific domain, such as law, while maintaining low costs, has been a significant challenge. Existing methods, such as fine-tuning or building mixture of experts (MoE) models, often struggle to balance model parameters, training costs, and domain-specific performance. Inspired by composition to augment language models (Bansal et al., 2024), we have developed Law-Neo, a novel model designed to enhance legal LLMs. This model significantly improves the model’s legal domain expertise at minimal training costs, while retaining the logical capabilities of a large-scale anchor model. Our Law-Neo model outperformed other models in comprehensive experiments on multiple legal task benchmarks, demonstrating the effectiveness of this approach.

1 Introduction

Large Language Models (LLMs) have shown significant capabilities, including commonsense and factual reasoning, world knowledge, and language generation. These abilities have been validated across various scientific fields such as finance, biochemistry, and medicine (Chen et al., 2023; Ren et al., 2023; Ferruz et al., 2022; Thirunavukarasu et al., 2023; Fan et al., 2024). However, the training cost escalates as the number of parameters in LLM increases when enhancing model’s domain-specific capabilities. This cost barrier is a significant challenge in developing domain-specific LLMs, such as those for the legal field.

To address these challenges, we propose the development of a comprehensive LLM-based legal assistance system.

Main Contributions In this paper, we present Law-Neo, a legal domain model trained at a rel-

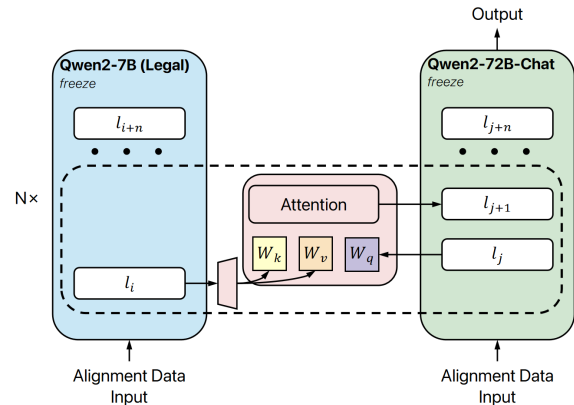


Figure 1: Architecture of Law-Neo. The Qwen2-72B-Chat model is enhanced with legal domain knowledge from Qwen2-7B (Legal) by sharing layer parameters. Both models remain unchanged during the composite training process, with a few additional parameters learned over their layer representations.

atively low cost while achieving satisfactory performance. This was accomplished by augmenting the Qwen2-72B-Chat model with a legal domain-specific model based on Qwen2-7B. We describe the training process of this composite model, which integrates multiple sub-models, each requiring different capability enhancements.

Our results on the Unified Qualification Exam for Legal Professionals and various downstream task benchmarks indicate that Law-Neo outperforms existing methods in several aspects. Qualitative analysis demonstrates that Law-Neo surpasses GPT-4 by 23 points in scoring on the Unified Qualification Exam for Legal Professionals, showcasing its robust legal consultation capabilities.

The data and training code used in this work are publicly available at <https://github.com/SkyFlap/Law-Neo>.

2 Related Work

Since the development of BERT (Devlin et al., 2019), significant efforts have been made to create

language models (LMs) specifically tailored for the legal domain. Initially, these models were small and followed the paradigm of pre-training followed by downstream task fine-tuning. Recent advancements have seen an increase in model size and the introduction of instruction fine-tuning, with evaluations extending across a broader spectrum of legal tasks. Most existing legal LLMs are text-based, with a focus on Chinese, English, or multi-language support (Chen et al., 2024).

2.1 Pre-Trained and Fine-Tuned PLMs

LegalBERT (Chalkidis, 2023) represents an early endeavor to develop pre-trained language models (PLMs) for legal tasks such as legal text classification (LTC). This model underwent additional pre-training on legal corpora and was subsequently fine-tuned with task-specific data. Lawformer (Xiao et al., 2021) is a Transformer-based model specifically pre-trained to manage lengthy legal texts, and it has been employed for tasks such as legal judgment prediction (LJP), legal reading comprehension (LRC), and legal question answering (LQA).

2.2 Pre-Trained and Fine-Tuned LLMs

In the realm of large language models (LLMs), models are pre-trained and fine-tuned specifically for legal tasks or datasets. These legal-specific LLMs often incorporate external knowledge bases and undergo extensive initial training to handle a wide range of legal data. Notable models include LexiLaw (Haitao, 2024), a fine-tuned Chinese legal model based on ChatGLM-6B (Zeng et al., 2024a), and Fuzi.mingcha (Deng et al., 2023), which is also based on ChatGLM-6B and fine-tuned on the CAIL2018 dataset (Xiao et al., 2018). Wisdom-Interrogatory (Wu et al.) builds upon Baichuan2-7B (Baichuan, 2023), and LawGPT-7B-beta1.0 (Nguyen, 2023) is pre-trained on 500,000 Chinese judgment documents, based on Chinese-LLaMA-7B (Cui et al., 2023). Additionally, HanFei (He et al., 2023) is a fully pre-trained and fine-tuned LLM with 7 billion parameters.

Further advancements in large-scale LLMs include LawyerLLaM (Huang et al., 2023), based on Chinese-LLaMA-13B (Cui et al., 2023) and fine-tuned with general and legal instructions, as well as ChatLaw-13B (Cui et al., 2024b), fine-tuned on Ziya-LLaMA-13B-v1 (Wang et al., 2022), and ChatLaw-33B (Cui et al., 2024b), fine-tuned on Anima-33B (Ogavinee and et al., 2022). Models in other languages have also emerged, such

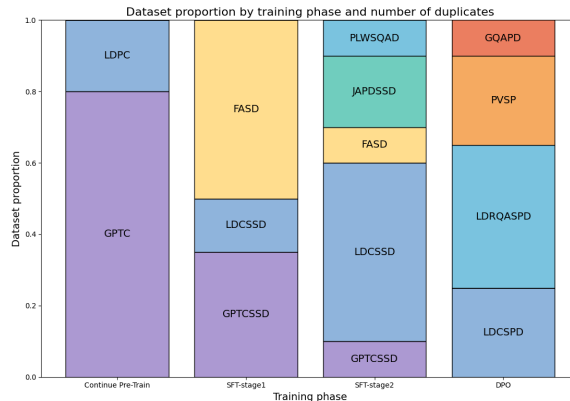


Figure 2: Proportions of Dataset Types Used at Each Step in Training Qwen2-7B (Legal). General Pre-training Corpus (GPTC), Legal Domain Pre-training Corpus (LDPTC), Foundational Abilities Supervised Data (FASD), Legal Domain Corpus Synthetic Supervised Data (LDCSSD), General Pre-training Corpus Synthetic Supervised Data (GPTCSSD), Public Legal Website Search and QA Data (PLWSQAD), Legal Domain Real QA Synthetic Preference Data (LDRQASPD), Public Video Case Synthetic Preference Data (PVSP), Legal Domain Corpus Synthetic Preference Data (LDCSPD), General QA Preference Data (GQAPD), and Judgment, Arbitration, and Prosecutorial Documents Synthetic Supervised Data (JAPDSSD). Homogeneous variant corpora were used at different stages to prevent catastrophic forgetting and capability degradation.

as SaulLM-7B (Colombo et al., 2024), based on Mistral-7B (Jiang et al., 2023), and JURU (Junior et al., 2024), the first LLM pre-trained for the Brazilian legal domain.

A recent innovation in this field is the introduction of ChatLaw-4x7B (Cui et al., 2024a), a mixture of experts model (MoE) designed to address hallucinations and insufficient domain expertise in LLMs. However, training MoE architectures presents significant challenges, particularly in balancing the training of expert models and sample load distribution (Zeng et al., 2024b; Pan et al., 2024).

These legal-specific LLMs, typically following an initial pre-training phase, are tailored to specific legal datasets and tasks. This tailoring enhances both the precision and practical applicability of legal NLP technologies.

3 Method

This section focuses on Law-Neo as illustrated in Fig.1. While our model comprises three main components—the **domain knowledge model** (Qwen2-

7B-Legal), the **anchor model** (Qwen2-72B-Chat Yang et al., 2024), and the **parameter-merging block**—this section will specifically discuss the **domain knowledge model** and the **parameter-merging block** in detail.

Our approach assumes: (i) **The model weights are fixed and unmodifiable**, reflecting the high computational cost of training or fine-tuning large LLMs from scratch in production environments. Pre-trained models are treated as fixed assets to ensure cost efficiency and stability. (ii) **We can access model weights, perform forward and backward passes, and retrieve intermediate representations**. This is feasible with many open-source LLMs, allowing us to use their parameters for inference and further training. (iii) **We lack access to the original training data, hyperparameters, or training states**, as open-source LLMs typically do not provide such information.

3.1 Legal Domain Model Qwen2-7B (Legal)

Here, we introduce the base model selection and provide more details about it. We’ve chosen Qwen2-7B-Base, which was released by the Qwen Team (Yang et al., 2024), is selected as the base model. We performed specified **Data Preparation** and **Model Training** upon the base model to better fit legal domain.

In June 2024, the Qwen Team open-sourced their Qwen2 series models. We used Qwen2-7B-Base as the base model. Fig.2 illustrates the processes applied to this base model, which involve two main steps: data preparation and model training.

Data Preparation: Following the data processing pipeline from the MAP-Neo (Zhang et al., 2024) technical report, we filtered legal-related training corpora from publicly available pre-training datasets (FinWeb, Penedo et al., 2024; Matrix, Zhang et al., 2024; etc.). We collected extensive current Chinese laws and regulations, including local and central regulations, multilateral and bilateral treaties involving China, and specific industry norms. The continued pre-training included over 10,000 manually collected books and papers, processed using methods from MAP-Neo. For post-training data preparation, we collected Chinese case law and synthesized supervised data using GLM4 (Zeng et al., 2024a). Inspired by CQIA (Bai et al., 2024), we gathered online case explanation videos, converted their audio to text, and generated preference data using GLM4. For specific details, please refer to the appendix A.

Model Training: We first performed full-parameter continued pre-training on the Qwen2-7B-Base model, utilizing both general pre-training corpora and the collected legal domain corpora. To ensure the LLM makes human-consistent judgments in the legal domain, we conducted supervised fine-tuning (SFT) in two phases. The first phase enhanced the model’s foundational abilities (e.g., code and math skills) using over 2 million instructional data points. The second phase focused on improving the model’s conversational abilities and legal judgment capabilities while retaining the foundational skills acquired in the first phase. We used the prepared legal domain SFT data and collected over 100,000 multi-turn conversation data from real user interactions. We then aligned the LLM using DPO.

3.2 Model Parameter-Merging Block

As illustrated in Fig.1, our approach involves concurrent operations on selected layers from two large language models (LLMs). Specifically, we introduce two sets of additional parameters over these layers: (1) a straightforward set of linear transformations, $f_{\text{proj}}(\cdot)$, which project an i^{th} layer representation from Qwen2-7B (Legal) to the dimensionality of representations from Qwen2-72B; and (2) a series of cross-attention layers, $f_{\text{cross}}(\cdot, \cdot)$, which perform cross-attention between this transformed layer representation and a j^{th} layer representation from Qwen2-72B. The output of the cross-attention is then added as a residual connection to the layer representations of Qwen2-72B. For specific details, please refer to the appendix D.

4 Experiments

We evaluated the performance of LawBench (Fei et al., 2023) and the Unified Qualification Exam for Legal Professionals. Additionally, we conducted benchmark testing for the LJP(Legal Judgment Prediction) task. Our primary focus, lies in LJP tasks utilizing fact-based articles from the CAIL2018 (Xiao et al., 2018) dataset.

4.1 Performance on LawBench

We evaluated our model on LawBench (Fei et al., 2023), a benchmark for the Chinese legal system assessing three cognitive levels: (1) Legal Knowledge Memory, (2) Legal Knowledge Understanding, and (3) Legal Knowledge Application.

As shown in Table 1, our **Law-Neo** model achieves an average score of 64.38, posi-

Model	LawBench Average Score	UQELP Average Score	CAIL2018 F-1
<i>GPT Series</i>			
GPT-3.5 (Brown et al., 2020)	42.15	78	0.29
GPT-4 (OpenAI et al., 2024)	52.35	103	0.52
<i>General LLMs</i>			
Baichuan2-7B (Baichuan, 2023)	38.08	61	–
ChatGLM2-6B (Zeng et al., 2024a)	29.88	34	–
InternLM2-7B (Cai et al., 2024)	43.78	41	–
Qwen2-72B-Chat (Yang et al., 2024)	56.26	–	–
<i>Legal LLMs</i>			
Fuzi-Mingcha-6B (Wang et al., 2022)	32.08	34	0.25
ChatLaw-13B (Cui et al., 2024b)	32.76	–	0.33
Wisdom-Interrogatory-7B (Wu et al.)	31.41	–	0.33
Chatlaw-MoE-4×7B (Cui et al., 2024a)	60.08	115	–
Qwen2-7B-Legal (ours)	51.25	84	0.39
Law-Neo (ours)	64.38	126	0.46

Table 1: Summary of LLM’s performance comparisons on benchmarks: We conducted experiments using three benchmark tests, namely LawBench, Unified Qualification Exam for Legal Professionals (UQELP), and CAIL2018.

tioned competitively between GPT-3.5 and GPT-4, which score 42.15 and 52.35, respectively. While Chatlaw-MoE scores higher at 60.08, our model significantly outperforms Legal LLMs like Fuzi-Mingcha (32.08) and General LLMs like InternLM2-7B (43.78), and also shows a marked improvement over models like Qwen2-7B-Legal (51.25). This demonstrates that our model parameter-merging training strategy is effective in achieving superior performance.

4.2 Performance on Unified Qualification Exam for Legal Professionals

We also assessed our model using China’s Unified Qualification Exam for Legal Professionals, which includes single-choice, multiple-choice, and uncertain-choice questions across various legal fields.

As indicated in Table 1, our **Law-Neo** model achieved an average score of 126, positioning it ahead of most models, including Chatlaw-MoE (115) and GPT-4 (103). Our model surpasses General LLMs such as Baichuan2-7B (61) and ChatGLM2-6B (34), as well as Legal LLMs like Fuzi-Mingcha (34). It also shows a significant improvement over Qwen2-7B-Legal (84), which further emphasizes the strength of our approach.

4.3 Performance on CAIL2018 Task

CAIL2018 (Xiao et al., 2018), a large-scale LJP task, includes over 2.6 million criminal cases from the Supreme People’s Court of China, annotated with applicable law articles, charges, and prison terms.

In Table 1, our **Law-Neo** model achieves an F-1 score of 0.46, showing strong performance. While GPT-4 scores higher at 0.52, our model outperforms GPT-3.5 (0.29) and General LLMs like Qwen2-7B-Chat (0.37). It also surpasses Legal LLMs such as Chatlaw-13B and Wisdom-Interrogatory (both 0.33). These results highlight **Law-Neo**’s robustness in legal language processing and its competitive edge in legal tasks, especially considering that it did not leverage a mixed expert model during training, unlike Chatlaw-MoE.

5 Training Overhead

Our training procedure was conducted on a GPU server equipped with eight 80GB A800 GPUs and an Intel Xeon 8470 processor. The entire training process took approximately 19.24 hours.

The comparative training regimen for ChatLaw was conducted on a GPU server equipped with eight 80GB A100 GPUs and two Intel Xeon 8358P processors. The entire training process was completed in approximately 23.14 hours, which exceeds our training duration by 3.9 hours.

6 Conclusion

In this paper, we introduced Law-Neo, an innovative approach to enhancing large language models (LLMs) for the legal domain by leveraging the concept of composition to augment existing models. Our methodology focused on integrating Qwen2-72B-Chat with a legal domain-specific model based on Qwen2-7B. Our comprehensive experiments, conducted on multiple legal benchmarks including LawBench, the Unified Qualification Exam for Legal Professionals (UQELP), and CAIL2018, demonstrate the efficacy of our approach. The Law-Neo model outperformed several existing models, including general-purpose LLMs and specialized legal LLMs. Our results indicate that integrating models through shared parameters can effectively enhance their specialized knowledge without sacrificing the foundational abilities of the base models.

7 Ethics Statement

The development and application of Law-Neo, an advanced legal large language model (LLM), bring forth significant ethical considerations, particularly regarding bias amplification, interpretability, accountability, and oversight. Law-Neo, like other LLMs, has been trained on extensive legal corpora, including laws, regulations, and judicial decisions. Despite efforts to ensure a balanced dataset, the model may still reflect and perpetuate biases found in the source material. This risk is especially concerning in the legal field, where unbiased and fair decision-making is crucial. Additionally, the complex decision-making process of these models is not easily transparent, making it difficult to scrutinize and understand their outputs fully, which can undermine trust in automated legal tools. Establishing clear guidelines and frameworks for the accountability and oversight of AI systems like Law-Neo is crucial. This includes defining the roles and responsibilities of developers, users, and regulatory bodies in monitoring the deployment and impact of these models. Regular audits, bias assessments, and updates should be conducted to ensure the model remains fair, transparent, and aligned with ethical standards.

References

Yuelin Bai, Xinrun Du, Yiming Liang, Yonggang Jin, Ziqiang Liu, Juntao Zhou, Tianyu Zheng, Xincheng Zhang, Nuo Ma, Zekun Wang, et al. 2024. [Coig-](#)

[cqia: Quality is all you need for chinese instruction fine-tuning](#). *Preprint*, arXiv:2403.18058.

Baichuan. 2023. [Baichuan 2: Open large-scale language models](#). *arXiv preprint arXiv:2309.10305*.

Rachit Bansal, Bidisha Samanta, Siddharth Dalmia, Nitish Gupta, Shikhar Vashishth, Sriram Ganapathy, Abhishek Bapna, Prateek Jain, and Partha Talukdar. 2024. [Llm augmented llms: Expanding capabilities through composition](#). *Preprint*, arXiv:2401.02412.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. 2024. [Internlm2 technical report](#). *Preprint*, arXiv:2403.17297.

Ilias Chalkidis. 2023. [Chatgpt may pass the bar exam soon, but has a long way to go for the lexglue benchmark](#). *Preprint*, arXiv:2304.12202.

Wei Chen, Qiushi Wang, Zefei Long, Xianyin Zhang, Zhongtian Lu, Bingxuan Li, Siyuan Wang, Jiarong Xu, Xiang Bai, Xuanjing Huang, et al. 2023. [Discfinllm: A chinese financial large language model based on multiple experts fine-tuning](#). *Preprint*, arXiv:2310.15205.

Zhiyu Zoey Chen, Jing Ma, Xinlu Zhang, Nan Hao, An Yan, Armineh Nourbakhsh, Xianjun Yang, Julian McAuley, Linda Petzold, and William Yang Wang. 2024. [A survey on large language models for critical societal domains: Finance, healthcare, and law](#). *Preprint*, arXiv:2405.01769.

Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, Andre F. T. Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, et al. 2024. [Saulmlm-7b: A pioneering large language model for law](#). *Preprint*, arXiv:2403.03883.

Jiayi Cui, Munan Ning, Zongjian Li, Bohua Chen, Yang Yan, Hao Li, Bin Ling, Yonghong Tian, and Li Yuan. 2024a. [Chatlaw: A multi-agent collaborative legal assistant with knowledge graph enhanced mixture-of-experts large language model](#). *Preprint*, arXiv:2306.16092.

Jiayi Cui, Munan Ning, Zongjian Li, Bohua Chen, Yang Yan, Hao Li, Bin Ling, Yonghong Tian, and Li Yuan. 2024b. [Chatlaw: Open-source legal large language model with integrated external knowledge bases](#). *Preprint*, arXiv:2306.16092.

Yiming Cui, Ziqing Yang, and Xin Yao. 2023. [Efficient and effective text encoding for chinese llama and alpaca](#). *arXiv preprint arXiv:2304.08177*.

- Wentao Deng, Jiahuan Pei, Keyi Kong, Zhe Chen, Furu Wei, Yujun Li, Zhaochun Ren, Zhumin Chen, and Pengjie Ren. 2023. [Syllogistic reasoning for legal judgment analysis](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13997–14009, Singapore. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Weijie Fan, Yi Yang, Jing Qi, Qichuan Zhang, Cuiwei Liao, Li Wen, Shuang Wang, Guangxian Wang, Yu Xia, Qihua Wu, et al. 2024. [A deep-learning-based framework for identifying and localizing multiple abnormalities and assessing cardiomegaly in chest x-ray](#). *Nature Communications*, 15(1):1347.
- Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Songyang Zhang, Kai Chen, Zongwen Shen, and Jidong Ge. 2023. [Lawbench: Benchmarking legal knowledge of large language models](#). *Preprint*, arXiv:2309.16289.
- Noelia Ferruz, Steffen Schmidt, and Birte Höcker. 2022. [Protgpt2 is a deep unsupervised language model for protein design](#). *Nature Communications*, 13(1):4348.
- C. Haitao. 2024. LexiLaw: A legal text processing toolkit. <https://github.com/CSHaitao/LexiLaw>. Accessed: 2024-04-29.
- Wanwei He, Jiabao Wen, Lei Zhang, Hao Cheng, Bowen Qin, Yunshui Li, Feng Jiang, Junying Chen, Benyou Wang, and Min Yang. 2023. [Hanfei-1.0](#).
- Quzhe Huang, Mingxu Tao, Chen Zhang, Zhenwei An, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. 2023. [Lawyer llama technical report](#). *Preprint*, arXiv:2305.15062.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Roseval Malaquias Junior, Ramon Pires, Roseli Romero, and Rodrigo Nogueira. 2024. [Juru: Legal brazilian large language model from reputable sources](#). *Preprint*, arXiv:2403.18140.
- Ha-Thanh Nguyen. 2023. [A brief report on lawgpt 1.0: A virtual legal assistant based on gpt-3](#). *Preprint*, arXiv:2302.05729.
- Vinayak Yogesh Ogavinee and et al. 2022. Anima: A comprehensive toolkit for medical image analysis. <https://github.com/lyogavin/Anima>. Accessed: 2024-07-12.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, et al. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Bowen Pan, Yikang Shen, Haokun Liu, Mayank Mishra, Gaoyuan Zhang, Aude Oliva, Colin Raffel, and Rameswar Panda. 2024. [Dense training, sparse inference: Rethinking training of mixture-of-experts language models](#). *Preprint*, arXiv:2404.05567.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). *Preprint*, arXiv:2406.17557.
- 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](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Ren Ren, Shenglin Cai, Xiaona Fang, Xiaoyi Wang, Zheng Zhang, Micol Damiani, Charlotte Hudlerova, Annachiara Rosa, Joshua Hope, Nicola J. Cook, et al. 2023. [Multiplexed detection of viral antigen and rna using nanopore sensing and encoded molecular probes](#). *Nature Communications*, 14(1):7362.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. [Large language models in medicine](#). *Nature Medicine*, 29(8):1930–1940.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Junjie Wang, Yuxiang Zhang, Lin Zhang, Ping Yang, Xinyu Gao, Ziwei Wu, Xiaoqun Dong, Junqing He, Jianheng Zhuo, Qi Yang, et al. 2022. [Fengshenbang 1.0: Being the foundation of chinese cognitive intelligence](#). *CoRR*, abs/2209.02970.
- Yiquan Wu, Yuhang Liu, Yifei Liu, Ang Li, Siying Zhou, and Kun Kuang. [wisdominterrogatory](#). Available at GitHub.
- Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. [Lawformer: A pre-trained language model for chinese legal long documents](#). *Preprint*, arXiv:2105.03887.
- Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei Han, Zhen Hu, Heng Wang, et al. 2018. [Cail2018: A large-scale legal dataset for judgment prediction](#). *Preprint*, arXiv:1807.02478.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Jingyu Sun, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. 2024a. *Chatglm: A family of large language models from glm-130b to glm-4 all tools*. *Preprint*, arXiv:2406.12793.

Zihao Zeng, Yibo Miao, Hongcheng Gao, Hao Zhang, and Zhijie Deng. 2024b. *Adamoe: Token-adaptive routing with null experts for mixture-of-experts language models*. *Preprint*, arXiv:2406.13233.

Ge Zhang, Scott Qu, Jiaheng Liu, Chenchen Zhang, Chenghua Lin, Chou Leuang Yu, Danny Pan, Esther Cheng, Jie Liu, Qunshu Lin, et al. 2024. *Map-neo: Highly capable and transparent bilingual large language model series*. *Preprint*, arXiv:2405.19327.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. *Llamafactory: Unified efficient fine-tuning of 100+ language models*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.

A Synthesized Prompt Templates

To create an effective training dataset, it is essential to ensure that the data are diverse and cover a wide range of types and tasks. Table 2 provides a detailed overview of the data sources.

During the various stages of training, we utilized homologous variant data, a significant portion of which was synthesized using GLM4. The detailed process, including the prompt templates and their effects, is illustrated in Table 3.

B Pre-Training

In order to continued pre-train the Qwen2-7B model, we adhere to the strategy it followed during the continued pre-training phase, which involves predicting the subsequent token based on the context provided by the preceding token. The context

length for our continued pre-training is set to 8192. For the creation of data batches, we shuffle and amalgamate the documents, subsequently truncating them to the aforementioned context lengths. To enhance computational efficiency and curtail memory consumption, we incorporate Flash Attention within the attention modules. The standard optimization algorithm employed for pretraining is AdamW. The hyperparameters are configured with $\beta_1 = 0.9$, $\beta_2 = 0.95$, and $\epsilon = 10^{-8}$. We utilize a cosine_with_warmup learning rate schedule, with a designated peak learning rate for each model size; the warmup steps are set to 3500. The learning rate is tapered down to a minimum of 10% of the peak learning rate, with the maximum learning rate established at 3×10^{-4} . The models are trained using BFloat16 mixed precision to ensure training stability.

C Post-Training

Consistent with pretraining, we also apply next-token prediction as the training task for SFT. We apply the loss masks for the system and user inputs. The model’s training process utilizes the AdamW optimizer, with the following hyperparameters: β_1 set to 0.9, β_2 set to 0.95 and ϵ set to 10^{-8} . The sequence length is limited to 8192, and the batch size is 64. The model undergoes a total of 10000 steps, with the learning rate gradually increased over 4096 steps, reaching a peak of 9×10^{-6} . To prevent over fitting, weight decay is applied with a value of 0.1, dropout is set to 0.1, and gradient clipping is enforced with a limit of 1.0

During the DPO training phase, we employed the LLaMa-Factory (Zheng et al., 2024) as an auxiliary tool, conducting a total of 5,000 training steps. The warmup_with_cosine strategy was utilized, wherein the learning rate gradually increased to reach its maximum value of 1.5×10^{-5} over the initial 2,237 steps, followed by a gradual decrease.

D Parameter-Merging

Compositional Layers: According to the technical report on the Qwen2 series models (Yang et al., 2024), Qwen2-7B (m_A) and Qwen2-72B (m_B) consist of 28 layers (N_A) and 80 layers (N_B), respectively. The hidden size of the two models is noted as 3,584 (D_A) for Qwen2-7B and 8,192 (D_B) for Qwen2-72B. Due to the significant difference in the number of layers between the two models, when selecting the combined layers \mathbb{L}_A and

Type	Description	Token
Legal Regulations	This category encompasses the Constitution, central-level regulations, local regulations at various levels, departmental rules and regulations, as well as bilateral and multilateral treaties, agreements, and other documents with the nature of treaties or agreements concluded by the People’s Republic of China with foreign countries.	7.5B
Case Documents	This includes judicial decisions from the courts, arbitration awards from arbitration institutions, and prosecutorial documents from the procuratorates.	1.33B
Legal Manuscripts	These are core journals indexed by CNKI (China National Knowledge Infrastructure), select theses and dissertations, and certain publications from legal publishing houses.	5.7B
Legal Q&A Data	Rich in legal Q&A data, including common legal questions and their corresponding answers. It covers multiple legal fields such as contract law, labor law, intellectual property, etc.	5.47B

Table 2: List of datasets used during training.

\mathbb{L}_B , we opted to include all layers from Qwen2-7B and a subset of layers from Qwen2-72B. The relationship between them is as follows:

$$N_A = |\mathbb{L}_A| = |\mathbb{L}_B|$$

$$l_{A,i} = l_{B,i} = n_{B,i+24}$$

where $l_{A,i}$ represents the i -th selected layer in the set of selected combined layers for the Qwen2-7B model, and $n_{B,i}$ represents the i -th model layer among all layers of the Qwen2-72B model. Further, $\mathbb{H}_A \in \{H_{A,1}, H_{A,2}, \dots, H_{A,28}\}$ denote the layer representations for the given input after each layer in \mathbb{L}_A .

Learned Projections: Next we map representations from Qwen2-7B to that of Qwen2-72B via a projection layer. In particular, for each layer in \mathbb{L}_A , we learn a projection function $f_{\text{proj}} : \mathbb{R}^{D_A} \rightarrow \mathbb{R}^{D_B}$, that projects representations from these layers to the desired representation size of Qwen2-72B. Let,

$$f_{\text{proj}}(\mathbb{H}_A) \leftarrow \{f_{\text{proj}}(H_{A,1}), \dots, f_{\text{proj}}(H_{A,28})\}$$

This transformation enables cross-attention across models, and also performs an alignment of representations from Qwen2-7B and Qwen2-72B despite frozen weights of the base models.

Cross-attention Layers: Similar to the multi-headed cross-attention in encoder-decoder models (for example Vaswani et al., 2017 and Raffel et al.,

2020)– we introduce cross-attention between representations of the Qwen2-72B and the Qwen2-7B model. In particular, we use $f_{\text{proj}}(H_{A,i})$ from the Qwen2-7B model as the *key* and *value* vectors for each head in cross-attention. We use the vector $H_{B,j}$ from the Qwen2-72B model as the *query* vector.

<i>↔Prompt Templates</i>	
<p>SFT Prompt for English Contents</p> <p># Legal Consultation Simulator</p> <p>You need to construct a pair of Q&A legal consultation data based on the legal regulations provided by the user. The questions should not directly ask about the content of the legal regulations but should conform to the user's normal questioning habits. The legal regulations provided by the user are: { Contents }</p>	<p>DPO Prompt for English Contents</p> <p># Legal Consultation Simulator</p> <p>You are an AI assistant proficient in law, and you need to use your legal knowledge to answer the questions posed by users. The user's question is: { Contents }</p>
<i>↔Demonstrations</i>	
<p>Before Synthesized</p> <p>Article 37: An applicant for a degree or a degree holder, who has one of the following circumstances during the process of pursuing the degree, may have the degree not conferred or revoked by the degree-granting institution upon the resolution of the Degree Evaluation Committee: (1) The thesis or practical results are identified as having academic misconduct such as ghostwriting, plagiarism, or forgery; (2) Misappropriating or impersonating someone else's identity to obtain admission qualifications, or obtaining admission qualifications or graduation certificates through other illegal means; (3) There are other serious illegal acts during the period of study that should not be granted a degree according to the law.</p>	<p>After Synthesized</p> <p>Question: What are the consequences if my thesis is found to be plagiarized during the period of pursuing a degree? Answer: According to the provisions of Article 37 of the relevant laws and regulations, if your thesis or practical results are identified as having academic misconduct such as ghostwriting, plagiarism, or forgery during the period of pursuing a degree, the Degree Evaluation Committee may decide not to confer the degree upon you or revoke the degree already granted. Please pay close attention to academic integrity and comply with the laws and regulations.</p>
<p>Chosen</p> <p>According to existing laws and regulations, an individual who maliciously attacks a public transport driver in the process of driving, or attempts to seize control, thereby disrupting the normal operation of the public transport vehicle and threatening public safety, may be sentenced to fixed-term imprisonment of not more than one year, detention, or control, and must also bear the responsibility of a fine. If the perpetrator's actions also violate other related criminal charges, such as seriously obstructing safe driving and causing significant personal injury and property damage, they should still be convicted and punished according to the more severe provisions. For example, they may be determined to have committed the crime of "endangering public safety by dangerous means." The minimum sentence for intentional crimes of this category starts at three years, and the maximum penalty can be the death penalty.</p>	<p>Rejected</p> <p>In China, criminal acts that obstruct safe driving, such as drunk driving, speeding, fatigue driving, and using a mobile phone, may be sentenced to detention, fines, or even traffic accident crimes according to the relevant provisions of the "Criminal Law of the People's Republic of China," based on the severity and consequences of the act. In serious cases, imprisonment may result. The specific judgment will also take into account factors such as the perpetrator's subjective malice, the dangerousness of the act, and the consequences caused. For more detailed legal advice or case analysis, it is recommended to consult a professional legal person.</p>

Table 3: The synthesized prompt templates with Demonstrations in English.