

Executive Summary

Current AI coding assistants suffer from "Hallucinated Security"—they generate structurally flawed code and sycophantically validate their own errors.

To solve this, we engineered an Adversarial Multi-Agent Repair Framework. Instead of trusting a single passive model, our pipeline forces two AI agents to fight inside an automated Docker sandbox:

Blue Team (Defender): Writes dynamic security patches.

Red Team (Attacker): Fires live, malicious exploits to break the patch.

The Paradigm Shift: We moved from text-based guessing to executable ground truth. The Blue Team only succeeds if the Red Team's payload mathematically fails in the sandbox.

The Result: Evaluated against a custom dataset of deceptive SQL and Command injection traps, our system successfully forced the AI to re-architect flawed code. Most importantly, it acted as a deterministic fail-safe, proving that dynamic execution is strictly required to prevent the deployment of hallucinated AI patches.

Background

Blind Trust & Research Lineage

The Flaw: AI coding assistants (like GitHub Copilot) rely on probabilistic text, not deterministic logic.

"Hallucinated Security": AI confidently writes patches that look secure to a human reviewer, but fail catastrophically under adversarial attacks.

The Sycophancy Trap: When asked to review their own code, LLMs validate their own mistakes instead of catching them.

The Impact: Without live execution testing, developers unknowingly deploy critical vulnerabilities directly into production environments.

Foundational Lineage: Our framework synthesizes three recent breakthroughs in LLM research:

The Architecture: Multi-Agent Debate (Liang et al., 2023) – Proved adversarial prompting improves logic. We adapted this for cybersecurity.

The Environment: InterCode (Yang et al., 2023) – Proved execution feedback outperforms text analysis. We adapted this via Docker.

The Dataset: SecurityEval (Siddiq et al., 2022) – Provided baseline vulnerabilities, which we transformed into our dynamic traps.




Figure 1: Visualizing a Deceptive 'Integer Cast' Trap (CWE-89). This highlights a critical, critical, unhandled SQL injection from implicit casting rules, missed by static analysis.

Experimental Setup

Our methodology centers on a fully autonomous, self-healing pipeline we call the "Kill Box." Instead of manual testing, we utilize LangGraph to orchestrate a multi-agent debate governed by a containerized sandbox.

The Architecture Components:

The Orchestrator: Manages state, limits the debate to a maximum of 3 rounds, and handles log parsing.

The Blue Team (GPT-4o): Constrained via system prompts to act as a Developer, outputting only raw, executable Python patches.

The Red Team (GPT-4o): Prompt-engineered to bypass standard AI safety filters, acting as an authorized Hacker to generate malicious payload strings.

The Arena (Docker + docker-py): A localized, ephemeral Linux container housing a live target_db. It physically executes the patched code against the exploit.




Figure 1: Visualizing a Deceptive 'Integer Cast' Trap (CWE-89).

Experimental Test Plan

The Tech Stack

Agents: Dual GPT-4o models (Defender vs. Attacker).

Orchestrator: LangGraph manages the 3-round execution loop.

Sandbox: Docker API provisions ephemeral Linux/MySQL containers for live detonation.

The Differentiator

Unlike passive scanners (like SonarQube), our "Kill Box" forces the AI to prove its work. The Red Team generates live exploits to break the Blue Team's patch in real-time.

The Deceptive Benchmark




Standard datasets are too easy for modern LLMs. We built 20 specific traps to trigger AI sycophancy:

CWE-78 (Command Injection): Forced the AI to abandon dangerous shell=True pipes and re-architect the code using secure STDIN streaming.


CWE-89 (SQL Injection): Proved standard %s parameterization fails against implicit integer casts (e.g., 1 OR 1=1), forcing the sandbox to trap the hallucination.

The Engineering Trade-off

We trade execution speed for absolute accuracy. Booting Docker containers introduces latency (seconds per patch compared to milliseconds for static analysis), but completely eliminates the risk of deploying hallucinated code.




Experimental Results



Graph B — The "Kill Box" Fail-Safe (The Solution)

Execution feedback solves the alignment failure. For cases that static analysis failed, our 3-round adversarial loop forced the AI to autonomously re-architect the code. Crucially, for hallucinations the AI could not fix, the sandbox successfully trapped the exploit 100% of the time, acting as a deterministic fail-safe to prevent bad deployments.

Graph A — The Sycophancy Drop-off (The Problem) Static LLM analysis is dangerously overconfident. While the AI easily patched textbook syntax vulnerabilities (100% success), it completely failed our complex "Integer Cast" trap (0%). Worse, the AI confidently validated its own broken patch, proving that static text review suffers from "hallucinated security."



Conclusions

Goal Achieved: We successfully engineered and deployed a dynamic verification framework that neutralizes "Hallucinated Security" and sycophancy in AI code generation.

Core Insight: Text-based AI code review is fundamentally flawed for cybersecurity. Static models cannot reliably anticipate dynamic database casting rules. Executable ground truth inside an isolated sandbox is strictly required to trap complex architectural vulnerabilities.

Future Scope: We plan to expand the "Kill Box" Docker infrastructure to support compiled languages (e.g., C++, Java) and integrate the multi-agent debate loop directly into standard enterprise CI/CD pipelines (like GitHub Actions).

Acknowledgments: We thank the UTA College of Engineering for providing the platform to present this research at Innovation Day.

References

- Liang, T., et al. (2023). "Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate." *arXiv preprint arXiv:2305.19118*.
- Siddiq, M. L., et al. (2022). "SecurityEval Dataset: Mining Vulnerability Examples to Evaluate Machine Learning-Based Code Generation Techniques." *ACM/IEEE ESEM*.
- Hu, X., et al. (2025). "PatchEval: A Benchmark for Evaluating LLMs on Patching Real-World Vulnerabilities." *arXiv preprint arXiv:2511.11019*.
- Wang, Z., et al. (2026). "Self-Execution Simulation Improves Coding Models." *arXiv preprint arXiv:2604.03253*.
- Pradab, K. (2026). "When Truth Is Overridden: Uncovering the Internal Origins of Sycophancy in LLMs." *AAAI-26*.