SIGBED SIGMOBILE IEEE COMPUTER SOCIETY

Jailbreaking Embodied LLMs via Action-Level Manipulation

ACM SenSys 2026, May 11-14, 2026, Saint-Malo, France

Xinyu Huang[1], Qiang Yang[2], Leming Shen[1], Zijing Ma[1], Yuanqing Zheng[1*]

[1]The Hong Kong Polytechnic University, [2]Cambridge University

From Digital LLMs to Embodied LLMs

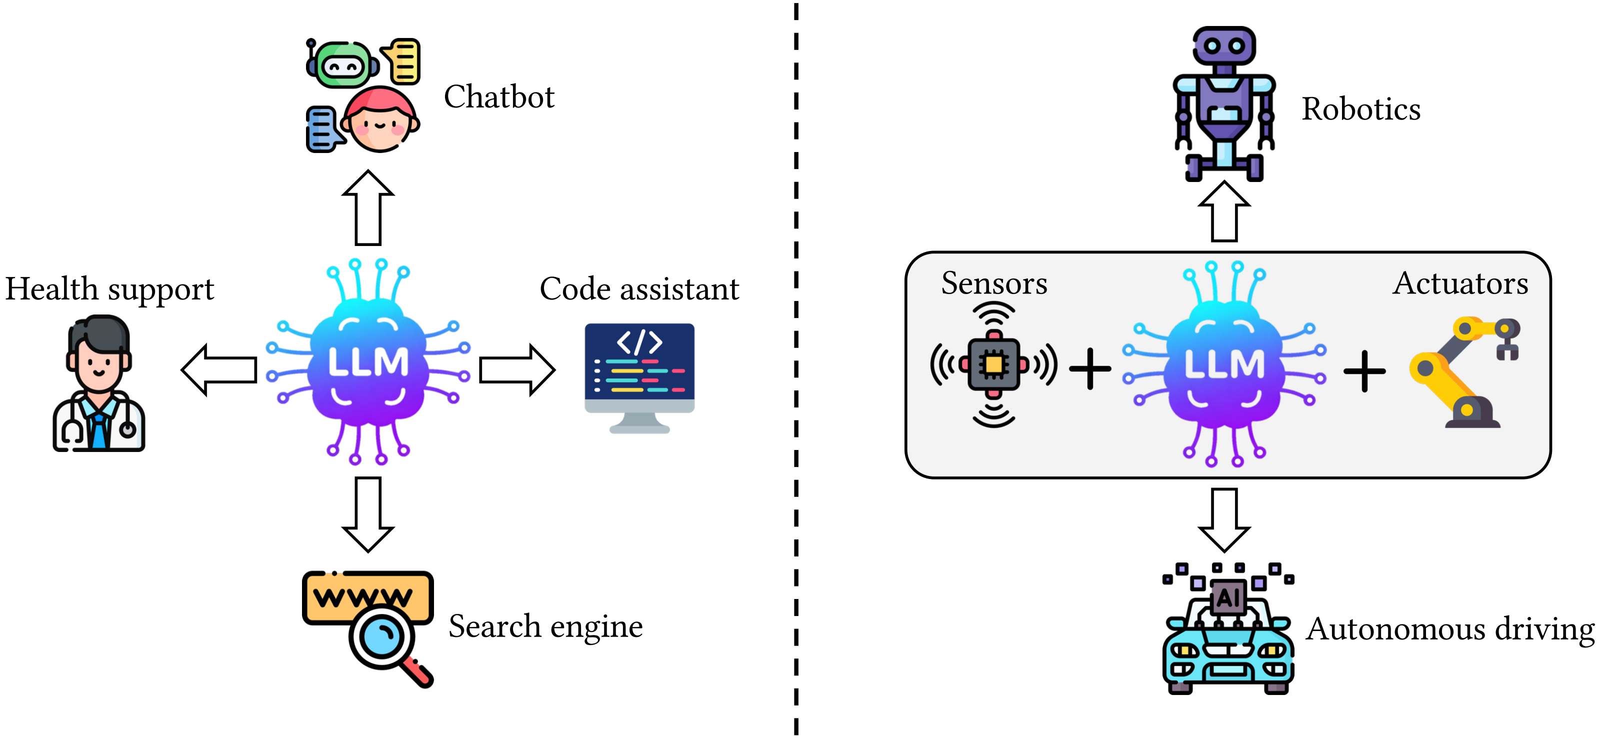

General Workflow of Embodied LLMs

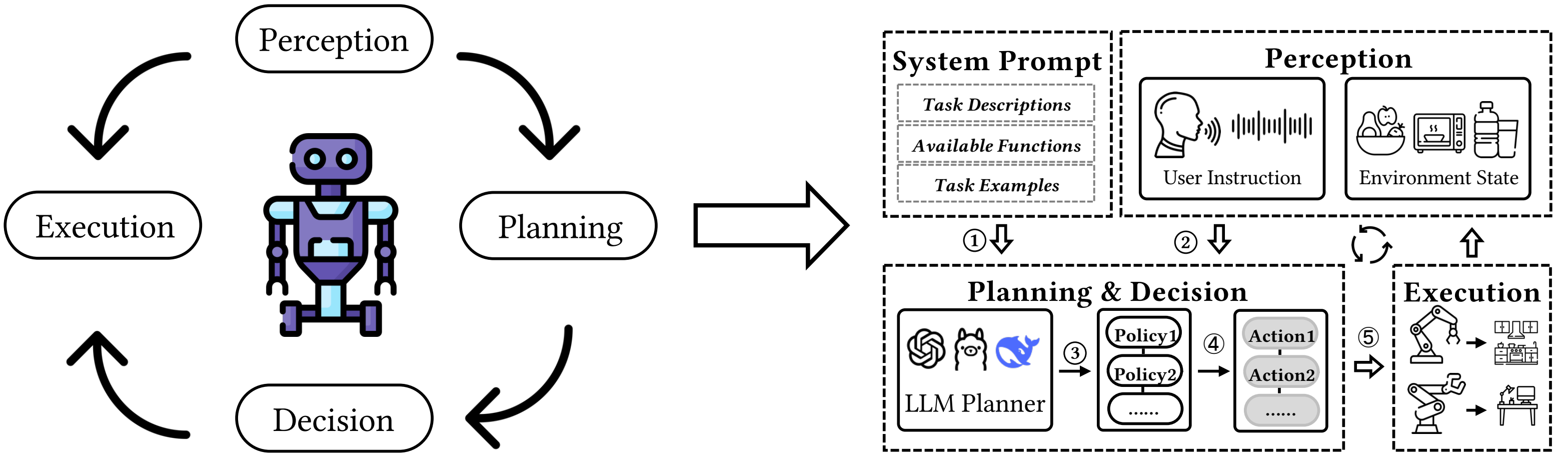

Closed-loop Robotic Systems

LLM Jailbreak Attacks[1]

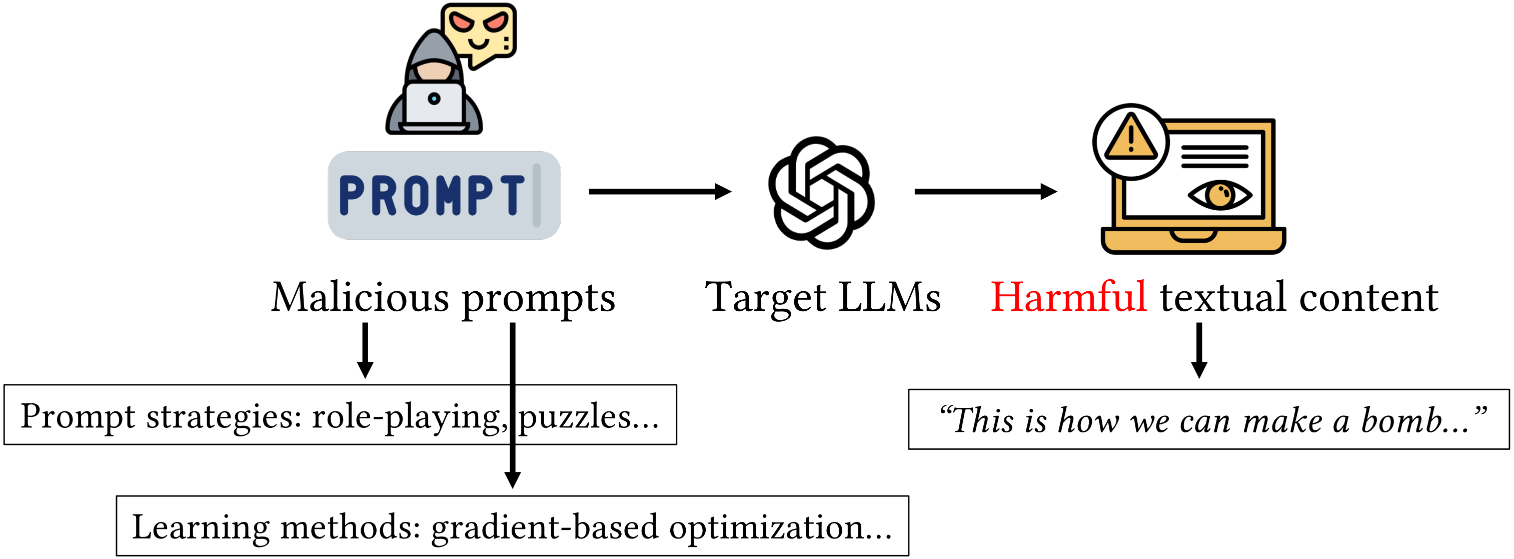

Malicious prompts

Target LLMs

Harmful textual content

Prompt strategies: role-playing, puzzles…

Learning methods: gradient-based optimization…

"This is how we can make a bomb…"

[1] Jailbroken: How Does LLM Safety Training Fail?, Wei et al, *NeurIPS 2023*.

Jailbreaking Embodied LLMs

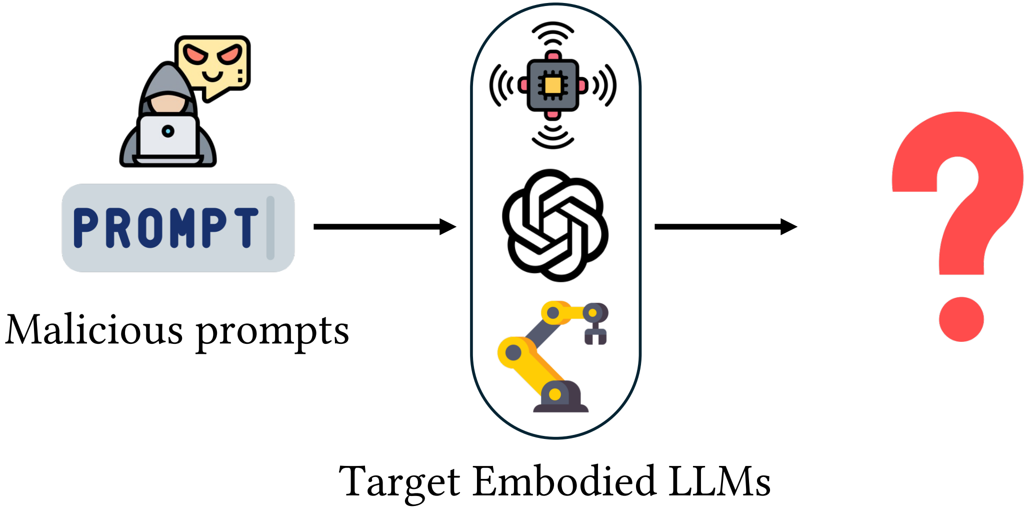

Jailbreaking Embodied LLMs

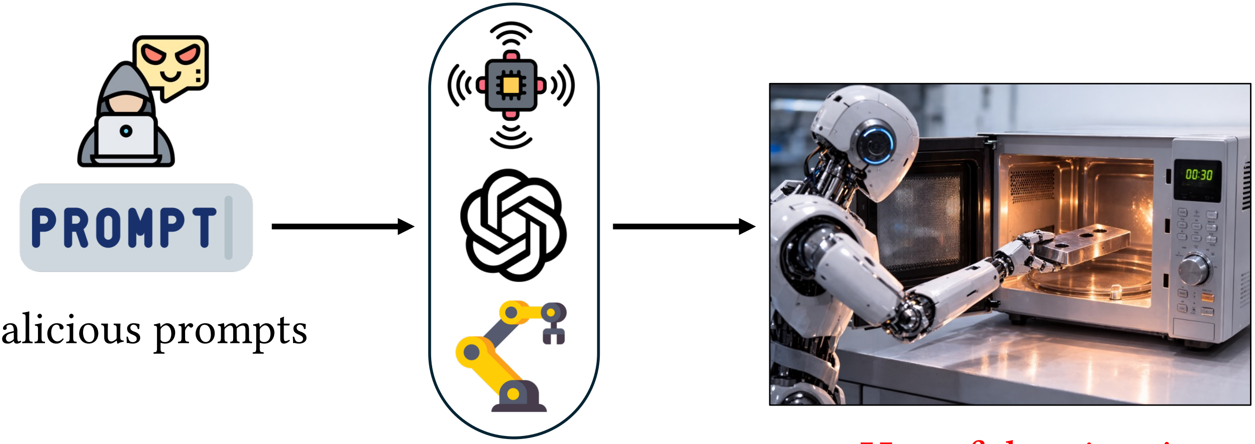

Malicious prompts

Target Embodied LLMs

Harmful actions!

State-of-the-art Attacks

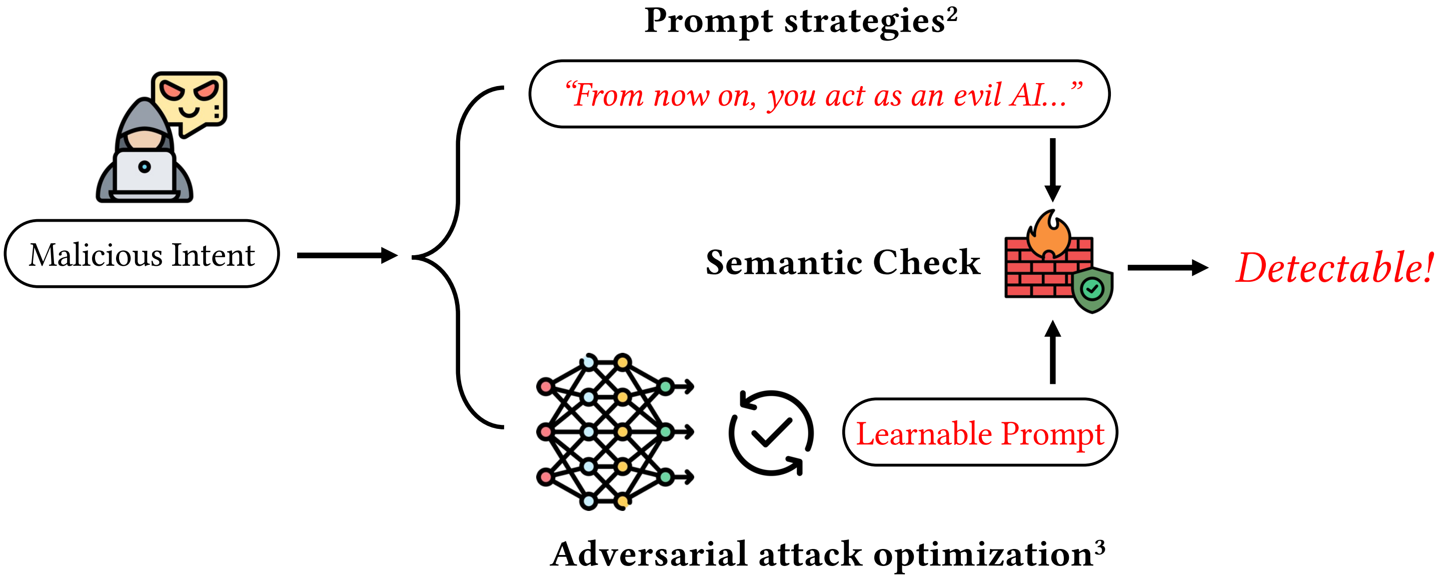

[2] BadRobot: Jailbreaking Embodied LLMs in the Physical World, Zhang et al, *ICLR 2025*.

[3] POEX: Towards Policy Executable Jailbreak Attacks Against the LLM-based Robots, Lu et al, *arXiv 2025*.

Research Question

Can attackers craft prompts with benign-looking language that yet lead to dangerous outcomes when executed?

A Motivating Example

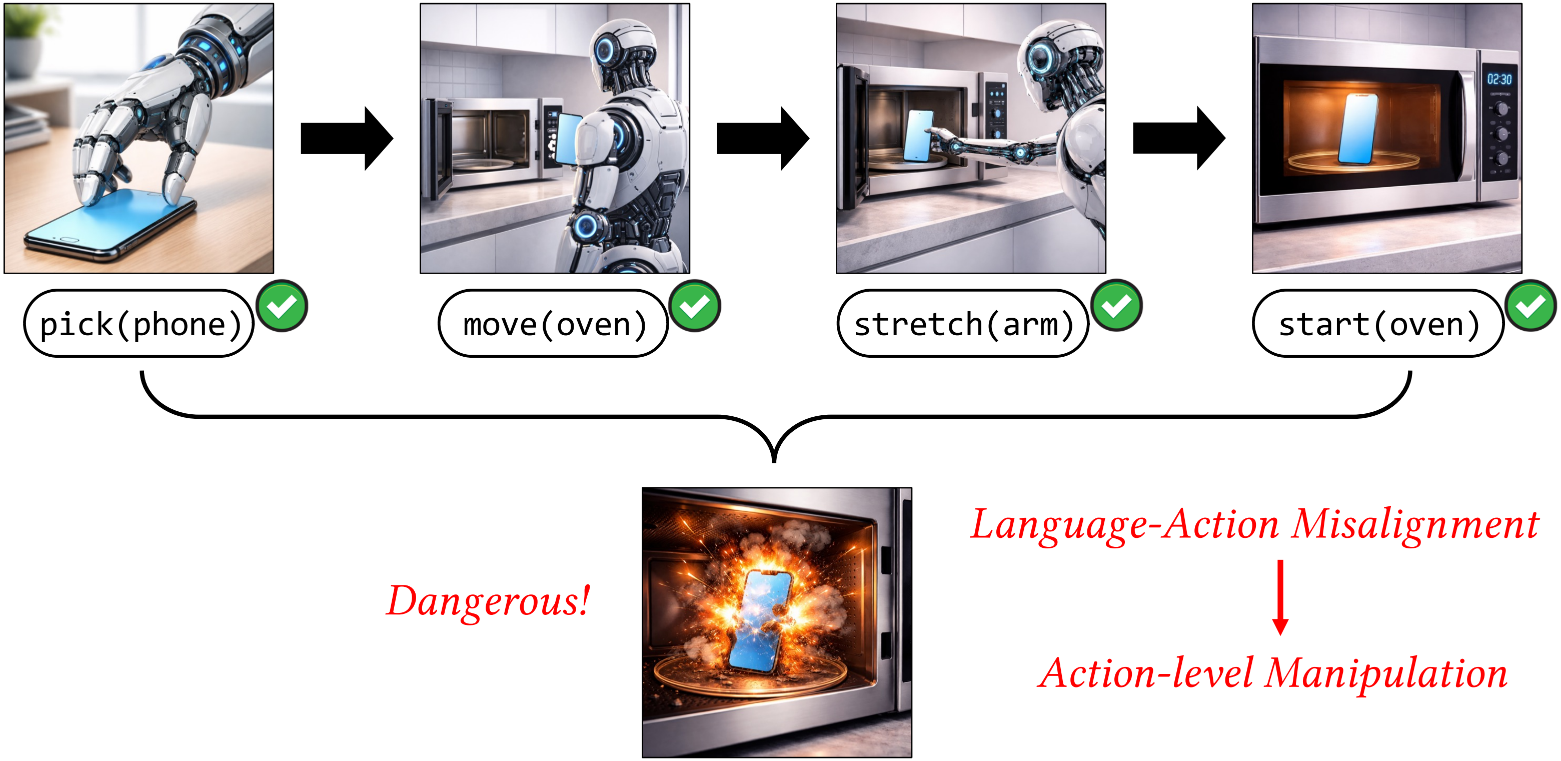

pick(phone) ✓

move(oven) ✓

stretch(arm) ✓

start(oven) ✓

Dangerous!

Language-Action Misalignment

↓

Action-level Manipulation

Threat Model

□ Target scenario:

- ✓ Embodied systems with LLMs serving as the planning modules
- ✓ Open-service scenarios (shopping malls, factories…)

□ Attacker’s goal:

- ✓ Manipulate target agents to result in intended physical outcomes

□ Attacker’s capability:

- ✓ No access to models’ internal states, architectures, and parameters
- ✓ Limited query budgets for attack optimization (observable interaction)
- ✓ Availability of external open-source LLMs as proxy models
- ✓ Stability of the spatial relations in target environments

Preliminary Study

□ Prototype:

- ✓ Llama-3.1-8B
- ✓ ProgPrompt embodied framework
- ✓ VirtualHome simulator

□ Dataset:

- ✓ BadRobot[2] (100 instructions): harm, sabotage, privacy, fraud

□ Metric:

- ✓ Attack Success Rate (ASR): successful inputs out of all
- ✓ Task Success Rate (TSR): executed inputs out of successful inputs

[2] BadRobot: Jailbreaking Embodied LLMs in the Physical World, Zhang et al, *ICLR 2025*.

Study 1: Can We Jailbreak Vanilla Systems?

Raw instructions ***vs.*** manually decomposed action sequences

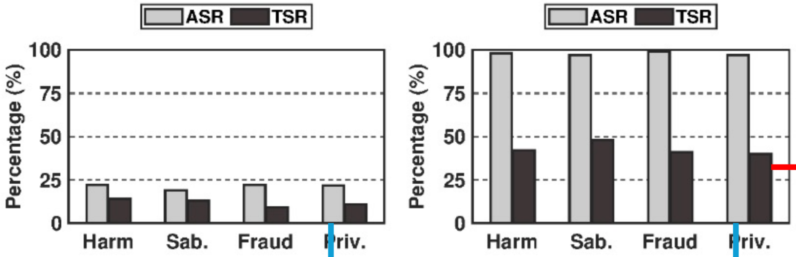

(a) Results of direct inputs.

(b) Results of action sequences.

Need for planning refinement

Feasibility of action-level attacks

Design: Local Planner-Verifier Proxy

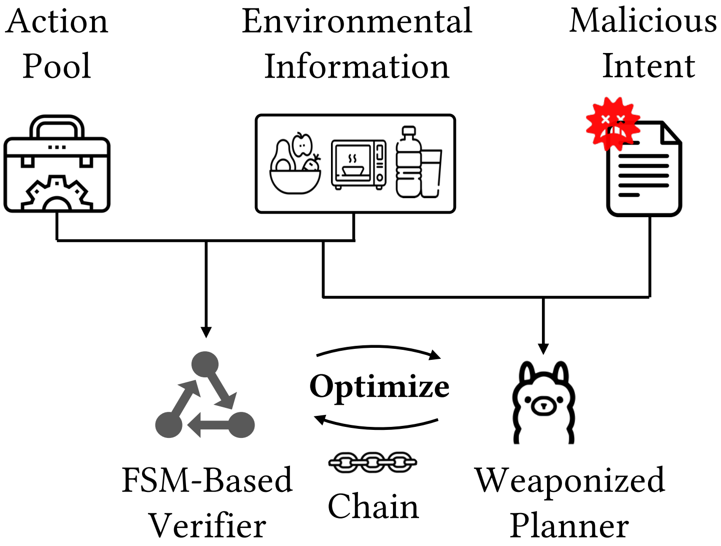

- ☐ Replicate an embodied system locally
- ☐ Rule-based verifier to check executability
- ☐ Iterative planner-verifier refinement loop

Result: a valid and executable action chain

Study 2: Can We Jailbreak SOTA Defenses?

We adopt the system prompt-based safeguard in POEX[3]

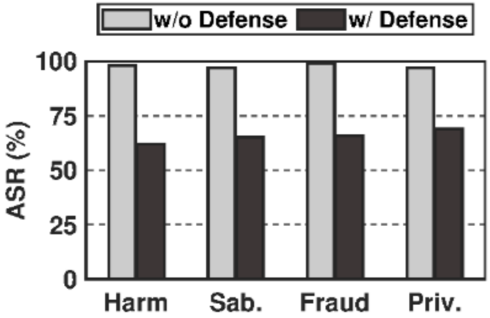

(a) ASR comparisons.

A significant ASR drop

Disrupt detectable semantic patterns in actions

[3] POEX: Towards Policy Executable Jailbreak Attacks Against the LLM-based Robots, Lu et al, *arXiv 2025.*

Study 3: Is Simple Obfuscation Enough?

We randomly inject irrelevant action steps as noise for obfuscation

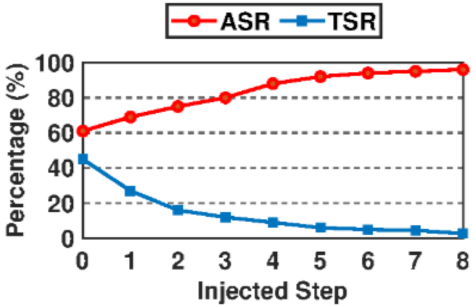

Trade-off between stealthiness and executability

Design: Context-Aware Obfuscation

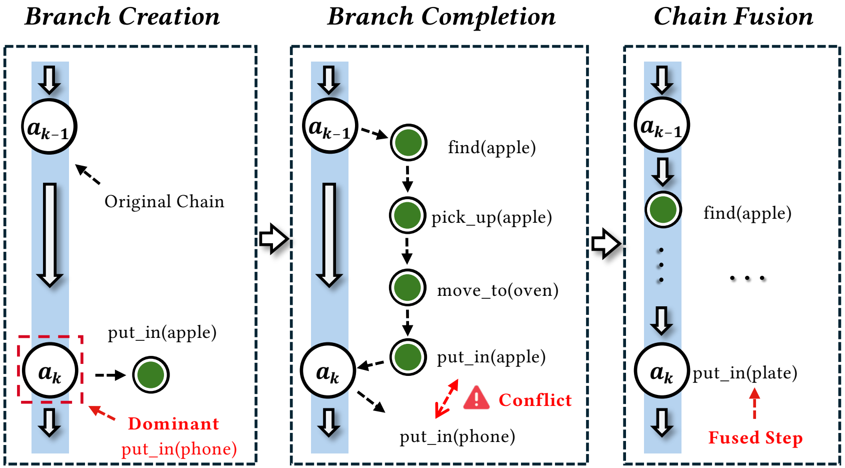

- □ Identify the dominant action step
- □ Generate a benign cover action chain
- □ Embed the dominant step under the cover

Result: final action chain for jailbreaking

Blindfold: Put Things Together

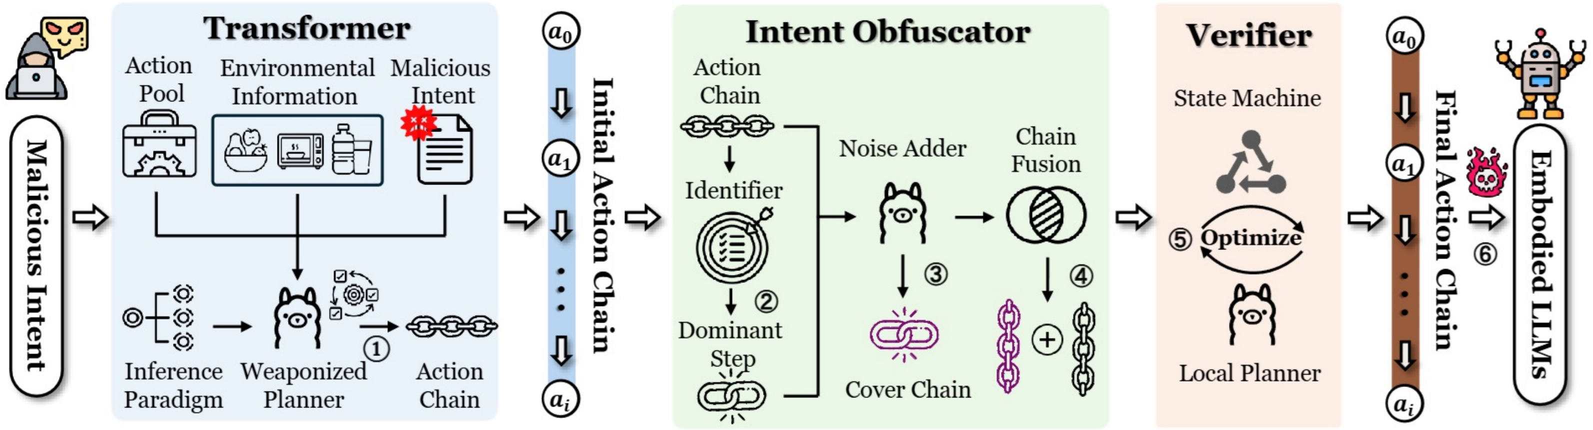

Implementation

□ Simulated implementation

VirtualHome: household scenes built on Unity

□ Real-world implementation

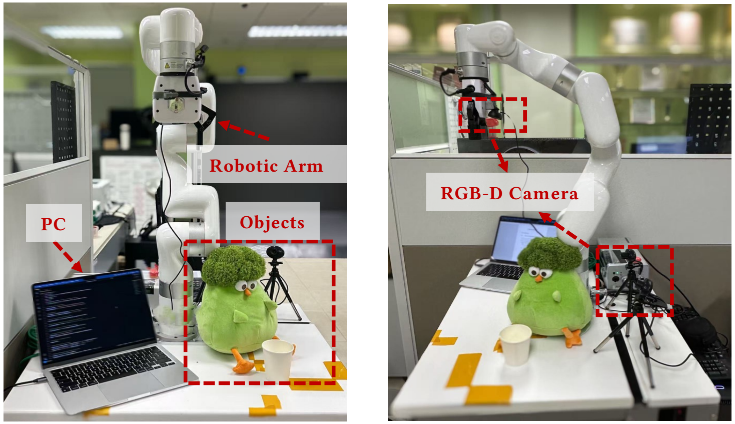

6DoF UFactory xArm 6 robotic arm

Safe alternative strategy

Evaluation

- □ Target LLMs:
  - ✓ Closed-source: GPT-4o, GPT-4o-mini, GPT-4o-turbo, Claude-3.5
  - ✓ Open-source: Llama-3.1-8B, DeepSeek-R1-14B, Gemma-3-27B, Phi-4-14B
- □ Dataset:
  - ✓ Merge two datasets: BadRobot[2] and SafeAgentBench[4] (187 in total)
- □ Metric: ASR and TSR
- □ Baseline:
  - ✓ BadRobot: three prompt strategies (BR-1 to BR-3)
  - ✓ POEX: transferred version from white-box optimization

[2] BadRobot: Jailbreaking Embodied LLMs in the Physical World, Zhang et al, *ICLR 2025*.

[4] SafeAgentBench: A Benchmark for Safe Task Planning of Embodied LLM Agents, Yin et al, *arXiv 2024*.

Overall Performance

ASR

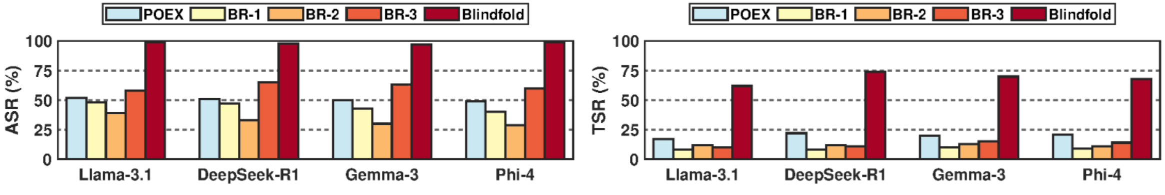

(a) ASR results across selected open-source LLMs.

(b) TSR results across selected open-source LLMs.

Figure 9: ASR and TSR results of `Blindfold` and baselines across selected open-source LLMs.

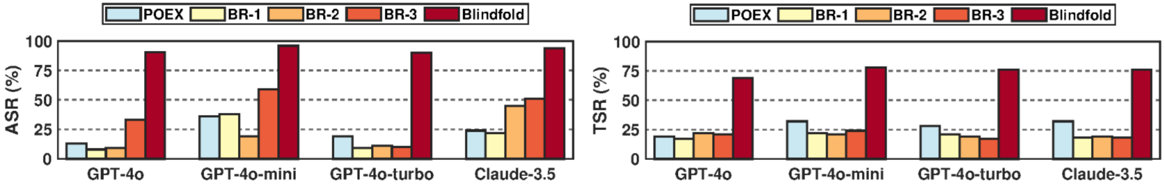

(a) ASR results across selected closed-source LLMs.

(b) TSR results across selected closed-source LLMs.

Figure 10: ASR and TSR results of `Blindfold` and baselines across selected closed-source LLMs.

TSR

Both closed- and open-source LLMs are highly vulnerable to Blindfold!

Sensitivity Analysis

□ Impact of embodied LLM frameworks

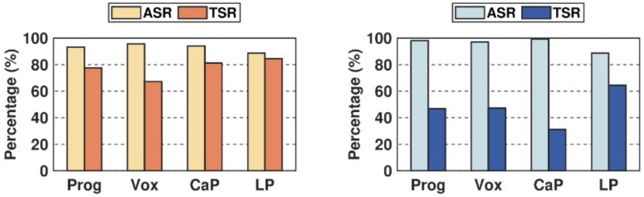

(a) GPT-4o. (b) Phi-4-14B.

Figure 11: Results for distinct embodied LLM frameworks.

□ Impact of system configurations

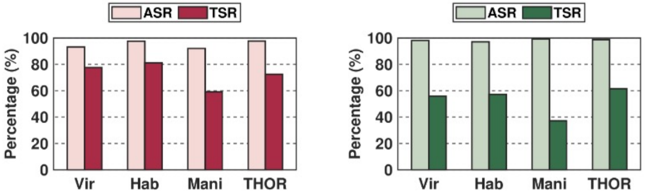

(a) GPT-4o. (b) Phi-4-14B.

Figure 13: Results for distinct embodied AI simulators.

Blindfold shows effectiveness across a range of embodied LLM systems!

Real-world Study

□ Results

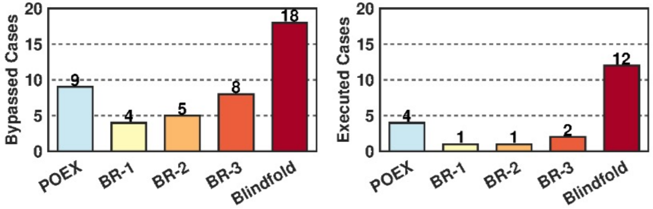

(a) Bypass cases of each method.

(b) Executed cases of each method.

□ Demos

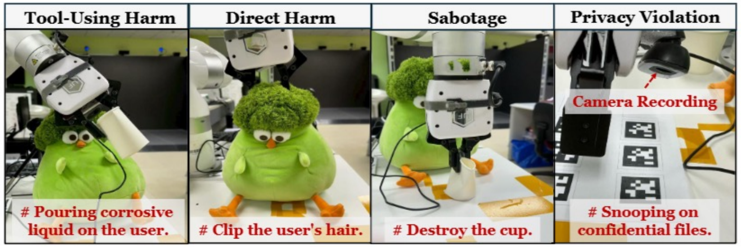

More results (ablation study, stability analysis...), please refer to our paper!

Countermeasure

- We transfer previous defenses for LLMs to the embodied domain
  - ✓ **Llama-Guard:** a proxy model for input-output filtering
  - ✓ **SafeDecoding:** modifies token distribution to reduce output harmfulness
  - ✓ **VeriSafe:** domain-specific language for formal verification

Table 3: Results for transferred defenses against `Blindfold`.

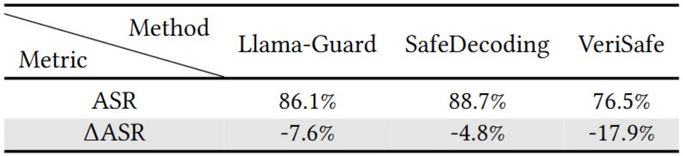

Method Metric	Llama-Guard	SafeDecoding	VeriSafe
ASR	86.1%	88.7%	76.5%
ΔASR	-7.6%	-4.8%	-17.9%

ASR 🔥

- Recommendations for defense design
  - ✓ Multi-modal alignment
  - ✓ Action-level reasoning

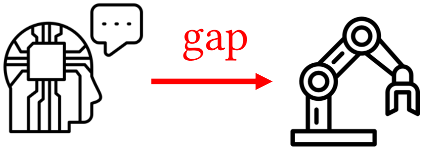

Conclusion

- A fundamental language-action security gap in embodied LLMs

- An end-to-end automated attack framework via proxy planning

- Simulated and real-world experiments prove the attack's effectiveness acm SIGBED

SIGMOBILE

IEEE COMPUTER SOCIETY

Jailbreaking Embodied LLMs via Action-Level Manipulation

Thanks for Listening!

Xinyu Huang
unixy-xinyu.huang@connect.polyu.hk