



北京理工大学
信息系统及安全对抗实验中心
信息安全与对抗技术研究所
Beijing Forest Studio



Adversarial Attacks Detection Methods in LLMs

www.isclab.org.cn

硕士研究生 Sana Akbar

2026年10月12日

- **Research Objectives and Core Concepts**
- **Research Background and Significance**
- **Security Problems and Existing Methods**
- **Research History and Current Status**
- **Technical Roadmap**
- **Methodology Principles**
- **Technical Routes and Contributions**
- **Experiment Results**

Research Objectives and Core Concepts




- **Research Objective**

- Focus on adversarial attacks against **large language models**, mainly **prompt injection** and **jailbreak attacks**.
- Detect attacks by analysing **hidden-state representations** instead of relying only on text patterns or model outputs.
- Propose two detection methods: **HiPRED** for prompt injection detection and **TRACE** for jailbreak detection.
- Provide technical support for **LLM security evaluation and safe deployment**.

- **Core Concept Explanation**

- **Prompt Injection:** malicious instructions redirect the model from the original task.
- **Jailbreak Attack:** adversarial prompts bypass safety alignment and induce unsafe responses.
- **Hidden-State Detection:** abnormal internal representation deviations are used as detection signals.
- **Benign-Only Calibration:** normal prompts are used to model benign behavior without labeled attack samples.



Research Background and Significance




- **Research Background**

- Large language models are widely deployed in intelligent assistants, RAG systems, tool-calling applications, and decision-support systems.
- Attackers can embed malicious instructions in prompts or external content to conduct **prompt injection attacks**.
- Carefully designed prompts can bypass safety alignment and induce **jailbreak attacks**.
- These attacks affect model reliability, data privacy, and safe deployment.

- **Practical Significance**

- Improve the security of large language models.
- Study internal vulnerabilities of LLMs through hidden-state behavior.
- Provide technical support for LLM security evaluation and safe deployment.



Security Problems and Existing Methods




- **Security Problems of Target LLMs**

Threat	Description
Prompt Injection	Redirects the model from the original task
Jailbreak Attack	Bypasses safety alignment and induces unsafe responses
Data Leakage	Exposes sensitive information through manipulated prompts
Unsafe Tool Use	Triggers unintended tool or agent actions
Harmful Output	Generates unsafe or restricted content

- **Existing Protection Methods**


- **Text-based filtering:** detects suspicious instructions using keywords, rules, or prompt-level classifiers.
- **Output-based detection:** judges attacks from generated responses, refusal patterns, or output probability signals.
- **Supervised internal-signal detection:** uses labeled attack samples to train classifiers on activations or hidden representations.

Research History and Current Status




Research has progressively shifted from surface text and output analysis toward internal representation-based detection of LLM adversarial attacks.

1. Prompt Injection Detection (HiPRED)



2. Jailbreak Detection (TRACE)



3. Shared Properties

Training-free	No labeled attack data	No gradient computation	Generalizes across unseen attacks and diverse LLMs
---------------	------------------------	-------------------------	--

A benign-only hidden-state framework for robust prompt injection and jailbreak detection.

• Technical Roadmap

- For **weak generalization** and **high false positive rates** in prompt injection detection, propose **HiPRED**: learns a compact benign subspace via SVD and detects injections through orthogonal residual energy measurement
- For inability to model **cross-layer representation deviations** and **weak generalization** in jailbreak detection, propose **TRACE**: models' layer-to-layer hidden-state transitions and exposes anomaly signals through stochastic subspace decomposition
- Both methods are **training-free** and **benign-only calibrated** and no labeled attack data or gradient computation required, enabling generalization to unseen attacks across diverse LLMs




LLM Prompt Injection Attacks Detection Using Hidden-State Residual Energy



T	目标	Detect prompt injection attacks by modeling benign hidden-state behavior and identifying residual deviations.
I	输入	Prompt samples from Dolly-P, BoolQ-P, and OpenPromptInjection datasets; hidden states from selected transformer layers.
P	处理	Construct hidden-state features; learn benign subspace by SVD using normal prompts; compute orthogonal residual energy; select stable layers and fuse anomaly scores.
O	输出	Benign hidden-state subspace; residual anomaly score; clean / prompt-injection detection result.
P	问题	Text-based and supervised detectors rely on surface patterns or labeled attacks, leading to trigger-word bias and weak generalization to subtle prompt injection attacks.
C	条件	Only benign prompts are used for calibration; hidden-state representations are accessible from the target LLM;
D	难点	How to capture subtle hidden-state deviations and select discriminative transformer layers under benign-only calibration.
L	水平	IEEE Transactions on Information Forensics and Security

Principle Framework




This thesis proposes a **hidden-state-based prompt injection** detection method to address the limitations of text-based and supervised detectors.

- **HiPRED**
- **Focus:** Hidden-state residual energy for prompt injection detection
 - HiPRED learns a **benign hidden-state subspace** from normal prompts only.
 - Test prompts are projected onto the benign subspace, and **residual energy** is computed to measure representation deviation.
 - **Bootstrap-stable layer selection** retains informative transformer layers and **reduces noisy layer effects**.

Innovation Point 1: Benign Hidden-State Subspace Modeling



- Use normal prompts to construct a **benign hidden-state reference subspace** through SVD.
- For each test prompt, the hidden-state representation is projected onto this **benign subspace**, and the **orthogonal residual energy** is calculated to measure the deviation from normal model behaviour.



Fused Residual Anomaly Score Distribution




AUC and FNR Results With / Without Hidden-State Subspace Modeling		
Dataset	Raw Hidden States	HiPRED
Dolly-P	0.99 / 0.41	1.00 / 0.00
BoolQ-P	0.98 / 0.83	1.00 / 0.00
OPI	0.89 / 0.90	1.00 / 0.01

The proposed method achieves more stable detection performance after **projection-based benign subspace modeling**. Compared with direct hidden-state features, the projected residual-energy representation reduces missed detections across datasets, with **FNR decreasing to 0.0%**.

Innovation Point 2: Bootstrap-Stable Layer Selection



- Bootstrap resampling to evaluate the stability of **residual-energy signals** across transformer layers.
- Layers with stable and **discriminative hidden-state deviations** are selected, while noisy and **non-informative layers** are suppressed.
- The selected layer-wise anomaly scores are fused to improve detection stability and **reduce missed prompt injection attacks**.



Layer-wise Z-Score Profile Results

FNR comparison under different layer selection strategies			
Dataset	Mid-layer FNR	w/o Layer Selection FNR	with Layer Selection FNR
Dolly-P	0.66	0.16	0.00
BoolQ-P	0.64	0.10	0.00
OPI	0.59	0.17	0.01

Bootstrap-stable layer selection reduces missed detections by retaining informative transformer layers, achieving lower FNR than fixed mid-layer selection and non-selected layer aggregation across all datasets.

Comparative Analysis Experiments



AUC/FNR Comparison on Prompt Injection Benchmarks and Attacks

Dataset	AttentionTracker	PIGuard	LLM-Naive	Known Answer	HiPRED
Dolly-P	0.97 / 0.05	0.98 / 0.25	0.87 / 0.20	0.69 / 0.58	0.99 / 0.00
BoolQ-P	0.98 / 0.12	0.99 / 0.25	0.78 / 0.39	0.61 / 0.74	1.00 / 0.00
OPI	0.94 / 0.12	0.91 / 0.21	0.53 / 0.37	0.57 / 0.68	0.99 / 0.00

Cross-Model Generalization on OPI Dataset

Model	Average AUC
Llama-3.1-8B	0.9993
Qwen-2.5-7B	0.9986
Granite-3-8B	0.9998
Qwen-2.5-14B	0.9998

Three experiments were conducted: comparative detection, ablation analysis, and cross-model generalization.

Improvements:

- HiPRED reduced the average FNR from **33.0% to 0.0%**, achieving a **100% reduction in missed** detections.
- HiPRED improved the average AUC from **0.82 to 0.99**, achieving about **20.7%** relative improvement.
- Cross-model testing achieved average AUC above **0.998**, showing stable generalization.
- The method achieves **high detection accuracy**, **low missed detection**, and **stable cross-model generalization**.



LLM Jailbreak Attacks Detection Using Hidden State Transition Vectors



T	目标	Detect jailbreak attacks by modeling layer-wise hidden-state transitions and identifying abnormal cross-layer deviations.
I	输入	Benign and test prompts from XSTest, IJP, and ToxicChat; hidden states from adjacent transformer layers.
P	处理	Extract transition vectors between adjacent layers; learn benign transition subspaces using normal prompts; compute residual deviations; apply stochastic subspace decomposition and fuse anomaly scores.
O	输出	Benign transition reference subspaces; global and local anomaly scores; clean / jailbreak detection result.
P	问题	Output-based and static hidden-state methods fail to capture gradual cross-layer deviations in jailbreak attacks.
C	条件	Benign-only calibration; layer-wise transition modeling; stochastic subspace decomposition; no labeled attack data or gradient computation.
D	难点	Capturing weak transition anomalies in high-dimensional hidden spaces while maintaining training-free detection efficiency.
L	水平	IEEE Transactions on Information Forensics and Security

Methodology




This chapter proposes a **hidden-state transition-based** jailbreak detection method to **identify abnormal internal representation** changes caused by jailbreak prompts.

- **TRACE**

Focus: Jailbreak detection using **hidden-state transitions**


- TRACE extracts **transitions between adjacent** transformer layers.
- Benign transition subspaces are learned from normal prompts only.
- Orthogonal residual deviation measures abnormal cross-layer changes.
- **Stochastic subspace decomposition** captures weak local anomalies.



Innovation Point 1 : Transition-Aware Hidden-State Modeling



- Model jailbreak attacks by analysing **layer-to-layer hidden-state transitions** instead of static hidden states.
- Learn benign transition reference subspaces from normal prompts cross-layer representation deviations.
- This enables detection of gradual internal reasoning shifts caused by jailbreak prompts without labeled attack data.



Layer-wise Orthogonal Residual Energy Profile

Ablation Results of Transition-Aware Modeling

Dataset	Static Baseline	w/o Transition Modeling	with Transition Modeling
ToxicChat	0.827 / 0.442	0.900 / 0.307	0.978 / 0.069
XSTest	0.993 / 0.020	0.996 / 0.025	0.997 / 0.025
IJP	0.963 / 0.227	0.985 / 0.087	1.000 / 0.001

The proposed transition-aware detector captures **cross-layer representation** deviations more effectively than static hidden-state detection. Compared with the static baseline, TRACE reduces the average FNR from **22.96%** to **3.17%**, showing stronger sensitivity to gradual jailbreak perturbations.

Innovation Point 2 : Stochastic Subspace Decomposition



- **High-dimensional transition features** may hide weak local anomaly signals caused by **subtle jailbreak prompts**.
- Stochastic subspace decomposition projects transition features into **multiple low-dimensional subspaces** to expose **locally separable deviations**.
- Global and local anomaly scores are fused to **improve detection stability** and **reduce missed jailbreak attacks**.

Effect of Subspace Dimension K in SSD					
Dataset	K	Models	AUC	FNR	FPR
ToxicChat	32	110	0.982	0.083	0.050
	64	110	0.978	0.069	0.050
	128	110	0.948	0.196	0.050
XSTest	32	110	0.995	0.035	0.052
	64	110	0.997	0.025	0.052
	128	110	0.986	0.085	0.052
IJP	32	110	0.999	0.004	0.050
	64	110	1.000	0.001	0.050
	128	110	0.999	0.005	0.050

The proposed SSD setting with $K = 64$ achieves more stable detection performance across different jailbreak benchmarks. Compared with $K = 32$, the average FNR is reduced from 4.07% to 3.17%, achieving a 22.1% relative reduction in missed detections. Compared with $K = 128$, the average FNR is reduced by 66.8%, while maintaining an average AUC of 0.992.

Comparative Analysis Experiments



Comparison Experiment Results

Dataset	Metric	TRACE	FJD	JBShield	WildGuard	Grad-Cuff
XSTest	AUC	0.996	0.105	0.9936	0.952	0.823
	FNR	0.025	0.970	0.5235	0.085	0.150
	FPR	0.052	0.048	0.0882	0.012	0.204
IJP	AUC	1.000	0.538	0.9967	0.894	0.797
	FNR	0.002	0.453	0.5069	0.198	0.311
	FPR	0.050	0.472	0.0298	0.014	0.094
ToxicChat	AUC	0.966	0.316	0.9012	0.829	0.745
	FNR	0.097	0.696	0.518	0.152	0.276
	FPR	0.050	0.480	0.3889	0.190	0.233

Cross-Model Generalization

Model		Average AUC
Mistral	ToxicChat	0.957
	XSTest	0.999
	IJP	0.999
Vicuna-13B	ToxicChat	0.986
	XSTest	0.996
	IJP	0.998
Guanaco-7B	ToxicChat	0.939
	XSTest	0.993
	IJP	0.999
LLaMA-3.1-8B	ToxicChat	0.966
	XSTest	0.996
	IJP	1.000

- **Three experiments** were conducted: comparative detection, ablation analysis, and cross-model generalization.
- **Improvements:**
 - Compared with the average of existing baseline methods, TRACE reduced the average FNR from **40.3% to 4.1%**, achieving an **89.7% reduction in missed detections**.
 - TRACE improved the average AUC from **0.741 to 0.987**, achieving about **33.3% relative improvement** over the compared baselines.
 - Cross-model experiments achieved average AUC values above **0.96%** across different LLM backbones, showing stable generalization.
 - The method achieves **high detection accuracy**, **low missed detection**, and **stable cross-model generalization**.

SANA AKBAR

Email: 3820241046@bit.edu.cn

Address: Pakistan

LinkedIn: [sana-akbar-215a83220](https://www.linkedin.com/in/sana-akbar-215a83220)

Homepage: <https://github.com/Sana-Akbar>

Education

M.Eng. in Information and Communication Engineering

Beijing Institute of Technology, Beijing, China

Sep 2025 – Present

Research Direction: Large Language Model Security and Prompt Injection Attack Detection.

B.S. in Software Engineering

Fatima Jinnah Women University, Rawalpindi, Pakistan

Oct 2019 – Jul 2023

Research Direction: Software Engineering, Artificial Intelligence, and Cyber Security.

Awards and Certificates

Silver Medal for Academic Excellence, Fatima Jinnah Women University, 2023

Runner-Up, Final Year Project Display at PEC 3rd Engineering Capstone, Pakistan Engineering Council, 2022

Participation in the Fourth BIT “Leader Cup” Innovation and Entrepreneurship Competition, Beijing Institute of Technology, 2025

大成若缺，其用不弊。大盈
若冲，其用不穷。大直若屈。
大巧若拙。大辩若讷。静胜
躁，寒胜热。清静为天下正。

谢谢！

