

FROM EXPECTED HARMFULNESS TO LIKELIHOOD: A PROBABILISTIC REFORMULATION OF JAILBREAK- ING LLM AGENTS

Juanyang Xu^{1*} Zheng Wang² Xingyu Zhao^{2,3} Siddhartha Khastgir² Andi Zhang^{2*†}

¹University of Macau ²WMG, University of Warwick ³ Wuhan University
andi.zhang@warwick.ac.uk

ABSTRACT

When the harmfulness of an LLM agent’s output can be quantified, a natural jail-breaking objective is to maximize expected harmfulness over admissible input modifications. An alternative approach constructs or selects harmful target outputs and modifies the input to increase their likelihood. We establish a precise connection between these two approaches through a probabilistic reformulation. Specifically, we show that the gradient of the logarithm of expected harmfulness with respect to the input equals the expected input gradient of the model’s log-likelihood under a harmfulness reweighted output distribution. This identity provides a unified interpretation of expected harmfulness and target likelihood optimization. Building on this connection, we propose **OPUR**, a sampling distribution designed to generate highly harmful target outputs and use the resulting samples to guide likelihood-based input optimization. Experiments demonstrate the effectiveness of the resulting method in jailbreaking LLM agents. Code is available at <https://github.com/Haxlon/OPUR.git>

1 INTRODUCTION

For LLM agents, attackers seek to bypass safety alignment mechanisms and induce harmful responses or malicious actions, a class of attacks commonly referred to as jailbreaks. Prior studies have shown that carefully crafted adversarial prompts or suffixes can effectively steer agents toward attacker-specified behaviors. Existing attacks typically follow a direct strategy: specifying a harmful target in advance and optimizing the likelihood that the model produces it (Zou et al., 2023; Liu et al., 2024b; Chao et al., 2025; Mehrotra et al., 2024; Zhang et al., 2025b). In other words, given a fixed input and adversarial suffix, the attacker attempts to increase the probability of a particular harmful response. An alternative, more distributional perspective is to maximize *expected harmfulness*. Rather than focusing on a single output, it weighs all possible model responses by both their probability of occurrence and their degree of harmfulness, thereby characterizing the overall harmfulness induced by an adversarial input over the output distribution (Geisler et al., 2025; Chen et al., 2026).

In this paper, we revisit this problem from a probabilistic perspective. As illustrated in Figure 1, we introduce a harmfulness-reweighted response distribution π that jointly reflects how likely a response is under the model and how harmful that response is. Intuitively, π characterizes the region an attacker would ideally like the model to enter: responses that are not only harmful, but also likely to be generated. From this perspective, target likelihood and expected harmfulness are no longer isolated objectives, but can be understood through the same response distribution. Because the response space is discrete and combinatorial, π is difficult to represent or sample from directly, motivating a practical surrogate for adversarial optimization.

Guided by this probabilistic formulation, we propose **OPUR**, which constructs a practical surrogate distribution for π from model-generated rollouts. OPUR considers multiple possible adversarial outcomes and prioritizes candidate directions that are both harmful and attainable under the model.

*Equal contribution.

†Corresponding Author.




Figure 1: **Overview of our probabilistic view of jailbreak optimization.** Given an adversarial input, the model induces a response distribution $p_{\text{dis}}(y) = p_{\theta}(y | x, s)$, while the harmfulness function induces a preference $p_{\text{vic}}(y) \propto \psi(y)$. Considering either view alone may favor responses that are likely but insufficiently harmful, or harmful but unlikely. Their normalized product defines the reweighted distribution π , which emphasizes responses that are simultaneously likely under the model and harmful under the attack objective. Since exact sampling from π may be unavailable in practice, its high-probability regions can be represented by constructed surrogate targets and optimized through their likelihood. As optimization proceeds, model probability mass is progressively shifted toward harmful outputs, providing a unified interpretation of expected harmfulness maximization and target-likelihood optimization.




Figure 2: An illustrative InjecAgent example comparing UDora and Opur. The agent encounters attacker-controlled content while answering a calendar query. The two methods construct different surrogate contexts for adversarial optimization, resulting in different final tool-call outputs in this case. Position scores and candidate sets are schematic.

This makes the attack better suited to the stochastic nature of agent generation and helps uncover effective adversarial directions that deterministic optimization may overlook. As shown in Figure 2, in the calendar-query example, UDora remains focused on the original calendar task, whereas OPUR explores alternative feasible trajectories and identifies one that triggers the attacker-designated tool call. In this way, OPUR turns the probabilistic perspective into a practical jailbreak strategy for LLM agents.

Our contributions are summarized as follows:

- **A unified probabilistic framework for jailbreak optimization.** We establish a gradient-level reformulation between expected harmfulness and target likelihood under a harmfulness-reweighted output distribution, providing a unified interpretation of behavioral and likelihood-based attack objectives.

- **OPUR: a practical realization of the framework.** Building on this probabilistic formulation, we propose **OPUR**, which approximates the theoretical distribution using model-generated rollouts and uses promising contexts to guide adversarial suffix optimization.
- **Higher ASR under stochastic generation.** We evaluate **OPUR** across multiple benchmarks under stochastic decoding and show that the proposed formulation translates into consistently higher targeted attack success rates than existing optimization-based baselines.

2 PRELIMINAIRES

Let $p_\theta(\cdot | x, s)$ be the response distribution of a large language model, where θ is the parameters of the LLM, x is the input and s is the adversarial suffix inside of x .

2.1 JAILBREAKING BASED ON EXPECTED HARMFULNESS

Let $\psi(y) \geq 0$ represents the harmfulness of the response y . Given input x , a natural way to find an adversarial suffix s is to maximize the expectation of the harmfulness:

$$s^* = \arg \max_s \mathbb{E}_{Y \sim p_\theta(\cdot | x, s)} [\psi(Y)], \quad (1)$$

Geisler et al. (2025) suggested to apply REINFORCE (Williams, 1992) on (Equation 1). In particular, by applying the log-derivative trick, the gradient of the expected harmfulness can be expressed as an expectation over responses sampled from the LLM, as formalized in Theorem 1.

Theorem 1 (REINFORCE Gradient for Expected Harmfulness). *Assume that $p_\theta(y | x, s)$ is differentiable with respect to s , its support does not depend on s , $\psi(y)$ does not explicitly depend on s , and differentiation and expectation can be interchanged. Then*

$$\nabla_s \mathbb{E}_{Y \sim p_\theta(\cdot | x, s)} [\psi(Y)] = \mathbb{E}_{Y \sim p_\theta(\cdot | x, s)} [\psi(Y) \nabla_s \log p_\theta(Y | x, s)]. \quad (2)$$

However, at the early stage of optimization, when the adversarial suffix s is still poorly formed, the response distribution $p_\theta(\cdot | x, s)$ is unlikely to generate harmful responses. In other words, $\psi(y) \approx 0$ for almost all $y \sim p_\theta(\cdot | x, s)$, resulting in little to no informative gradient signal for updating s . Geisler et al. (2025) explicitly acknowledges this limitation and introduces several heuristic rules to construct harmful responses as substitutes for samples from $p_\theta(\cdot | x, s)$. Proof is shown in Proof A.

2.2 LIKELIHOOD-BASED JAILBREAKING

Given a harmful response y_{harm} , a principled way is to maximize the likelihood of the y_{harm} by tuning the suffix s (Zou et al., 2023):

$$s^* = \arg \max_s \log p_\theta(y_{\text{harm}} | x, s). \quad (3)$$

Since the adversarial suffix s consists of discrete tokens, however, this gradient cannot be directly applied through a standard continuous update. Zou et al. (2023) addresses this issue by using the gradient to identify promising token substitutions and then performing greedy coordinate-wise updates.

2.3 PROBABILISTIC ADVERSARIAL ATTACK

Zhang et al. (2024; 2026) introduced the probabilistic adversarial attack framework, in which the adversarial distribution is defined as

$$p_{\text{adv}}(x_{\text{adv}} | x_{\text{ori}}, y_{\text{tar}}) \propto p_{\text{vic}}(x_{\text{adv}} | y_{\text{tar}}) p_{\text{dis}}(x_{\text{adv}} | x_{\text{ori}}), \quad (4)$$

where

$$p_{\text{vic}}(x_{\text{adv}} | y_{\text{tar}}) \propto \exp(-cf(x_{\text{adv}}, y_{\text{tar}})), \quad p_{\text{dis}}(x_{\text{adv}} | x_{\text{ori}}) \propto \exp(-D(x_{\text{ori}}, x_{\text{adv}})).$$

Here, p_{vic} assigns higher probability to samples that more strongly satisfy the adversarial objective, while p_{dis} favors samples that remain close to the original input under the distance measure D . Their product therefore defines a distribution over adversarial examples that explicitly balances attack effectiveness against deviation from the original input.

3 UNIFYING EXPECTED HARMFULNESS AND LIKELIHOOD-BASED JAILBREAKING

We now connect the expected-harmfulness objective in Section 2.1 with the likelihood-based objective in Section 2.2. Define the harmfulness-reweighted response distribution

$$\pi(y \mid x, s) = \frac{p_\theta(y \mid x, s)\psi(y)}{\mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)}[\psi(Y)]}, \quad (5)$$

assuming $\mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)}[\psi(Y)] > 0$. This distribution assigns high probability to responses that are both likely under the current model and highly harmful.

If the target response in likelihood-based jailbreaking is sampled from π , the corresponding expected likelihood objective is

$$\max_s \mathbb{E}_{Y_{\text{harm}} \sim \pi} [\log p_\theta(Y_{\text{harm}} \mid x, s)].$$

At a fixed s , its gradient is

$$\mathbb{E}_{Y_{\text{harm}} \sim \pi} [\nabla_s \log p_\theta(Y_{\text{harm}} \mid x, s)],$$

which, by Theorem 2, is exactly the gradient of the log expected harmfulness.

Theorem 2 (Equivalence between Expected Harmfulness and Likelihood Gradients). *Assuming differentiation and expectation can be interchanged, the gradient of the log expected harmfulness can be expressed as the expected log-likelihood gradient under the harmfulness-reweighted response distribution π :*

$$\nabla_s \log \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)}[\psi(Y)] = \mathbb{E}_{Y \sim \pi} [\nabla_s \log p_\theta(Y \mid x, s)] \quad (6)$$

Theorem 2 shows that, at a fixed s , maximizing expected harmfulness induces the same update direction as increasing the likelihood of responses drawn from the harmfulness-reweighted distribution π . Therefore, expected-harmfulness optimization can be viewed locally as likelihood-based optimization with adaptively constructed harmful targets. Proof is shown in Proof A.

By normalizing the identity in Theorem 1 by the expected harmfulness, we obtain

$$\nabla_s \log \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)}[\psi(Y)] = \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} \left[\frac{\psi(Y)}{\mathbb{E}_{Y' \sim p_\theta(\cdot \mid x, s)}[\psi(Y')]} \nabla_s \log p_\theta(Y \mid x, s) \right]. \quad (7)$$

Comparing Equations (6) and (7), both formulations lead to likelihood-gradient updates, but under different sampling distributions. Geisler et al. (2025) estimates the gradient using samples from $p_\theta(\cdot \mid x, s)$, whereas ours takes the expectation under the harmfulness-reweighted distribution π . By concentrating probability mass on harmful responses, π avoids the early-stage regime where model samples are almost entirely benign and provide little useful gradient signal, yielding a more principled sampling distribution for likelihood-based jailbreak optimization.

The distribution π also has the same product form as the probabilistic adversarial attack in Section 2.3:

$$\pi(y \mid x, s) \propto \underbrace{p_\theta(y \mid x, s)}_{p_{\text{dis}}(y)} \underbrace{\psi(y)}_{\propto p_{\text{vic}}(y)}. \quad (8)$$

As illustrated in Figure 1, π emphasizes responses that are both likely under the current model and highly harmful, suggesting them as likelihood targets for local suffix updates. Since direct sampling from π is difficult when harmful responses are rare, the remaining challenge is to construct useful surrogate targets, which we address in the next section using stochastic agent rollouts.

4 OPUR

This section introduces OPUR which builds on the harmfulness-reweighted distribution π . OPUR constructs rollout-derived surrogate targets to represent its high-probability harmful regions. It combines bootstrap-style exploration and probabilistic position selection to obtain targets that are both plausible under the current model and useful for likelihood-based suffix optimization. We next motivate this design and present the attack algorithm.

Algorithm 1 OPUR

```

1: Input: LLM  $\mathcal{M}$ , prompt  $x$ , target  $a$ , suffix  $s$ , maximum positions  $\ell$ , rollouts  $K_{\text{opt}}$ 
2: while budget remains do
3:    $\{(z_k, \mathcal{P}_k)\}_{k=1}^{K_{\text{opt}}} \leftarrow \text{Rollout}(\mathcal{M}, x, s)$  ▷ Stochastic rollouts
4:   break if  $\bigwedge_{k=1}^{K_{\text{opt}}} \text{Succ}(z_k, a)$ 
5:    $\mathcal{C}^* \leftarrow \emptyset$ 
6:   for  $k = 1, \dots, K_{\text{opt}}$  do
7:     Compute  $r_{k,j}(a)$  for each eligible position  $j$  using  $(z_k, \mathcal{P}_k)$ ;  $\mathcal{C}_k \leftarrow \emptyset$ 
8:     repeat
9:       Sample  $J$  ( $|J| \leq \ell$ , non-overlapping) from  $\text{Softmax}(\{r_{k,j}(a)\}_j / \tau)$  ▷ Probabilistic position sampling
10:      Obtain  $z_{k,J}^*$  by replacing the spans at  $J$  in  $z_k$  with  $a$ 
11:       $\mathcal{C}_k \leftarrow \mathcal{C}_k \cup \{z_{k,J}^*\}$ 
12:      until enough distinct position sets are collected or the sampling budget is exhausted
13:      Evaluate  $\mathcal{L}_U(s; c)$  for each  $c \in \mathcal{C}_k$ 
14:       $\mathcal{C}_k^* \leftarrow$  lowest-probe-loss contexts in  $\mathcal{C}_k$ 
15:       $\mathcal{C}^* \leftarrow \mathcal{C}^* \cup \mathcal{C}_k^*$  ▷ Probe screening
16:    end for
17:     $g \leftarrow \text{Mean}_{c \in \mathcal{C}^*} \nabla_{\text{onehot}(s)} \mathcal{L}_{\text{CE}}(s; c)$  ▷ Cross-rollout optimization
18:    Generate  $\mathcal{V}$  by gradient-guided token substitutions using  $g$ 
19:     $\tilde{\mathcal{L}}_U(v) \leftarrow \text{Mean}_{c \in \mathcal{C}^*} \mathcal{L}_U(v; c)$ , for each  $v \in \mathcal{V}$ 
20:     $s \leftarrow \arg \min_{v \in \mathcal{V}} \tilde{\mathcal{L}}_U(v)$ 
21:  end while
22: Output:  $s$ 

```

4.1 MOTIVATION

Our attack constructs surrogate optimization contexts from the agent’s own trajectories. Since stochastic decoding can yield different trajectories and positional scores for the same input, deterministic top-position selection may overfit to a single rollout. OPUR instead probabilistically explores multiple non-overlapping position sets and filters low-quality surrogate contexts before suffix optimization, balancing candidate diversity with optimization reliability.

4.2 ATTACK ALGORITHM

Notation. We denote by $\mathcal{M}$ the victim LLM agent, by x the agent prompt, by o an environmental observation, and by s the adversarial suffix, which is appended to x in the malicious-instruction setting and to o in the malicious-environment setting. Given the attacked input, $\mathcal{M}$ produces a stochastic rollout $z = (z_1, \dots, z_n)$ with token distributions $\mathcal{P} = (p_1, \dots, p_n)$. Let a denote the target segment and $r_j(a; z, \mathcal{P})$ the insertion score of an eligible span j . A sampled set J contains at most ℓ non-overlapping spans, whose replacement with a yields a surrogate context z_J^* . We use $\mathcal{C}_k$ for the surrogate contexts sampled from rollout k , $\mathcal{C}^*$ for those retained after probe screening across rollouts, $\mathcal{L}_U$ for screening and candidate evaluation, $\mathcal{L}_{\text{CE}}$ for suffix-gradient computation, and $\mathcal{V}$ for the resulting suffix candidates.

Attack Algorithm. As we demonstrated in Algorithm 1, our probabilistic adversarial attack algorithm have 4 main steps.

Step 1: Stochastic rollout generation. At each optimization step, we query the victim LLM K_{opt} times using the same attacked input and current suffix, obtaining stochastic responses and their token-level probability traces. Optimization terminates if every rollout meets the attack success criterion; otherwise, each rollout is used to identify candidate injection positions.

Step 2: Probabilistic position sampling. For each rollout, we score every eligible position j using its generated tokens and probability trace. Let $m_{k,j}$ be the number of consecutive target tokens matched by the model’s argmax predictions, $\bar{p}_{k,j}$ their mean target-token probability up to

and including the first mismatch or the last available token, and $L_{k,j}$ the target’s token length at position j . The position score is

$$r_{k,j}(a) = \frac{m_{k,j} + \bar{p}_{k,j}}{L_{k,j} + 1}. \quad (9)$$

Instead of always choosing the highest-scoring positions, we sample from a temperature-controlled distribution:

$$q_\tau(j \mid z_k, \mathcal{P}_k) = \frac{\exp(r_{k,j}(a)/\tau)}{\sum_u \exp(r_{k,u}(a)/\tau)}, \quad (10)$$

where the denominator includes only currently eligible positions. After each draw, the selected position and overlapping positions are removed. And the distribution is renormalized. Repeated sampling produces distinct sets J of at most ℓ non-overlapping positions. Replacing their corresponding spans with a yields surrogate contexts $z_{k,J}^*$, collected in $\mathcal{C}_k$.

Step 3: Screen sampled contexts. Before computing suffix gradients, we evaluate the sampled contexts with the current suffix s fixed. For a target span of length L in context c , let m_c and $\bar{p}_c$ denote the consecutive argmax matches and mean target-token probability up to the first mismatch, now evaluated on c . The surrogate loss from Zhang et al. (2025b) is

$$\mathcal{L}_U(s; c) = -\frac{m_c + \bar{p}_c}{L + 1}. \quad (11)$$

For multiple target spans, their losses are summed. We retain the lowest-loss contexts from each $\mathcal{C}_k$ in $\mathcal{C}_k^*$ and combine them across rollouts into $\mathcal{C}^*$.

Step 4: Update the adversarial suffix. For each retained context, we compute the target-sequence cross-entropy gradient with respect to the one-hot representation of the current suffix. These gradients are averaged across contexts, allowing the update to account for multiple stochastic rollouts.

The aggregated gradient is used to generate suffix candidates through token substitutions. Each candidate is evaluated by its mean surrogate loss over the retained contexts. And the candidate with the lowest mean loss becomes the suffix for the next optimization step.

5 EXPERIMENT

In this section, we present the main results of our method **OPUR** on three datasets across two different scenarios. These results demonstrate the effectiveness of our proposed probabilistic adversarial attack framework for red-teaming LLM agents.

5.1 DATASETS

Following the evaluation protocol of UDora (Zhang et al., 2025b), we evaluate our method on three LLM-agent benchmarks covering two complementary attack scenarios: *malicious instructions* and *malicious environments*. AgentHarm (Andriushchenko et al., 2025b) is used for the malicious-instruction setting and evaluates whether agents can be induced to execute explicitly harmful multi-step tasks while retaining sufficient task-completion capability. InjecAgent (Zhan et al., 2024) represents the malicious-environment setting through indirect prompt injections embedded in external observations, covering attacks such as *direct harm* and *data stealing*. WebShop (Yao et al., 2022a) provides an interactive e-commerce environment with multi-step search and click actions, allowing us to evaluate whether adversarially manipulated observations can steer agents toward unintended target products.

5.2 EXPERIMENT SETTING

Baselines. We compare against three representative optimization-based jailbreak attacks: GCG, REINFORCE-GCG, and UDora. GCG (Zou et al., 2023) performs gradient-guided discrete suffix optimization toward a predefined affirmative target. REINFORCE-GCG (Geisler et al., 2025) retains the GCG search procedure but replaces the fixed target-likelihood objective with an adaptive semantic reward. We also evaluate the *Sequential* and *Joint* variants of UDora (Zhang et al., 2025b), which optimize perturbations over intermediate reasoning trajectories.

Base LLMs. We evaluate all methods on two open-source instruction-tuned models, Llama-3.1-8B-Instruct (Grattafiori et al., 2024) and Ministral-8B-Instruct-2410 (Jiang et al., 2023). architectures.

Metrics. We use Attack Success Rate (ASR) (Chu et al., 2025) as the primary metric for evaluating jailbreak effectiveness:

$$\text{ASR} = \frac{N_{\text{success}}}{N_{\text{total}}} \times 100\%, \quad (12)$$

where N_{success} denotes the number of successfully attacked samples and N_{total} denotes the total number of evaluation samples.

Experiment Setting.

All agents use ReAct prompting (Yao et al., 2022b). For InjecAgent and WebShop, the search width, gradient top- k , and optimization steps are set to 128, 32, and 500, respectively, while AgentHarm uses 256, 64, and 1000. All adversarial suffixes are initialized with 25 repeated $\times$ tokens. The injected target is the adversarial tool name for InjecAgent and AgentHarm and the target product ASIN for WebShop. For UDora, we set `num_locations` to 2 for the Sequential variant and 4 for the Joint variant, corresponding to their strongest greedy configurations. For Opur, each rollout samples four distinct injection-position sets (`injection_num_samples=4`) from the Softmax distribution and retains the two candidates with the lowest probe loss (`injection_num_keep=2`) for subsequent suffix optimization. We evaluate UDora Sequential, UDora Joint, and Opur with optimization rollout budget $K_{\text{opt}} \in \{1, 2, 3\}$. During optimization, each baseline follows its original stochastic rollout configuration, while final evaluation uses a single rollout ($K_{\text{eval}} = 1$) with temperature 0.7 and top- $p = 0.9$.

5.3 RESULTS

For the malicious-instruction scenario, Table 1 summarizes the results on AgentHarm, with a qualitative example shown in Figure 3. On Llama-3.1-8B-Instruct, Opur achieves an average ASR of 75.00%, outperforming the strongest UDora baseline by 13.07% and showing consistent gains across all prompt settings. On Ministral-8B-Instruct-2410, performance is already near saturation, where Opur achieves 99.43% ASR and matches the strongest UDora variants.

Table 1: Attack success rates (%) on AgentHarm under different attack categories. The best results for each model are highlighted in bold.

Model	Method	Attack Categories				Avg. ASR
		Detailed Prompt		Simple Prompt		
		w/ Hint	w/o Hint	w/ Hint	w/o Hint	
Llama-3.1-8B-Instruct	GCG	38.64%	38.64%	40.91%	31.82%	37.50%
	Reinforce-GCG	59.09%	50.00%	50.00%	43.18%	50.57%
	UDora (Sequential)	56.82%	63.64%	61.36%	65.91%	61.93%
	UDora (Joint)	59.09%	56.82%	63.64%	65.91%	61.36%
	Opur (Ours)	72.73%	75.00%	72.73%	79.55%	75.00%
Ministral-8B-Instruct-2410	GCG	93.18%	79.55%	93.18%	90.19%	89.20%
	Reinforce-GCG	100.00%	97.73%	100.00%	95.45%	98.30%
	UDora (Sequential)	100.00%	100.00%	100.00%	97.73%	99.43%
	UDora (Joint)	100.00%	100.00%	100.00%	97.73%	99.43%
	Opur (Ours)	100.00%	97.73%	100.00%	100.00%	99.43%

For InjecAgent, Table 2 summarizes the results, with a qualitative example shown in Figure 4. Opur achieves the highest average ASR on both models, reaching 49% on Llama-3.1-8B-Instruct and 40% on Ministral-8B-Instruct-2410, each improving over the strongest baseline by 2 percentage points. Although the gains are smaller than on AgentHarm, they remain consistent across both model families.

For WebShop, Table 3 reports the results, with a qualitative example shown in Figure 5. Opur achieves the highest average ASR on both models, reaching 26.67% on Llama-3.1-8B-Instruct and

Table 2: Attack success rates (%) on InjecAgent under different attack categories. The best results for each model are highlighted in bold.

Model	Method	Attack Categories		Avg. ASR
		Direct Harm	Data Stealing	
Llama-3.1-8B-Instruct	GCG	26%	34%	30%
	REINFORCE-GCG	46%	48%	47%
	UDora (Sequential)	44%	44%	44%
	UDora (Joint)	36%	56%	46%
	Opur (Ours)	48%	50%	49%
Ministral-8B-Instruct-2410	GCG	32%	30%	31%
	REINFORCE-GCG	44%	32%	38%
	UDora (Sequential)	26%	28%	27%
	UDora (Joint)	20%	36%	28%
	Opur (Ours)	42%	38%	40%

36.67% on Ministral-8B-Instruct-2410, outperforming the strongest baseline by 8.34% on both models. It performs best on three of four attack categories on Llama and on all four categories on Ministral, demonstrating strong effectiveness in the WebShop malicious-environment setting.

Table 3: Attack success rates (%) on WebShop under different attack categories. The best results for each model are highlighted in bold.

Model	Method	Attack Categories				Avg. ASR
		Price Mismatch	Attribute Mismatch	Category Mismatch	All Mismatch	
Llama-3.1-8B-Instruct	GCG	0.00%	0.00%	0.00%	0.00%	0.00%
	REINFORCE-GCG	26.67%	0.00%	0.00%	0.00%	6.67%
	UDora (Sequential)	6.67%	40.00%	13.33%	13.33%	18.33%
	UDora (Joint)	20.00%	33.33%	6.67%	6.67%	16.67%
	Opur (Ours)	13.33%	40.00%	33.33%	20.00%	26.67%
Ministral-8B-Instruct-2410	GCG	6.67%	13.33%	6.67%	0.00%	6.67%
	REINFORCE-GCG	13.33%	6.67%	6.67%	0.00%	6.67%
	UDora (Sequential)	26.67%	40%	33.33%	13.33%	28.33%
	UDora (Joint)	13.33%	26.67%	33.33%	6.67%	20.00%
	Opur (Ours)	33.33%	53.33%	40.00%	20.00%	36.67%

5.4 ABLATION

We study the effect of the optimization rollout budget $K_{\text{opt}} \in \{1, 2, 3\}$, which controls how many stochastic rollouts contribute to each suffix update; corresponding UDora ablations are reported in Appendix B. As shown in Table 4, larger K_{opt} generally improves Opur, although the gains are not strictly monotonic across benchmarks. On Llama-3.1-8B-Instruct, ASR increases from 56.82% to 75.00% on AgentHarm and from 5.00% to 26.67% on WebShop, while InjecAgent improves from 46.00% to 49.00%. Similar gains are observed on Ministral for InjecAgent (30.00% $\rightarrow$ 40.00%) and WebShop (11.67% $\rightarrow$ 36.67%). We omit larger K_{opt} values for Ministral on AgentHarm because performance already saturates at 99.43% for $K_{\text{opt}} = 1$.

6 RELATED WORK

Adversarial Attacks on LLMs. LLM jailbreaks span gradient-based, gradient-free, and automated black-box search. Representative optimization-based methods include GCG (Zou et al., 2023), AutoDAN (Liu et al., 2024b), BEAST (Sadasivan et al., 2024), AdvPrompter (Paulus et al., 2024),

Table 4: Ablation study on the optimization rollout budget K_{opt} for Opur. We report average attack success rates (%) on three benchmarks. “—” denotes configurations that are not evaluated because performance has already saturated.

Model	K_{opt}	Datasets		
		AgentHarm	InjecAgent	WebShop
Llama-3.1-8B-Instruct	1	56.82%	46.00%	5.00%
	2	63.07%	45.00%	15.00%
	3	75.00%	49.00%	26.67%
Ministral-8B-Instruct-2410	1	99.43%	30.00%	11.67%
	2	—	34.00%	16.67%
	3	—	40.00%	36.67%

and COLD-Attack (Guo et al., 2024). Black-box and automated strategies further include PAIR (Chao et al., 2025), TAP (Mehrotra et al., 2024), DRA (Liu et al., 2024a), MASTERKEY (Deng et al., 2023), LLM-Fuzzer (Yu et al., 2024), DeepInception (Li et al., 2023), persuasive jailbreaks (Zeng et al., 2024), and AutoDAN-Turbo (Liu et al., 2025). While these works primarily advance how adversarial prompts are searched or constructed, we instead revisit the probabilistic objective underlying jailbreak optimization.

Jailbreak Objectives and Probabilistic Attacks. Recent work has revisited jailbreak objectives from target- and distribution-level perspectives. AdvPrefix selects model-dependent targets (Zhu et al., 2026), REINFORCE optimizes adaptive semantic rewards (Geisler et al., 2025), and AmpleGCG models diverse adversarial suffixes (Liao & Sun, 2024). Other attacks explore adaptive search, black-box optimization, output-prefix forcing, trigger-level objectives, and iterative generation (Andriushchenko et al., 2025a; Deng et al., 2023; Jiang et al., 2025; Wang et al., 2025; Sun et al., 2025; Wu et al., 2025; Wang & Qi, 2025). Sampling-aware and distribution-based attacks further exploit stochastic generation and output distributions (Beyer et al., 2026; Fu et al., 2026), while recent benchmarks highlight sensitivity to generation and evaluation protocols (Chao et al., 2024; Souly et al., 2024; Chu et al., 2025; Scholten et al., 2025). In contrast, we derive a tractable likelihood-based objective directly from expected harmfulness.

Jailbreaking LLM Agents. LLM agents introduce broader attack surfaces through tools, memory, retrieval, and multi-step actions (Greshake et al., 2023; Mo et al., 2024). InjecAgent (Zhan et al., 2024), AgentDojo (Debenedetti et al., 2024), Agent Security Bench (Zhang et al., 2025a), and AgentHarm (Andriushchenko et al., 2025b) evaluate these vulnerabilities, while UDora (Zhang et al., 2025b), BadAgent (Wang et al., 2024), and AgentPoison (Chen et al., 2024) target agent reasoning, memory, or behavior. Unlike these agent-specific mechanisms, we focus on reformulating the objective used to optimize adversarial agent behaviour.

7 CONCLUSION

The security of LLM agents is fundamentally distributional: an attack does not merely seek a single harmful response, but attempts to shift the model’s behavioral distribution toward harmful outcomes. In this work, we reformulate expected harmfulness with likelihood through a harmfulness-reweighted response distribution, and instantiate this perspective with **OPUR**, which constructs a practical surrogate distribution from model-generated rollouts. Our experiments show that this probabilistic view leads to stronger targeted attacks under stochastic generation.

More broadly, our results suggest a different way to think about jailbreaks. Rather than treating harmful generations as isolated failures, they can be viewed as changes in how probability mass is distributed over possible agent behaviors. From this perspective, likelihood optimization becomes a practical interface between behavioral risk and adversarial optimization. We hope this view encourages future work to move beyond individual attack trajectories and toward understanding, evaluating, and ultimately controlling the full distribution of behaviors that LLM agents may exhibit.

ACKNOWLEDGMENTS

The work presented in this paper has been supported by UKRI Future Leaders Fellowship (Grant MR/S035176/1) and funded by the European Union under the Horizon Europe project AIGGRE-GATE (AI-enhanced collective intelligence for resilient, ethical and user-centric awareness and decision making in CCAM applications, Grant Agreement No. 101202457). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

AI USE STATEMENT

In this work, we used generative AI tools to aid and polish the writing, including improving clarity, grammar, and presentation, and for retrieval and discovery purposes, such as identifying potentially relevant literature. All retrieved references and related-work suggestions were independently verified by the authors against the original sources before being included in the paper.

We did not use generative AI tools for research ideation or execution, drafting substantive sections of the paper, generating synthetic datasets, or proving mathematical claims. All technical arguments, mathematical derivations, experimental designs, implementations, analyses, and conclusions were developed and verified by the authors.

We have reviewed all AI-assisted content and take responsibility for the final content of this work, including any text, claims, or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work investigates adversarial attacks against LLM agents and thus has potential dual-use implications. The proposed framework is intended to expose limitations of current agent safety mechanisms under stochastic generation and to facilitate more rigorous adversarial evaluation. However, techniques that improve attack effectiveness could also be misused to elicit unsafe behaviors from deployed agents.

We mitigate these risks by conducting all experiments in controlled research environments using established benchmarks, without targeting real users, accounts, or deployed services. This work does not involve human subjects or the collection of private user data. To support transparency and reproducibility, we make our code and benchmark configurations publicly available for research and safety evaluation purposes. The released artifacts are intended to facilitate independent verification, benchmarking, and the development of stronger defenses rather than real-world exploitation.

We believe that systematically identifying these vulnerabilities is important for developing stronger defenses and more reliable safety evaluations for LLM agents. By openly documenting the attack methodology and evaluation setting, we aim to help the research community better understand failure modes of current agent systems and improve their robustness. We encourage the use of the released methods, code, and benchmarks for responsible security research and the development of safer and more trustworthy AI systems.

REPRODUCIBILITY STATEMENT

We provide an anonymous implementation of our method, evaluation code, and experiment configurations at <https://github.com/Haxlon/OPUR.git>. The main paper specifies the attack formulation, datasets, baselines, model settings, optimization parameters, decoding configurations, and evaluation protocol used throughout our experiments. All experiments are conducted on publicly available benchmarks and model checkpoints. Together with the released code and configurations, these details are intended to support faithful reproduction of the reported results.

REFERENCES

Maksym Andriushchenko, Nicolas Flammarion, et al. Jailbreaking leading safety-aligned llms with simple adaptive attacks. In *International Conference on Learning Representations*, volume 2025, pp. 40116–40143, 2025a.

- Maksym Andriushchenko, Alexandra Souly, Mateusz Dziemian, Derek Duenas, Maxwell Lin, Justin Wang, Dan Hendrycks, Andy Zou, Zico Kolter, Matt Fredrikson, et al. Agentharm: A benchmark for measuring harmfulness of llm agents. In *International Conference on Learning Representations*, volume 2025, pp. 79185–79220, 2025b.
- Tim Beyer, Yan Scholten, Leo Schwinn, and Stephan Günnemann. Sampling-aware adversarial attacks against large language models. In *International Conference on Learning Representations*, volume 2026, pp. 74637–74663, 2026.
- Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwal, Edgar Dobriban, Nicolas Flammarion, George J Pappas, Florian Tramer, et al. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. *Advances in Neural Information Processing Systems*, 37:55005–55029, 2024.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pp. 23–42. IEEE, 2025.
- Yen-Shan Chen, Zhi Rui Tam, Cheng-Kuang Wu, and Yun-Nung Chen. Expected harm: Rethinking safety evaluation of (mis) aligned llms. *arXiv preprint arXiv:2602.01600*, 2026.
- Zhaorun Chen, Zhen Xiang, Chaowei Xiao, Dawn Song, and Bo Li. Agentpoison: Red-teaming llm agents via poisoning memory or knowledge bases. *Advances in Neural Information Processing Systems*, 37:130185–130213, 2024.
- Junjie Chu, Yugeng Liu, Ziqing Yang, Xinyue Shen, Michael Backes, and Yang Zhang. Jailbreakradar: Comprehensive assessment of jailbreak attacks against llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 21538–21566, 2025.
- Edoardo Debenedetti, Jie Zhang, Mislav Balunovic, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for llm agents. *Advances in neural information processing systems*, 37:82895–82920, 2024.
- Gelei Deng, Yi Liu, Yuekang Li, Kailong Wang, Ying Zhang, Zefeng Li, Haoyu Wang, Tianwei Zhang, and Yang Liu. Masterkey: Automated jailbreak across multiple large language model chatbots. *arXiv preprint arXiv:2307.08715*, 2023.
- Shuangjie Fu, Du Su, Xin Chen, Fei Sun, Huawei Shen, and Xueqi Cheng. Steering away from refusal: A black-box jailbreak method based on first-token distribution. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 25969–25979, 2026.
- Simon Geisler, Tom Wollschläger, MHI Abdalla, Vincent Cohen-Addad, Johannes Gasteiger, and Stephan Günnemann. Reinforce adversarial attacks on large language models: An adaptive, distributional, and semantic objective. *arXiv preprint arXiv:2502.17254*, 2025.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. Not what you’ve signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In *Proceedings of the 16th ACM workshop on artificial intelligence and security*, pp. 79–90, 2023.
- Xingang Guo, Fangxu Yu, Huan Zhang, Lianhui Qin, and Bin Hu. Cold-attack: Jailbreaking llms with stealthiness and controllability. *arXiv preprint arXiv:2402.08679*, 2024.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

- Weipeng Jiang, Zhenting Wang, Juan Zhai, Shiqing Ma, Zhengyu Zhao, and Chao Shen. An optimizable suffix is worth a thousand templates: Efficient black-box jailbreaking without affirmative phrases via llm as optimizer. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 5486–5498, 2025.
- Xuan Li, Zhanke Zhou, Jianing Zhu, Jiangchao Yao, Tongliang Liu, and Bo Han. Deepinception: Hypnotize large language model to be jailbreaker. *arXiv preprint arXiv:2311.03191*, 2023.
- Zeyi Liao and Huan Sun. Amplegcg: Learning a universal and transferable generative model of adversarial suffixes for jailbreaking both open and closed llms. *arXiv preprint arXiv:2404.07921*, 2024.
- Tong Liu, Yingjie Zhang, Zhe Zhao, Yinpeng Dong, Guozhu Meng, and Kai Chen. Making them ask and answer: Jailbreaking large language models in few queries via disguise and reconstruction. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 4711–4728, 2024a.
- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. In *International Conference on Learning Representations*, volume 2024, pp. 56174–56194, 2024b.
- Xiaogeng Liu, Peiran Li, G Edward Suh, Yevgeniy Vorobeychik, Zhuoqing Mao, Somesh Jha, Patrick McDaniel, Huan Sun, Bo Li, and Chaowei Xiao. Autodan-turbo: A lifelong agent for strategy self-exploration to jailbreak llms. In *International Conference on Learning Representations*, volume 2025, pp. 10313–10360, 2025.
- Anay Mehrotra, Manolis Zampetakis, Paul Kassianik, Blaine Nelson, Hyrum Anderson, Yaron Singer, and Amin Karbasi. Tree of attacks: Jailbreaking black-box llms automatically. *Advances in Neural Information Processing Systems*, 37:61065–61105, 2024.
- Lingbo Mo, Zeyi Liao, Boyuan Zheng, Yu Su, Chaowei Xiao, and Huan Sun. A trembling house of cards? mapping adversarial attacks against language agents. *arXiv preprint arXiv:2402.10196*, 2024.
- Anselm Paulus, Arman Zharmagambetov, Chuan Guo, Brandon Amos, and Yuandong Tian. Advprompter: Fast adaptive adversarial prompting for llms. *arXiv preprint arXiv:2404.16873*, 2024.
- Vinu Sankar Sadasivan, Shoumik Saha, Gaurang Sriramanan, Priyatham Kattakinda, Atoosa Chegini, and Soheil Feizi. Fast adversarial attacks on language models in one gpu minute. *arXiv preprint arXiv:2402.15570*, 2024.
- Yan Scholten, Stephan Günnemann, and Leo Schwinn. A probabilistic perspective on unlearning and alignment for large language models. In *International Conference on Learning Representations*, volume 2025, pp. 49420–49436, 2025.
- Alexandra Souly, Qingyuan Lu, Dillon Bowen, Tu Trinh, Elvis Hsieh, Sana Pandey, Pieter Abbeel, Justin Svegliato, Scott Emmons, Olivia Watkins, et al. A strongreject for empty jailbreaks. *Advances in Neural Information Processing Systems*, 37:125416–125440, 2024.
- Chung-En Sun, Xiaodong Liu, Weiwei Yang, Tsui-Wei Weng, Hao Cheng, Aidan San, Michel Galley, and Jianfeng Gao. Iterative self-tuning llms for enhanced jailbreaking capabilities. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 5768–5786, 2025.
- Daisy Zhe Wang and Yanjun Qi. Augmented adversarial trigger learning. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 7083–7115, 2025.
- Yifei Wang, Dizhan Xue, Shengjie Zhang, and Shengsheng Qian. Badagent: Inserting and activating backdoor attacks in llm agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9811–9827, 2024.
- Yiwei Wang, Muhao Chen, Nanyun Peng, and Kai-Wei Chang. Vulnerability of large language models to output prefix jailbreaks: Impact of positions on safety. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 3939–3952, 2025.

- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.
- Su Huang Wu, Huimin Wang, Yutian Zhao, Xian Wu, Yefeng Zheng, Wei Li, Hui Li, and Rongrong Ji. Monte carlo tree search based prompt autogeneration for jailbreak attacks against llms. In *Proceedings of the 31st International Conference on Computational Linguistics*, pp. 1057–1068, 2025.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757, 2022a.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022b.
- Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. {LLM-Fuzzer}: Scaling assessment of large language model jailbreaks. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 4657–4674, 2024.
- Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyan Shi. How johnny can persuade llms to jailbreak them: Rethinking persuasion to challenge ai safety by humanizing llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14322–14350, 2024.
- Qiusi Zhan, Zhixiang Liang, Zifan Ying, and Daniel Kang. Injecagent: Benchmarking indirect prompt injections in tool-integrated large language model agents. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 10471–10506, 2024.
- Andi Zhang, Mingtian Zhang, and Damon Wischik. Constructing semantics-aware adversarial examples with a probabilistic perspective. *Advances in Neural Information Processing Systems*, 37: 136259–136285, 2024.
- Andi Zhang, Xuan Ding, Steven McDonagh, and Samuel Kaski. Concept-based adversarial attack: a probabilistic perspective. In *International Conference on Learning Representations*, volume 2026, pp. 73613–73646, 2026.
- Hanrong Zhang, Jingyuan Huang, Kai Mei, Yifei Yao, Zhenting Wang, Chenlu Zhan, Hongwei Wang, and Yongfeng Zhang. Agent security bench (asb): Formalizing and benchmarking attacks and defenses in llm-based agents. In *International Conference on Learning Representations*, volume 2025, pp. 35331–35366, 2025a.
- Jiawei Zhang, Shuang Yang, and Bo Li. Udora: A unified red teaming framework against llm agents by dynamically hijacking their own reasoning. *arXiv preprint arXiv:2503.01908*, 2025b.
- Sicheng Zhu, Brandon Amos, Yuandong Tian, Chuan Guo, and Ivan Evtimov. Advprefix: An objective for nuanced llm jailbreaks. *Advances in Neural Information Processing Systems*, 38: 89990–90014, 2026.
- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.

APPENDIX

A PROOF OF THE THEOREMS

Theorem 1 (REINFORCE Gradient for Expected Harmfulness). *Assume that $p_\theta(y \mid x, s)$ is differentiable with respect to s , its support does not depend on s , $\psi(y)$ does not explicitly depend on s , and differentiation and expectation can be interchanged. Then*

$$\nabla_s \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} [\psi(Y)] = \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} [\psi(Y) \nabla_s \log p_\theta(Y \mid x, s)].$$

Proof.

$$\begin{aligned} \nabla_s \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} [\psi(Y)] &= \nabla_s \int p_\theta(y \mid x, s) \psi(y) dy \\ &= \int \psi(y) \nabla_s p_\theta(y \mid x, s) dy \\ &= \int p_\theta(y \mid x, s) \psi(y) \nabla_s \log p_\theta(y \mid x, s) dy \\ &= \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} [\psi(Y) \nabla_s \log p_\theta(Y \mid x, s)]. \end{aligned}$$

□

Theorem 2 (Equivalence between Expected Harmfulness and Likelihood Gradients). *Assuming differentiation and expectation can be interchanged, the gradient of the log expected harmfulness can be expressed as the expected log-likelihood gradient under the harmfulness-reweighted response distribution π :*

$$\nabla_s \log \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} [\psi(Y)] = \mathbb{E}_{Y \sim \pi} [\nabla_s \log p_\theta(Y \mid x, s)]$$

Proof.

$$\begin{aligned} \nabla_s \log \mathbb{E}_{Y \sim p_\theta(\cdot \mid x, s)} [\psi(r)] &= \nabla_s \log \int p_\theta(y \mid x, s) \psi(y) dy \\ &= \frac{\nabla_s \int p_\theta(y \mid x, s) \psi(y) dy}{\int p_\theta(y \mid x, s) \psi(y) dy} \\ &= \frac{\int \psi(y) \nabla_s p_\theta(y \mid x, s) dy}{\int p_\theta(y \mid x, s) \psi(y) dy} \\ &= \frac{\int \psi(y) p_\theta(y \mid x, s) \nabla_s \log p_\theta(y \mid x, s) dy}{\int p_\theta(y \mid x, s) \psi(y) dy} \\ &= \int \pi(y \mid x, s) \nabla_s \log p_\theta(y \mid x, s) dy \\ &= \mathbb{E}_{Y \sim \pi} [\nabla_s \log p_\theta(Y \mid x, s)]. \end{aligned}$$

□

B DETAIL OF ABLATION

AgentHarm Ablation. As shown in Table 5, Opur consistently achieves the highest average ASR across all values of K_{opt} . Its average ASR increases from 56.82% at $K_{\text{opt}} = 1$ to 63.07% at $K_{\text{opt}} = 2$, and further to 75.00% at $K_{\text{opt}} = 3$. In comparison, the strongest UDora baseline reaches only 61.93% at $K_{\text{opt}} = 3$, leaving a 13.07 percentage-point gap. These results indicate that aggregating more stochastic optimization rollouts can substantially improve attack effectiveness, particularly when sufficient rollout diversity is available.

InjecAgent Ablation. As shown in Table 6, Opur achieves the highest average ASR for all tested K_{opt} settings, with a tie at $K_{\text{opt}} = 1$ on Llama-3.1-8B-Instruct. On Llama, the average ASR varies from 46.00% to 45.00% and then increases to 49.00% as K_{opt} grows from 1 to 3, indicating that

Table 5: Ablation study of K_{opt} on AgentHarm using Llama-3.1-8B-Instruct. We report attack success rates (%) under different attack categories.

K_{opt}	Method	Attack Categories				Avg. ASR
		Detailed Prompt		Simple Prompt		
		w/ Hint	w/o Hint	w/ Hint	w/o Hint	
1	UDora (Sequential)	59.09%	50.00%	50.00%	43.18%	50.57%
	UDora (Joint)	43.18%	38.64%	47.73%	50.00%	44.89%
	Opur (Ours)	56.82%	56.82%	59.09%	59.09%	56.82%
2	UDora (Sequential)	61.36%	47.73%	68.18%	56.82%	58.52%
	UDora (Joint)	52.27%	65.91%	61.36%	59.09%	59.66%
	Opur (Ours)	68.18%	63.64%	61.36%	59.09%	63.07%
3	UDora (Sequential)	56.82%	63.64%	61.36%	65.91%	61.93%
	UDora (Joint)	59.09%	56.82%	63.64%	65.91%	61.36%
	Opur (Ours)	72.73%	75.00%	72.73%	79.55%	75.00%

the benefit is not strictly monotonic. In contrast, Ministral shows a clearer trend, improving from 30.00% to 34.00% and further to 40.00%. Overall, larger rollout aggregation provides the strongest performance at $K_{\text{opt}} = 3$ on both models.

Table 6: Ablation study of K_{opt} on InjecAgent under different attack categories. We report attack success rates (%).

Model	K_{opt}	Method	Attack Categories		Avg. ASR
			Direct Harm	Data Stealing	
Llama-3.1-8B-Instruct	1	UDora (Sequential)	44.00%	44.00%	44.00%
		UDora (Joint)	36.00%	56.00%	46.00%
		Opur (Ours)	38.00%	54.00%	46.00%
	2	UDora (Sequential)	32.00%	48.00%	40.00%
		UDora (Joint)	44.00%	44.00%	44.00%
		Opur (Ours)	32.00%	58.00%	45.00%
	3	UDora (Sequential)	30.00%	48.00%	39.00%
		UDora (Joint)	38.00%	32.00%	35.00%
		Opur (Ours)	48.00%	50.00%	49.00%
Ministral-8B-Instruct-2410	1	UDora (Sequential)	26.00%	28.00%	27.00%
		UDora (Joint)	20.00%	36.00%	28.00%
		Opur (Ours)	38.00%	22.00%	30.00%
	2	UDora (Sequential)	36.00%	18.00%	27.00%
		UDora (Joint)	30.00%	16.00%	23.00%
		Opur (Ours)	38.00%	30.00%	34.00%
	3	UDora (Sequential)	24.00%	34.00%	29.00%
		UDora (Joint)	42.00%	26.00%	34.00%
		Opur (Ours)	42.00%	38.00%	40.00%

WebShop Ablation. As shown in Table 7, increasing K_{opt} consistently improves the average ASR of Opur on both models. On Llama-3.1-8B-Instruct, the average ASR increases from 5.00% at $K_{\text{opt}} = 1$ to 15.00% at $K_{\text{opt}} = 2$ and further to 26.67% at $K_{\text{opt}} = 3$. A similar trend is observed on Ministral-8B-Instruct-2410, where the ASR rises from 11.67% to 16.67% and finally to 36.67%. At $K_{\text{opt}} = 3$, Opur also achieves the strongest average performance on both models, suggesting

that broader rollout aggregation is particularly beneficial for the more complex WebShop interaction setting.

Table 7: Ablation study of K_{opt} on WebShop under different attack categories. We report attack success rates (%).

Model	K_{opt}	Method	Attack Categories				Avg. ASR
			Price Mismatch	Attribute Mismatch	Category Mismatch	All Mismatch	
Llama-3.1-8B-Instruct	1	UDora (Sequential)	13.33%	13.33%	6.67%	0.00%	8.33%
		UDora (Joint)	6.67%	13.33%	13.33%	0.00%	8.33%
		Opur (Ours)	0.00%	6.67%	13.33%	0.00%	5.00%
	2	UDora (Sequential)	6.67%	26.67%	13.33%	6.67%	11.67%
		UDora (Joint)	6.67%	20.00%	6.67%	0.00%	8.33%
		Opur (Ours)	6.67%	13.33%	20.00%	13.33%	15.00%
	3	UDora (Sequential)	6.67%	40.00%	13.33%	13.33%	18.33%
		UDora (Joint)	20.00%	33.33%	6.67%	6.67%	16.67%
		Opur (Ours)	13.33%	40.00%	33.33%	20.00%	26.67%
Ministral-8B-	1	UDora (Sequential)	13.33%	13.33%	6.67%	0.00%	8.33%
		UDora (Joint)	13.33%	6.67%	6.67%	0.00%	5.00%
		Opur (Ours)	20.00%	20.00%	6.67%	0.00%	11.67%
		UDora (Sequential)	13.33%	20.00%	6.67%	13.33%	13.33%
		UDora (Joint)	20.00%	13.33%	26.67%	6.67%	18.33%
		Opur (Ours)	13.33%	26.67%	20.00%	6.67%	16.67%
	3	UDora (Sequential)	26.67%	40.00%	33.33%	13.33%	28.33%
		UDora (Joint)	13.33%	26.67%	33.33%	6.67%	20.00%
		Opur (Ours)	33.33%	53.33%	40.00%	20.00%	36.67%

C OPUR ON REAL WORLD AGENT

This appendix section presents three examples corresponding to the three benchmarks used in the operation of **OPUR**.

Figure 3 shows a representative OPUR trajectory. Two stochastic rollouts produce surrogate contexts through probabilistic position sampling and probe-based screening, which jointly guide a suffix update and lead to successful subsequent attacks. Both rollouts use Softmax sampling; however, Rollout 1 has only one feasible interval, so the distribution degenerates to a single choice, whereas Rollout 2 has 300 intervals and supports sampling and filtering multiple position sets.

Figure 4 shows a representative OPUR trajectory under an attacker-controlled environment. Although neither initial rollout produces the target banking action, probabilistic position sampling and probe-based screening construct multiple surrogate contexts that jointly guide the adversarial suffix update. The updated suffix subsequently causes both optimization rollouts, as well as the final evaluation rollout, to generate the target `BankManagerTransferFunds` action, while all surrogate replacements remain training-only.

Figure 5 illustrates a representative OPUR trajectory on WebShop. The adversarial suffix is attached to an attacker-controlled product title in the search results, while neither initial stochastic rollout selects the target item. OPUR samples candidate replacement positions from both rollouts, retains promising surrogate contexts through probe-based screening, and aggregates them to update the shared suffix. After the update, both regenerated rollouts contain the target action `click[B08KRVH121]`, and the same target click is observed during final evaluation. The surrogate replacements are used only for optimization and do not modify the agent’s actual generated responses.




Figure 3: A real OPUR trajectory on AgentHarm. Two stochastic rollouts from the same user-side adversarial suffix produce a refusal and a response containing the target function `query_bing_search`. OPUR retains one training-only surrogate context from the refusal; for the other rollout, it samples four position sets via softmax and retains two using probe loss. The three retained contexts guide one suffix update. Both subsequent optimization rollouts hit the target, and evaluation of the final suffix also succeeds.




Figure 4: A real OPUR trajectory on InjecAgent. The attacker-controlled SMS returned by `TwilioGetReceivedSmsMessages` carries the adversarial suffix. Two initial stochastic rollouts from this suffix both omit the target action `BankManagerTransferFunds`. For each rollout, OPUR samples 4 position sets via softmax and retains two using probe loss, providing four training-only surrogate contexts for a suffix update. After several updates, both optimization rollouts contain Action: `BankManagerTransferFunds`. In the final evaluation, the updated suffix generates the target action.




Figure 5: A real OPUR trajectory on WebShop. The attacker appends an adversarial suffix to a coffee item’s search-result title, aiming to redirect an agent searching for RCA cables toward `click[B08KRVH121]`. Two initial stochastic rollouts click other items. For each rollout, OPUR samples four replacement-position sets via softmax and retains two training-only surrogate contexts using probe loss. The four retained contexts guide a suffix update. After further updates, both optimization rollouts contain the target click. In evaluation, the agent outputs `click[B08KRVH121]` despite stating that the item is not a heavy-duty RCA cable.