

Enabling Discriminative Reasoning in LLMs for Legal Judgment Prediction

Chenlong Deng, Kelong Mao, Yuyao Zhang, Zhicheng Dou*
Gaoling School of Artificial Intelligence, Renmin University of China
{dengchenlong, dou}@ruc.edu.cn

Abstract

Legal judgment prediction is essential for enhancing judicial efficiency. In this work, we identify that existing large language models (LLMs) underperform in this domain due to challenges in understanding case complexities and distinguishing between similar charges. To adapt LLMs for effective legal judgment prediction, we introduce the *Ask-Discriminate-PredicT* (ADAPT) reasoning framework inspired by human judicial reasoning. ADAPT involves decomposing case facts, discriminating among potential charges, and predicting the final judgment. We further enhance LLMs through fine-tuning with multi-task synthetic trajectories to improve legal judgment prediction accuracy and efficiency under our ADAPT framework. Extensive experiments conducted on two widely-used datasets demonstrate the superior performance of our framework in legal judgment prediction, particularly when dealing with complex and confusing charges.

1 Introduction

Legal judgment prediction (LJP) is a key research area within the legal natural language processing (NLP) community, aiming to provide automated reference judgments to help judges and other professionals manage cases more efficiently (Luo et al., 2017; Chalkidis et al., 2020; Niklaus et al., 2021). The main challenges in enhancing judgment prediction systems are twofold: understanding case facts and distinguishing between similar charges. Understanding case facts involves extracting key information from complex descriptions (Yue et al., 2021), while distinguishing between charges requires identifying the correct labels among confusing options (Xu et al., 2020). To tackle these issues, researchers have explored the use of advanced language models, the incorporation of external legal knowledge, and the reference to precedents to enhance model performance (Zhao et al., 2022).

Recently, large language models (LLMs) have achieved state-of-the-art performance across a range of tasks due to their expanded parameter scales and training data (Zhao et al., 2023; Zhu et al., 2023; Wang et al., 2024; Jiao et al., 2023). These models also display emergent abilities, such as instruction following and in-context learning, allowing them to quickly adapt to specific tasks based on minimal instructions or examples (Brown et al., 2020; OpenAI, 2023). Although preliminary evaluations of mainstream LLMs in legal judgment prediction have been conducted, results indicate that their performance still lags behind many traditional supervised methods, suggesting that the existing LLMs are still far from being good judgment predictors (Shui et al., 2023; Vats et al., 2023; Jiang and Yang, 2023).

Through our experiments, we find that current LLMs still struggle significantly more with distinguishing “confusing charges”, which refer to charges that have similar or even overlapping key behaviors with other charges. This highlights a critical aspect of the LJP task: *certain criminal behaviors can satisfy parts or even all of the conditions for multiple charges, creating ambiguity*. Due to the lack of extensive domain knowledge and reasoning training in the legal context, existing LLMs exhibit insufficient reasoning capability to effectively differentiate between similar charges.

To adapt LLMs for effective legal judgment prediction, we introduce the *Ask-Discriminate-PredicT* (ADAPT) reasoning framework in this paper. Our framework is inspired by the thought process of human judges, who use legal knowledge to navigate between facts and norms, as described by the classic phrase: “*the gaze shuttles back and forth between facts and norms*” (Rüthers et al., 2013). Specifically, In the first step, *Ask*, we decompose the noisy case description into multiple aspects under legal theory to clarify the key criminal facts. In the second step, *Discriminate*, the

*Corresponding author.

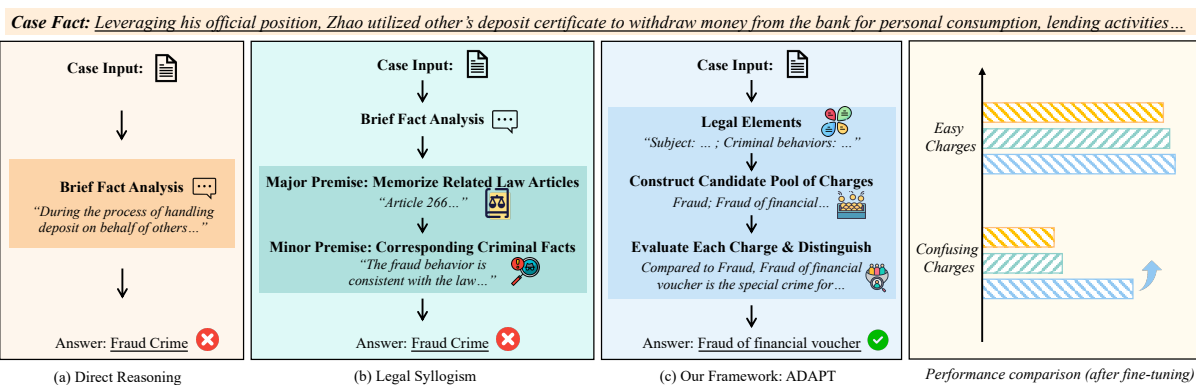


Figure 1: Comparison of our framework with direct reasoning and legal syllogism. We notice that our approach improves the performance on confusing charges more obviously after fine-tuning.

model uses its parameterized knowledge to generate a candidate pool of the most probable charges and relevant law articles. Within this pool, the model further differentiates among the candidates, assessing the degree of alignment between each candidate and the criminal facts. Finally, in the *Predict* step, the model synthesizes the previous reasoning process to provide the final prediction.

Experimental results show that our ADAPT prompting framework outperforms traditional prompting methods, such as direct reasoning and legal syllogism (Figure 1), in leveraging LLMs for legal judgment prediction. However, LLMs still struggle to consistently generate accurate ADAPT reasoning trajectories so as to finally make correct predictions. We hypothesize that this limitation is due to their narrow legal-specific knowledge and lack of familiarity with our specialized ADAPT reasoning patterns. Furthermore, we find that current LLMs frequently avoid providing sentencing ranges because of strict safety alignment protocols, which hinders their ability to complete this essential task in legal judgment prediction.

To address these issues, we further propose fine-tuning an enhanced LLM within our ADAPT framework for more comprehensive, efficient, and effective legal judgment prediction. Specifically, we strengthen the LLM by incorporating additional context labels—such as discriminative labels, charges, legal articles, and sentencing ranges—and prompt it to generate high-quality synthetic reasoning trajectories tailored to our ADAPT framework. We then use these multi-task synthetic trajectories to fine-tune a smaller LLM, enabling it to perform accurate reasoning under our ADAPT framework.

We conduct extensive experiments on two datasets, CAIL2018 and MultiLJP, which belong to single-defendant and multi-defendant scenarios,

respectively. The results show that our approach achieves new state-of-the-art in both scenarios, especially on the most challenging set of charges.

Our contributions are summarized as:

(1) We pinpoint that the underperformance of LLMs in legal judgment prediction primarily stems from their difficulty in distinguishing between confusing charges.

(2) We propose the ADAPT reasoning framework to emulate human judicial reasoning, which guides LLM to navigate between legal facts and norms to improve the accuracy of legal judgments.

(3) We fine-tune an enhanced LLM using knowledge distillation on multi-task synthetic trajectories to achieve more comprehensive, efficient, and effective legal judgment prediction under our ADAPT framework.

2 Related Work

Legal judgment prediction Legal judgment prediction is a long-standing legal NLP task. The evolution of this task’s technology has transitioned through various phases: initially relying on rule-based approaches (Nagel, 1963; Segal, 1984), advancing to statistical machine learning techniques (Katz et al., 2017; Sulea et al., 2017), and currently dominated by deep learning methodologies (Xu et al., 2020; Yue et al., 2021; Zhang et al., 2023a). Additionally, incorporating domain-specific legal knowledge (Zhao et al., 2022) or precedents (Zhang and Dou, 2023; Wu et al., 2023) is also an important direction in existing research. With the continuous advancement of methods, this task has expanded from a simplified multi-class classification problem to complex scenarios that mirror real-world situations, such as dealing with multiple defendants (Lyu et al., 2023) and multiple law articles (Liu et al., 2023). In this context, we

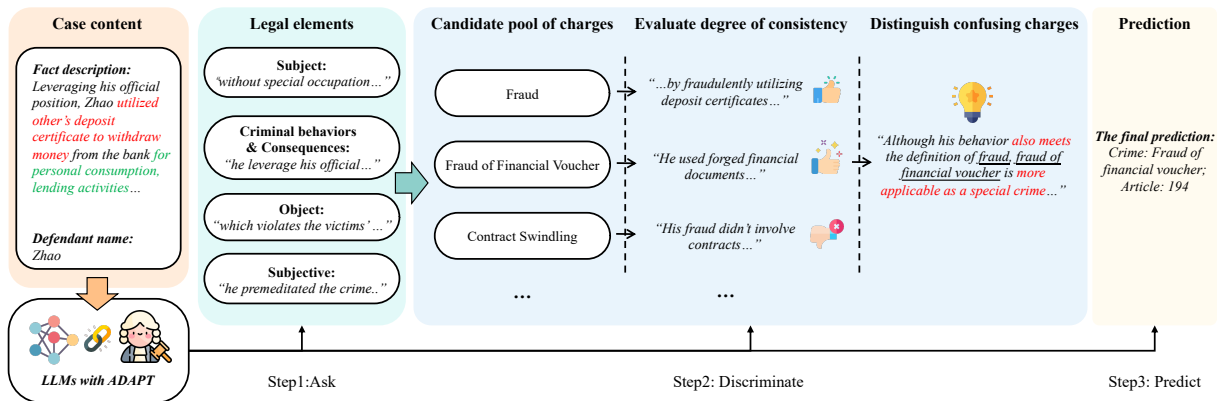


Figure 2: Overview of our framework. The final judgment is predicted based on three different reasoning steps.

explore the use of large language models as the foundation model and conduct robust reasoning under the setting of multi-label classification.

Reasoning skills in language models Recent studies have shown that effective reasoning can be achieved in LLMs by using prompting techniques (Wei et al., 2022; Press et al., 2023; Deng et al., 2024b). In the legal context, previous study shows that legal syllogism can enhance the performance of LJP (Jiang and Yang, 2023). Furthermore, some research attempts focus on distilling the reasoning processes of large models into smaller models, thereby achieving approximate reasoning capabilities at a lower cost (Ho et al., 2023; Mukherjee et al., 2023). Our approach synthesizes the trajectory of the ADAPT framework to fine-tune a 7B model, enabling it to achieve both robust and efficient reasoning.

3 Methodology

Existing LLMs mostly rely on direct answering or judicial syllogism for reasoning, which requires the model to directly provide the correct law articles and charges. This exceeds the capabilities of these language models and even human experts, leading to bad performance. In this work, we enable LLM to emulate the reasoning pattern of real-world judges to conduct discriminative reasoning for accurate legal judgment prediction.

3.1 Preliminaries

We first formally describe the workflow of legal judgment prediction. Given the criminal fact f from a real case, the name of one of the defendants, a set of charges \mathcal{C} , and a set of law articles \mathcal{A} , our task is to predict the applicable subset of charges \mathcal{C}_d , the relevant subset of law articles \mathcal{A}_d , and the

term of imprisonment for the defendant d . To align with legal practice, we follow recent studies that treat charge and law article prediction as multi-label classification tasks, and term of imprisonment prediction as a multi-class classification task.

3.2 Discriminative Reasoning Framework

We propose a discriminative reasoning framework, called ADAPT, to guide LLM to gradually deduce the most appropriate charges and law articles step by step, including *Ask*, *Discriminate*, and *Predict*. As shown in Figure 2, in the first step, *Ask*, the model is prompted to identify the key legal elements that constitute crimes through a question-answering approach. Then, in the second step, *Discriminate*, we prompt LLM to utilize the extracted key elements from the *Ask* step to initially identify the top- K most possible charges, and subsequently evaluate the consistency between each charge and the established facts. Finally, in the *Predict* step, LLM integrates the reasoning signals from the previous two steps to identify the most suitable charges and law articles. We describe the details of these three steps in the following.

Step1: Ask The objective of the *Ask* step is to clarify the key elements that constitute crimes. We use legal theory (Rüthers et al., 2013) to guide LLM to summarize four aspects from the facts (Deng et al., 2024a): (1) *Subject*, which refers to the defendant’s occupation and identity characteristics, such as state officials. (2) *Criminal behaviors and consequences*, which contains the defendant’s specific actions and resulting harm. (3) *Object*, which is the entities or legal interests violated by the criminal acts. (4) *Subjective aspect*, which is the psychological state of the defendant, such as direct purpose, negligence, and so on. Inspired by (Zhang et al., 2023b), we prompt in question-answering form for

accurate summarization.

Step2: Discriminate To avoid invalid reasoning caused by selecting from incorrect confusing charges, it is necessary to carefully distinguish candidate charges before making predictions. Specifically, we prompt the LLM to first provide several most likely candidate charges based on its parametric knowledge. Based on these candidates, the model then evaluates the consistency of each candidate with the key elements and distinguishes the main differences between these charges.

Step3: Predict By contextualizing the reasoning trajectories of the initial two steps, the LLM is prompted to predict the final judgment result.

3.3 Improving ADAPT with Fine-tuning

While our ADAPT prompting framework outperforms traditional prompting methods such as direct reasoning or legal syllogism, we find that current general LLMs still struggle to consistently generate fully accurate reasoning trajectories. This limitation arises from their restricted legal-specific knowledge and lack of exposure to our specialized ADAPT reasoning patterns. Additionally, we find that current LLMs often reject giving the sentencing ranges, likely because of their strict safety alignment. To address these issues, we propose fine-tuning a better LLM under our ADAPT framework for more comprehensive, efficient, and effective legal judgment prediction.

Synthetic trajectories generation we first generate synthetic ground-truth reasoning trajectories for the three steps of ADAPT using a larger model, specifically a 72B parameter LLM, which is provided with refined instructions and additional context labels, such as ground-truth discriminative labels, charges, and relevant legal articles. By incorporating these additional context labels, we find that the 72B LLM is able to generate highly accurate reasoning trajectories for each step. Subsequently, we use these high-quality synthetic reasoning trajectories to fine-tune a smaller LLM.

Multi-task instruction tuning We have five specific tasks in our fine-tuning uniformly using the language modeling loss function:

$$\mathcal{L}_{\text{task}} = -\frac{1}{T_{\text{task}}} \sum_{t=1}^{T_{\text{task}}} \log P_{\theta}(y_t^{\text{task}} | \mathcal{F}_{\text{task}}(x_{\text{task}}), y_{<t}^{\text{task}}),$$

where \mathcal{F} is the task-specific prompting function for formulating the input instruction. x , y , and T

are the input, target response, and the number of tokens in the response, respectively.

Specifically, the first two tasks are *ask* and *discriminate*, corresponding to the first two steps of ADAPT. The input x_{ask} consists of the criminal fact f and the specified defendant d . The input x_{disc} additionally contains y^{ask} , which is the target output of the *Ask* step generated by the 72B LLM.

The third task is *sentencing*, which is to improve the model’s perception of the sentencing factors. Its input x_{sent} consists of a set of charges C_d against the defendant d ; The fourth task is *article*, which is to improve the model comprehension of the correspondence between the case facts and the specified law articles. The model learns to recite the content of the given article numbers as well as explain in detail how the defendant’s actions align with these articles. Its input x_{article} contains the criminal fact f , the specified defendant d , and the article number. The training targets of these two tasks are also generated by employing the 72B LLM with additional context labels, including the articles and the sentencing ranges.

Finally, the last task, *predict_all*, is to contextualize all of the previous reasoning results and predict the final charges, articles, and sentencing ranges just in one prompt. Its input $x_{\text{predict_all}}$ contains the criminal fact f and the specified defendant d . Its training target $y_{\text{predict_all}}$ is the concatenation of the synthetic reasoning trajectories of the three steps of ADAPT and the sentence range labels.

For clarity, we show the task-specific prompting functions and all synthetic prompts of different tasks in Appendix A. We equally mix the training samples of different tasks to perform multi-task fine-tuning.

4 Experiments

4.1 Datasets and Evaluation Metrics

We conduct extensive experiments in both single-defendant and multi-defendant scenarios to comprehensively evaluate our method’s performance in real-world applications. In the single-defendant context, we employ the widely-used CAIL2018 dataset (Xiao et al., 2018). For the multi-defendant case, we select the MultiLJP (Lyu et al., 2023) dataset whose labels are verified by human experts. Both datasets are divided into 11 intervals to convert prison term prediction to a multi-class classification task. Detailed statistics of both datasets are provided in Table 1. For evaluation metrics, we

Dataset	CAIL2018	MultiLJP
# Train cases	118,399	18,968
# Test cases	1,120	2,370
# Charges	191	23
# Articles	162	22
# Intervals of prison term	11	11
# Average criminal per case	1	3.71
Average length per case	440.9	3,040.8

Table 1: Basic statistics of the datasets.

follow previous works to adopt Accuracy (Acc.), Macro Precision (Ma-P), Macro Recall (Ma-R), and Macro F1 (Ma-F) across all sub-tasks.

4.2 Baselines

Fine-tuning methods. We categorize the fine-tuning baselines according to their characteristics as follows: (1) *Topological Relationships*: **Top-Judge** (Zhong et al., 2018) explicitly models the dependency relationships among the three sub-tasks in the prediction workflow. (2) *Graph-related Modeling*: **LADAN** (Xu et al., 2020) designs a graph distillation module to distinguish confusing law articles. (3) *Fact Decomposition*: **NeurJudge** (Yue et al., 2021) decomposes textual fact into different representations for each sub-task. (4) *Different Pre-trained Language Models*: For the “Text-to-Class” style, we select **BERT** (Devlin et al., 2019) and **Lawformer** (Xiao et al., 2021), while for the “Text-to-Text” style, we choose **mT5** (Xue et al., 2021) as the backbone model. (5) *Hierarchical Reasoning*: **HRN** (Lyu et al., 2023) improves predictions by learning intermediate reasoning steps, but this also restricts its evaluation on the MultiLJP dataset with corresponding annotations. (6) *LLM-based Fine-tuning*: **Vanilla-SFT** processes training data into a unified chat template for fine-tuning. **Finetune-CoT** (Ho et al., 2023) initially generates Chain-of-Thought trajectories for each training data, then finetune the base model with the synthesized data. In addition to the above methods, we also consider approaches such as **CL4LJP** (Zhang et al., 2023a) and **CECP** (Zhao et al., 2022). However, these methods are excluded from our main experiment due to their limited adaptability to multi-label classification.

Prompting setting We employ two types of models (1) *Law Specific LLMs*: We select **Disc-LawLLM** (Yue et al., 2023) as the representative, which undergo supervised fine-tuning with high-quality task data from both legal and general sce-

narios. (2) *General-Purpose LLMs*: We choose **Qwen2-[7B, 72B]-Instruct** (Bai et al., 2023) to investigate performance across different models and scales. For each model, we evaluate the performance of different prompting methods under both zero-shot and few-shot settings.

4.3 Implementation Details

For all LLM-based fine-tuning methods, we utilize Qwen2-7B as the foundation model. LoRA is adopted for parameter-efficient fine-tuning of the large language model. We apply LoRA to all linear modules of the model, with both alpha and rank set to 32. The language modeling head is also unfrozen to enhance learning. Our model is fine-tuned by 10 epochs, with a learning rate of $5e-5$ and a batch size of 64. The total number of training reasoning trajectories for CAIL2018 and MultiLJP is 80,141 and 157,763, respectively. Greedy decoding is used for all generative models to enhance the stability of results. For those generated charges that are not present in the label pool, we use BGE (Xiao et al., 2023) to map them to the closest charge in the pool based on their representations.

4.4 Evaluation on the Fine-tuning Setting

The results on the fine-tuning setting are presented in Table 2. We have the following findings:

(1) Our method outperforms all baselines across all metrics on both datasets. Overall, the LLM-based approaches show superior performance, indicating that large causal language models can adapt effectively to the tasks of legal judgment prediction with targeted fine-tuning. On the CAIL2018 dataset, our approach achieves relative accuracy and Ma-F improvements of 4.1% and 4.7%, respectively, over Finetune-CoT. This demonstrates that our synthetic data significantly enhances the effectiveness of legal judgment prediction.

(2) The ADAPT framework achieves a more significant advantage in charge and law article prediction. These two tasks are both multi-label classification tasks, meaning a successful shot satisfies the predicted set to be completely consistent with the label set. Our improvements suggest that discriminative reasoning offers more evidence for the precise inference of the target set. And on the other side, the performance across methods is relatively close in prison term prediction. We hypothesize that this is due to the uncertainty introduced by real-world judges’ discretion, which cannot be fully mitigated by classification metrics based on rigid

Methods	CAIL2018						MultiLJP					
	Charge		Law Article		Prison Term		Charge		Law Article		Prison Term	
	Acc.	Ma-F	Acc.	Ma-F	Acc.	Ma-F	Acc.	Ma-F	Acc.	Ma-F	Acc.	Ma-F
TopJudge	65.5	74.1	68.2	74.3	32.1	32.4	67.6	55.7	73.9	54.1	36.1	33.1
LADAN	63.1	71.8	62.5	71.0	30.1	31.2	60.4	43.2	68.2	49.0	35.1	34.6
NeurJudge	65.7	71.4	67.4	70.9	29.6	33.2	64.8	51.2	71.8	55.7	33.9	32.0
BERT	64.6	74.6	68.3	73.5	31.7	33.5	66.3	54.2	73.6	54.0	35.6	32.9
Lawformer	66.2	73.1	67.5	74.4	30.4	30.7	68.1	53.8	76.2	53.8	36.1	34.7
mT5	72.3	77.5	73.2	74.4	33.9	30.8	78.4	44.6	82.9	44.1	30.7	20.3
HRN	-	-	-	-	-	-	83.5	60.9	84.3	62.1	34.3	33.4
<i>LLM-based Fine-tuning</i>												
Vanilla-SFT	74.1	78.6	74.0	75.5	32.0	31.3	85.4	65.2	87.7	63.5	32.0	31.3
Finetune-CoT	74.8	79.3	75.6	77.7	31.5	31.9	86.2	66.7	88.0	64.8	32.4	32.7
ADAPT (Ours)	77.9	83.0	78.3	80.0	37.9	35.8	90.3	73.1	91.1	75.4	37.3	35.2

Table 2: Experimental results on the fine-tuning setting. The best results are in bold.

Model	Params.	Demos.	Charge				Law Article			
			Acc.	Ma-P	Ma-R	Ma-F	Acc.	Ma-P	Ma-R	Ma-F
<i>In-domain LLMs</i>										
Disc-LawLLM	13B	✗	44.2	59.7	61.8	56.6	55.0	54.5	70.5	57.7
<i>General-purpose LLMs</i>										
Qwen2-7B	7B	✗	41.7	<u>56.3</u>	58.6	<u>53.0</u>	<u>50.4</u>	<u>51.9</u>	<u>64.7</u>	<u>52.8</u>
+ Few-shot	-	✓	44.4	55.1	56.9	50.7	48.7	49.8	60.6	48.7
+ CoT	-	✓	45.7	54.7	<u>58.7</u>	52.1	49.4	47.6	60.2	47.5
+ ADAPT	-	✓	<u>45.0</u>	59.2	59.4	55.2	53.3	56.0	64.8	55.6
Qwen2-72B	72B	✗	56.2	60.9	<u>72.4</u>	<u>63.2</u>	57.7	57.2	<u>71.3</u>	59.4
+ Few-shot	-	✓	<u>57.4</u>	61.0	69.8	61.4	58.6	54.3	68.5	57.9
+ CoT	-	✓	56.9	<u>61.4</u>	71.4	62.2	62.9	<u>58.6</u>	73.2	<u>60.3</u>
+ ADAPT	-	✓	58.4	62.3	73.3	65.0	<u>59.7</u>	60.3	70.4	61.5

Table 3: Performance on the prompting setting. The best results and the second-best results of each setting are in bold and underlined, respectively.

interval divisions.

(3) Fine-tuning with discriminative reasoning trajectories further enhances the language models’ performance. Finetune-CoT similarly utilizes synthetic reasoning data to promote autoregressive learning. However, the synthesized CoT data is observed to result in sub-optimal performance. We argue that this is because the “step-by-step reasoning” generated by the language model does not contribute to more supporting logic. Instead, it merely discusses the plain facts in most cases.

4.5 Evaluation on the Prompting Setting

We evaluate under the prompting setting using the CAIL2018 dataset, as the MultiLJP dataset contains only 23 charges and can not reflect the effect of confusing charges in real-world scenarios.

During our experiments, we observed that the in-domain model fails to effectively adhere to instructions for few-shot learning, so we only tested its zero-shot ability. Moreover, we discard the prison term prediction task because LMs typically refuse to predict the accurate terms. The results are reported in Table 3, from which we can observe:

(1) Our approach generally demonstrates superior performance across models of different sizes, suggesting that prompts can still activate discriminative reasoning to some extent. Moreover, we also observe that the improvements are not pronounced as in the fine-tuning setting. This indicates that fine-tuning can further enhance the model’s capabilities within our framework.

(2) Larger models typically bring better performance, but they still exhibit a notable disparity

Model	Charge				Law Article				Prison Term			
	Acc.	Ma-P	Ma-R	Ma-F	Acc.	Ma-P	Ma-R	Ma-F	Acc.	Ma-P	Ma-R	Ma-F
<i>w/o Ask</i>	76.3	81.4	84.5	81.5	76.4	78.5	81.4	79.2	35.7	33.6	35.3	33.8
<i>w/o Disc</i>	76.8	81.2	85.0	81.8	76.6	78.2	82.1	79.0	35.3	33.0	35.7	33.5
<i>w/o Article</i>	75.7	80.3	81.9	80.4	76.0	78.1	80.7	78.5	34.9	32.6	33.9	32.4
<i>w/o Sentencing</i>	77.6	82.0	86.4	82.7	78.0	79.6	83.0	79.7	36.0	34.0	35.6	34.1
<i>ADAPT→Refine</i>	70.2	74.4	75.3	74.8	69.5	72.1	77.1	74.7	28.9	29.4	30.7	29.8
ADAPT (Ours)	77.9	82.3	86.9	83.0	78.3	79.6	83.5	80.0	37.9	35.6	37.2	35.8

Table 4: Ablation results on CAIL2018. The best results are in bold.




Figure 3: Fine-tuning performance of each sub-task with epochs on the CAIL2018 dataset.

when compared to models trained for LJP specifically. This might be because recent leading open-source models have been trained on extensive data from major domains. This suggests that fine-tuning remains a valuable strategy for LLMs to adapt to the requirements of LJP tasks

4.6 Ablation Study

We investigate the effects of various reasoning data on the final LJP tasks within the fine-tuning context. The specific ablation strategies are described as follows: (1) *w/o Ask*: The task of summarizing legal elements from facts is removed, and y^{Ask} is excluded from y^{All} . (2) *w/o Disc*: After the *Ask* step, the language model must directly predict the charges and law articles. (3) *w/o Article*: We remove the law article-related reasoning trajectories in the training data. (4) *w/o Sentencing*: The language model no longer analyzes sentencing factors before determining the prison term prediction. (5) *ADAPT→Refine*: We construct candidate items for charges and law articles and provide them to the large language model, requiring it to refine them and determine the final prediction. During inference, the candidates are derived from the top-k items in the probability distribution of BERT, while during training, the correct labels are ensured to be included among these candidates.

The ablation results are shown in Table 4. Firstly, we can observe the removal of each sub-task leads to a performance decline, indicating that each type

of synthetic data makes a positive contribution. Additionally, we notice that the ablation associated with law articles causes obvious impact. This impact may be attributed to the better alignment between legal provisions and facts of real cases, which enhances legal reasoning abilities in other forms. Finally, it is observed that the ablation related to sentencing almost exclusively affects prison term prediction. This can likely be attributed to the fact that sentencing factors and conviction factors are orthogonal in most cases.

Moreover, we find that *ADAPT→Refine* results in a significant performance decline. This is because the top-k candidates provided by the domain model (i.e., BERT) during inference often do not include the ground truth labels, thus the LLM can only select relatively close labels in such cases. We think that the fine-tuned LLMs are inherently capable of generating high-quality candidates and don't require assistance from smaller external models.

4.7 Effect of Training Epochs

Instruction tuning has been observed to lead to performance degradation with an increasing number of training epochs. However, previous studies demonstrate that it often requires many epochs (e.g., 20 epochs) for model performance to coverage in LJP tasks. We investigate the effects of training epochs for the ADAPT framework in this section. The results, presented in Figure 3, reveal that:

- (1) In the predictions of charges and law articles,




Figure 4: Performance in different charge subgroups of various methods. Intervals of larger numbers (e.g., 75%-100%) indicate greater difficulty of the corresponding subgroup.

metrics gradually increase with the number of training epochs and stabilize after reaching a peak at a certain epoch. This indicates that the model can continuously learn effective features from the data even after the initial complete iteration.

(2) The metric curves for different sub-tasks exhibit notable distinctions. Unlike the other two tasks, most of the metrics for prison term prediction show a marked decline after the third epoch. This divergence highlights the inherent difficulty in predicting prison terms. Judicial discretion in sentencing can span multiple pre-defined intervals, thereby increasing the risk of the model overfitting when learning from repeated samples.

(3) In multi-label classification tasks (i.e., charge and law article prediction), the model consistently exhibits higher performance in the Macro-Recall compared to the other three metrics. This indicates that the fine-tuned language model tends to identify more possible positive candidates. We believe this also suggests the potential for further refining results in our proposed reasoning framework.

4.8 Performance on Different Difficulty

Exploring performance across different charges can help us better understand the detailed improvements achieved by our method. We first calculate the F1 scores of the finetuned BERT model for each charge in the CAIL2018 dataset and rank these charges from highest to lowest. For clearer visualization, we categorize the ranked charges into four sets and then evaluate the macro-F1 scores of all finetuned models on the charge prediction task. Generally, the charges in the higher quartiles (e.g., 75%-100%) exhibit greater prediction difficulty, as evidenced by the poor performance of mainstream "Text-to-Label" style models on these charges. The experimental results are shown in Figure 4, from which we can observe the following findings:

(1) Our ADAPT framework achieves a more sig-

nificant improvement on difficult sets. Specifically, in the 75%-100% interval, ADAPT achieves relative improvements of 15.1% and 16.7% compared to mT5 and Vanilla-SFT, respectively. Conversely, in the 0%-25% interval, the relative improvements are merely 6.5% and 1.4%. This suggests that our discriminative reasoning approach effectively delineates subtle differences between various charges, thereby significantly enhancing prediction accuracy for more confusing charges.

(2) The marginal benefits of employing simple instruction fine-tuning on larger models are limited. We can observe that despite Vanilla-SFT leveraging a language model with 7B parameters, its improvements over mT5 are not substantial. Notably, in cases involving more difficult charges, Vanilla-SFT is even likely to demonstrate a slight decline in performance. This finding highlights the importance of identifying an appropriate reasoning pattern to enhance the effectiveness of large language models in legal judgment prediction.

5 Conclusion

In this paper, we presented a novel framework to enable discriminative reasoning in LLMs for legal judgment prediction. Our ADAPT framework effectively distinguishes confusing charges by determining the degree of alignment between each candidate charge and the criminal facts before the final prediction. Furthermore, we utilize multi-task instruction tuning on synthetic data to enhance the language model's comprehension of this reasoning pattern. Extensive experiments demonstrate the effectiveness of our approach, particularly its robustness in handling confusing charges. We believe that our work will improve the integration of LLMs in legal judgment prediction and contribute to the community's understanding of reasoning patterns in specific tasks.

Limitations

Despite the promising results that have been demonstrated in our framework, several limitations must be acknowledged:

Cost of Synthesized Data. Our method requires the synthesis of reasoning trajectories from existing judgment data to facilitate fine-tuning, potentially leading to increased computational expenses or API costs. Fortunately, the cost of utilizing large language models is rapidly decreasing. For instance, recent inference services such as Qwen and DeepSeek require less than \$0.0001 per 1,000 tokens. Therefore, we believe that generating reasoning data at this scale is entirely acceptable.

Limited Scope of Open Datasets. Our framework demonstrates strong generalization capabilities for the case types it was trained on. However, the most diverse dataset we employed, CAIL2018, encompasses fewer than 200 distinct charges, potentially limiting its applicability in all real-world scenarios. Consequently, we recommend further synthesizing more comprehensive reasoning trajectories using public or private domain data to meet the specific needs of real-world applications.

Potential Dataset Leakage Risks. Although the large language models utilized in our experiments are open-source, the datasets employed during their training are not entirely transparent, potentially posing risks of data leakage. Our solution is to evaluate various methods on the same foundation model to ensure fair comparisons. The relative improvements in various settings can prove our advantage.

Ethical Discussion

Privacy and Data Security. Legal data often includes sensitive and confidential information about individuals and entities. Mishandling this data can lead to serious privacy breaches. The two datasets adopted in our experiments are robustly anonymized to protect this information.

Potential Bias in Training Data. Large language models may learn bias from judgments in the training set. In a real-world case, the final judgment can be affected by some factors like social comments, the judge's style, etc. We need to test and identify possible biases before application.

Legal and Ethical Compliance. Adhering to existing legal and ethical standards is essential when

deploying LLMs for legal judgment prediction. We advise users to critically evaluate the model's predictions and make independent decisions about their adoption, rather than uncritically accepting the machine's reasoning.

Acknowledgement

Zhicheng Dou is the corresponding author. This work was supported by the National Science and Technology Major Project No. 2022ZD0120103, National Natural Science Foundation of China No.62272467, the fund for building world-class universities (disciplines) of Renmin University of China, Public Computing Cloud of RUC, and School of Interdisciplinary Studies of RUC. The work was partially done at the Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE.

References

- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen technical report](#). *CoRR*, abs/2309.16609.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. [LEGAL-BERT: "preparing the muppets for court"](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 2898–2904. Association for Computational Linguistics.

- Chenlong Deng, Zhicheng Dou, Yujia Zhou, Peitian Zhang, and Kelong Mao. 2024a. [An element is worth a thousand words: Enhancing legal case retrieval by incorporating legal elements](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 2354–2365. Association for Computational Linguistics.
- Chenlong Deng, Kelong Mao, and Zhicheng Dou. 2024b. [Learning interpretable legal case retrieval via knowledge-guided case reformulation](#). *CoRR*, abs/2406.19760.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023. [Large language models are reasoning teachers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 14852–14882. Association for Computational Linguistics.
- Cong Jiang and Xiaolei Yang. 2023. [Legal syllogism prompting: Teaching large language models for legal judgment prediction](#). In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law, ICAIL 2023, Braga, Portugal, June 19-23, 2023*, pages 417–421. ACM.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. [Is chatgpt A good translator? A preliminary study](#). *CoRR*, abs/2301.08745.
- Daniel Martin Katz, Michael J Bommarito, and Josh Blackman. 2017. A general approach for predicting the behavior of the supreme court of the united states. *PloS one*, 12(4):e0174698.
- Yifei Liu, Yiquan Wu, Yating Zhang, Changlong Sun, Weiming Lu, Fei Wu, and Kun Kuang. 2023. [ML-LJP: multi-law aware legal judgment prediction](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*, pages 1023–1034. ACM.
- Bingfeng Luo, Yansong Feng, Jianbo Xu, Xiang Zhang, and Dongyan Zhao. 2017. [Learning to predict charges for criminal cases with legal basis](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2727–2736. Association for Computational Linguistics.
- Youngang Lyu, Jitai Hao, Zihan Wang, Kai Zhao, Shen Gao, Pengjie Ren, Zhumin Chen, Fang Wang, and Zhaochun Ren. 2023. [Multi-defendant legal judgment prediction via hierarchical reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 2198–2209. Association for Computational Linguistics.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive learning from complex explanation traces of GPT-4](#). *CoRR*, abs/2306.02707.
- Stuart S Nagel. 1963. Applying correlation analysis to case prediction. *Tex. L. Rev.*, 42:1006.
- Joel Niklaus, Ilias Chalkidis, and Matthias Stürmer. 2021. [Swiss-judgment-prediction: A multilingual legal judgment prediction benchmark](#). In *Proceedings of the Natural Legal Language Processing Workshop 2021, NLLP@EMNLP 2021, Punta Cana, Dominican Republic, November 10, 2021*, pages 19–35. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A. Smith, and Mike Lewis. 2023. [Measuring and narrowing the compositionality gap in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 5687–5711. Association for Computational Linguistics.
- Bernd Rüthers, Christian Fischer, and Axel Birk. 2013. *Rechtstheorie mit juristischer Methodenlehre*. Beck.
- Jeffrey A Segal. 1984. Predicting supreme court cases probabilistically: The search and seizure cases, 1962-1981. *American Political Science Review*, 78(4):891–900.
- Ruihao Shui, Yixin Cao, Xiang Wang, and Tat-Seng Chua. 2023. [A comprehensive evaluation of large language models on legal judgment prediction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 7337–7348. Association for Computational Linguistics.
- Octavia-Maria Sulea, Marcos Zampieri, Shervin Malmasi, Mihaela Vela, Liviu P. Dinu, and Josef van Genabith. 2017. [Exploring the use of text classification in the legal domain](#). In *Proceedings of the Second Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 16th International Conference on Artificial Intelligence and Law (ICAIL 2017), London, UK, June 16, 2017*, volume 2143 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Shaurya Vats, Atharva Zope, Somsubhra De, Anurag Sharma, Upal Bhattacharya, Shubham Kumar Nigam,

- Shouvik Kumar Guha, Koustav Rudra, and Kripabandhu Ghosh. 2023. [Llms - the good, the bad or the indispensable?: A use case on legal statute prediction and legal judgment prediction on indian court cases](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 12451–12474. Association for Computational Linguistics.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024. [A survey on large language model based autonomous agents](#). *Frontiers Comput. Sci.*, 18(6):186345.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Yiquan Wu, Siying Zhou, Yifei Liu, Weiming Lu, Xiaozhong Liu, Yating Zhang, Changlong Sun, Fei Wu, and Kun Kuang. 2023. [Precedent-enhanced legal judgment prediction with LLM and domain-model collaboration](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12060–12075. Association for Computational Linguistics.
- Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. [Lawformer: A pre-trained language model for chinese legal long documents](#). *AI Open*, 2:79–84.
- Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei Han, Zhen Hu, Heng Wang, and Jianfeng Xu. 2018. [CAIL2018: A large-scale legal dataset for judgment prediction](#). *CoRR*, abs/1807.02478.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighof. 2023. [C-pack: Packaged resources to advance general chinese embedding](#). *CoRR*, abs/2309.07597.
- Nuo Xu, Pinghui Wang, Long Chen, Li Pan, Xiaoyan Wang, and Junzhou Zhao. 2020. [Distinguish confusing law articles for legal judgment prediction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3086–3095. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mt5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 483–498. Association for Computational Linguistics.
- Linan Yue, Qi Liu, Binbin Jin, Han Wu, Kai Zhang, Yanqing An, Mingyue Cheng, Biao Yin, and Dayong Wu. 2021. [Neurjudge: A circumstance-aware neural framework for legal judgment prediction](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 973–982. ACM.
- Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, Chenchen Shen, Shujun Liu, Yuxuan Zhou, Yao Xiao, Song Yun, Xuanjing Huang, and Zhongyu Wei. 2023. [Disc-lawllm: Fine-tuning large language models for intelligent legal services](#).
- Han Zhang and Zhicheng Dou. 2023. [Case retrieval for legal judgment prediction in legal artificial intelligence](#). In *Chinese Computational Linguistics - 22nd China National Conference, CCL 2023, Harbin, China, August 3-5, 2023, Proceedings*, volume 14232 of *Lecture Notes in Computer Science*, pages 434–448. Springer.
- Han Zhang, Zhicheng Dou, Yutao Zhu, and Ji-Rong Wen. 2023a. [Contrastive learning for legal judgment prediction](#). *ACM Trans. Inf. Syst.*, 41(4):113:1–113:25.
- Kai Zhang, Bernal Jiménez Gutiérrez, and Yu Su. 2023b. [Aligning instruction tasks unlocks large language models as zero-shot relation extractors](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 794–812.
- Jie Zhao, Ziyu Guan, Cai Xu, Wei Zhao, and Enze Chen. 2022. [Charge prediction by constitutive elements matching of crimes](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 4517–4523. ijcai.org.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.
- Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. 2018. [Legal judgment prediction via topological learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3540–3549. Association for Computational Linguistics.
- Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Zhicheng Dou, and Ji-Rong Wen. 2023. [Large language models for information retrieval: A survey](#). *CoRR*, abs/2308.07107.

Task 1: Ask
Please use legal theory to analyze the defendant `{{defendant}}` in the above case in terms of the subject of the behavior, criminal behavior and results, object of the behavior, and subjective.

Task 2: Discriminate
First, you need to give a set of possible charges the defendant `{{defendant}}` would be accused of. Then, you need to carefully distinguish the differences among these charges and conclude that the most appropriate charges.

Task 3: Article
Please list the specific content of the legal provisions applicable to the defendant `{{defendant}}` in the above cases, and explain the reasons for applying these law articles.

Task 4: Sentencing
Please analyze all the possible sentencing factors of the defendant `{{defendant}}` in the above case.

Task 5: Predict_all
Please predict the charges, applicable law articles, and prison term using the ADAPT framework.

Figure 5: Prompt for each task of our multi-task instruction tuning.

Appendix

A Prompts

In this section, we provide prompts for multi-task instruction tuning and synthesizing data, respectively. The text in pink denotes the input information.

System Prompt:
You are a helpful assistant.

Context:
`{{fact description}}`

User Instruction:
Please answer the following four questions about the defendant `{{defendant}}`:

Q1. Subject of behavior: Please infer and describe whether the defendant has a special professional identity, such as being a manager of a company, a public official, etc. If there is no special professional identity, you should also make it clear that the person did not mention a special professional identity.

Q2. Criminal Behavior and results: Please summarize the specific behavior of the defendant that constitutes the crime and the results caused.

Q3. Object of behavior: Please infer and describe the object of the criminal behavior, which can be the rights and interests of the state, society, collective or individual. If it is an individual, you need to describe the specific characteristics of the person violated and the specific rights violated. For example, the object of the crime of trafficking in women and children must be a woman or a child. Different crimes violate different objects of behavior. For example, the crime of theft violates property ownership, and the crime of intentional injury violates personal rights.

Q4. Subjective: Please infer and describe the psychological state of the criminal when committing the crime, mainly including intent and negligence. Intention is divided into direct intent (hoping for the result to happen) and indirect intent (letting the result happen). Negligence is divided into negligent negligence (failure to foresee the consequences due to negligence) and overconfident negligence (foreseen the consequences but believed they could be avoided).

Output Format:

1. **Subject of behavior:** [The answer of Q1]
2. **Criminal Behavior and results:** [The answer of Q2]
3. **Object of behavior:** [The answer of Q3]
4. **Subjective:** [The answer of Q4]

System Prompt:
You are a helpful assistant.

Context:
`{{fact description}}`
`{{related law articles and legal provisions}}`

User Instruction:
First, you need to analyze if the defendant `{{defendant}}` would be accused of each charges in `{{all possible charges}}` based on the above law articles and legal provisions. Then, you need to carefully distinguish the differences among these charges and conclude that the most appropriate charges to be accused are `{{charge labels}}`

Figure 7: Prompt for synthesizing the trajectory of the step *Discriminate*.

System Prompt:
You are a helpful assistant.

Context:
{{fact description}}
{{related law articles and legal provisions}}

User Instruction:
Please analyze why the defendant {{defendant}} can be accused by using the law article {{law article number}}.

Figure 8: Prompt for synthesizing the trajectory of the task *Article*.

System Prompt:
You are a helpful assistant.

Context:
{{fact description}}
{{collected legal provisions about sentencing}}
{{charge labels}}

User Instruction:
Please analyze the sentencing factors can be applied to the defendant {{defendant}} according to the above information.

Figure 9: Prompt for synthesizing the trajectory of the task *Sentencing*.