

KnowProxy: Adapting Large Language Models by Knowledge-guided Proxy

Gukhyeon Lee*, Yeachan Kim*, SangKeun Lee

(* Equal contribution)

ICLR 2026



Introduction

Adapting large language models (LLMs) is expensive or impossible.

[Open-source LLMs]



> **Massive** computational **costs**

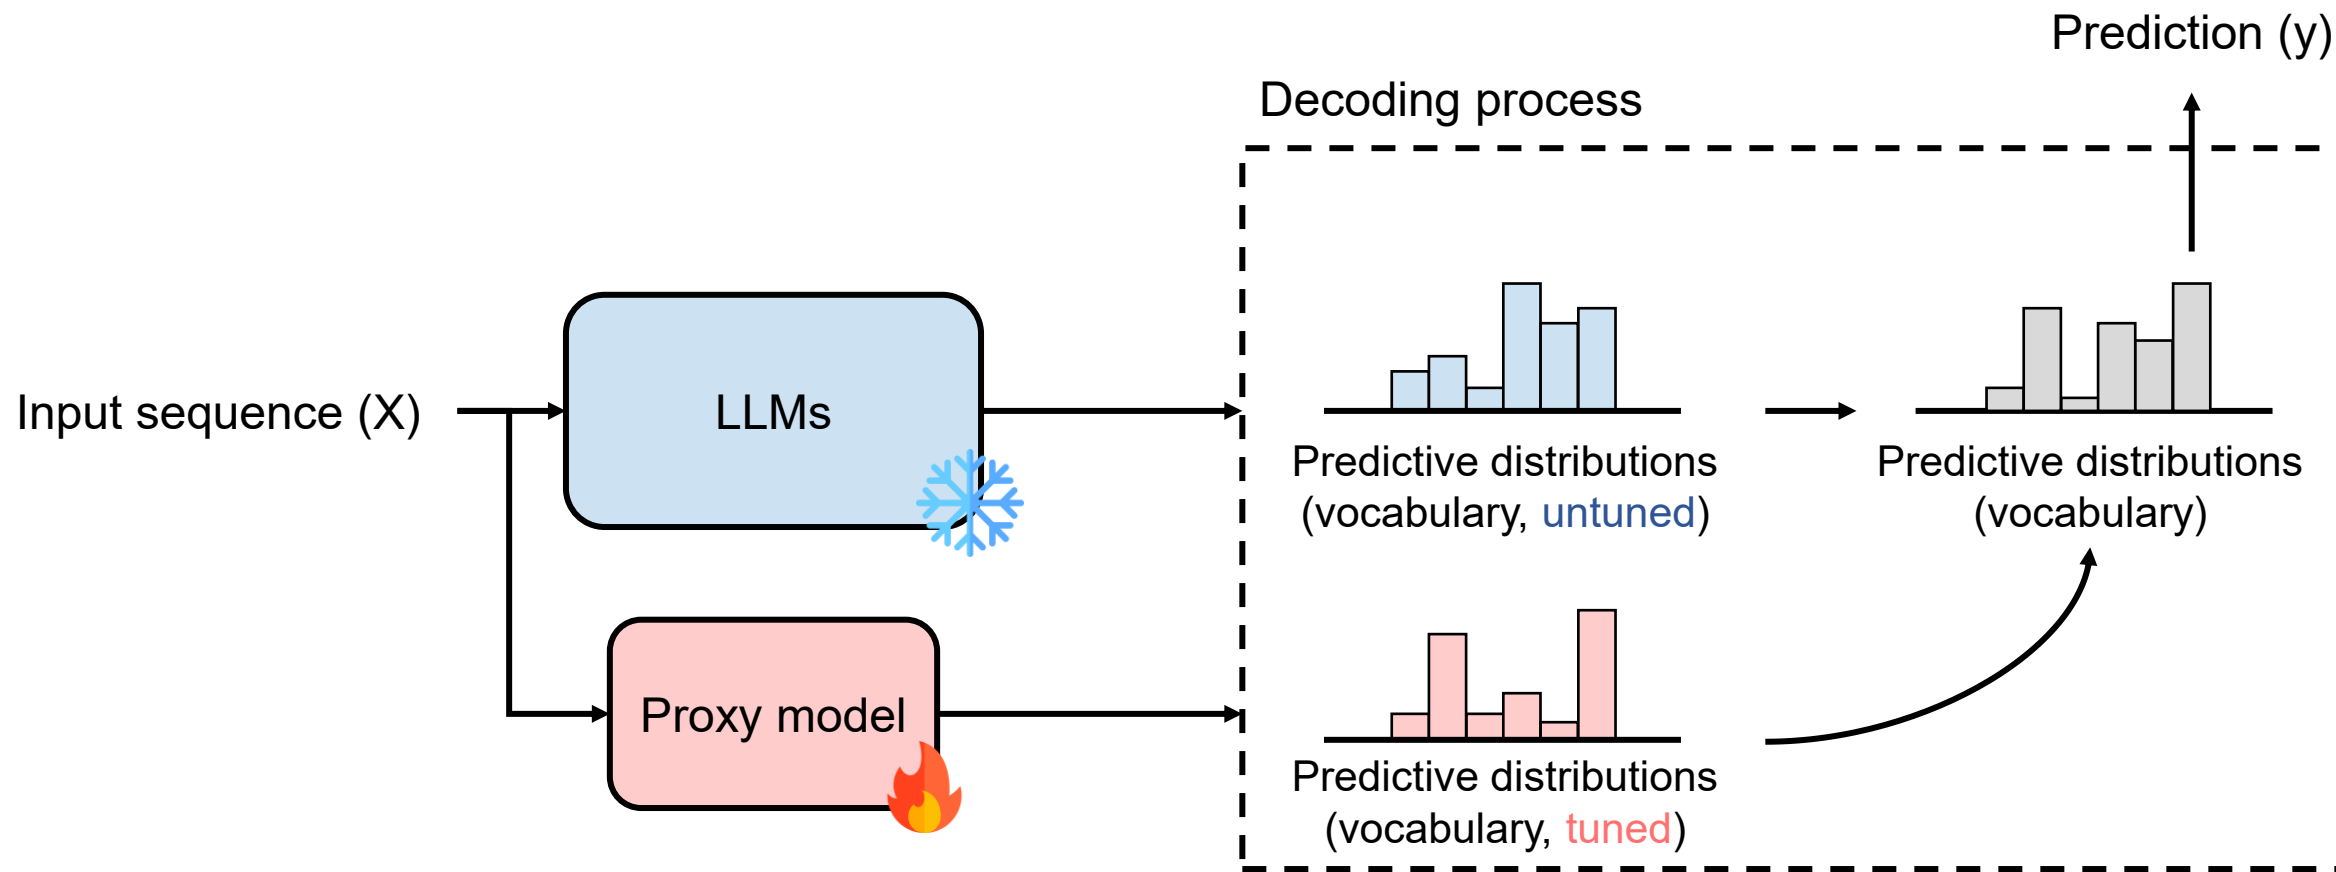
[API-based LLMs]



> **Inaccessible** to black-box LLMs

Related Works: Proxy-based training

Fine-tuning approach based on the predictive distributions from LLMs^[1, 2].



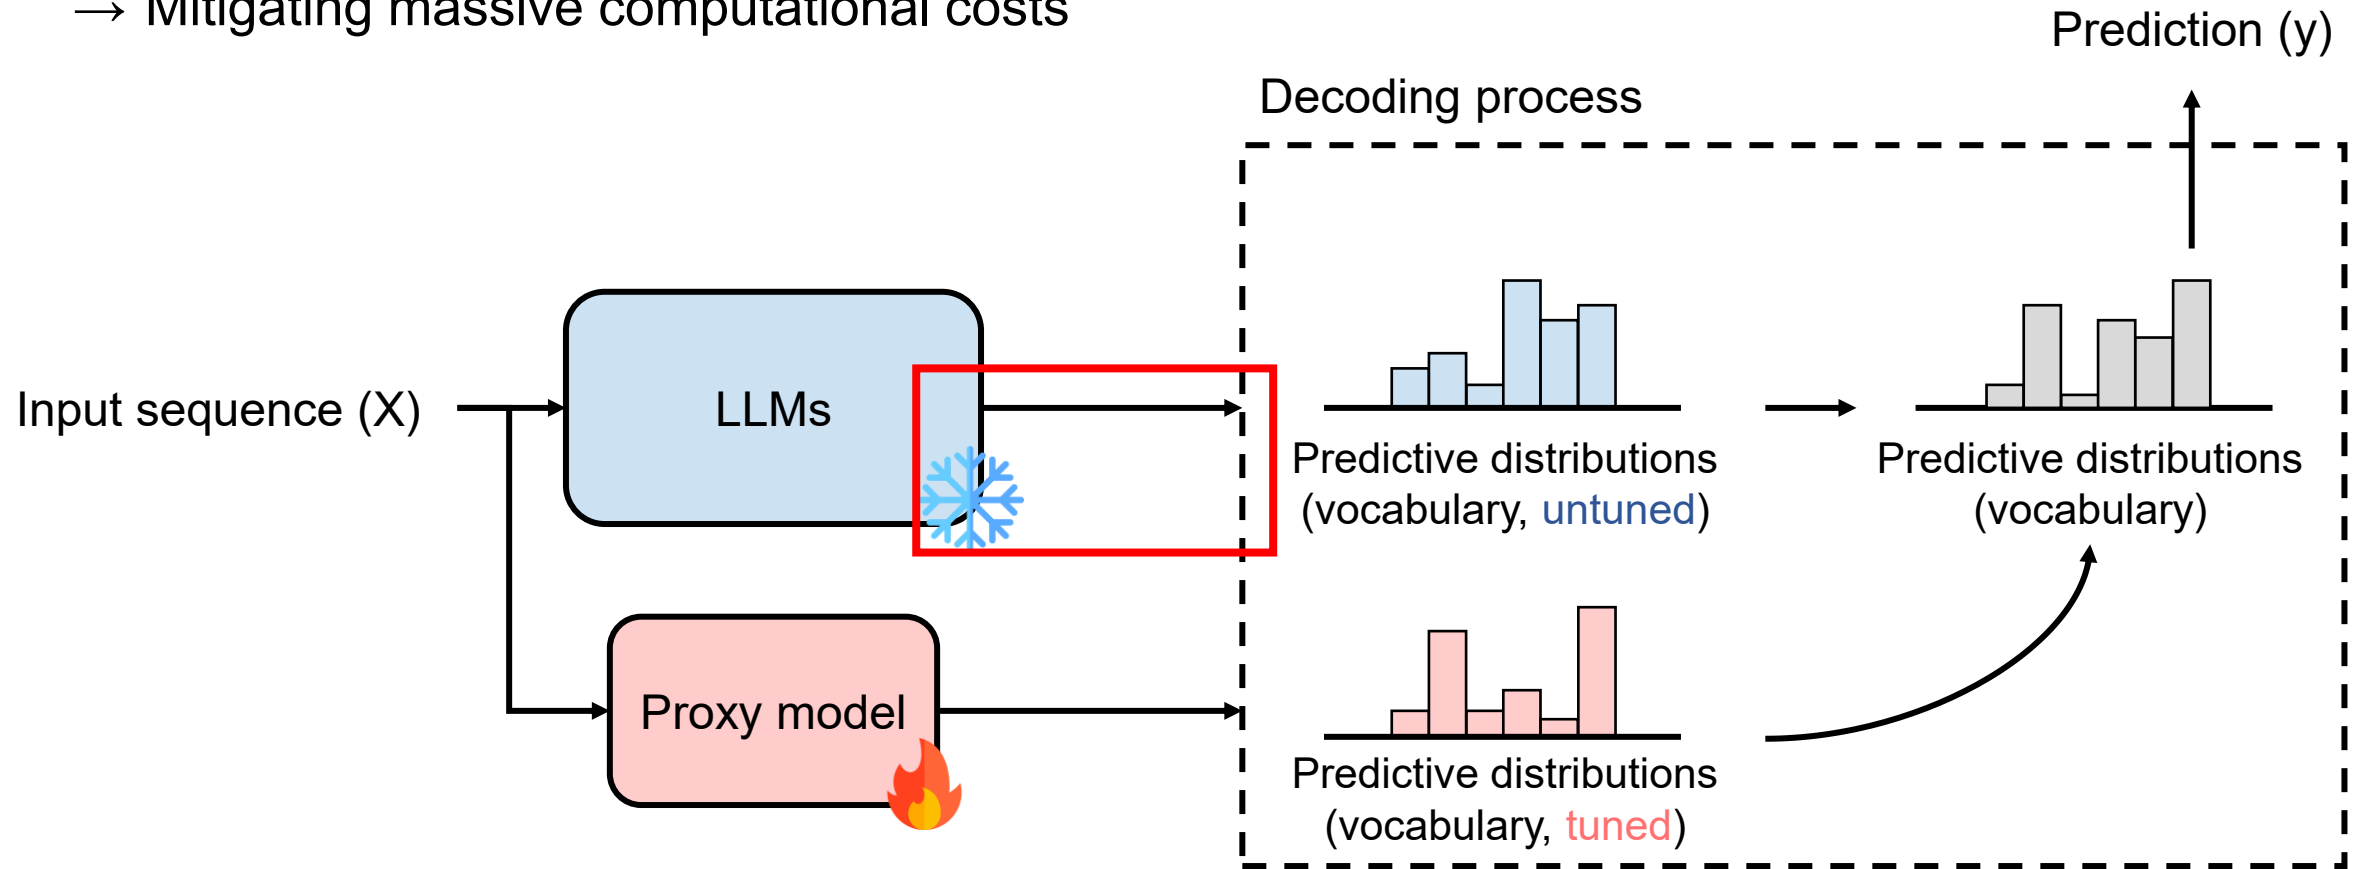
[1] Liu et al., Tuning Language Models by Proxy, COLM 2024.

[2] Ormazabal et al., CombLM: Adapting Black-Box Language Models through Small Fine-Tuned Models, EMNLP 2023.

Related Works: Proxy-based training

Fine-tuning approach based on the predictive distributions from LLMs^[1, 2].

- **No back-propagation** process through LLMs
→ Mitigating massive computational costs



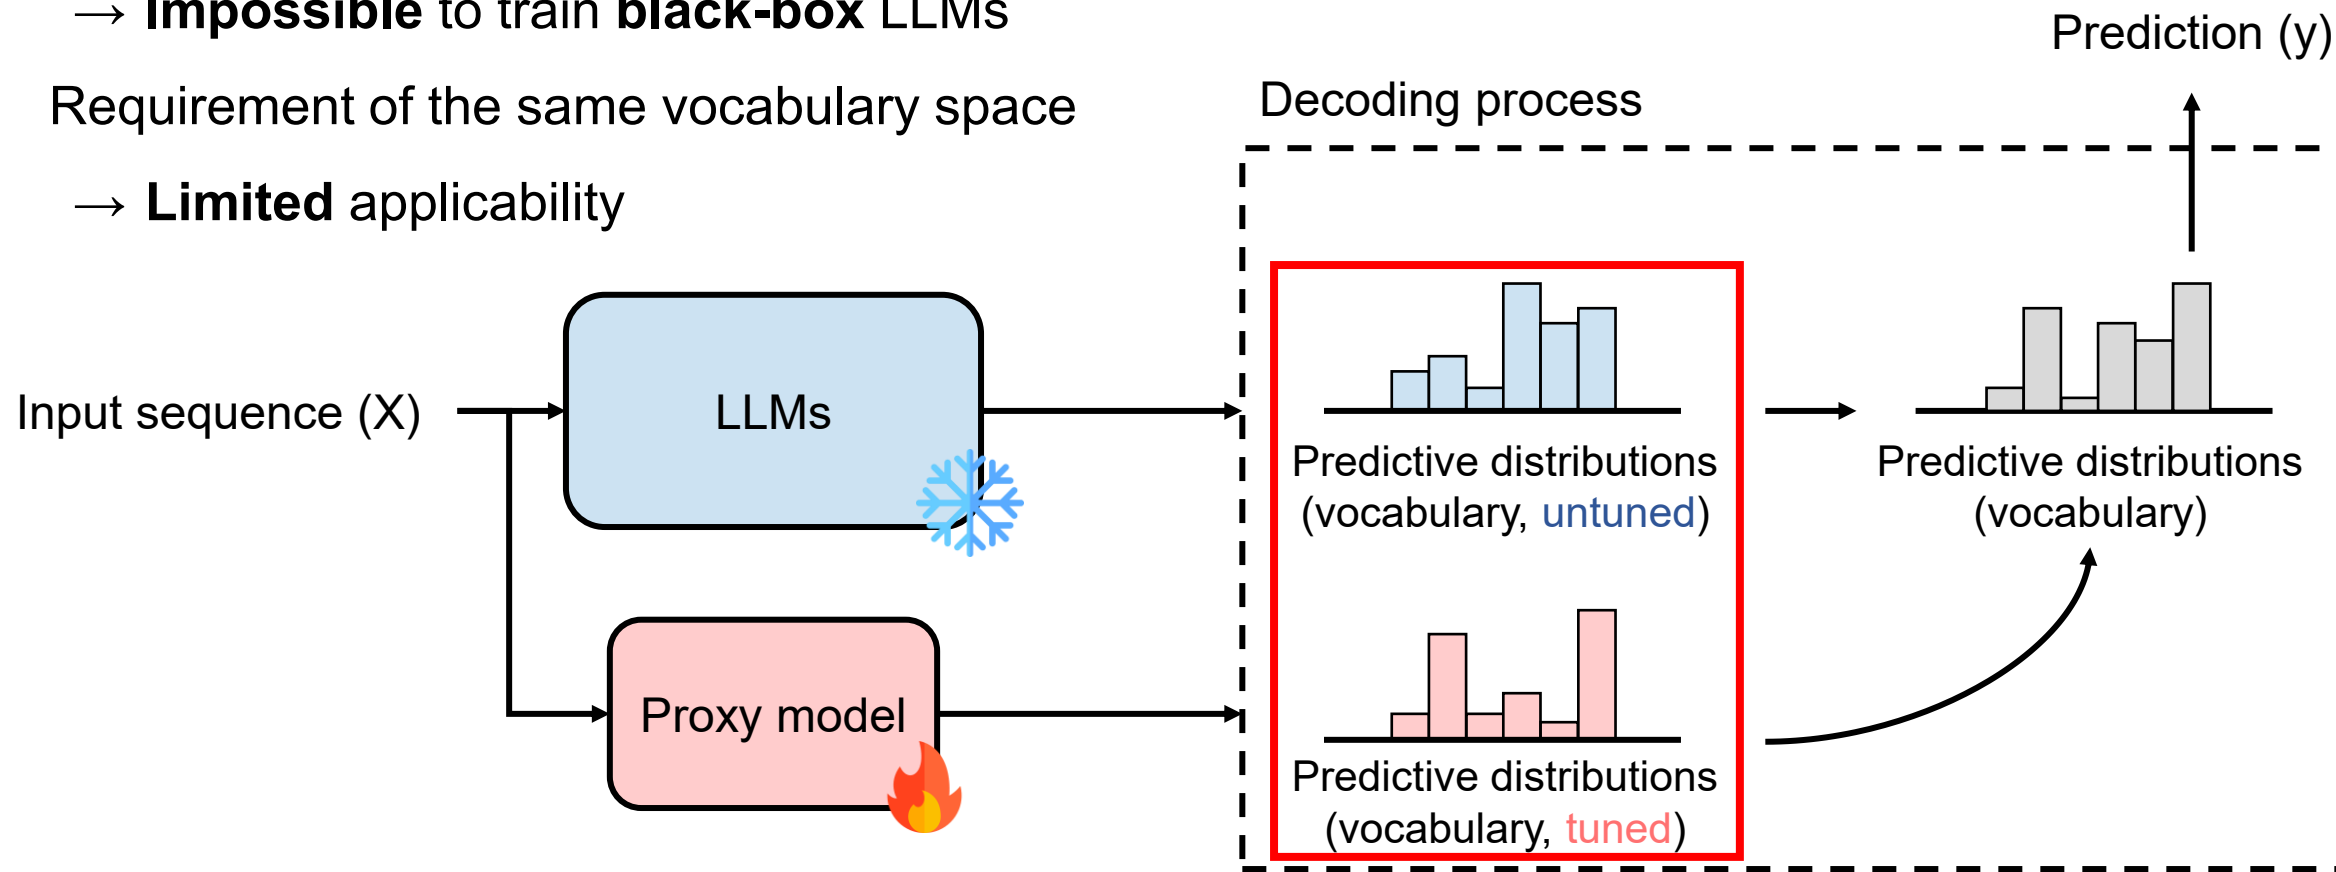
[1] Liu et al., Tuning Language Models by Proxy, COLM 2024.

[2] Ormazabal et al., CombLM: Adapting Black-Box Language Models through Small Fine-Tuned Models, EMNLP 2023.

Related Works: Proxy-based training

Fine-tuning approach based on the predictive distributions from LLMs^[1, 2].

- Access to the predictive distributions of LLMs
 - **Impossible** to train **black-box** LLMs
- Requirement of the same vocabulary space
 - **Limited** applicability



[1] Liu et al., Tuning Language Models by Proxy, COLM 2024.

[2] Ormazabal et al., CombLM: Adapting Black-Box Language Models through Small Fine-Tuned Models, EMNLP 2023.

Research Question

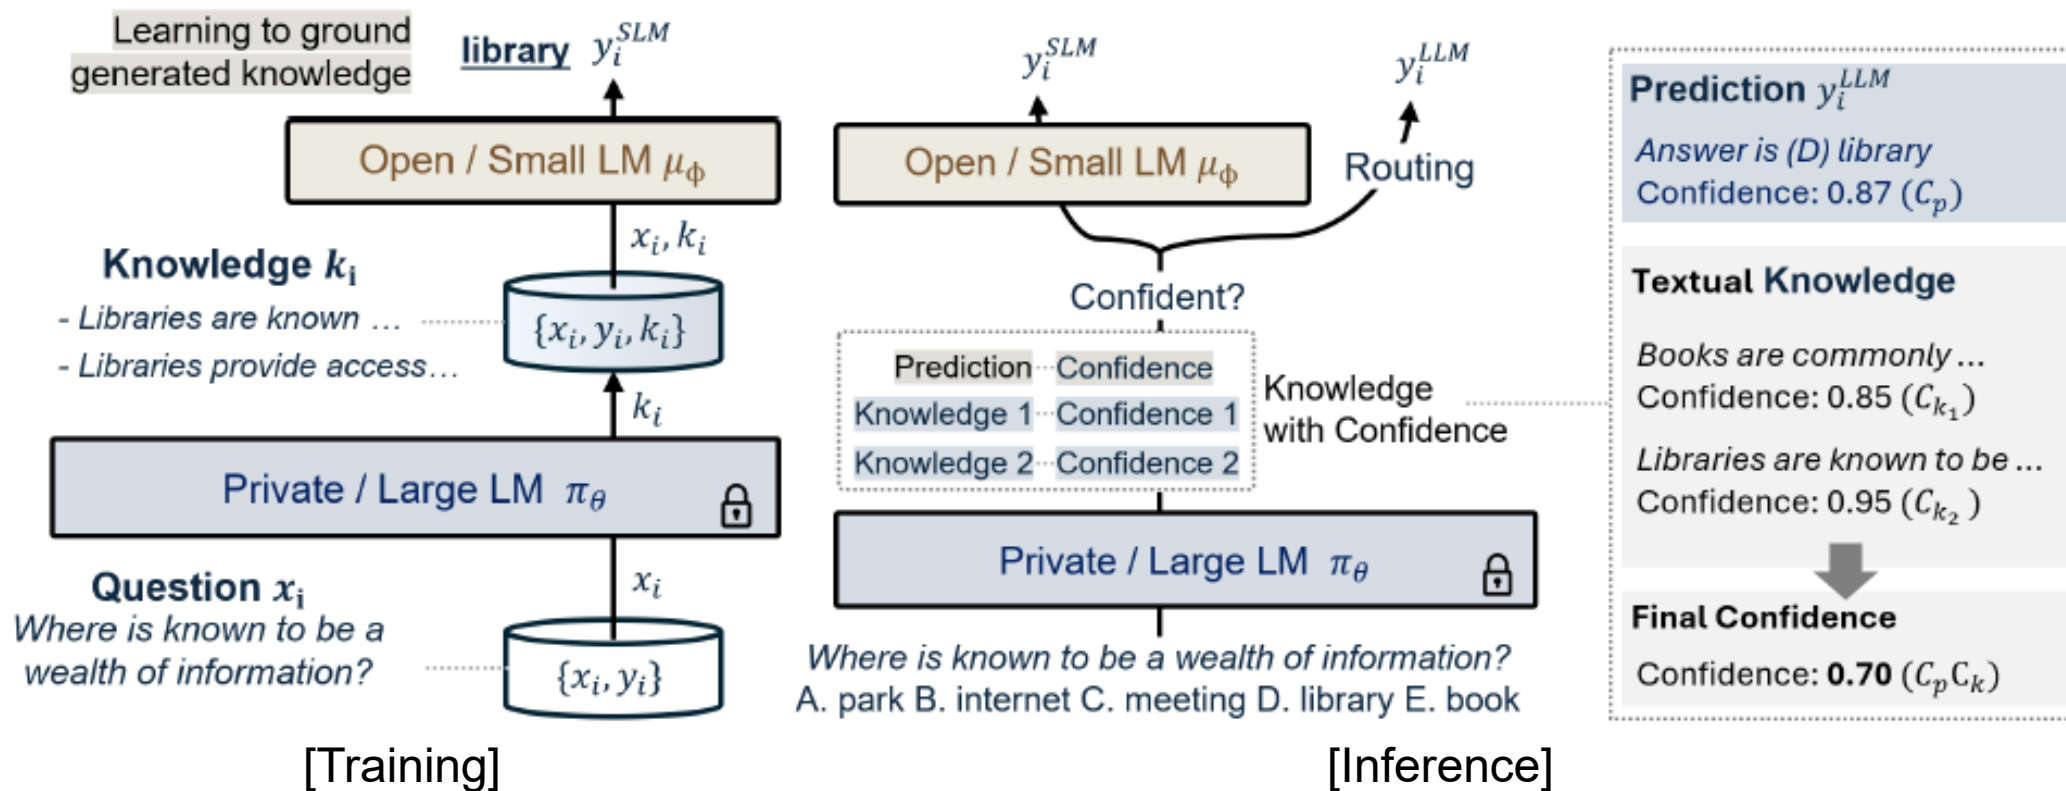
How can we adapt large language models using smaller models **without** relying on their **predictive distributions**?

- ✓ Applicable to all LLMs
- ✓ No need to match vocabulary space

Our Approach: KnowProxy

Proxy-based training through **knowledge adaptation**.

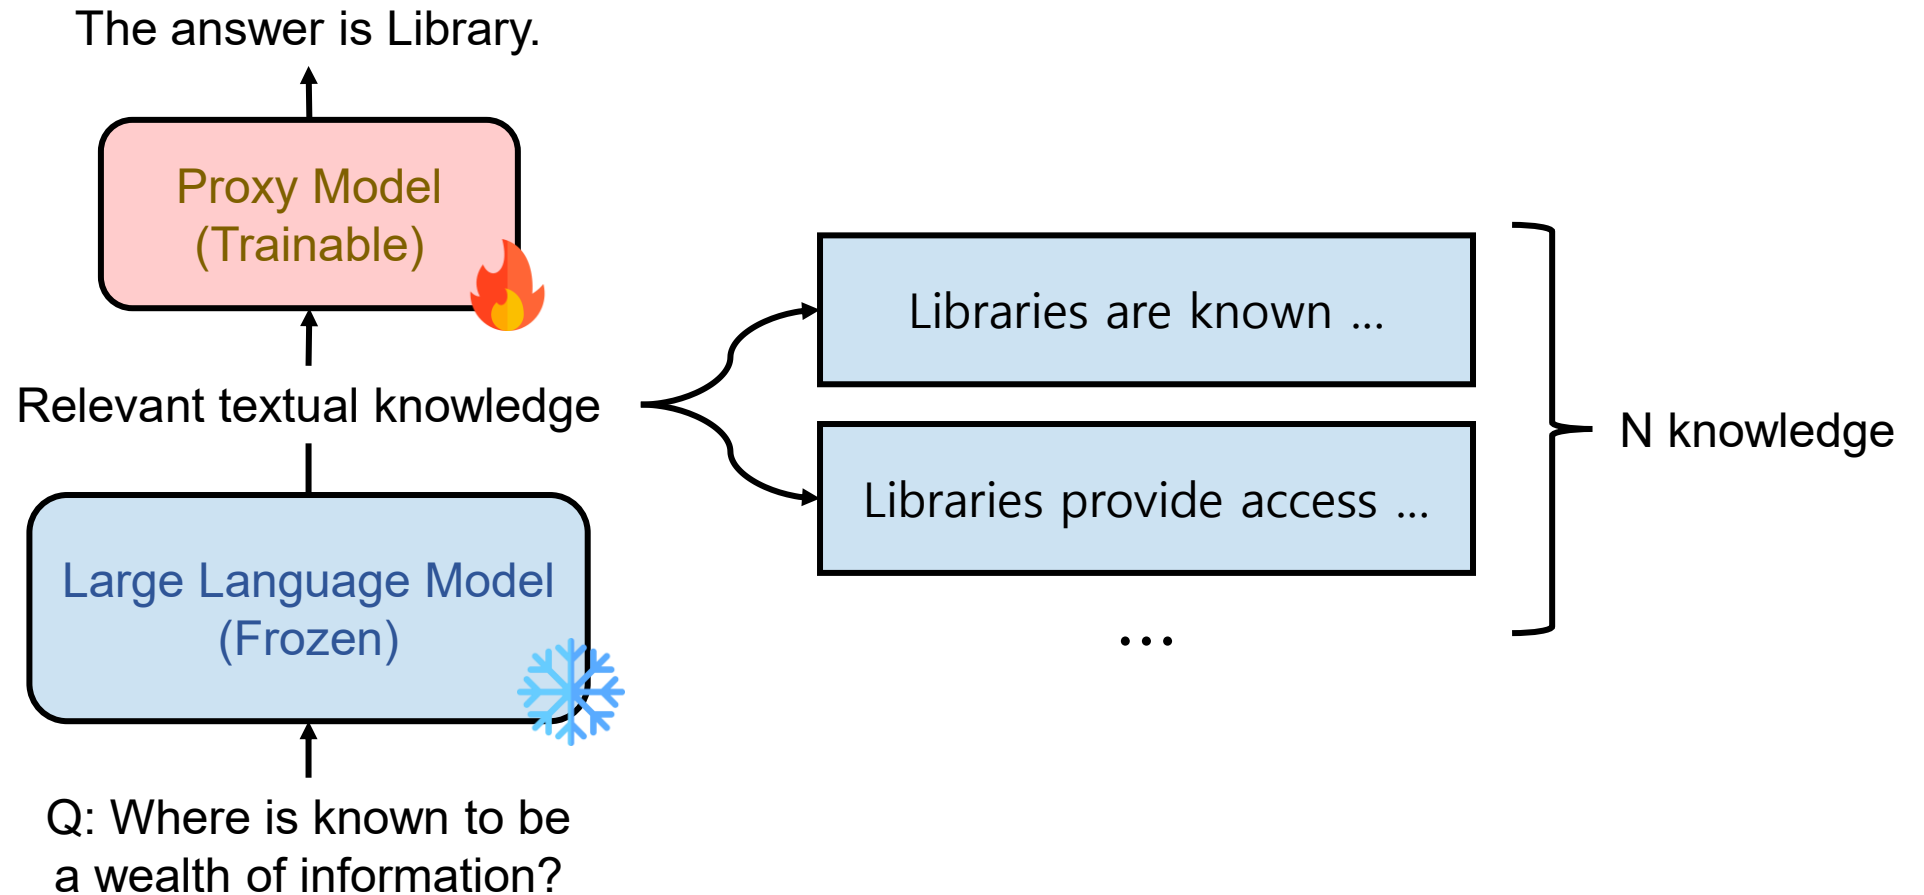
- **LLM**: generates **textual knowledge** relevant to the input sequence.
- **SLM**: learns a **mapping** from LLM-generated knowledge to the ground-truth.



Our Approach: KnowProxy

Proxy-based training through **knowledge adaptation**.

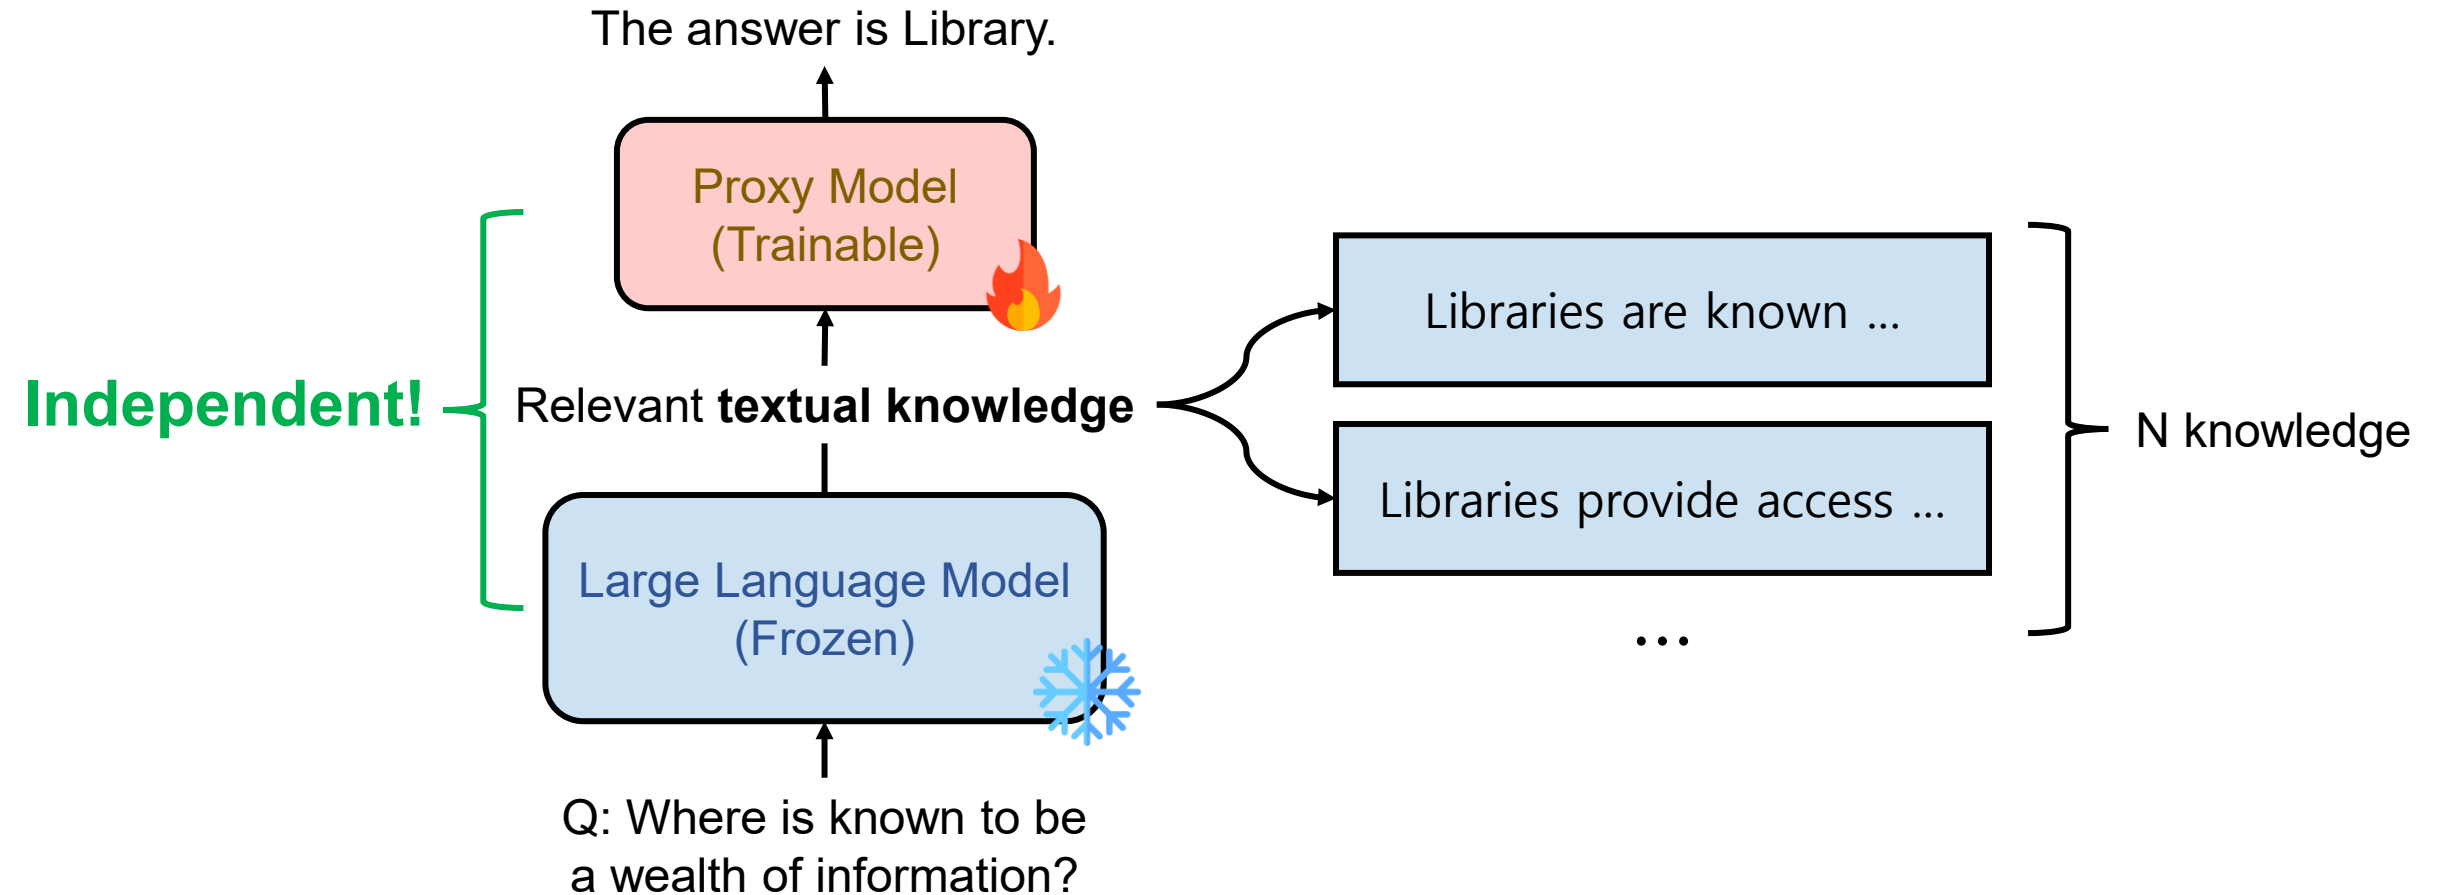
- **LLM**: generates **textual knowledge** relevant to the input sequence.
- **SLM**: learns a **mapping** from LLM-generated knowledge to the ground-truth.



Our Approach: KnowProxy

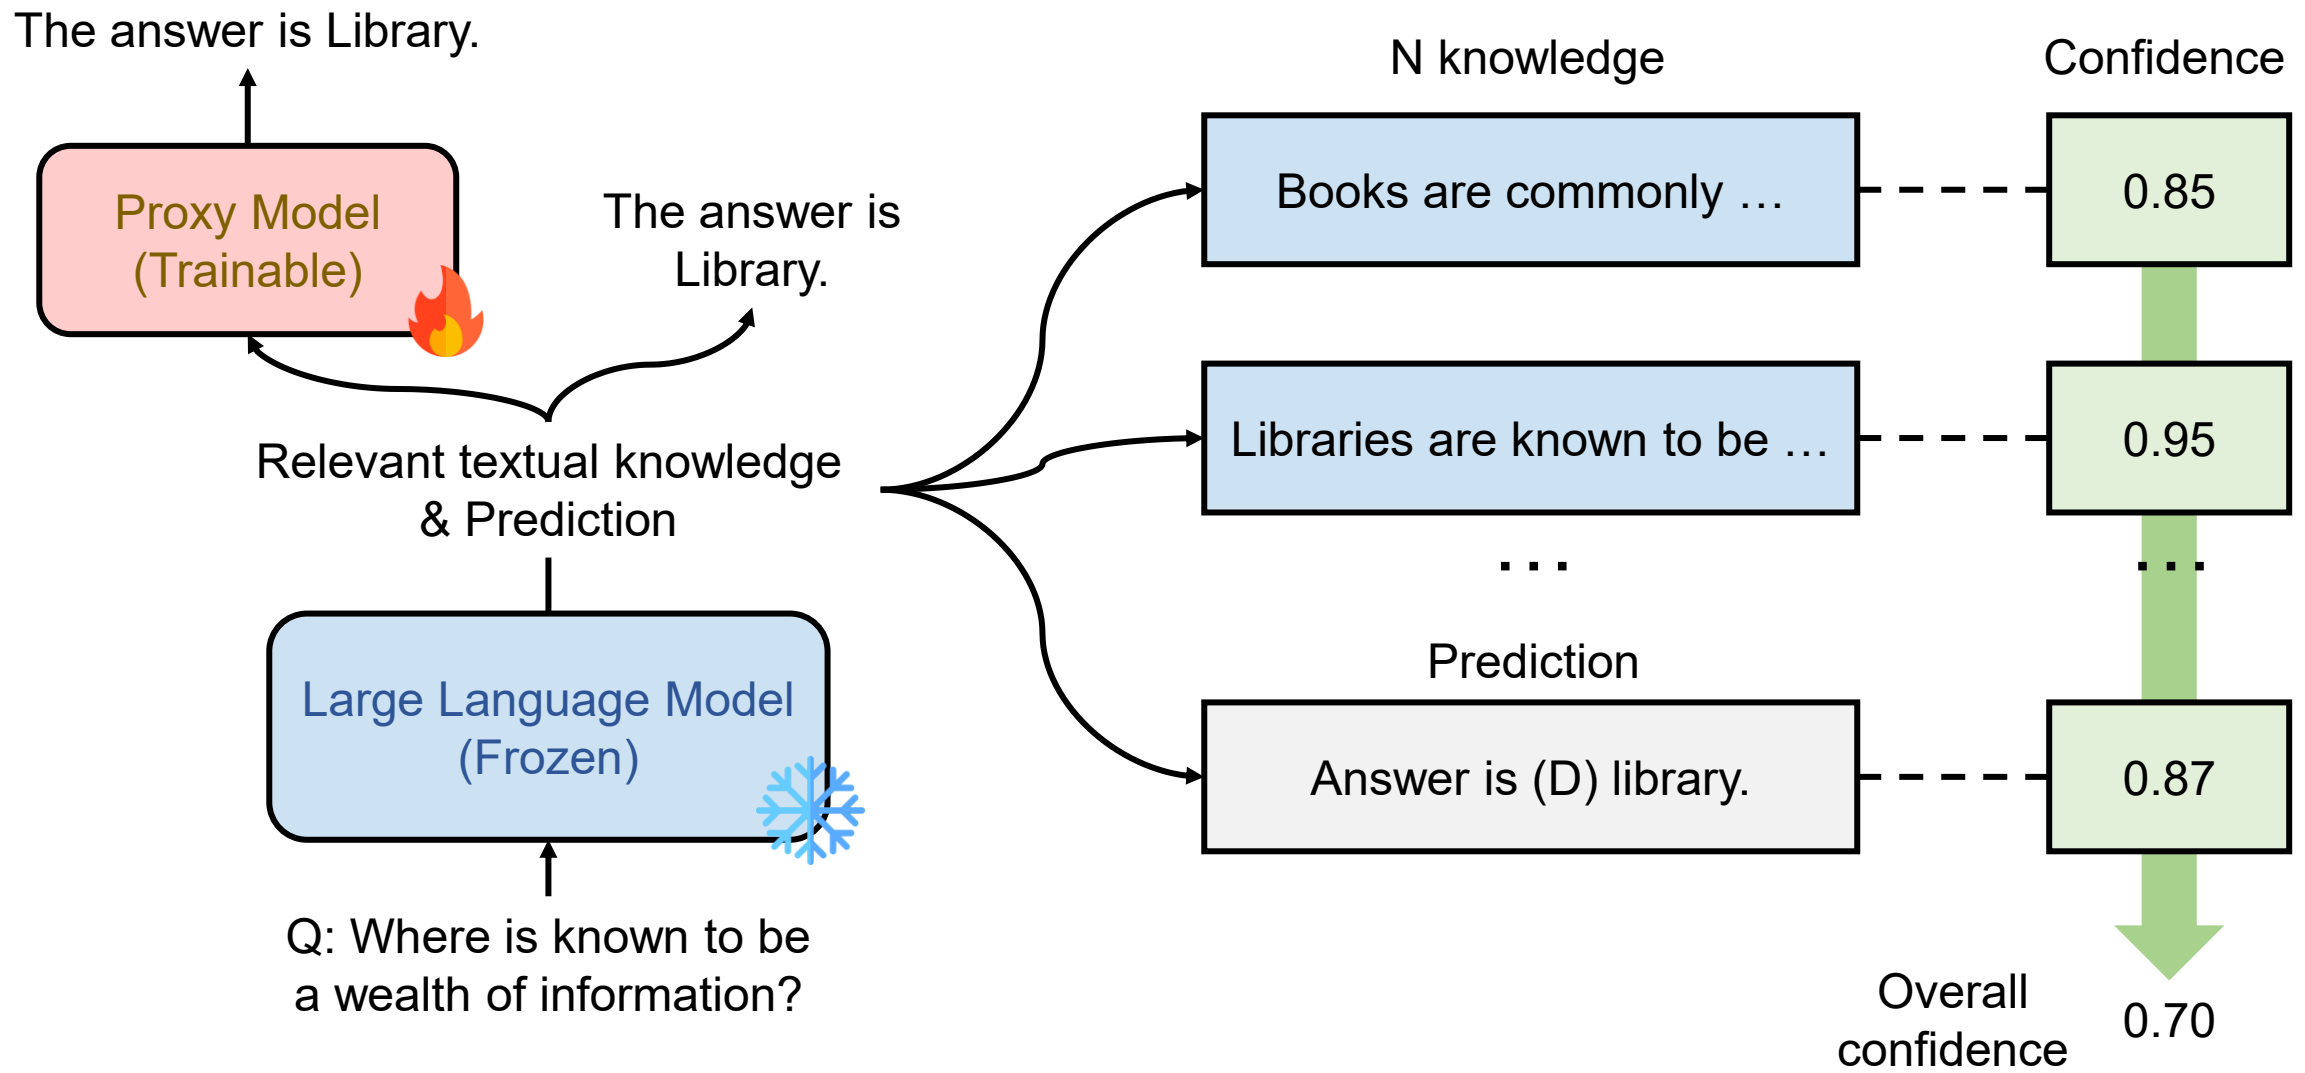
Proxy-based training through **knowledge adaptation**.

- **LLM**: generates **textual knowledge** relevant to the input sequence.
- **SLM**: learns a **mapping** from LLM-generated knowledge to the ground-truth.



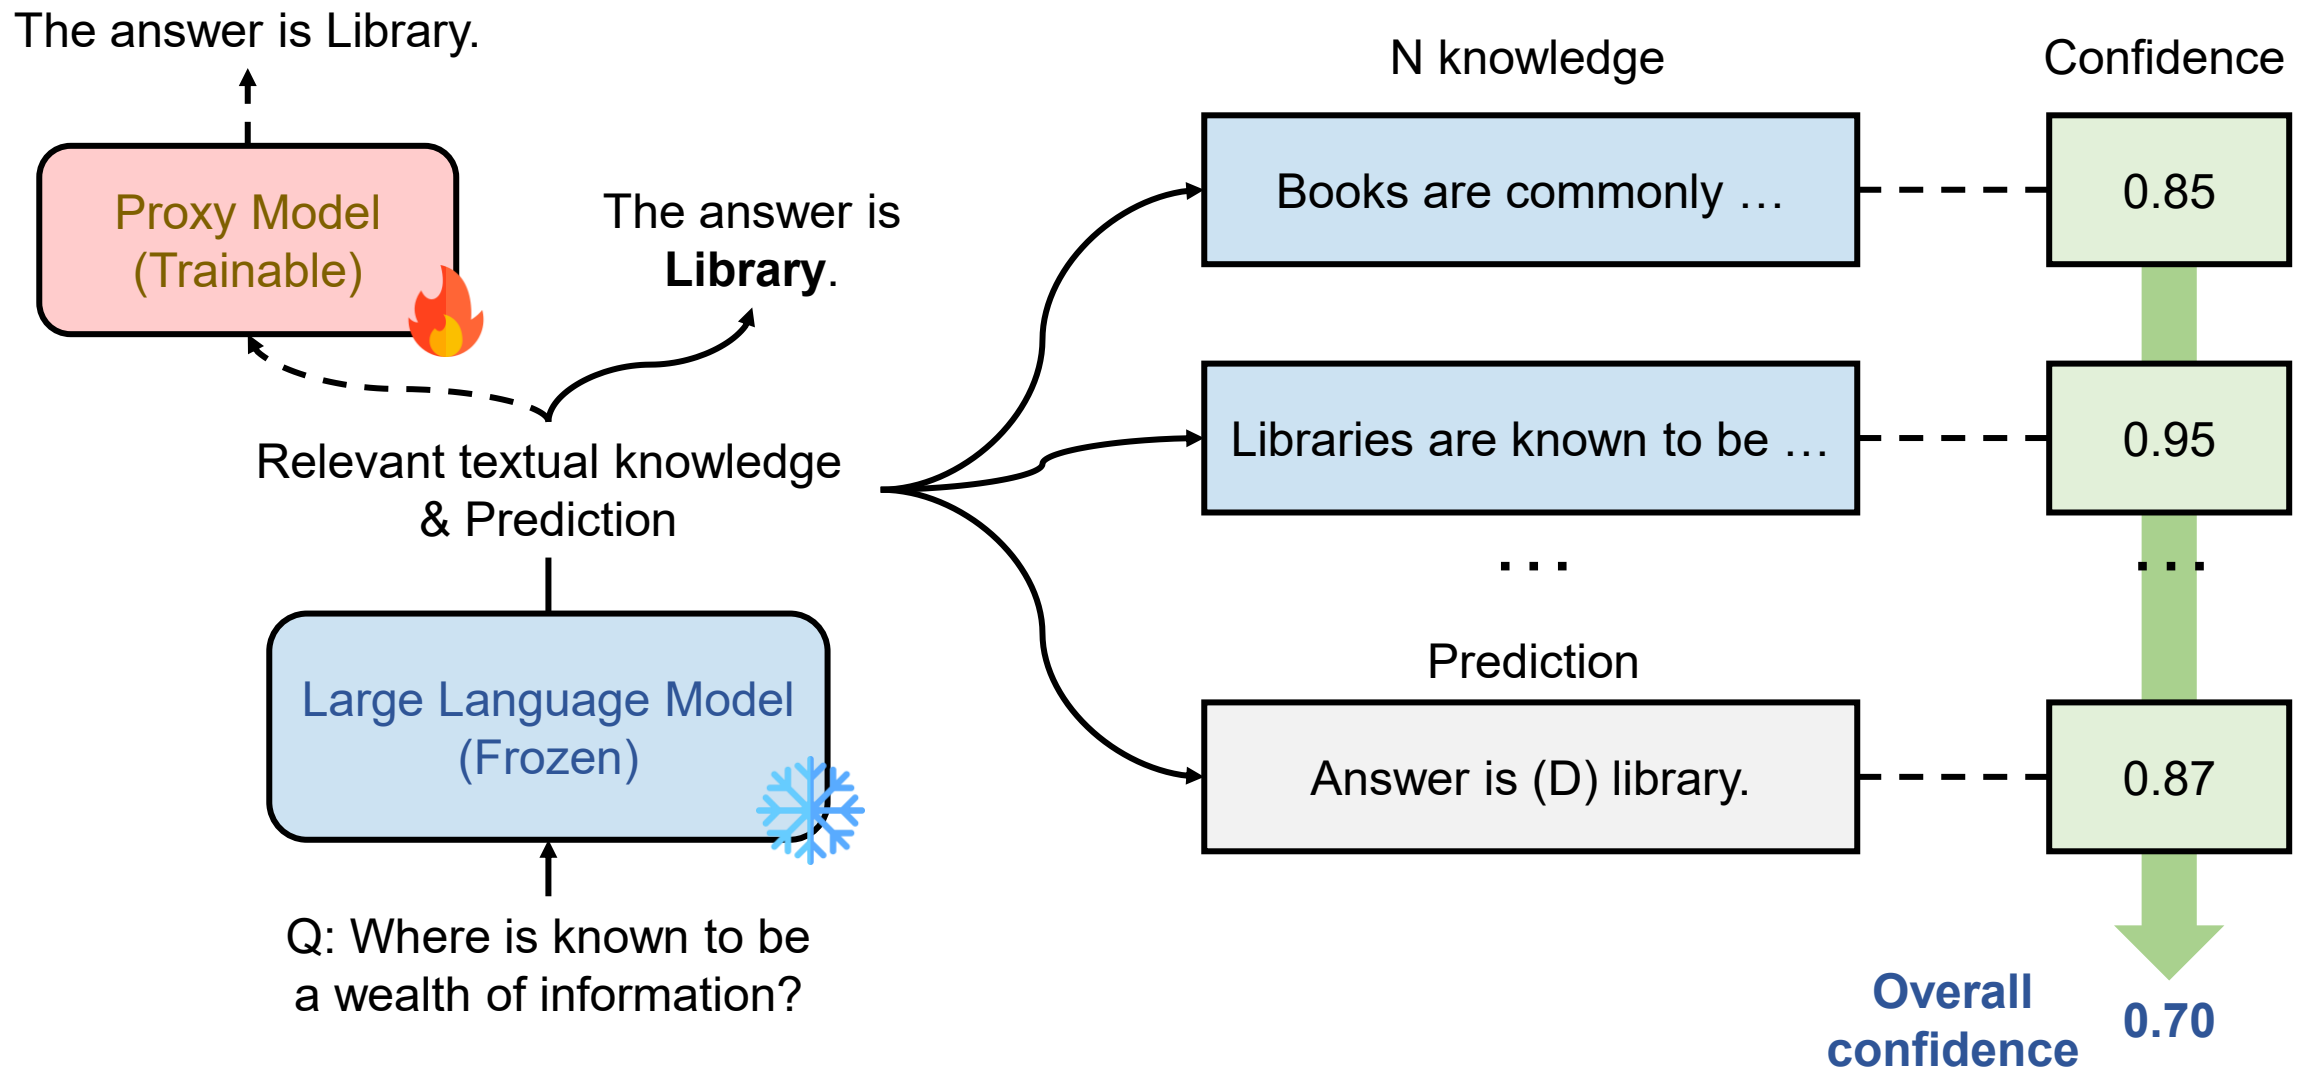
Our Approach: KnowProxy

Dynamic routing inference based on LLM's uncertainty.



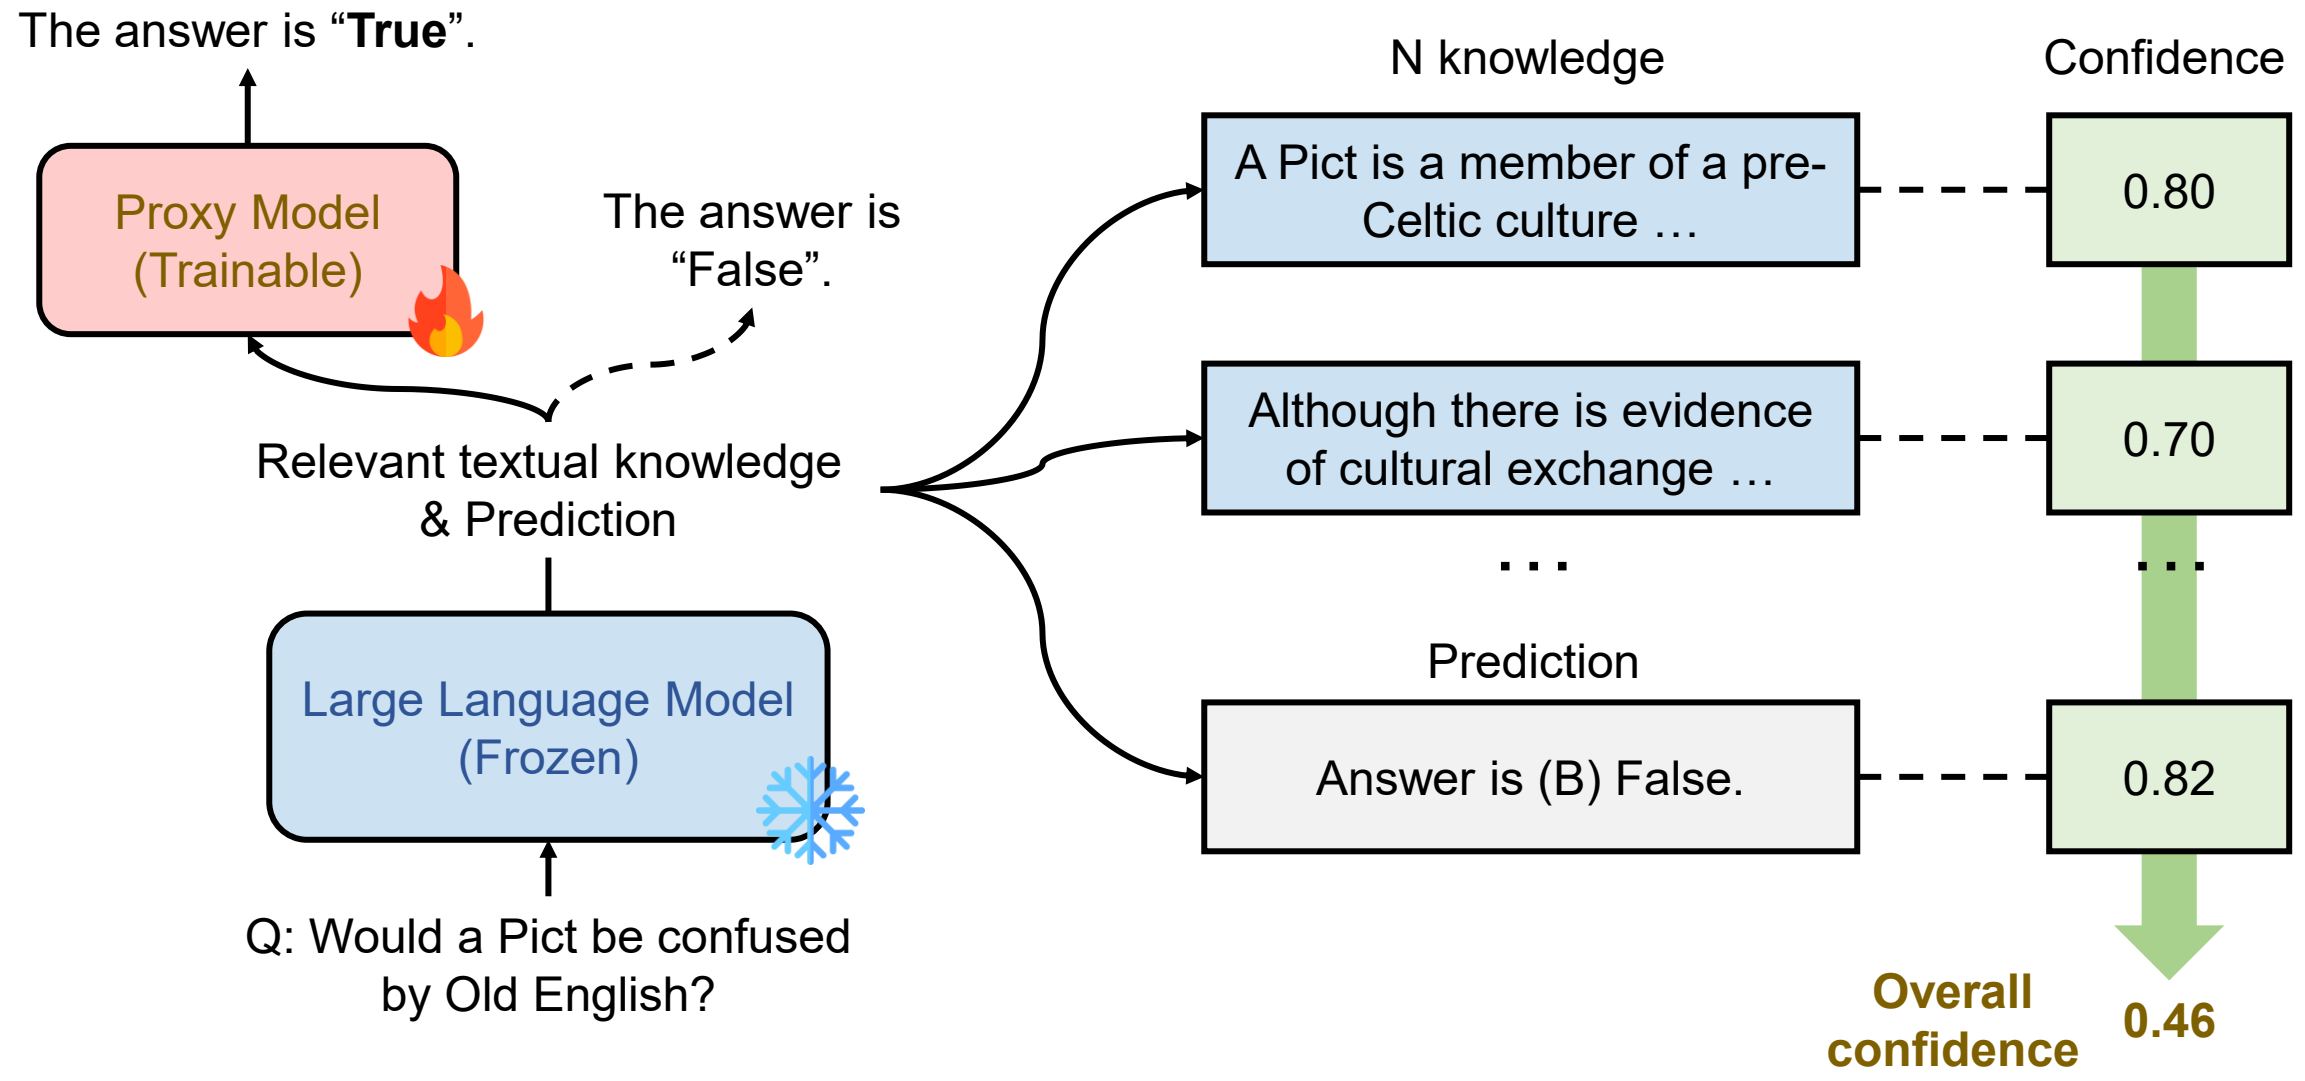
Our Approach: KnowProxy

Dynamic routing inference based on LLM's uncertainty.




Our Approach: KnowProxy

Dynamic routing inference based on LLM's uncertainty.



Results

Finding 1: KnowProxy demonstrates effectiveness across diverse fine-tuning scenarios.

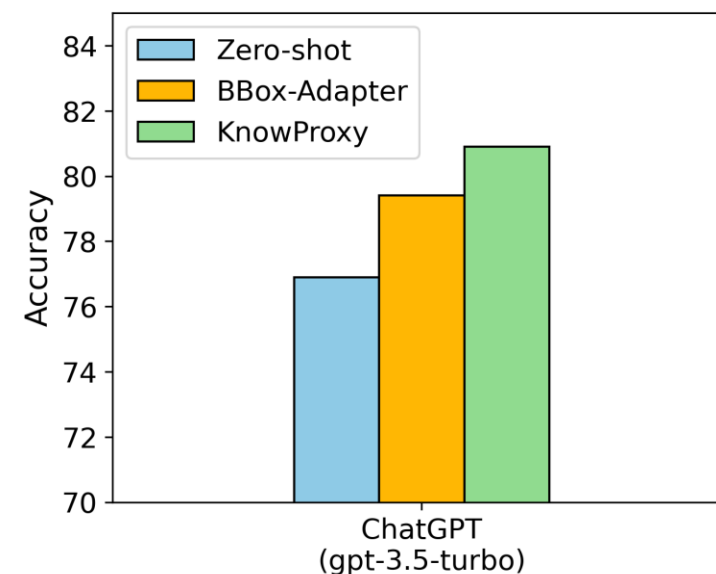
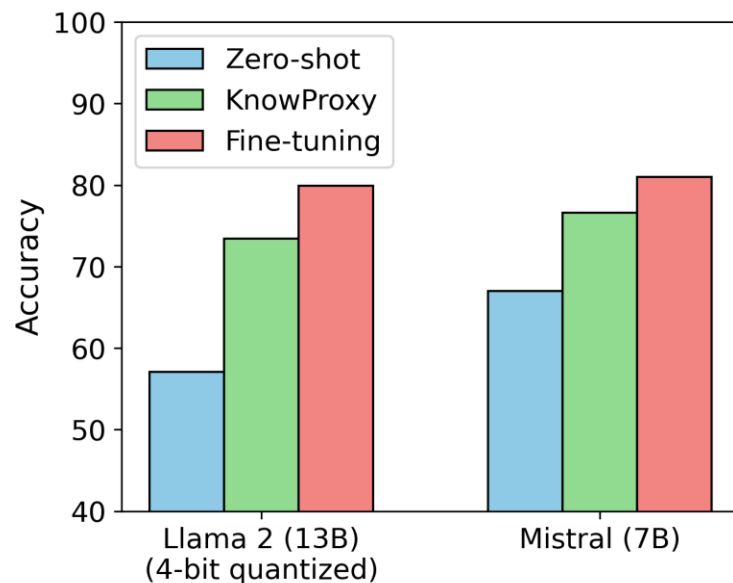
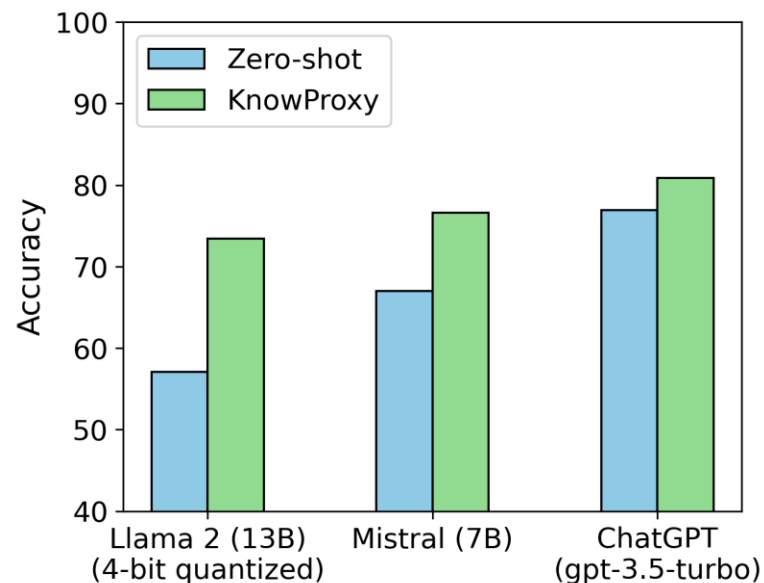


Method	OBQA	ARC _h	PIQA	CSQA	QASC	SIQA	WNGR	StrategyQA	BoolQ	Avg.
Fine-tuning LLM	82.2	76.2	87.7	79.5	82.9	80.5	87.3	71.5	86.9	81.6
Fine-tuning SLM	73.2	60.9	80.3	72.0	68.0	74.9	75.4	66.5	85.4	73.0
<i>Advanced Zero-shot Reasoning (Frozen LLM)</i>										
Zero-shot	72.2	68.6	75.8	67.7	75.9	65.3	53.6	60.9	78.3	68.7
Self-Talk	74.8	74.2	75.9	72.3	80.5	67.0	54.6	57.6	<u>78.5</u>	70.6
Chain-of-Thought	77.6	80.0	75.6	73.1	79.0	68.6	57.8	<u>69.0</u>	<u>76.7</u>	73.1
Plan-and-Solve	76.6	<u>75.3</u>	74.3	<u>73.7</u>	<u>79.8</u>	66.4	58.2	66.8	73.9	71.7
<i>Proxy-based Training (Frozen LLM + SLM)</i>										
Proxy-tuning	77.2	69.6	80.1	70.8	69.9	72.6	65.7	64.6	76.2	71.9
CombLM	78.6	72.6	<u>81.1</u>	72.5	76.9	<u>73.7</u>	<u>69.3</u>	67.2	76.8	74.3
BBox-Adapter	76.2	68.6	<u>73.8</u>	73.3	73.8	<u>72.7</u>	<u>53.7</u>	69.0	70.5	70.2
KNOWPROXY (ours)	80.2	75.2	83.4	75.0	78.1	76.3	77.8	72.9	85.1	78.2

KnowProxy achieves strong performance compared to existing proxy-based fine-tuning approaches.

Results

Finding 1: KnowProxy demonstrates effectiveness across diverse fine-tuning scenarios.

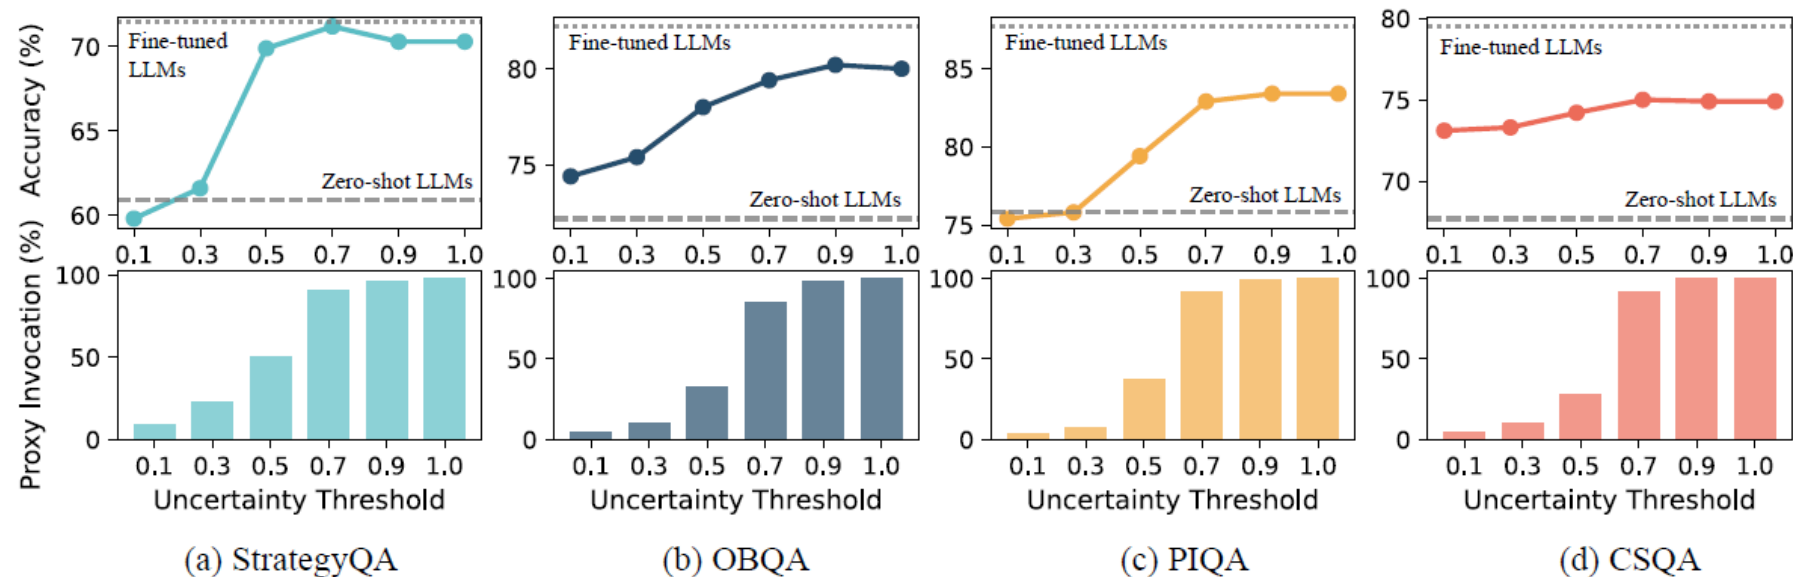


KnowProxy effectively enhances the reasoning performance in diverse types of LLMs.

- ✓ Quantized LLMs
- ✓ Black-box LLMs

Analysis

Finding 2: KnowProxy improves efficiency while preserving performance.



Method	OBQA	PIQA	StrategyQA	SIQA
KNOWPROXY	85.0	87.2	74.7	77.0
w/o routing	82.0	87.2	74.7	76.8
w/o filtering	85.0	86.2	72.1	76.0
w/o adaptation	80.6	85.1	59.4	75.3
w/ LLM answer	76.8	83.7	72.9	76.4

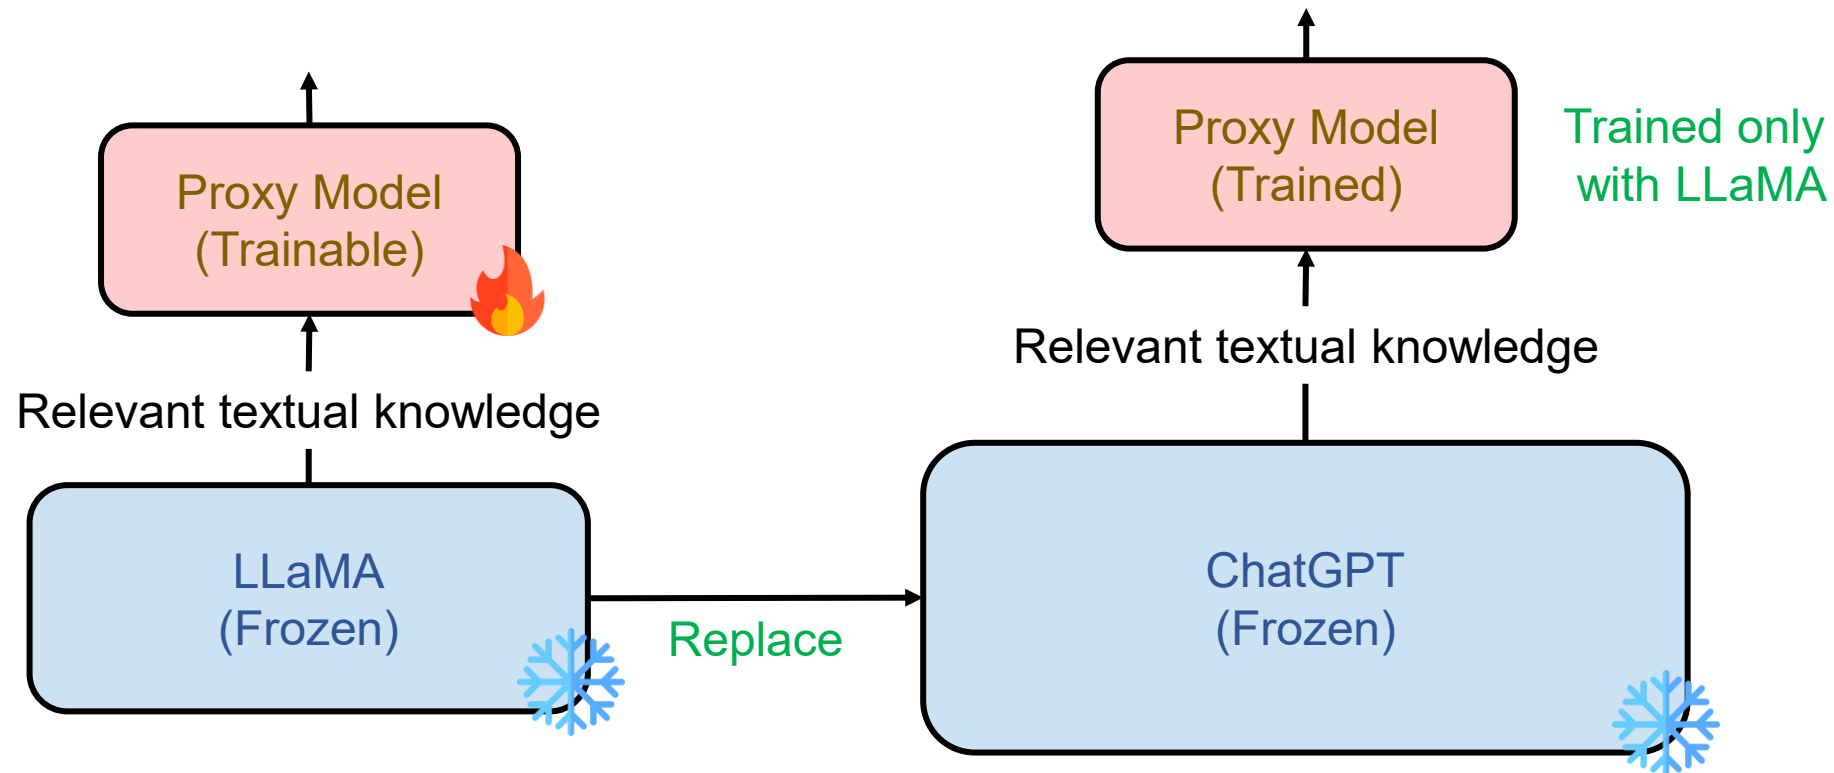
by activating the proxy model
only when necessary.

Analysis

Finding 3: The proxy learns to leverage textual knowledge in a model-agnostic way.

Method	Trained Knowledge	OBQA	ARC _h	PIQA	CSQA	StrategyQA	QASC	Avg.
Zero-shot	-	78.8	81.2	82.8	76.3	68.1	79.0	77.7
KNOWPROXY	ChatGPT	85.0	83.9	87.2	78.1	74.7	80.2	81.5
	Llama 3.2	81.0	83.5	85.1	76.7	69.9	78.4	79.1
	Mistral-v0.2	82.2	84.3	86.1	76.4	71.6	80.2	80.1

e.g.,



Conclusion

- I. **KnowProxy** is a novel **proxy-based** fine-tuning framework that adapts LLMs **through their textual outputs**, enabling applicability to black-box settings.

- II. **KnowProxy** adaptively activates the proxy model **only when necessary**, improving efficiency while preserving performance.

- III. **KnowProxy** demonstrates **effectiveness** across diverse benchmarks and training setups, highlighting **its practical value** in fine-tuning scenarios.