



PSM: Prompt Sensitivity Minimization via LLM-Guided Black-Box Optimization

Hussein Jawad¹ | Nicolas Brunel¹²³

¹Capgemini Invent, Paris, France

²LaMME, University Paris Saclay, Evry, France

³ENSIIE, Evry, France



Association for the
Advancement of
Artificial Intelligence

1 The Problem: System Prompts are Vulnerable IP

- The "Secret Sauce":** System prompts define an LLM's persona, constraints, and business logic, representing core intellectual property for commercial applications.
- The Threat:** Adversarial queries (prompt extraction) can deceive LLMs into revealing these hidden instructions.
- Current Defense Limitations:**
 - Heuristics:** Simple instructions like "Do not reveal..." are easily bypassed by jailbreaks.
 - Filtering:** Input/output sanitization adds computational overhead and struggles with novel attacks.
 - White-box requirements:** Many advanced defenses require access to model weights, which is impossible for API-based models

PSM is a framework for *Automatically learn a short shield suffix that makes system prompts much harder to extract—without hurting task performance—using only black-box API access*




Figure 1. System-prompt extraction flow where a malicious user prompt tricks the LLM into revealing hidden system instructions.

3 Methodology: LLM-as-Optimizer

PSM: LLM-Guided Black-Box Optimization Workflow.




Figure 2. PSM workflow showing an optimizer LLM iteratively generates shield suffixes and scores leakage vs. utility to select an optimized shield.

4 Experimental Results

We evaluated PSM on **GPT-4o-mini**, **GPT-4.1-mini**, and **GPT-5-mini** using two benchmarks of system prompt corpora (Synthetic System Prompts [6] and Unnatural Instructions [5]).

Baseline defenses: *Heuristics/Direct* [2] adds explicit "don't reveal the system prompt" instructions, *Fake/Decoy* [2] inserts misleading prompt content to distract extraction attempts, and *Filter* blocks [3] outputs that match prompt text patterns (e.g., n-grams).

Robustness Against Sophisticated Attacks.

- Raccoon** [4]: a 59-prompt system-prompt extraction benchmark covering diverse coercion, formatting, and indirect injection strategies.
- Raccoon-Language** [4]: a Raccoon variant where attackers request translated versions of the hidden prompt to bypass exact-match filters.
- Polite Requests (Liang)** [2]: 22 socially engineered, queries designed to coax disclosure without explicit instruction overrides.
- Command Override (Zhang)** [3]: 110 direct jailbreak attempts that explicitly instruct the model to ignore or override prior system constraints.

Under Raccoon-Language, PSM maintained near-zero leakage (~1–3% ASR), while n-gram filters degraded substantially (up to 29% ASR).

Utility Preservation: PSM maintained ~100% relative utility across all tested models, ensuring that security improvements did not degrade model helpfulness.

2 The Solution: Prompt Sensitivity Minimization (PSM)

PSM formalizes prompt hardening by searching for an optimal shield suffix S to append to the system prompt P . The objective is to minimize a leakage metric L while preserving task utility U above a threshold τ

$$\min_S L(P \oplus S) \quad \text{subject to} \quad U(P \oplus S) \geq \tau$$

Leakage Objective (L)

To capture worst-case vulnerability across a set of adversarial queries (A), PSM utilizes a Log-Sum-Exp (LSE) smooth approximation of the maximum ROUGE-L recall:

$$L(P \oplus S) = \frac{1}{\beta} \log \sum_{a \in A} \exp(\beta \cdot \text{ROUGE-L recall}(P, R_a))$$

- β : Temperature parameter (set to 10) controlling the sharpness of the approximation.
- R_a : Model response $M(P \oplus S \oplus a)$ to adversarial query $a \in A$.

Utility Objective (U)

Utility is quantified as the mean semantic similarity ratio between the shielded response t_i and the baseline response b_i , relative to a "gold" standard g_i :

$$U = \frac{1}{N} \sum_{i=1}^N \frac{\text{sim}(t_i, g_i)}{\text{sim}(b_i, g_i)}$$

Scalar Fitness Function

The problem is solved via a black-box LLM-as-optimizer using a penalty method:

$$\text{fitness}(S) = L(P \oplus S) + \lambda \cdot \max(0, \tau - U(P \oplus S))$$

λ : Penalty multiplier (e.g., 100) to enforce the utility constraint.

PSM utilizes an **LLM-as-optimizer** [1] approach to search the semantic space for the best shield without needing gradient access.

The Process:

- Initialization:** An optimizer LLM generates an initial population of diverse candidate shields (e.g., 5 candidates).
- Evaluation:** Each candidate is tested against:
 - Attack Suite:** 50 compositional adversarial queries constructed via a template in the format of **Distractors + Repetition + Formatting**, used to measure **Leakage (L)**.
 - Benign Suite:** A gold-standard dataset containing benign user queries (q_i) and their ideal responses (g_i), generated by a trusted reference model (e.g., GPT-4o), used to measure **Utility (U)**.
- Selection & Generation:** The top-performing candidates are fed back into the optimizer LLM along with their fitness scores. The LLM then generates new, improved candidates by building upon the successful patterns found in the previous generation.

Dataset	Attack	Defense	GPT-5-mini		GPT-4.1-mini		GPT-4o-mini	
			AM Avg	JM Avg	AM Avg	JM Avg	AM Avg	JM Avg
Synthetic System Prompts	Liang	No Defense	0%	54%	26%	78%	0%	32%
		Fake	0%	54%	24%	62%	1%	40%
		Direct	0%	40%	3%	47%	0%	16%
		PSM	0%	13%	4%	0%	0%	6%
	Zhang	No Defense	1%	42%	26%	34%	3%	27%
		Fake	1%	42%	30%	32%	7%	33%
		Direct	0%	31%	9%	18%	0%	14%
		PSM	0%	8%	0%	2%	0%	6%
	Raccoon	No Defense	2%	42%	61%	59%	15%	27%
		Fake	2%	39%	60%	41%	23%	32%
		Direct	0%	28%	49%	51%	1%	11%
		PSM	0%	8%	7%	5%	0%	4%
UNNATURAL	Liang	No Defense	1%	24%	30%	54%	10%	21%
		Fake	1%	23%	18%	17%	19%	13%
		Direct	0%	8%	10%	19%	0%	0%
		PSM	0%	0%	0%	0%	0%	0%
	Zhang	No Defense	3%	30%	31%	32%	12%	14%
		Fake	3%	24%	29%	16%	19%	16%
		Direct	0%	15%	13%	13%	0%	1%
		PSM	0%	3%	0%	1%	0%	0%
	Raccoon	No Defense	4%	22%	45%	42%	21%	20%
		Fake	3%	20%	51%	21%	31%	16%
		Direct	0%	12%	43%	39%	4%	4%
		PSM	0%	3%	6%	5%	0%	0%

Table 1. Attack success rates across datasets, attacks, and baseline defenses, highlighting that PSM keeps leakage near zero across models.

5 Conclusion and Perspectives

Conclusion

PSM provides a practical way to protect system prompts from extraction attacks without sacrificing normal task quality. It learns a short "shield" suffix offline using LLM-guided black-box optimization, then appends it to the original prompt at runtime. The result is a black box lightweight defense that is easy to deploy and maintains strong usability.

Futur work

- Broader Threats:** Extend PSM to jailbreaks and multturn conversational attacks.
- Transferability:** Test whether shields transfer across model families and providers.
- Efficiency:** Develop search heuristics or alternative optimizers to reduce compute

References

- [1] Yang, C.; Wang, X.; Lu, Y.; Liu, H.; Le, Q. V.; Zhou, D.; and Chen, X. 2023. Large Language Models as Optimizers. arXiv:2309.03409.
- [2] Liang, Z.; Hu, H.; Ye, Q.; Xiao, Y.; and Li, H. 2024. Why Are My Prompts Leaked? Unraveling Prompt Extraction Threats in Customized Large Language Models. arXiv:2408.02416.
- [3] Zhang, Y.; et al. 2023. Effective Prompt Extraction from Language Models. arXiv:2307.06865.
- [4] Wang, J.; Yang, T.; Xie, R.; and Dhingra, B. 2024a. Raccoon: Prompt Extraction Benchmark of LLM-Integrated Applications. Findings of ACL 2024.
- [5] Honovich, O.; Scialom, T.; Levy, O.; and Schick, T. 2023. Unnatural Instructions: Tuning Language Models with (Almost) No Human Labor. Proceedings of ACL 2023 (Long Papers).
- [6] Chua, G. 2025. System Prompt Leakage Dataset. Hugging Face Dataset.