

# BBox-Adapter

Lightweight Adapting for Black-Box Large Language Models

Haotian Sun<sup>1\*</sup>, Yuchen Zhuang<sup>1\*</sup>, Wei Wei<sup>2</sup>, Chao Zhang<sup>1</sup>, Bo Dai<sup>1</sup>

<sup>1</sup> Georgia Institute of Technology

<sup>2</sup>Accenture



Paper



Code




Website




- Adapting Blackbox LLMs *only through* their APIs raises problems with **privacy, transparency, and cost.**



*uploading training data  
via APIs...*




*a restricted set of adjustable  
hyperparameters...*

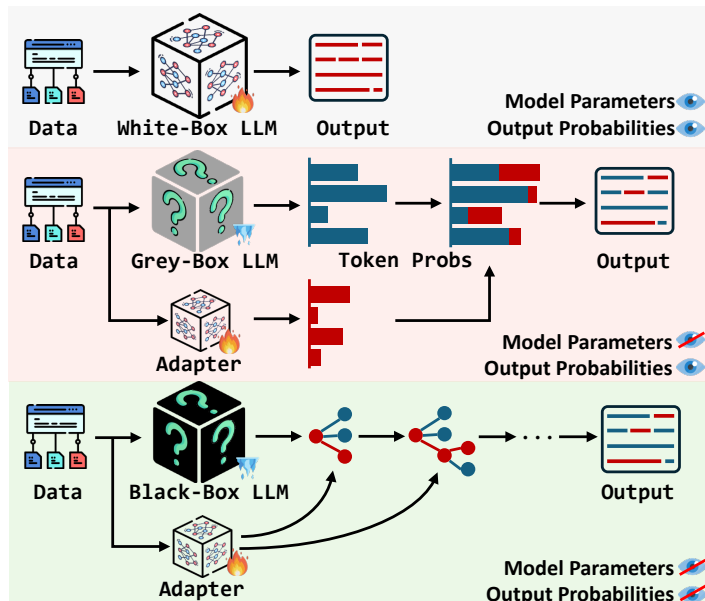


*fine-tuning APIs much higher  
compared to inference...*

- Adapting Blackbox LLMs *only through* their APIs raises problems with **privacy, transparency, and cost**.
- Existing methods fail to support this real black-box LLM setting, where neither model parameters nor output probabilities can be accessed in the most recent LLM APIs.

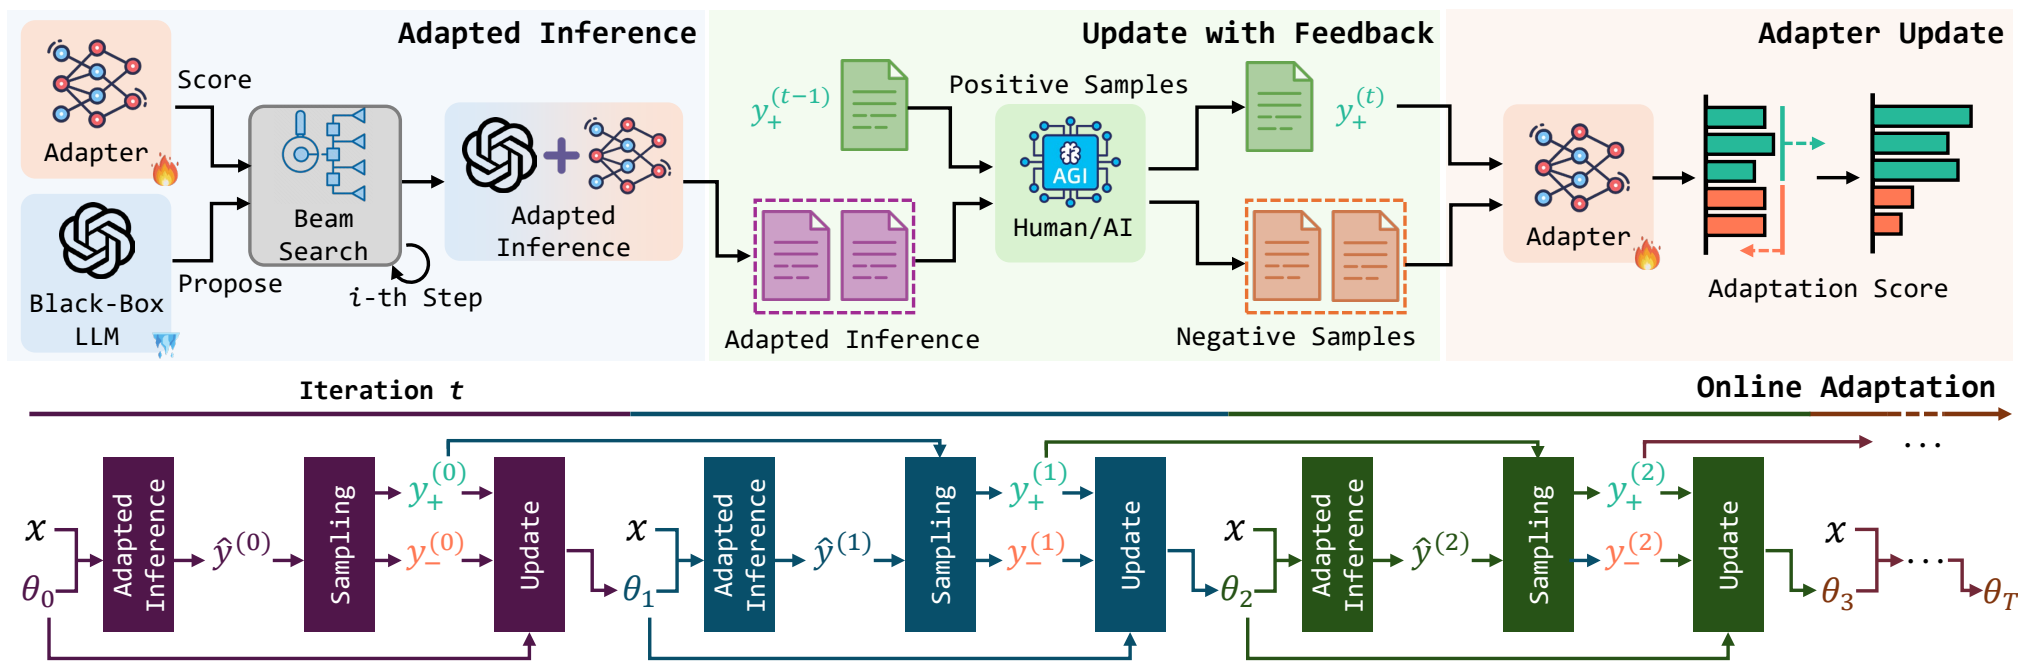


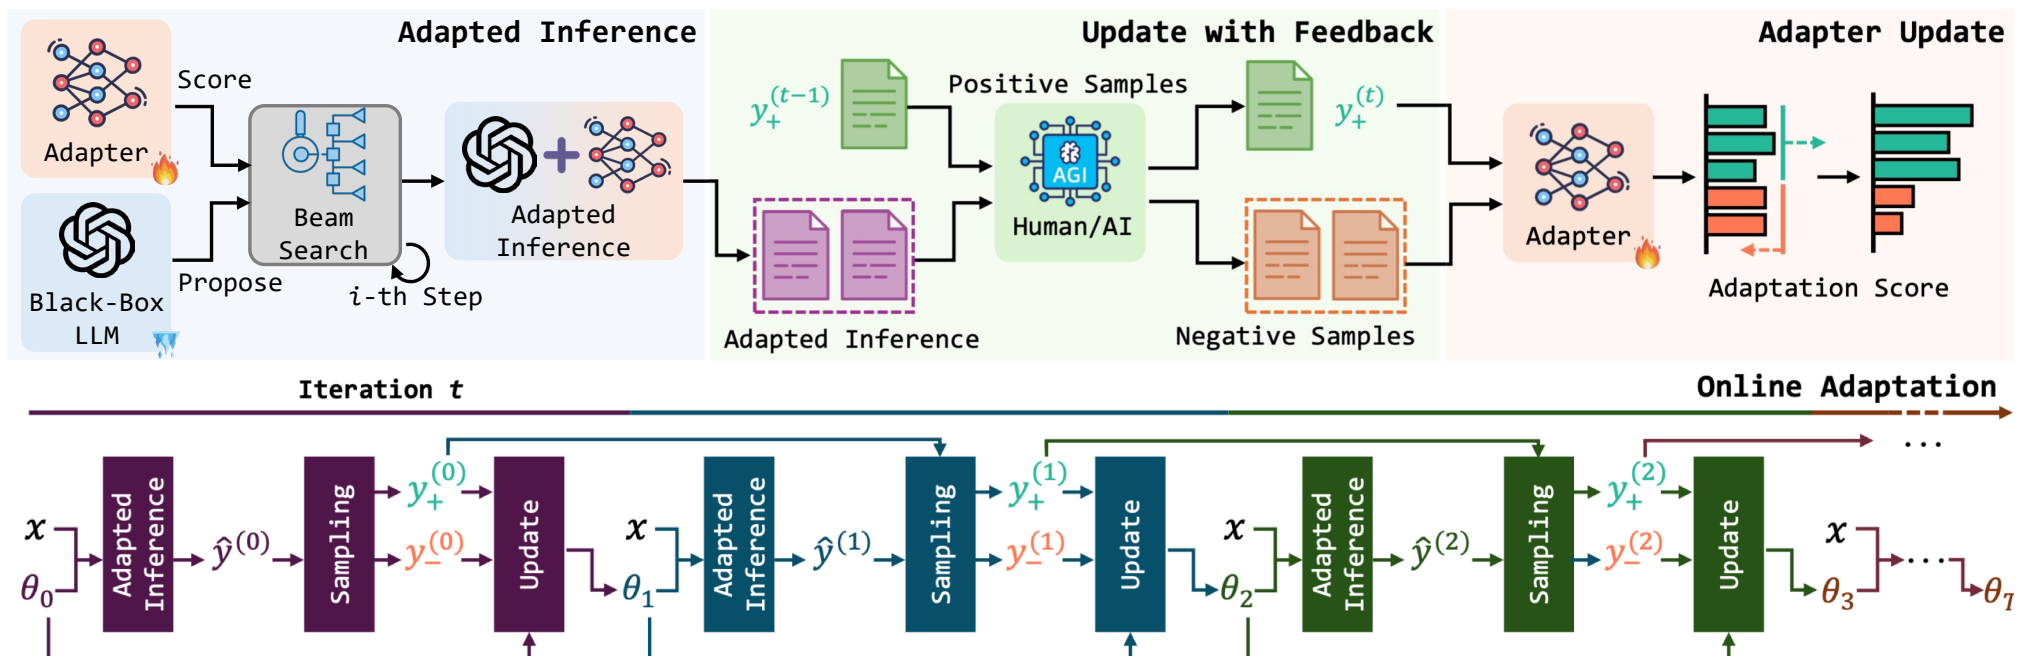
- Adapting Blackbox LLMs *only through* their APIs raises problems with **privacy, transparency, and cost.**
- Existing methods fail to support this real black-box LLM setting, where neither model parameters nor output probabilities can be accessed in the most recent LLM APIs.



Methods	w/o Model Parameters	w/o High-Dimensional Representation	w/o Token Probabilities	w/o Retrieval Corpus	w/ Smaller Adapter
<i>White-Box LLM Fine-Tuning</i>					
Fine-Tuning (Devlin et al., 2019)	✗	✗	✗	✓	✗
Instruction-Tuning (Wei et al., 2021)	✗	✗	✗	✓	✗
Continual Pre-Training (Gururangan et al., 2020)	✗	✗	✗	✓	✗
Adapter (Houlsby et al., 2019)	✗	✗	✗	✓	✓
Prefix-Tuning (Liu et al., 2022)	✗	✗	✗	✓	✓
LoRA (Hu et al., 2021)	✗	✗	✗	✓	✓
<i>Grey-Box LLM Adaptation</i>					
LMaaS (Sun et al., 2022)	✓	✗	✗	✓	✓
kNN-Adapter (Huang et al., 2023)	✓	✓	✗	✗	✓
CombLM (Ormazabal et al., 2023)	✓	✓	✗	✓	✓
IPA (Lu et al., 2023)	✓	✓	✗	✓	✓
Proxy-Tuning (Liu et al., 2024)	✓	✓	✗	✓	✓
<i>Black-Box LLM Adaptation</i>					
<b>BBOX-ADAPTER (Ours)</b>	✓	✓	✓	✓	✓

**BBox-Adapter** offers a novel lightweight adaption solution for customizing commercial black-box LLMs with *only* APIs.





## Adaptation as Energy-based Models

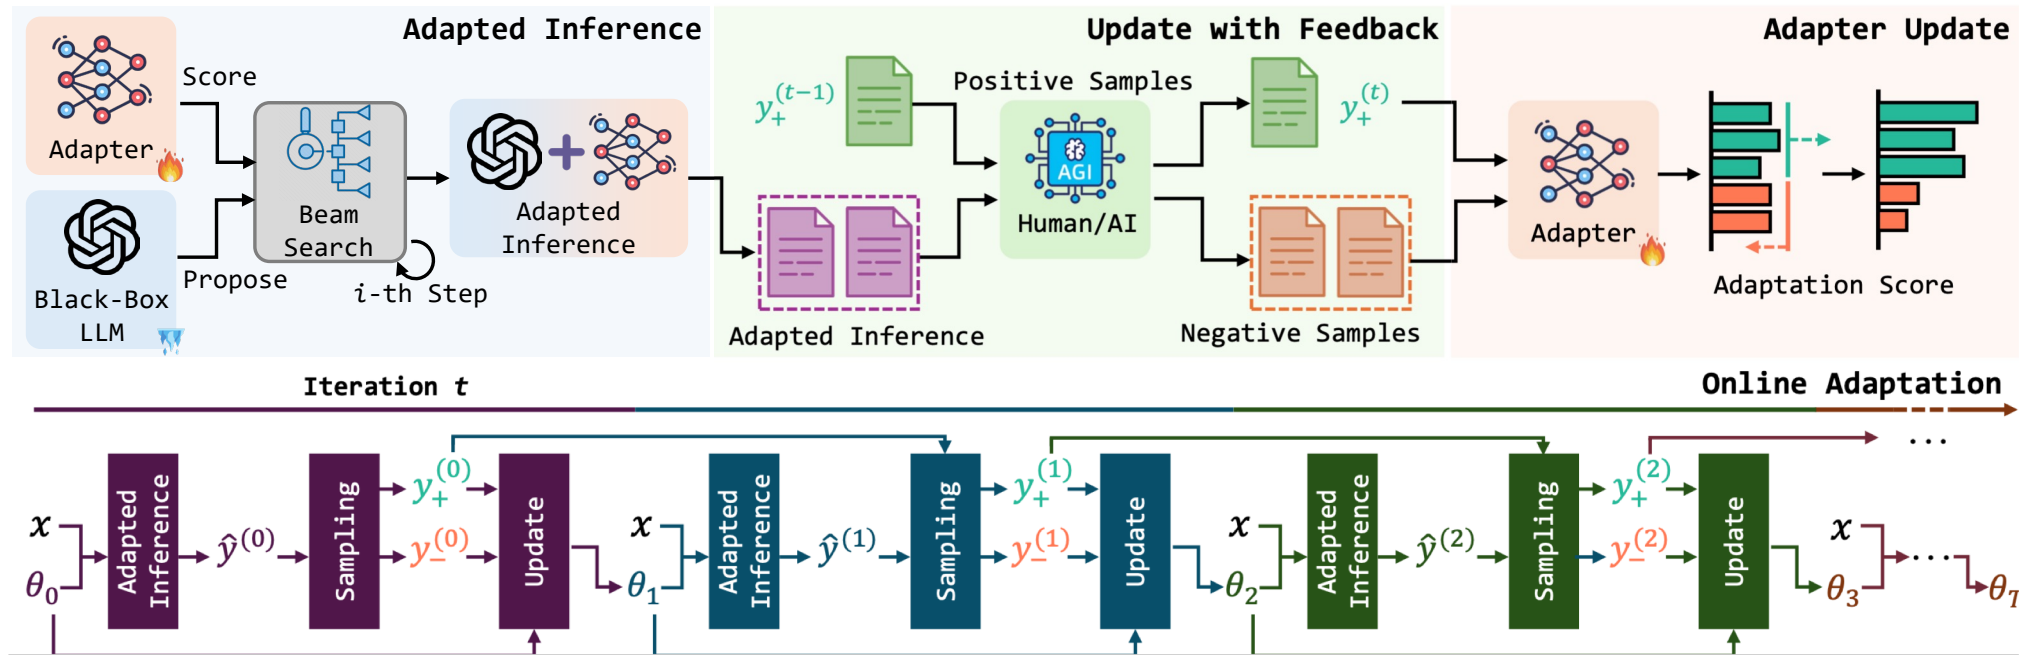
$$p_{\theta}(\mathbf{y}|\mathbf{x}) = p_{\text{LLM}}(\mathbf{y}|\mathbf{x}) \frac{\exp(g_{\theta}(\mathbf{x}, \mathbf{y}))}{Z_{\theta}(\mathbf{x})}$$

$$Z_{\theta}(\mathbf{x}) = \int p_{\text{LLM}}(\mathbf{y}|\mathbf{x}) \exp(g_{\theta}(\mathbf{x}, \mathbf{y})) d\mathbf{y}$$



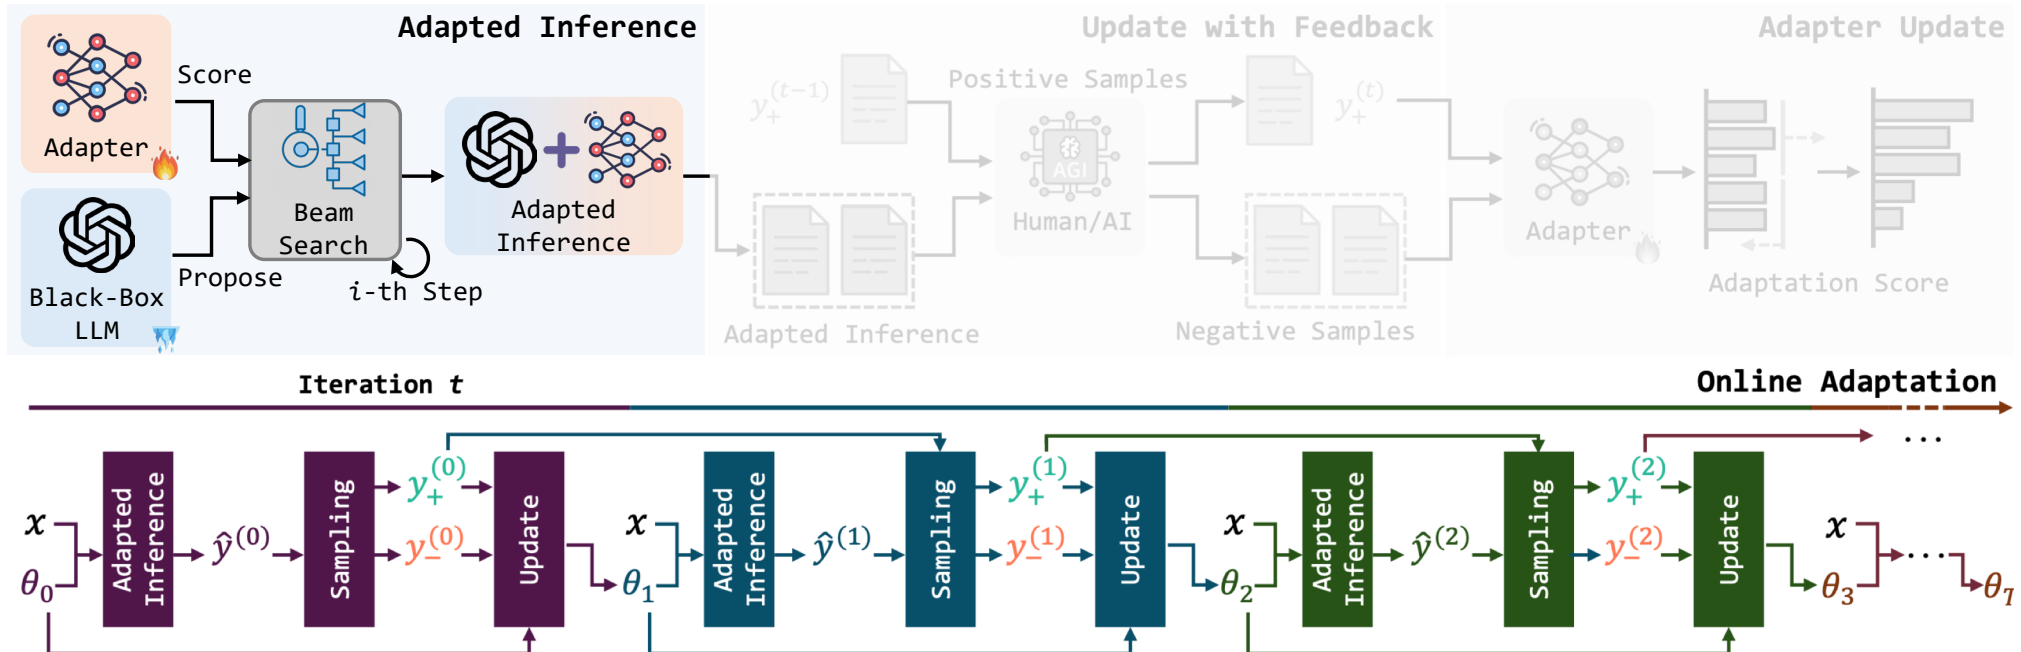
### NCE loss

$$\nabla_{\theta} \ell(\theta) = \nabla_{\theta} \{ -\mathbb{E}_{\mathbf{y}_+ \sim p_{\text{data}}(\mathbf{y}|\mathbf{x})} [g_{\theta}(\mathbf{x}, \mathbf{y}_+)] + \alpha \mathbb{E}[g_{\theta}(\mathbf{x}, \mathbf{y}_+)^2] + \mathbb{E}_{\mathbf{y}_- \sim p_{\theta}(\mathbf{y}|\mathbf{x})} [g_{\theta}(\mathbf{x}, \mathbf{y}_-)] + \alpha \mathbb{E}[g_{\theta}(\mathbf{x}, \mathbf{y}_-)^2] \}$$



**Online adaptation** draws training samples from dynamic distributions to optimize its backend adapter.

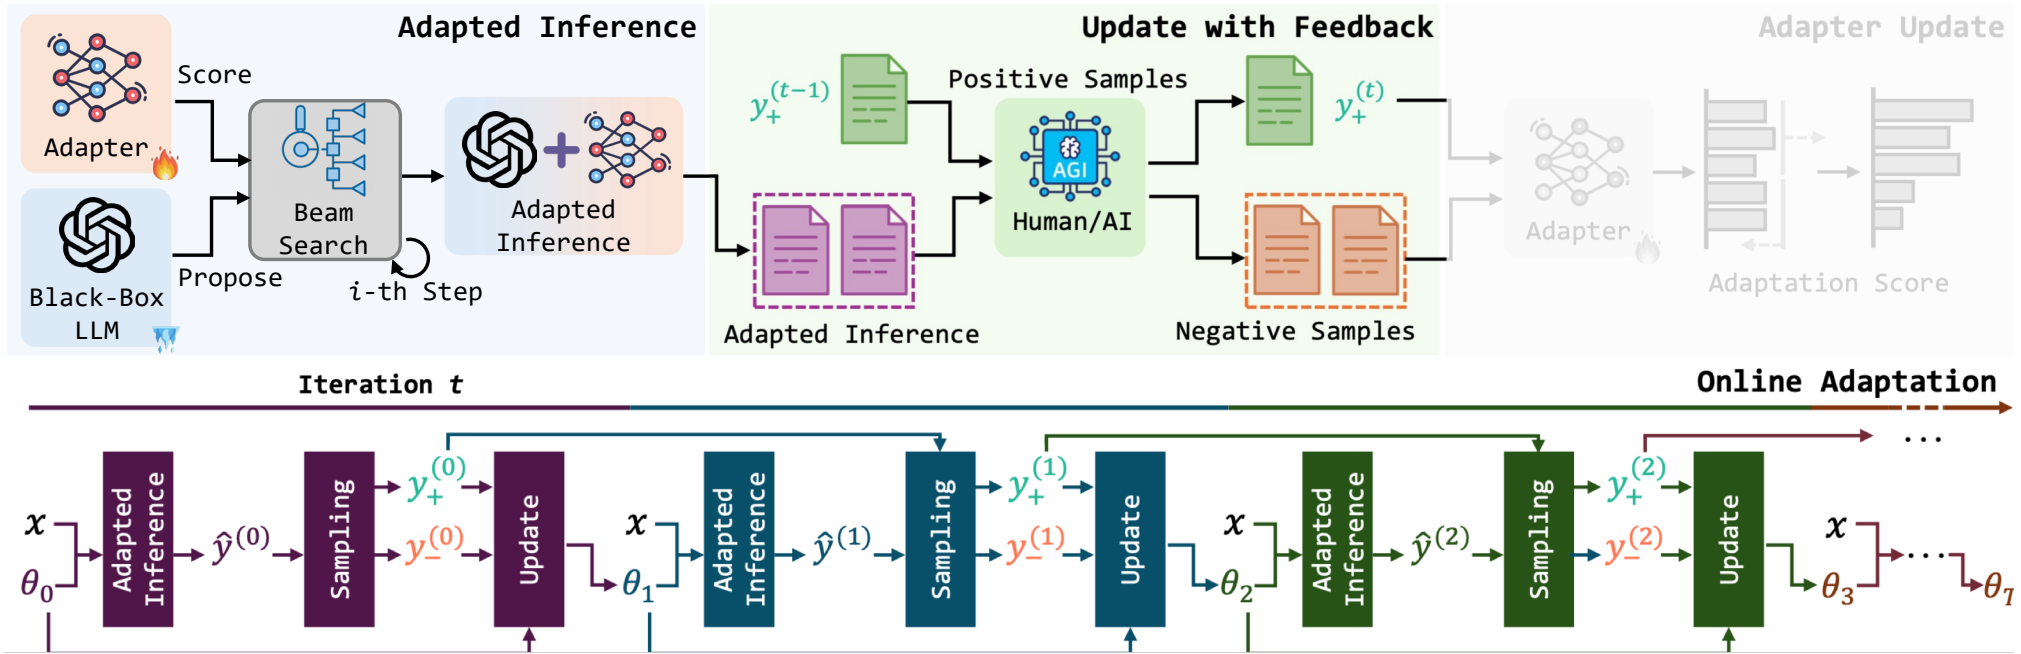
- 1) *Sampling from Adapted Inference.*
- 2) *Updating Training Data with Feedback.*
- 3) *Update Adapter Parameters.*



**Adapted Inference** employs a black-box Language Model (LM) to generate proposals and an adapter for evaluation. The process simplifies complex tasks into sentence-level searches by selecting and evaluating candidate sequences.

$$\begin{aligned}
 p_{\theta}(\mathbf{y}|\mathbf{x}) &= p_{\theta}(\mathbf{s}^{1:L}|\mathbf{x}) = p_{\text{LLM}}(\mathbf{s}^{1:L}|\mathbf{x}) \exp(g_{\theta}(\mathbf{s}^{1:L}, \mathbf{x})) \\
 &= \exp(g_{\theta}(\mathbf{s}^{1:L}, \mathbf{x})) \prod_l p_{\text{LLM}}(s^l|\mathbf{x}, \mathbf{s}^{1:l-1}).
 \end{aligned}$$

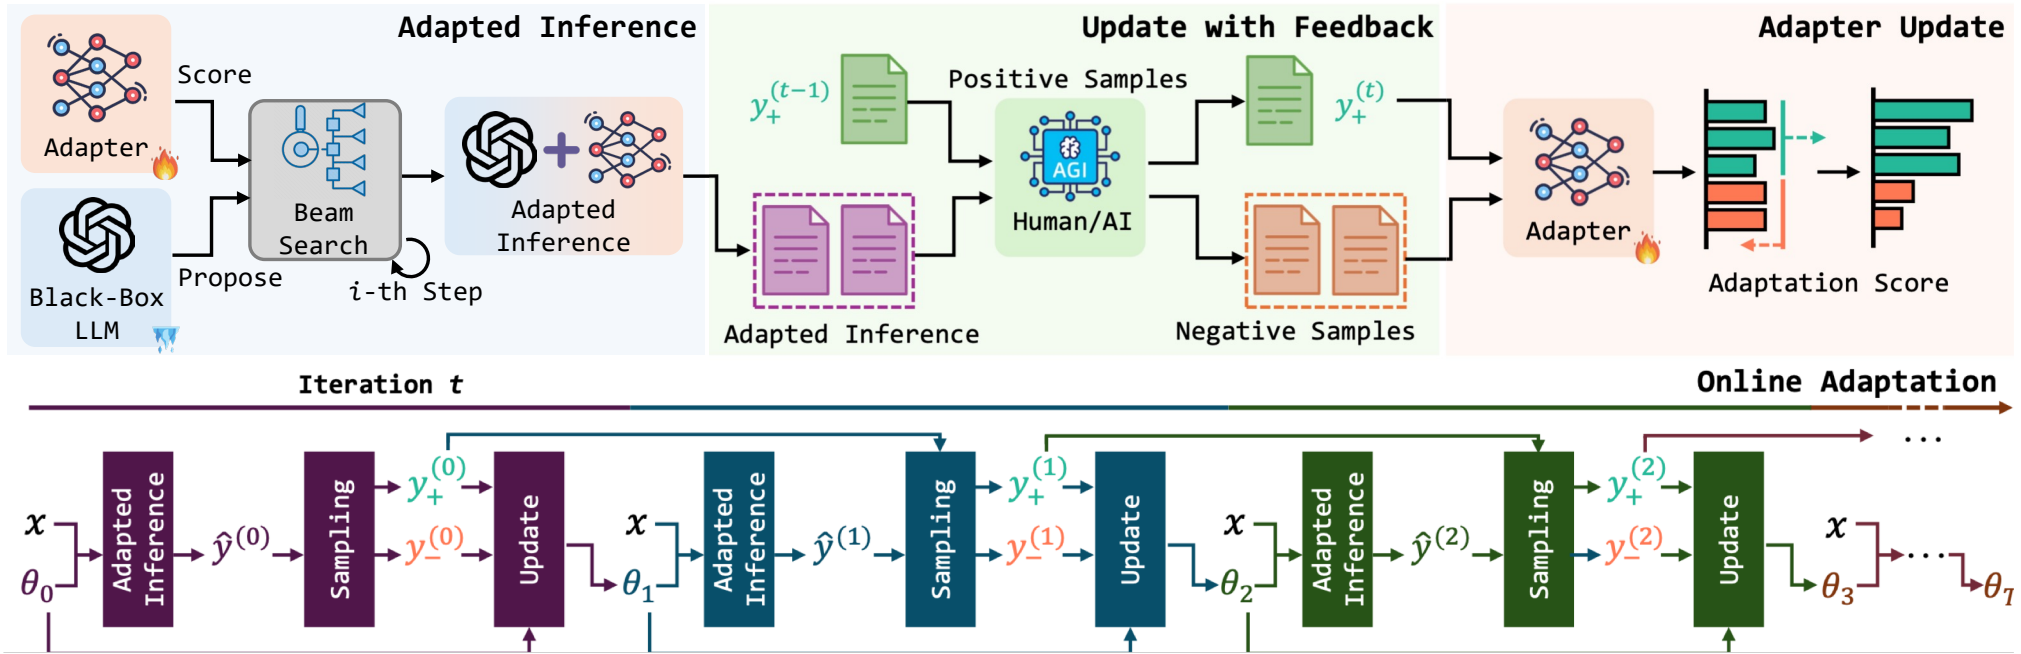




## Updating Training Data with Feedback

Positive samples are selected based on ground-truth solutions or AI feedback from advanced LLMs like GPT-4, which simulate human judgment. We also employ outcome supervision to categorize further adapted inferences into correct outcomes (positive) and incorrect (negative) and update the respective sets accordingly.

$$\begin{aligned}
 \mathbf{y}_{i+}^{(t)} &= \text{SEL}(\mathbf{y}_{i+}^{(t-1)}, \{\hat{\mathbf{y}}_{i,m}\}_{m=1}^M). \\
 \mathbf{y}_{i-}^{(t)} &= \{\hat{\mathbf{y}}_{i,m} | \hat{\mathbf{y}}_{i,m} \neq \mathbf{y}_{i+}^{(t)}\}_{m=1}^M.
 \end{aligned}$$



**Adapter Update.** With the updated positive and negative samples, we update the adapter parameters using the NCE loss, which minimizes energy for positive samples and increases it for negative ones.

$$\nabla_{\theta} \ell(\theta) = \nabla_{\theta} \left\{ -\mathbb{E}_{\mathbf{y}_+ \sim p_{\text{data}}(\mathbf{y}|\mathbf{x})} [g_{\theta}(\mathbf{x}, \mathbf{y}_+)] + \alpha \mathbb{E}[g_{\theta}(\mathbf{x}, \mathbf{y}_+)^2] + \mathbb{E}_{\mathbf{y}_- \sim p_{\theta}(\mathbf{y}|\mathbf{x})} [g_{\theta}(\mathbf{x}, \mathbf{y}_-)] + \alpha \mathbb{E}[g_{\theta}(\mathbf{x}, \mathbf{y}_-)^2] \right\}.$$

Dataset (→)	StrategyQA		GSM8K		TruthfulQA		ScienceQA	
	Acc. (%)	Δ (%)	Acc. (%)	Δ (%)	True + Info (%)	Δ (%)	Acc. (%)	Δ (%)
gpt-3.5-turbo (OpenAI, 2022)	66.59	-	67.51	-	77.00	-	72.90	-
Azure-SFT (Peng et al., 2023)	76.86	+10.27	69.94	+2.43	95.00	+18.00	79.00	+6.10
<b>BBOX-ADAPTER (Ground-Truth)</b>	71.62	+5.03	73.86	+6.35	79.70	+2.70	78.53	+5.63
<b>BBOX-ADAPTER (AI Feedback)</b>	69.85	+3.26	73.50	+5.99	82.10	+5.10	78.30	+5.40
<b>BBOX-ADAPTER (Combined)</b>	<b>72.27</b>	<b>+5.68</b>	<b>74.28</b>	<b>+6.77</b>	<b>83.60</b>	<b>+6.60</b>	<b>79.40</b>	<b>+6.50</b>

- Empirical experiments show that BBox-Adapter consistently outperforms gpt-3.5-turbo by an average of **6.39%** across all datasets, highlighting its efficacy in adapting black-box LLMs to various tasks.

Plugger (→)	BBOX-ADAPTER (gpt-3.5-turbo)							
Dataset (→)	StrategyQA		GSM8K		TruthfulQA		Average	
Black-Box LLMs (↓) / Metrics (→)	Acc. (%)	Δ (%)	Acc. (%)	Δ (%)	True + Info (%)	Δ (%)	Acc. (%)	Δ (%)
davinci-002	44.19	-	23.73	-	31.50	-	33.14	-
<b>davinci-002 (Plugged)</b>	<b>59.61</b>	<b>+15.42</b>	<b>23.85</b>	<b>+0.12</b>	<b>36.50</b>	<b>+5.00</b>	<b>39.99</b>	<b>+6.85</b>
Mixtral-8×7B	59.91	-	47.46	-	40.40	-	49.26	-
<b>Mixtral-8×7B (Plugged)</b>	<b>63.97</b>	<b>+4.06</b>	<b>47.61</b>	<b>+0.15</b>	<b>49.70</b>	<b>+9.30</b>	<b>53.76</b>	<b>+4.50</b>

- **Plug-and-Play**. Our trained adapter demonstrates an average performance improvement of **6.85%** and **4.50%** across all datasets.

Dataset (→)	StrategyQA			GSM8K		
	Acc.(%)	Training Cost (\$)	Inference Cost (\$)/1k Q	Acc.(%)	Training Cost (\$)	Inference Cost (\$)/1k Q
gpt-3.5-turbo	66.59	-	0.41	67.51	-	1.22
Azure-SFT (Peng et al., 2023)	76.86	153.00	7.50	69.94	216.50	28.30
BBOX-ADAPTER (Single-step)	69.87	2.77	2.20	71.13	7.54	3.10
BBOX-ADAPTER (Full-step)	71.62	3.48	5.37	74.28	11.58	12.46

- **Cost-efficiency**. BBox-Adapter achieves **5.90%** improvement over the base model with **31.30** times less training cost and **1.84** times less inference cost than the official SFT service.

- Case study on GSM8K

Q: An airport has only 2 planes that fly multiple times a day. Each day, the first plane goes to Greece for three-quarters of its flights, and the remaining flights are split equally between flights to France and flights to Germany. The other plane flies exclusively to Poland, and its 44 trips only amount to half the number of trips the first plane makes throughout each day. How many flights to France does the first plane take in one day?

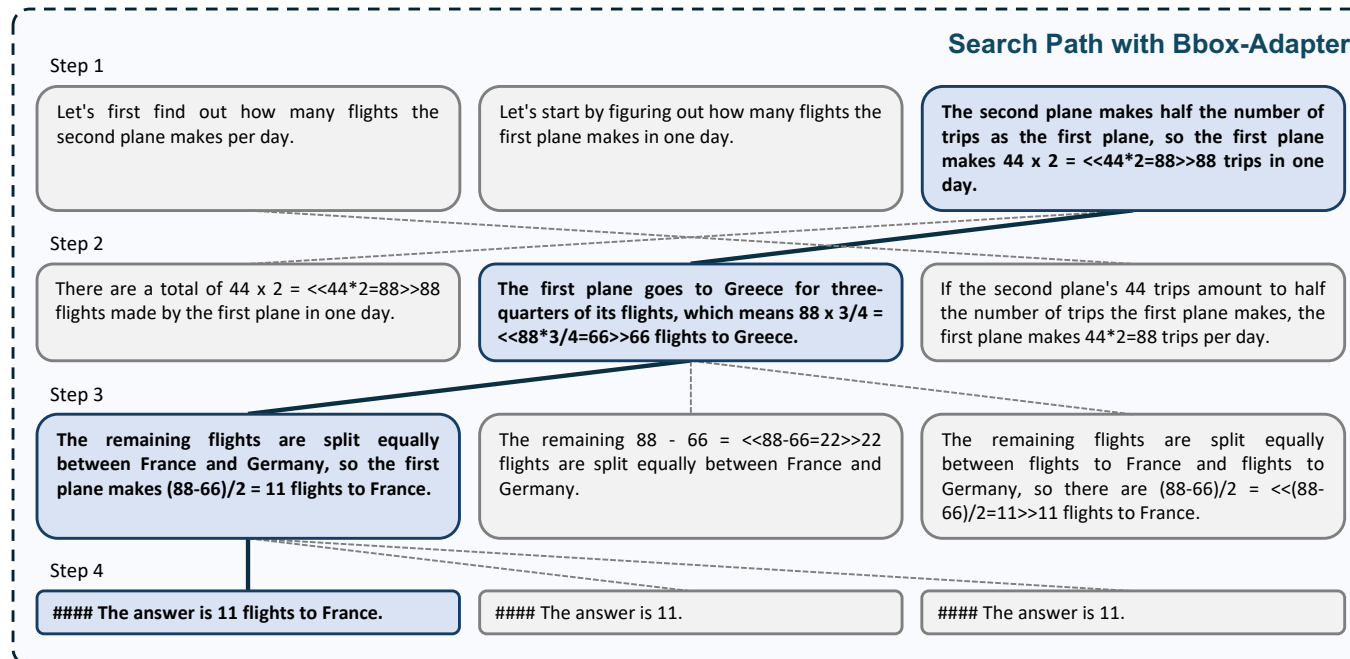
The first plane makes 4 trips to Greece (3/4 of its flights), and the remaining 1/4 is split equally between France and Germany. So, it makes  $1/4 * 4 = 1$  trip to France.  
The second plane makes 44 trips, which is half the number of trips the first plane makes. So, the first plane makes  $44 * 2 = 88$  trips in total.  
Therefore, the first plane takes 1 trip to France in one day.

#### The answer is 1 trip to France. ❌

CoT Result

The second plane makes half the number of trips as the first plane, so the first plane makes  $44 * 2 = 88$  trips in one day. The first plane goes to Greece for three-quarters of its flights, which means  $88 * 3/4 = 66$  flights to Greece. The remaining flights are split equally between France and Germany, so the first plane makes  $(88-66)/2 = 11$  flights to France. #### The answer is 11 flights to France. ✅

Adapted Result



# *Thank You!*

## BBox-Adapter: Lightweight Adapting for Black-Box Large Language Models

Haotian Sun<sup>1\*</sup>, Yuchen Zhuang<sup>1\*</sup>, Wei Wei<sup>2</sup>, Chao Zhang<sup>1</sup>, Bo Dai<sup>1</sup>

<sup>1</sup> Georgia Institute of Technology

<sup>2</sup>Accenture



Paper



Code



Website

