

WenyanGPT: A Large Language Model for Classical Chinese Tasks

Xinyu Yao¹, Mengdi Wang¹, Bo Chen^{1,2*}, Xiaobing Zhao^{1,2}

¹School of Information Engineering, Minzu University of China

²National Language Resource Monitoring and Research Center of Minority Languages
{xinyu230101, wang_mengdi_wmd}@163.com, chenbomuc@muc.edu.cn, nmzxb_cn@163.com

Abstract

Classical Chinese, as the core carrier of Chinese culture, plays a crucial role in the inheritance and study of ancient literature. However, existing natural language processing models primarily optimize for Modern Chinese, resulting in inadequate performance on Classical Chinese. This paper presents a comprehensive solution for Classical Chinese language processing. By continuing pre-training and instruction fine-tuning on the LLaMA3-8B-Chinese model, we construct a large language model, WenyanGPT, which is specifically designed for Classical Chinese tasks. Additionally, we develop an evaluation benchmark dataset, WenyanBENCH. Experimental results on WenyanBENCH demonstrate that WenyanGPT significantly outperforms current advanced LLMs in various Classical Chinese tasks. We make the model’s training data, instruction fine-tuning data, and evaluation benchmark dataset publicly available to promote further research and development in the field of Classical Chinese processing.

1 Introduction

Classical Chinese is an important component of Chinese culture, with a long history and profound cultural heritage. It is essential for understanding traditional Chinese culture. With the rapid development of artificial intelligence technology, intelligent processing of ancient Chinese texts offers a new solution for the preservation and inheritance of Classical Chinese. Modern technologies, such as digitalization and natural language processing (NLP), can efficiently preserve and spread traditional culture while fostering deeper and more innovative academic research. Maximizing the potential of AI in processing Classical Chinese texts has become an urgent need for cultural inheritance and academic development.

Early research in Classical Chinese language processing focused on tasks such as punctuation, word segmentation, part-of-speech tagging (POS), named entity recognition (NER) and translation. These tasks initially relied on traditional machine learning methods, such as Hidden Markov

*Corresponding author.

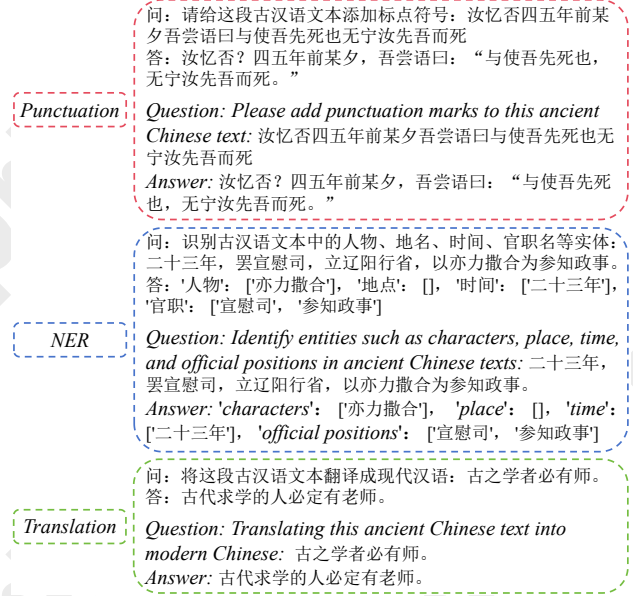


Figure 1: Examples of tasks from WenyanGPT. The model demonstrates advanced knowledge in Classical Chinese and shows strong performance in both Classical Chinese understanding and generation tasks.

Models (HMMs) for POS [Huang *et al.*, 2002], Conditional Random Fields (CRFs) for punctuation [Huang *et al.*, 2010] and NER [Yuan *et al.*, 2019; Li, 2018]. In deep learning, RNNs, LSTMs, GRUs, and Attention mechanisms have been applied to various tasks, including couplet generation and classical poetry generation [Yan *et al.*, 2016; Yi *et al.*, 2017], punctuation and POS using BiLSTM-CRF models [Wang *et al.*, 2019; Cheng *et al.*, 2020; Zhang *et al.*, 2023b; Chang *et al.*, 2024]. As Transformer architectures emerged [Vaswani *et al.*, 2017], studies began using large-scale parallel corpora to train models for translating Classical Chinese into Modern Chinese [Liu *et al.*, 2018] and for generating classical poetry [Huang *et al.*, 2020]. The introduction of pre-trained models including BERT [Kenton and Toutanova, 2019] and GPT [Radford and Narasimhan, 2018] provided new opportunities for intelligent Classical Chinese processing. Some research integrated ancient Chinese texts into the training data of general pre-trained models, improving the

processing performance of Classical Chinese compared to typical pre-trained models [Tian *et al.*, 2020; Wang *et al.*, 2022; Wang *et al.*, 2023a; Liu *et al.*, 2023a]. Other studies used Classical Chinese corpus to continue pre-training and fine-tuning large language models, aiming to build conversational models for Classical Chinese [Zhang *et al.*, 2024; Yang *et al.*, 2024b; Cao *et al.*, 2023; Cao *et al.*, 2024].

However, challenges remain in Classical Chinese processing. Different tasks require training specialized models, and no effective universal model exists. Additionally, there is a lack of standardized evaluation benchmarks in this field, with existing evaluation tasks, datasets, and metrics being inconsistent, making it difficult to perform cross-task comparisons and systematic assessments of model performance.

To address these issues, we propose WenyanGPT, a large language model for Classical Chinese. Some examples of WenyanGPT are shown in Figure 1. We also construct the largest pre-training corpus for continued pre-training, enhancing the model’s domain adaptability. Additionally, we propose a framework for generating domain-specific instruction data for supervised fine-tuning in the development of WenyanGPT. To promote research in the intelligent processing of Classical Chinese, we build the WenyanBENCH evaluation dataset and conduct extensive experiments for detailed analysis. The main contributions are as follows:

- We propose WenyanGPT, a large language model focused on Classical Chinese. It demonstrates superior performance and wide applicability in tasks such as punctuation, POS, translation, etc.
- We release pre-training and instruction fine-tuning datasets, along with a novel method for constructing domain-specific fine-tuning data, providing valuable resources for future research.
- We introduce WenyanBENCH, an evaluation benchmark for Classical Chinese tasks, with extensive experiments verifying WenyanGPT’s leading performance across multiple tasks.

2 Related Work

2.1 PLMs

In 2017, Google introduced a new neural network architecture, Transformer. It utilizes self-attention mechanisms to better handle long-distance dependencies and significantly improves training efficiency through parallel computation. Based on Transformer, various LLMs have been proposed. BERT employs an encoder-only Transformer architecture and is pre-trained using masked language modeling and next sentence prediction tasks. The GPT series, uses a decoder-only Transformer architecture and an autoregressive language model (ALM). Over the course of the GPT series, the model size has steadily increased, from the original GPT to subsequent iterations, including GPT-2 [Radford *et al.*, 2019], GPT-3 [Brown *et al.*, 2020], and GPT-4 [Achiam *et al.*, 2023], with continuous improvements in performance. PaLM [Chowdhery *et al.*, 2023] uses the standard Transformer architecture in a decoder-only model with a modified SwiGLU activation function. This model, with 540 billion parameters,

achieved human-level performance in 1-shot learning on the BIG-bench [Srivastava *et al.*, 2023] dataset. In 2023, Meta AI released the LLaMA model [Touvron *et al.*, 2023]. This model also follows a decoder-only Transformer architecture and excels in various NLP tasks after large-scale training. In 2024, LLaMA 3 [Dubey *et al.*, 2024] was released, including a pre-trained version with 405 billion parameters and a post-training version, alongside the LLaMA Guard 3 model for input-output safety. Pre-trained language models have seen rapid development, with Transformer-based models becoming the mainstream technology in NLP.

2.2 Classical Chinese PLMs

Pre-trained language models have achieved widespread success in the field of NLP. However, studies show that general-domain models often lack specialized knowledge for tasks in specific domains. Models pre-trained with domain-specific data tend to perform better for specialized tasks [Ke *et al.*, 2023; Gupta *et al.*, 2023; Ibrahim *et al.*, 2024; Taylor *et al.*, 2022; Lehman *et al.*, 2023; Liu *et al.*, 2020]. In the field of Classical Chinese, several studies have extended models such as BERT, RoBERTa, and GPT by incorporating Classical Chinese corpora for pre-training, resulting in specialized models like AnchiBERT [Tian *et al.*, 2020], SikuBERT and SikuRoBERTa [Wang *et al.*, 2022], GujiBERT and GujiGPT series [Wang *et al.*, 2023a], and SikuGPT [Liu *et al.*, 2023a]. These models show improved performance over general pre-trained models in Classical Chinese tasks. Instruction fine-tuning is another effective strategy. Using Supervised Fine-Tuning (SFT) can activate LLMs’ ability to understand and answer questions in a specific domain [Liu *et al.*, 2023b; Xiong *et al.*, 2023; Wang *et al.*, 2023b; Yue *et al.*, 2023; Huang *et al.*, 2023; Cui *et al.*, 2023; Yang *et al.*, 2023b; Zhang *et al.*, 2023a; Dan *et al.*, 2023]. Classical Chinese LLMs [Zhang *et al.*, 2024; Yang *et al.*, 2024b; Cao *et al.*, 2023; Cao *et al.*, 2024] are in the early stages of development. For example, the "Xunzi"¹ Classical Chinese large language model was trained on Classical Chinese-related corpora based on general models such as Qwen2.5 [Yang *et al.*, 2024a], Baichuan2 [Yang *et al.*, 2023a], and GLM-4 [Zeng *et al.*, 2024]. It has shown excellent performance in intelligent tagging, information extraction, and other tasks. TongGu [Cao *et al.*, 2024], through two-stage instruction fine-tuning, is capable of Classical Chinese punctuation, translation, and appreciation tasks. In this paper, WenyanGPT is fine-tuned on higher-quality pre-training data and a larger, more diverse instruction dataset, showing superior and more comprehensive task handling capabilities.

3 WenyanGPT

In order to obtain the WenyanGPT Classical Chinese model, we first construct a Classical Chinese pre-training corpus and continue pre-training based on LLaMA3-8B-Chinese (Section 3.1). Then, we propose a method to construct domain instruction data (Section 3.2). In our framework, instruction generation is manually constructed, guided by LLMs, and

¹<https://github.com/Xunzi-LLM-of-Chinese-classics/XunziALLM>.

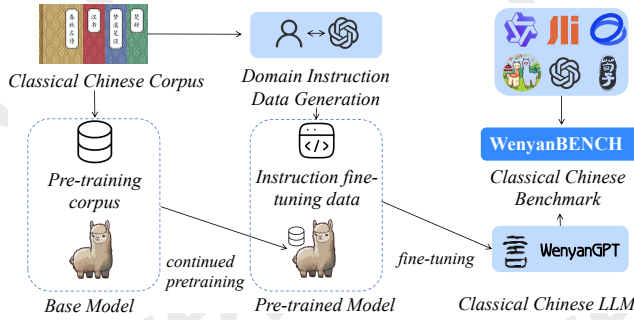


Figure 2: Overall Training Framework of WenyanGPT.

Source	Scale	Source	Scale
Daizhige	5.2G	Poetry-master	323M
wenyanguji.com	1.6G	PoetrySlider	16M
network resource	1.1G	TCM-Ancient-Books	322M
chinese-novel-master-long	294M	text-to-picture-sidamingzhi	6.7M
chinese-gushiwen	23M	chtxt-main	88M
Classical-Chinese	208M	chinese-poetry	115M
Classical-Modern	853M	guner2023	63M
core-books-main	752M	kangxi-master	37M
core-texts	232M	poems-db	660M
GuWen-master	2.5M	scripta-sinica	3.7G

Table 1: Sources and Scale of Classical Chinese Pre-training Corpus.

tested to ensure the high quality of the fine-tuning data. The complete training process is shown in Figure 2.

3.1 Pre-training

The corpus used in the pre-training phase is sourced from authoritative websites such as Daizhige, Wenyanguji, and various Classical Chinese-related data collected and organized from GitHub. The detailed data sources and scale are shown in Table 1. We uniformly format and store data from these different sources, removing redundant information, errors, special symbols, and invalid characters. As a result, we obtain a clean, large-scale, high-quality Classical Chinese corpus of approximately 16GB. This corpus covers the Four Books and Five Classics, including Confucian classics, historical records, works of various philosophers, poetry, essays, drama, novels, miscellanies, and other literary genres. It also encompasses diverse fields such as local gazetteers, genealogies, religious texts, agricultural, law, medicine, astronomy, geography, craft books, and military texts. The corpus integrates both simplified and traditional Chinese characters, spanning from the Pre-Qin period to the Republic of China, providing rich materials for deep learning and research on ancient Chinese texts. We select LLaMA3-8B-Chinese as the base model and use the bfloat16 data format during training to improve efficiency. The hyper-parameter settings in pre-

Hyper parameter	Value
per device train batch size	16
gradient accumulation steps	1
learning rate	1.0e-4
num train epochs	1
lr scheduler type	cosine
warmup ratio	0.1

Table 2: Hyper-parameter Settings in Continue Pre-training.

training are shown in Table 2.

3.2 Supervised Fine-Tuning

Based on continued pre-training, we perform supervised fine-tuning to better adapt the model to specific tasks and instructions. We use a high-quality set of instruction fine-tuning data, which we have previously collected and organized, to trigger the knowledge the model acquires during pre-training. The detailed process for constructing the instruction fine-tuning data is shown in Figure 3.

Data Selection and Initial Organization. We select relevant data from the Classical Chinese corpus, including three major categories: question-answer pairs, parallel corpora, and tagged corpora. Parallel corpora are used for translation tasks, while tagged corpora support fine-grained tasks such as punctuation and POS. In cases where the corpus lacks clear question-answer pairs, we generate supplementary data. During the selection phase, we prioritize high-quality data with clear content, standardized semantics, and task relevance to build the initial high-quality input-output pairs.

Manual Design of Task Instructions and Model Expansion. We design initial task instruction templates manually based on high-quality input and output, covering tasks such as Classical Chinese punctuation, and translating Classical Chinese into Modern Chinese. Then, we use LLMs, such as GPT and Qwen series to expand the task instructions. On the one hand, we generate diverse instructions from existing ones; on the other hand, we allow LLMs to perform reverse reasoning from the existing high-quality input-output pairs to generate new instructions, ensuring diversity in the instructions. After the expansion, we conduct an initial screening of the generated instructions, removing those with unclear or unreasonable semantics, resulting in a seed instruction set.

Testing the Instruction Set and Optimizing Fine-Tuning Data. We randomly select high-quality input-output pairs and combine them with the seed instruction set to evaluate the model’s adherence to instructions across different task scenarios. We analyze the test results to refine the instructional design and identify the task instructions that produce higher-quality outputs. Then, we combine the optimized instruction set with high-quality input-output pairs to construct fine-tuning data for specific instructions, providing reliable data support for subsequent model training.

Generation and Supplementation of Specific Task Data. Due to the lack of specific task instruction data in the corpus, we manually design diverse initial instructions and inputs.

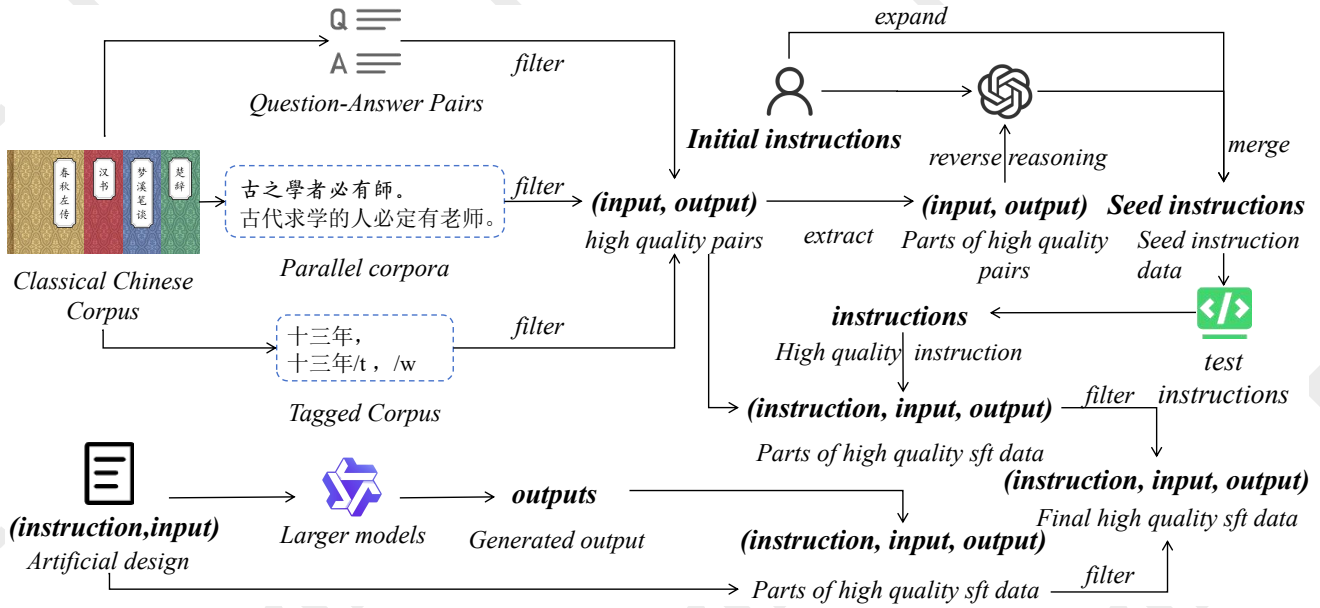


Figure 3: Instruction Fine-Tuning Data Construction Process.

Task	Data Source	Num
Punctuation	Daizhige	107,3017
POS	evahan	9,952
NER	Self-built	29923
Translation	classical-modern	222,700
	wenyanguji.com	302,724
Word explanation	gushiwen.com	31,088
Reverse dictionary	chinese-dictionary	39,708
	chinese-xinhua	138,810
Total		1,847,922

Table 3: Sources and Scale of Instruction Fine-Tuning Data.

Hyper parameter	Value
per device train batch size	8
gradient accumulation steps	2
learning rate	1.0e-4
num train epochs	1
lr scheduler type	cosine
warmup ratio	0.1

Table 4: Hyper-parameter Settings in Fine-Tuning.

Then we use LLMs including Qwen2.5-14B and Qwen2.5-72B to generate high-quality answers. We remove low-quality or irrelevant content through manual selection and automatic quality checks, forming another portion of the instruction data.

Data Integration and Output Validation. We integrate the data generated from the corpus and the supplementary data from the LLMs to form a complete instruction dataset. During integration, we ensure consistency between instruc-

Task	Num	Task	Num
Punctuation	7,559	Translation	5,013
POS	1,247	Word Explanation	3,931
NER	3,741	Reverse Dictionary	4,462
Total	25,953		

Table 5: Data Sources and Detailed Statistics of WenyanBench.

tions and outputs, covering multiple Classical Chinese task scenarios. Through comprehensive data validation and optimization, we create a high-quality instruction dataset. In the end, we obtain about 1.85 million instruction fine-tuning data. The detailed data sources and counts are shown in Table 3. We use these data to fine-tune the pre-trained model. The hyper-parameter settings in fine-tuning are shown in Table 4.

4 Benchmarking Classical Chinese Tasks

WenyanBench. In order to evaluate the model’s performance on Classical Chinese tasks, we devise a benchmark known as WenyanBench. WenyanBench shares the same data sources as the instruction fine-tuning data and has undergone duplicate data removal, as well as validation by both manual and LLMs. For quality control, we sample a subset of the data. The distribution and detailed statistics of WenyanBench are shown in Table 5.

Tasks. Our benchmark includes six tasks related to Classical Chinese. Among them, we subdivide 14 types of punctuation marks in the punctuation task, divide Classical Chinese word classes into 17 categories in the POS task, such as nr for personal nouns, ns for place nouns, and define 4 categories for the NER task.

Metrics. For the WenyanBench benchmark, different evaluation metrics are used for different types of tasks (under-

Model	Punctuation			POS			NER		
	P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)
Qwen2.5-7B-Instruct	54.34	53.31	53.82	51.25	48.16	49.65	66.05	46.55	54.61
Baichuan2-7B-Chat	51.05	21.03	29.79	47.11	30.97	37.37	35.10	10.58	16.26
GLM-4-9B-Chat	52.39	55.00	53.66	49.90	54.98	52.32	52.19	45.42	48.57
Meta-Llama-3-8B-Instruct	55.05	22.41	31.85	25.73	17.06	20.52	47.50	57.48	52.01
Llama3-8B-Chinese-Chat	45.76	38.07	41.56	21.34	19.34	20.29	46.85	66.69	55.04
Xunzi-Qwen-1.5-7B-Chat	52.08	47.19	49.51	77.54	78.07	77.81	49.79	51.21	50.49
GPT-4o	52.00	50.70	51.34	82.41	81.11	<u>81.75</u>	61.58	76.97	<u>68.42</u>
Deepseek-V3	56.33	61.94	<u>59.01</u>	79.12	79.18	79.15	56.83	79.75	66.36
WenyanGPT	76.84	74.52	75.66	89.66	88.54	89.1	92.14	90.19	91.16

Table 6: Results for understanding tasks (Punctuation, POS, NER) on WenyanBench. The results underlined represent the second-best model’s F1 score.




Figure 4: The F1 scores of WenYanGPT for the subcategories of understanding tasks (including Punctuation, POS, and NER) on WenyanBench.

standing tasks and generation tasks). For understanding tasks, evaluation primarily relies on Precision, Recall, and F1-Score. For generation tasks, BLEU and BERT-Score are used as evaluation metrics. BLEU measures the N-gram overlap between generated content and reference answers, while BERT-Score better captures the semantic similarity between the generated content and reference answers.

Evaluation Method. To efficiently assess model performance, we design a set of scripted tools to automatically compute BLEU, BERT-Score, and other metrics. These tools quickly and accurately quantify model outputs, providing clear feedback for model optimization. This automated evaluation approach improves evaluation efficiency and ensures the consistency and comparability of the results.

5 Experiments

The experiments evaluate WenyanGPT’s performance on understanding and generation tasks in Classical Chinese.

5.1 Experimental Setup

Baselines. The baselines include general-domain and Classical Chinese domain LLMs. The general-domain LLMs are Qwen2.5-7B-Instruct, Baichuan2-7B-Chat, GLM-4-9B-Chat, Meta-Llama-3-8B-Instruct, Llama3-8B-Chinese-Chat, GPT-4o [Hurst *et al.*, 2024], and Deepseek-V3 [Liu *et al.*, 2024]. The Classical Chinese domain LLM is Xunzi-Qwen1.5-7B-Chat.

Data and Evaluation. We use the WenyanBench benchmark for testing. The understanding tasks, including punctuation, POS, and NER, are evaluated by Precision, Recall, and F1-Score. The generation tasks include word explanation, translation, and reverse dictionary, where BLEU is used for word explanation and translation, and BERT-Score is used for reverse dictionary.

5.2 Experimental Analysis

WenyanGPT demonstrates a significant lead in language understanding tasks. The experimental results for understanding tasks are presented in Table 6. In the NER task, WenyanGPT’s precision, recall, and F1 score all exceed 90%, while the second-best model, GPT-4o, fails to surpass 77% in any of these metrics. In the punctuation task, WenyanGPT’s F1 score is 16.65% higher than that of the second-best model, Deepseek-V3, reaching 75.66%. Additionally, in the POS task, WenyanGPT’s F1 score is 7.35% higher than that of the second-best model, GPT-4o. These results highlight WenyanGPT’s overwhelming superiority in understanding tasks, particularly in NER and punctuation. This performance is attributed to the model’s extensive pre-training on Classical Chinese data, enabling it to handle complex linguistic phenomena and ensuring higher accuracy and stability in fundamental language understanding tasks, such as POS and NER.

WenyanGPT demonstrates excellent capabilities and high efficiency in subcategories of Classical Chinese compre-

Model	Translation				Word explanation				Reverse dictionary		
	BLEU1	BLEU2	BLEU3	BLEU4	BLEU1	BLEU2	BLEU3	BLEU4	P(%)	R(%)	F1(%)
Qwen2.5-7B-Instruct	0.37	0.23	0.17	0.14	0.16	0.09	0.07	0.05	68.43	68.99	68.66
Baichuan2-7B-Chat	0.33	0.20	0.14	0.11	0.14	0.08	0.05	0.04	64.81	66.21	65.42
GLM-4-9B-Chat	0.34	0.21	0.15	0.12	0.15	0.09	0.06	0.05	65.58	68.04	66.69
Meta-Llama-3-8B-Instruct	0.16	0.09	0.06	0.05	0.11	0.06	0.05	0.04	59.41	64.13	61.48
Llama3-8B-Chinese-Chat	0.26	0.15	0.10	0.08	0.11	0.06	0.04	0.03	61.8	65.18	63.28
Xunzi-Qwen1.5-7B-Chat	0.22	0.15	0.11	0.09	0.11	0.08	0.06	0.05	66.47	68.45	67.35
GPT-4o	0.41	0.27	0.19	0.14	0.19	0.13	0.09	0.07	64.96	66.76	65.81
Deepseek-V3	0.30	0.19	0.13	0.10	0.20	0.14	0.11	0.08	71.93	71.84	71.88
WenyanGPT	0.47	0.33	0.24	0.19	0.35	0.31	0.27	0.23	75.51	75.31	75.39

Table 7: Results for generation tasks (Translation, Word explanation, Reverse dictionary) on WenyanBench. The results underlined represent the second-best model’s BLEU1 score and BERT-Score-F1 score.

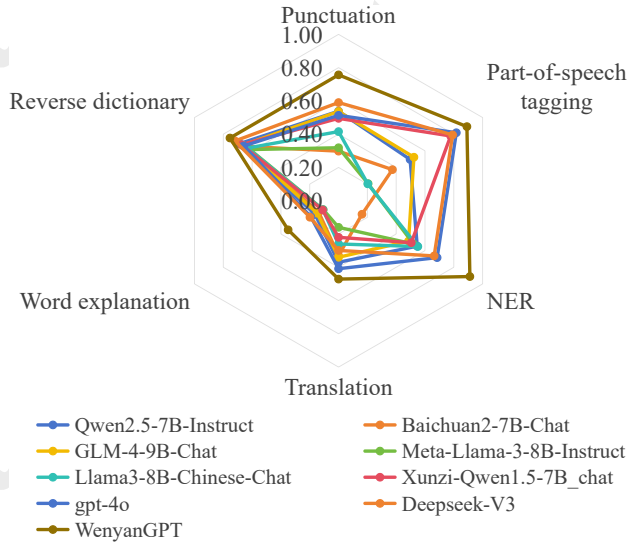


Figure 5: Radar plot showing model performance on WenyanBench, with values normalized to a 0-1 scale.

hension tasks. As shown in Figure 4, the F1 scores of our model in the three task’s subcategories are generally stable and high. Specifically, in the NER task, WenyanGPT’s F1 score remains above 80%, demonstrating its strong ability to correctly identify entities such as historical figures, place, and proper nouns in ancient Chinese texts. This performance shows that it has strong accuracy and robustness when processing ancient Chinese texts, and can effectively capture complex contextual relationships and word meaning changes. Overall, WenyanGPT’s high F1 score in ancient Chinese comprehension tasks not only reflects its high efficiency in basic tasks, but also demonstrates its advantages and potential in processing fine-grained tasks in ancient Chinese.

WenyanGPT can significantly enhance the quality of generated content in generation tasks. The experimental results for generation tasks are displayed in Table 7. WenyanGPT exhibits strong semantic preservation and contextual consistency, demonstrates superior and more stable performance in translation and word explanation tasks, with

consistently high BLEU scores (BLEU1-BLEU4). In classical Chinese translation, the model’s BLEU1 score of 0.47 exceeds that of the second-best model by 0.06. Furthermore, WenyanGPT’s F1 score in the reverse dictionary task is 3.47% higher than the second best model. Through meticulous pre-training and multitask training, WenyanGPT develops robust contextual consistency, enabling it to generate content that accurately reflects the intended meaning and context. This capability ensures high-quality semantic transmission and content coherence in both long-text generation and complex tasks, such as word explanation.

WenyanGPT can achieve better performance than the previous open source LLMs in Classical Chinese tasks. WenyanGPT, through large-scale pre-training and multitask joint optimization, demonstrates substantial capabilities in Classical Chinese processing tasks, establishing a clear lead over existing mainstream LLMs. This validates the effectiveness of the proposed approach. The performance of various LLMs on WenyanBench is shown in Figure 5. It is evident that WenyanGPT has achieved the highest score in six Classical Chinese tasks. WenyanGPT’s multitask training strategy, which integrates various Classical Chinese processing tasks, enhances the model’s ability to learn across tasks. The mutual reinforcement of shared linguistic features and semantic information between tasks notably boosts the model’s generalization capabilities. This inter-task synergy not only improves performance on individual tasks but also enables WenyanGPT to effectively handle multiple tasks simultaneously, maintaining stable performance, especially in complex tasks, and demonstrating superior adaptability across tasks.

5.3 Case Study

We provide response examples for understanding tasks (POS) and generation tasks (word explanation) from five LLMs: WenyanGPT, Deepseek-V3, GPT-4o, Qwen2.5-7B-Instruct, and Xunzi-Qwen1.5-7B-Chat, as shown in Table 8 and Table 9. By analyzing these examples, we can see the performance differences of different LLMs in ccessing tasks, especially in the understanding and generation of classical Chinese.

Table 8 illustrates some typical errors in the POS task. Specifically, GPT-4o makes mistakes in tagging time words and proper nouns. For instance, it incorrectly tags the time

Case 1	Please segment the following Classical Chinese content and accurately tag the parts of speech: 四年春，衛州吁弑桓公而立。
Ground Truth	四年/t 春/n， /w 衛州吁/nr 弑/v 桓公/nr 而/c 立/v。 /w
WenyanGPT	四年/t 春/n， /w 衛州吁/nr 弑/v 桓公/nr 而/c 立/v。 /w
Deepseek-V3	四/m 年/t 春/t， /w 衛/ns 州吁/nr 弑/v 桓公/nr 而/c 立/v。 /w
GPT-4o	四年/n 春/n， /w 衛州/n 吁/v 弑/v 桓公/n 而/c 立/v。 /w
Qwen2.5-7B-Instruct	四年/t 春/w， /w 卫州吁/nr 射/v 桓公/nr 而/c 立/v。 /w
Xunzi-Qwen1.5-7B-Chat	四年春/t， /w 衛州吁/nr 杀/v 桓公/nr 而/c 立/v。 /w

Table 8: Response examples for the POS task from different LLMs. Errors in POS are marked in red, while text errors are highlighted in blue.

expression "四年" (the fourth year) as a common noun, and it also struggles with recognizing the proper noun "衛州吁" (Wei Zhou Xu). Qwen2.5-7B-Instruct errors in tagging the part of speech for "春" (spring) and sometimes replaces characters in the original text with inappropriate alternatives. Xunzi-Qwen1.5-7B-Chat primarily fails in distinguishing between time words and nouns, mistakenly using the simplified character "杀" instead of the correct "弑". These errors highlight the model's difficulty in making precise distinctions between similar words.

For the **word explanation** task, WenyanGPT's response is the closest to the Ground Truth. Although WenyanGPT adds some modifiers, it maintains simplicity and refinement in describing the color, which matches the tone of the poem. Deepseek-V3, GPT-4o, Qwen2.5-7B-Instruct, and Xunzi-Qwen1.5-7B-Chat do not specify the subject as "海棠花" (begonia flowers) and instead over-explained the moistness of the petals and the process of spreading color, deviating from the standard, concise description of "晕红" (blushing red). Specifically, the responses from Qwen2.5-7B-Instruct and Xunzi-Qwen1.5-7B-Chat extend the impact of the external environment, which do not fully align with the original text's expression. This over-expansion shows that these LLMs fail to capture the intended simplicity and depth of the original description.

WenyanGPT demonstrates strong semantic understanding and retention capabilities in handling Classical Chinese tasks. Its ability to accurately tag parts of speech and generate concise, poetic word explanations demonstrates a deep grasp of the nuances of Classical Chinese. In contrast, other LLMs struggle with both understanding the subtle distinctions of Classical Chinese and producing responses that remain faithful to the content and meaning of the original text. This positions WenyanGPT as a powerful tool for handling complex tasks in the domain of Classical Chinese, far surpassing the other LLMs in terms of both accuracy and literariness.

Case 2	Explain the meaning of "晕红" in the line "一片晕红才著雨，几丝柔绿乍和烟" from "浣溪沙·谁道飘零不可怜".
Ground Truth	形容海棠花的色泽。 Describes the color of begonia flowers.
WenyanGPT	形容海棠花色泽鲜艳而浓丽。 Describes the bright and rich color of the begonia flower.
Deepseek-V3	指雨后花朵泛起的朦胧红晕。 It refers to the hazy red glow that appears on flowers after the rain.
GPT-4o	指花瓣在雨水滋润下呈现出的一种柔和而略带湿润光泽的红色。 Refers to the soft and slightly moist red color that the petals appear to have when nourished by rain.
Qwen2.5-7B-Instruct	指花瓣上晕染开的淡红色，适合颜色柔和、渐变，类似于涂抹或被雨水滋润后散开的红晕。 Refers to the light red color that spreads across the petals. It is suitable for soft and gradual colors, similar to the blush that spreads after being smeared or moistened by rain.
Xunzi-Qwen1.5-7B-Chat	指花瓣上的红色晕染，因为花瓣在雨中被淋湿，所以红色晕染更加明显。 Refers to the red smudge on the flower petals. Because the petals got wet in the rain, the red smudge is more obvious.

Table 9: Response examples for the word explanation task from different LLMs.

6 Conclusion

We propose a comprehensive solution for Classical Chinese language processing challenges, including the development of WenyanGPT, a large language model focused on the Classical Chinese domain, and WenyanBENCH, an evaluation benchmark dataset for Classical Chinese tasks. We release pre-training and instruction fine-tuning datasets and describe the method for constructing the instruction fine-tuning dataset. Through systematic experiments and analysis, we demonstrate the significant impact of domain-specific pre-training and multi-task instruction fine-tuning on improving Classical Chinese processing capabilities. Our model outperforms existing mainstream LLMs in various downstream tasks. In future, we attend to explore the potential of multimodal models by combining Classical Chinese texts with image data (such as inscriptions and manuscripts) to enhance processing capabilities.

Acknowledgements

We would like to sincerely thank all reviewers for their valuable feedback and constructive comments. This work is supported by the National Social Science Fund of China (Grant No. 22&ZD035).

References

- [Achiam *et al.*, 2023] Josh Achiam, Steven Adler, Sandhini Agarwal, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [Cao *et al.*, 2023] Jiahuan Cao, Dezhi Peng, Yongxin Shi, Zongyuan Jiang, and Lianwen Jin. Translating ancient chinese to modern chinese at scale: A large language model-based approach. In *International Conference on Algorithmic Learning Theory*, 2023.
- [Cao *et al.*, 2024] Jiahuan Cao, Dezhi Peng, Peirong Zhang, Yongxin Shi, et al. Tonggu: Mastering classical chinese understanding with knowledge-grounded large language models. In *Conference on Empirical Methods in Natural Language Processing*, 2024.
- [Chang *et al.*, 2024] Bolin Chang, Yiguo Yuan, Bin Li, et al. Automatic word segmentation and part-of-speech tagging for classical chinese based on radicals. *Data Analysis and Knowledge Discovery*, 8(11):102–113, 2024.
- [Cheng *et al.*, 2020] Ning Cheng, Bin Li, Liming Xiao, et al. Integration of automatic sentence segmentation and lexical analysis of ancient chinese based on bilstm-crf model. In *Proceedings of LT4HALA 2020-1st Workshop on Language Technologies for Historical and Ancient Languages*, pages 52–58, 2020.
- [Chowdhery *et al.*, 2023] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- [Cui *et al.*, 2023] Jiayi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. Chatlaw: Open-source legal large language model with integrated external knowledge bases. *CoRR*, 2023.
- [Dan *et al.*, 2023] Yuhao Dan, Zhikai Lei, Yiyang Gu, et al. Educhat: A large-scale language model-based chatbot system for intelligent education, 2023.
- [Dubey *et al.*, 2024] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [Gupta *et al.*, 2023] Kshitij Gupta, Benjamin Th’erien, Adam Ibrahim, et al. Continual pre-training of large language models: How to (re)warm your model? *ArXiv*, abs/2308.04014, 2023.
- [Huang *et al.*, 2002] Liang Huang, Yinan Peng, Huan Wang, and Zhenyu Wu. Statistical part-of-speech tagging for classical chinese. In *Proceedings of the 5th International Conference on Text, Speech and Dialogue*, TSD ’02, page 115–122, Berlin, Heidelberg, 2002. Springer-Verlag.
- [Huang *et al.*, 2010] Hen-Hsen Huang, Chuen-Tsai Sun, and Hsin-Hsi Chen. Classical chinese sentence segmentation. In *CIPS-SIGHAN Joint Conference on Chinese Language Processing, Beijing, China, August 28–29, 2010*, 2010.
- [Huang *et al.*, 2020] Chuen-Min Huang, Kuo-Lin Lu, Yi-Ying Cheng, and Yu-Chen Peng. Generating chinese classical poetry with quatrain generation model (qgm) using encoder-decoder lstm. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 5700–5702. IEEE, 2020.
- [Huang *et al.*, 2023] Quzhe Huang, Mingxu Tao, Zhenwei An, Chen Zhang, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. Lawyer llama technical report. 2023.
- [Hurst *et al.*, 2024] Aaron Hurst, Adam Lerer, Adam P Goucher, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- [Ibrahim *et al.*, 2024] Adam Ibrahim, Benjamin Th’erien, Kshitij Gupta, Mats L. Richter, et al. Simple and scalable strategies to continually pre-train large language models. *ArXiv*, abs/2403.08763, 2024.
- [Ke *et al.*, 2023] Zixuan Ke, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bin Liu. Continual pre-training of language models. In *International Conference on Learning Representations*, 2023.
- [Kenton and Toutanova, 2019] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, page 2. Minneapolis, Minnesota, 2019.
- [Lehman *et al.*, 2023] Eric Lehman, Evan Hernandez, Diwakar Mahajan, et al. Do we still need clinical language models? *ArXiv*, abs/2302.08091, 2023.
- [Li, 2018] N Li. Automatic extraction of alias in ancient local chronicles based on conditional random fields. *J. Chin. Inf. Process*, 32:41, 2018.
- [Liu *et al.*, 2018] Dayiheng Liu, Jiancheng Lv, Kexin Yang, and Qian Qu. Ancient-modern chinese translation with a new large training dataset. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 19:1 – 13, 2018.
- [Liu *et al.*, 2020] Zhuang Liu, Degen Huang, Kaiyu Huang, Zhuang Li, and Jun Zhao. Finbert: A pre-trained financial language representation model for financial text mining. In *International Joint Conference on Artificial Intelligence*, 2020.
- [Liu *et al.*, 2023a] Chang Liu, Dongbo Wang, Zhixiao Zhao, et al. Sikugpt: A generative pre-trained model for intelligent information processing of ancient texts from the perspective of digital humanities. *ArXiv*, abs/2304.07778, 2023.
- [Liu *et al.*, 2023b] June M. Liu, Donghao Li, He Cao, Tianhe Ren, Zeyi Liao, and Jiamin Wu. Chatcounselor: A

- large language models for mental health support. *ArXiv*, abs/2309.15461, 2023.
- [Liu *et al.*, 2024] Aixin Liu, Bei Feng, Bing Xue, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.
- [Radford and Narasimhan, 2018] Alec Radford and Karthik Narasimhan. Improving language understanding by generative pre-training. 2018.
- [Radford *et al.*, 2019] Alec Radford, Jeff Wu, Rewon Child, David Luan, et al. Language models are unsupervised multitask learners. 2019.
- [Srivastava *et al.*, 2023] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Trans. Mach. Learn. Res.*, 2023, 2023.
- [Taylor *et al.*, 2022] Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, et al. Galactica: A large language model for science. *ArXiv*, abs/2211.09085, 2022.
- [Tian *et al.*, 2020] Huishuang Tian, Kexin Yang, Dayiheng Liu, and Jiancheng Lv. Anchibert: A pre-trained model for ancient chinese language understanding and generation. *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2020.
- [Touvron *et al.*, 2023] Hugo Touvron, Thibaut Lavril, Gautier Izacard, et al. Llama: Open and efficient foundation language models. *ArXiv*, abs/2302.13971, 2023.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, et al. Attention is all you need. In *Neural Information Processing Systems*, 2017.
- [Wang *et al.*, 2019] Hongbin Wang, Haibing Wei, Jianyi Guo, and Liang Cheng. Ancient chinese sentence segmentation based on bidirectional lstm+ crf model. *Journal of advanced computational intelligence and intelligent informatics*, 23(4):719–725, 2019.
- [Wang *et al.*, 2022] Dongbo Wang, Chang Liu, Ziheng Zhu, Jiangfeng Liu, et al. Sikubert and sikuroberta: Construction and application of pre-trained models for the siku quanshu in the field of digital humanities. *Library Tribune*, 42(06):31–43, 2022.
- [Wang *et al.*, 2023a] Dongbo Wang, Chang Liu, Zhixiao Zhao, et al. Gujibert and gujigpt: Construction of intelligent information processing foundation language models for ancient texts. *arXiv preprint arXiv:2307.05354*, 2023.
- [Wang *et al.*, 2023b] Hao Wang, Chi-Liang Liu, Nuwa Xi, Zewen Qiang, Sendong Zhao, Bing Qin, and Ting Liu. Huatuo: Tuning llama model with chinese medical knowledge. *ArXiv*, abs/2304.06975, 2023.
- [Xiong *et al.*, 2023] Honglin Xiong, Sheng Wang, Yitao Zhu, Zihao Zhao, Yuxiao Liu, Linlin Huang, Qian Wang, and Dinggang Shen. Doctorglm: Fine-tuning your chinese doctor is not a herculean task. *ArXiv*, abs/2304.01097, 2023.
- [Yan *et al.*, 2016] Rui Yan, Cheng-Te Li, Xiaohua Hu, and Ming Zhang. Chinese couplet generation with neural network structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2347–2357, 2016.
- [Yang *et al.*, 2023a] Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, et al. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*, 2023.
- [Yang *et al.*, 2023b] Yi Yang, Yixuan Tang, and Kar Yan Tam. Investlm: A large language model for investment using financial domain instruction tuning, 2023.
- [Yang *et al.*, 2024a] An Yang, Baosong Yang, Beichen Zhang, et al. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.
- [Yang *et al.*, 2024b] Songhua Yang, Hanjie Zhao, Senbin Zhu, et al. Zhongjing: Enhancing the chinese medical capabilities of large language model through expert feedback and real-world multi-turn dialogue. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19368–19376, 2024.
- [Yi *et al.*, 2017] Xiaoyuan Yi, Ruoyu Li, and Maosong Sun. Generating chinese classical poems with rnn encoder-decoder. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data: 16th China National Conference, CCL 2017, and 5th International Symposium, NLP-NABD 2017, Nanjing, China, October 13-15, 2017, Proceedings 16*, pages 211–223. Springer, 2017.
- [Yuan *et al.*, 2019] Y Yuan, D Wang, S Huang, and B Li. The comparative study of different tagging sets on entity extraction of classical books. *Data Analysis and Knowledge Discovery*, 3(03):57–65, 2019.
- [Yue *et al.*, 2023] Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, et al. Disc-lawllm: Fine-tuning large language models for intelligent legal services. *ArXiv*, abs/2309.11325, 2023.
- [Zeng *et al.*, 2024] Team Glm Aohan Zeng, Bin Xu, Bowen Wang, et al. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *ArXiv*, abs/2406.12793, 2024.
- [Zhang *et al.*, 2023a] Xuanyu Zhang, Qing Yang, and Dongliang Xu. Xuanyuan 2.0: A large chinese financial chat model with hundreds of billions parameters. *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2023.
- [Zhang *et al.*, 2023b] Yiqin Zhang, Sanhong Deng, Qi Zhang, Dongbo Wang, and Hongcun Gong. Comparative analysis of language models for linguistic examination of ancient chinese classics: A case study of zuozhuan corpus. In *2023 International Conference on Asian Language Processing (IALP)*, pages 154–161. IEEE, 2023.
- [Zhang *et al.*, 2024] Jundong Zhang, Songhua Yang, Jiangfeng Liu, and Qi Huang. Aigc empowering the revitalization of ancient books on traditional chinese medicine: building the huang-di large language model. *Library Tribune*, 44(10):103–112, 2024.