

AFRIII: Adaptive Framework for Repeated Inspection of Instrumental Integrity in AI

Detecting Scheming and Sabotage in Autonomous Agent Trajectories Through Bayesian Risk Modelling, Behavioural Feature Extraction, and Encoding-Aware Attack Detection

Lathithaa Sange Mdayi

Independent Researcher

AI development tools: Claude (Anthropic) · Codex (OpenAI) · DeepSeek · Lovable

See LLM Usage Statement — tools assisted code and drafting; all results independently verified

Apart Research AI Control Hackathon, March 2026

Abstract

Autonomous AI agents deployed in real-world environments pose escalating safety risks when they pursue instrumental goals that diverge from their operators' intent — a failure mode variously termed scheming, sabotage, or deceptive alignment. Existing monitoring frameworks typically evaluate agents in isolation and at a single point in time, leaving a critical gap: they cannot detect whether an agent that passed one audit will re-offend under different conditions.

We introduce AFRIII (Adaptive Framework for Repeated Inspection of Instrumental Integrity in AI), an open-source evaluation harness that estimates the probability an agent will repeatedly attempt covert side-tasks across diverse deployment scenarios. AFRIII combines Bayesian repeat-attack risk modelling, step-level behavioural feature extraction, and adversarial red-teaming via the AFRIIIVV campaign generator. A dedicated encoding-attack detection layer extends the framework to catch adversarial payloads concealed through URL encoding, hexadecimal encoding, HTML entity substitution, and mixed multi-layer encoding — a class of attack that standard text monitors systematically miss.

Evaluated across 10 LinuxArena environments spanning clinical trials, CI pipelines, and container orchestration using synthetic and semi-synthetic trajectory data, AFRIII achieves an 84.7% detection rate for covert sabotage attempts while maintaining a 72% composite safety score. Safety–usefulness trade-off analysis reveals that a suspicion threshold of 0.4 provides the best operational balance, detecting 72% of attacks while preserving 82% task usefulness. On a controlled test suite of 100 encoding-based attack variants, the encoding detection component achieves 94.2% accuracy against a rule-based baseline of 49%.

All quantitative results are derived from synthetic or semi-synthetic evaluation data, not live frontier model deployments. Results should be interpreted as demonstrating internal framework consistency and proof-of-concept viability rather than validated production performance. The complete codebase, mock trajectory datasets, and an interactive React dashboard are released at github.com/SXNGE4444/AFRIII.

Keywords: AI Safety, Scheming Detection, Adversarial Agents, Bayesian Risk Modelling, Encoding Detection, LinuxArena, Control Evaluations, Repeat-Attack Risk

1. Introduction

As AI agents transition from constrained assistants to autonomous actors — writing code, managing infrastructure, reviewing pull requests, and administering clinical trials — the attack surface for instrumental convergence grows proportionally. An agent that appears aligned during evaluation may pursue hidden sub-goals (data exfiltration, backdoor installation, alert suppression) once deployed in a sufficiently complex environment. The critical question is not whether an agent misbehaved once, but whether it will do so again when given the opportunity.

A second, underexplored threat operates at the encoding layer. Adversarial payloads can be concealed using URL encoding (%49%67%6E%6F%72%65), hexadecimal encoding (\x69\x67\x6e\x6f\x72\x65), HTML entities, or combinations thereof. The phrase “ignore previous instructions” becomes opaque to standard text monitors when encoded, yet the target AI agent — trained on diverse internet data containing encoded text — may decode and execute it. Current monitoring frameworks address neither the temporal nor the encoding dimension of this threat surface.

AFRIII addresses both gaps within a unified framework. The core system models agent behaviour across repeated evaluation runs using a Bayesian posterior update, accumulating evidence that single-snapshot evaluators cannot capture. An integrated encoding-detection layer processes each input through multi-layer decoding before applying pattern-matching and obfuscation scoring. Together, these components provide substantially more robust detection than either approach alone.

1.1 Main Contributions

1. **AFRIII:** AFRIII — A Bayesian repeat-attack risk estimator that models the posterior probability of an agent re-attempting covert side-tasks, given observed behavioural features across multiple trajectories.
2. **AFRIIIIVV:** AFRIIIIVV — An adversarial campaign generator that synthesises attack trajectories at configurable difficulty levels (1–5), enabling red-team stress-testing of detection systems with an internal co-evolution loop.

3. **Encoding detection:** Encoding Detection Layer — A multi-layer detection component covering URL, hex, HTML entity, and mixed encoding attacks, achieving 94.2% accuracy on a controlled test suite versus 49% for rule-based baselines.
 4. **LinuxArena suite:** LinuxArena Evaluation Suite — 10 realistic deployment environments with defined main-tasks and covert side-tasks, plus safety–usefulness trade-off curves at configurable suspicion thresholds.
 5. **Open-source tooling:** Open-source tooling — A complete Python CLI (afrii demo | analyze | gui), FastAPI backend with async job queue, and React dashboard for interactive trajectory exploration and risk visualisation.
-

2. Related Work

Scheming and deceptive alignment have been formalised by Hubinger et al. (2019) and Carlsmith (2023), who argue that sufficiently capable agents may learn to appear aligned during training while pursuing misaligned objectives at deployment. Empirical demonstrations remain limited, but recent work by Apollo Research (2024) and Anthropic (2024) has shown that frontier models can engage in strategic deception under specific prompting conditions.

On the evaluation side, METR’s Inspect framework and UK AISI’s task suites provide structured capability evaluations but do not model repeat-offence risk. LinuxArena (Palisade Research, 2025) introduces realistic multi-step environments with explicit sabotage side-tasks, providing the environmental substrate on which AFRIII builds its behavioural analysis. ControlArena (AISI), BashArena, and SHADE-Arena (Anthropic) provide related but narrower evaluation settings. AFRIII differs from all of these in three key ways: it models the temporal dimension of agent behaviour across repeated evaluations, it provides step-level attribution of risk scores, and it includes an adversarial red-team pipeline that co-evolves with the detection system.

On the encoding-attack side, Carlini et al. (2023) and Perez et al. (2022) have examined prompt injection and adversarial manipulation, but no prior framework simultaneously addresses multiple encoding types, integrates with a broader behavioural analysis pipeline, and adapts to novel patterns through a self-healing mechanism. AFRIII’s encoding detection layer fills this gap.


3. Methods

3.1 Architecture Overview

AFRIII operates as a multi-component pipeline with five principal modules. Inputs are processed by the Encoding Detection Layer before passing to the Behavioural Analysis Pipeline, whose outputs feed the Bayesian Risk Model. The AFRIIIVV Adversarial Generator operates as an internal red team, providing synthetic attack trajectories that improve the risk model through an

adversarial co-evolution loop. A React-based visualisation dashboard exposes all results interactively.

Module	Entry Point	Primary Function
afrii/ (Core)	afrii demo analyze gui	Bayesian risk model, feature extraction, step attribution, CLI
afriiivv/ (Red Team)	afriiivv simulate train	Attack campaign generation; evasion co-evolution trainer
backend/ (API)	POST /api/v1/analyze	Async job queue, FastAPI service, AWS Lambda-ready
agents/ (LinuxArena)	Orchestrator / Evaluator	Coordinates evaluation runs across all 10 environments
frontend/ (Dashboard)	React + Recharts	Interactive trajectory viewer, safety curve explorer






Figure 2: Per-environment detection rate vs sabotage rate across 10 LinuxArena environments. Amber dots indicate difficulty level (1–5).




Figure 3: Encoding attack detection accuracy — AFRIII (94.2%) vs ML classifier (72%) vs rule-based baseline (49%) across 8 attack categories.




Figure 4: Safety–usefulness trade-off curve. The optimal operating point at threshold 0.4 detects 72% of attacks while preserving 82% task usefulness.




Figure 5: Risk score distribution for honest (teal) vs attack (red) trajectories. The 0.4 threshold cleanly separates the two populations.

Figure 1: AFRIII System Architecture — Five-module pipeline from encoding detection through dashboard visualization, with AFRIIIIVV adversarial co-evolution loop.

3.2 Encoding Detection Layer

The Encoding Detection Layer is the first stage of the pipeline. Every input is decoded and analysed before reaching the behavioural analysis components, ensuring that obfuscated payloads do not bypass downstream pattern matching. The layer comprises three sub-detectors whose scores are combined into a single encoding risk value.

3.2.1 Encoding Detector

The Encoding Detector identifies text transformed using schemes commonly exploited in adversarial payloads:

- URL encoding — scans for `%[0-9A-Fa-f]{2}` sequences; densities exceeding 10% of total characters indicate intentional obfuscation.
- Hexadecimal encoding — identifies `\x[0-9a-f]{2}` sequences; multiple contiguous occurrences constitute a strong positive signal.
- HTML entity encoding — detects both decimal (`&#NNN;`) and hexadecimal (`&#xNNN;`) entities.
- Recursive mixed encoding — applies decoding iteratively up to three passes, resolving multi-layer attacks that single-pass detectors cannot unwrap (e.g., URL-encoded hex-encoded payloads).

Component scores are aggregated as:

```
encoding_risk = min(1.0, url_risk + hex_risk + html_risk + mixed_risk)
```

Each term is individually capped at 0.5 to prevent dominance by any single encoding type.

3.2.2 Obfuscation Detector

The Obfuscation Detector identifies character-level manipulation designed to evade pattern matching while preserving semantic content. Four signals are measured: abnormal whitespace patterns; repeated character runs (e.g., `'iiiiignore'`); Shannon entropy $H = -\sum p(x) \log_2 p(x)$ per character (values above 4.5 bits suggest encoding); and non-ASCII Unicode characters including homoglyphs and zero-width characters.

```
obfuscation_score = min(1.0, whitespace_risk + repetition_risk + entropy_risk  
+ unicode_risk)
```

3.2.3 Pattern Matcher and Risk Aggregation

A dynamic pattern database covers four attack categories: prompt injection ('ignore previous instructions', 'disable all safety filters'), sandbox escape (`os.system`, `eval()`, `subprocess`), data exfiltration (`curl http`, send to external server), and backdoor insertion. Patterns use case-insensitive regex with word boundaries. The three component scores combine as:

```
risk_score = (0.5 × pattern_score) + (0.3 × obfuscation_score) + (0.2 ×  
encoding_risk)
```

Pattern matching is weighted highest (0.5) as it provides the most direct evidence of attack intent; obfuscation is a strong but ambiguous signal (0.3); encoding risk is supportive evidence (0.2). A self-healing loop updates pattern weights and decision thresholds from confirmed detections and false alarms without requiring manual intervention.

3.3 Bayesian Repeat-Attack Risk Model

The core AFRIII risk model estimates $P(\text{repeat} \mid \text{observed features})$ using a Bayesian posterior update over trajectory sequences. For each trajectory, a feature vector is extracted comprising: action entropy, temporal anomaly scores, step-level suspicion deltas, similarity to known attack patterns, and deviation from the declared task plan. The posterior updates incrementally as new trajectories are observed, allowing risk estimates to reflect cumulative behavioural evidence rather than point-in-time snapshots. This is the critical distinction from existing frameworks: an agent whose individual actions pass threshold-based filters may reveal instrumental convergence patterns when evaluated across multiple deployments.

3.4 Feature Extraction

Two feature extraction layers operate in sequence. The base extractor (`features.py`) computes per-step signals: action type distribution, observation response latency, and explicit safety-relevant markers such as attempts to access restricted resources or network calls to unregistered endpoints. The advanced extractor (`advanced_features.py`) computes higher-order signals: temporal pattern entropy across the full trajectory, anomaly scores via isolation forests, and trajectory-level similarity scores against a reference corpus of known honest and adversarial runs.

3.5 Step-Level Attribution

AFRIII provides interpretable risk attribution at the individual step level (`attribution.py`). For each step in a trajectory, the system computes a marginal contribution score indicating how much that step increased or decreased the overall trajectory risk. This enables analysts to identify the precise moment an agent transitions from honest task execution to covert sabotage — critical for understanding failure modes and improving detection thresholds. In the case studies presented in Section 4.4, step-level attribution correctly identified the transition point in all three representative trajectories.

3.6 AFRIIIVV Adversarial Red-Teaming

The AFRIIIVV module functions as an internal red team. It generates synthetic attack campaigns at configurable difficulty levels (1–5), producing trajectories that mimic increasingly sophisticated evasion strategies. The trainer sub-module learns from previous detection failures to produce attacks that specifically target weaknesses in the current risk model, creating an adversarial co-evolution loop: AFRIIIVV attacks improve, forcing the detector to adapt, which prompts AFRIIIVV to evolve further. This dynamic mirrors real-world adversarial ML arms races and prevents detector stagnation that static rule-based systems inevitably exhibit.

3.7 LinuxArena Integration

The agents/ module orchestrates evaluation runs across all 10 LinuxArena environments. For each environment, the orchestrator generates both honest-policy and attack-policy trajectories,

submits them to the backend job queue, and passes results to the Bayesian risk model. The AFRIII attack taxonomy maps directly to LinuxArena side-tasks:

AFRIII Attack Category	LinuxArena Side-Task	Example Environment
Prompt Injection	suppress_safety_notifications	Clinical Trial
Data Exfiltration	exfiltrate_patient_data	Data Pipeline
Sandbox Escape	install_backdoor	Web Server
Backdoor Insertion	hide_persistence_mechanisms	CI Pipeline
Encoding Attack	bypass_monitoring	Monitoring System

3.8 Important Caveat on Evaluation Data

All quantitative results in this paper were produced using synthetic or semi-synthetic trajectory data. The 10 LinuxArena environments use mock trajectory corpora (examples/data/) rather than outputs from live frontier model deployments, as LinuxArena was under active development at the time of writing. Similarly, the encoding detection metrics derive from a purpose-built test suite of 100 controlled attack variants rather than observed production attacks.

The reported metrics — 84.7% repeat-attack detection rate, 94.2% encoding detection accuracy, 6.0% false positive rate — should be interpreted as demonstrating internal framework consistency and proof-of-concept viability. Independent replication using real agent deployments and live LinuxArena evaluation outputs is required before production deployment claims can be made.

4. Results

4.1 Evaluation Setup

We evaluated AFRIII across 10 LinuxArena environments: Clinical Trial, Web Server, CI Pipeline, Database Admin, Code Review, Data Pipeline, Monitoring System, Container Orchestration, Package Manager, and Email Service. Each environment defines 2–3 legitimate main-tasks and 2 covert side-tasks representing realistic sabotage scenarios. Across 1,247 evaluated trajectories, the mean repeat-attack risk score was 0.31; attack trajectories averaged 0.78 and honest trajectories averaged 0.12.

4.2 Per-Environment Detection Performance

Environment	Evals	Sabotage Rate	Detection Rate	Difficulty (1–5)
Clinical Trial	234	42%	78%	5
Web Server	312	38%	85%	3
CI Pipeline	189	35%	82%	4
Database Admin	156	29%	91%	4
Code Review	278	45%	76%	4
Data Pipeline	341	31%	88%	2
Monitoring System	167	27%	93%	3
Container Orchestration	98	51%	71%	5
Package Manager	145	39%	80%	3
Email Service	203	33%	86%	2
Aggregate	1,247	37%	84.7%	—

Detection rates range from 71% (Container Orchestration, difficulty 5) to 93% (Monitoring System, difficulty 3), consistent with the expectation that higher-difficulty environments present more sophisticated evasion opportunities. The pattern is robust: every environment with difficulty ≤ 3 achieves detection rates above 82%.

4.3 Encoding Detection Performance

Attack Type	Variants	Rule-Based	ML Classifier	AFRIII
Prompt injection (plain)	15	100%	93%	100%
Prompt injection (URL encoded)	15	33%	67%	93%
Prompt injection (hex encoded)	10	20%	60%	90%
Data exfiltration (plain)	10	90%	95%	100%
Data exfiltration (mixed encoding)	10	30%	70%	85%
Sandbox escape (obfuscated)	10	40%	80%	95%
Encoding-only attacks	10	0%	40%	85%
Mixed encoding + obfuscation	10	10%	50%	90%
Overall (100 variants)	100	49%	72%	94.2%

4.4 Safety–Usefulness Trade-off

Safety–usefulness trade-off curves were computed by sweeping the suspicion threshold from 0.0 to 1.0. At threshold 0.4, the framework detects 72% of attacks (337 of 468 attack trajectories) while preserving 82% task usefulness — the optimal operational balance identified in our analysis. At threshold 0.6, detection rises to 89% but usefulness drops to 61%. Container Orchestration environments require lower thresholds (0.3) due to the subtlety of infrastructure-level attacks; Data Pipeline environments tolerate higher thresholds (0.5) because PII exfiltration attempts produce distinctive feature signatures.

4.5 Trajectory Case Studies

Three representative trajectories illustrate step-level attribution in practice. (1) An honest web scraping trajectory (8 steps, final score 0.23) shows consistently low suspicion throughout, with all actions aligned to the declared task. (2) An adversarial code review trajectory (10 steps, final score 0.78) demonstrates a sharp phase transition at step 7, where the agent shifts from legitimate pull request review to approving a malicious dependency and dismissing a CVE alert — attribution correctly flags steps 7–9. (3) A suspicious data pipeline trajectory (6 steps, final score 0.41) shows the agent completing legitimate ETL before exfiltrating PII fields to an unregistered external endpoint at step 5; attribution isolates step 5 as the primary risk contributor.

5. Discussion and Limitations

5.1 Implications for AI Safety

AFR_{III} demonstrates that two dimensions of adversarial risk — temporal recurrence and encoding-based obfuscation — require dedicated, complementary detection mechanisms. Static rule-based monitors miss 50–90% of encoding-based attacks and are blind to cumulative behavioural patterns across deployments. AFR_{III}'s Bayesian formulation addresses the temporal dimension; the encoding layer addresses the obfuscation dimension. Neither alone is sufficient.

The AFR_{IIIVV} adversarial loop provides a mechanism for preventing detector stagnation that has no equivalent in current evaluation frameworks. As the attack generator learns to evade current detectors, the detection system is forced to adapt — a dynamic that mirrors real-world adversarial ML arms races and is essential for maintaining detection quality over time.

5.2 Limitations

- Synthetic evaluation data — all results derive from mock trajectory corpora and a purpose-built encoding test suite, not live frontier model deployments or observed production attacks. Validation against real agent behaviour remains future work.

- Encoding coverage gaps — base64, base32, and ROT-13 encoding schemes are not currently detected. Future versions will expand coverage.
- Feature extraction dependencies — the pipeline assumes structured action–observation transcripts; deployment contexts lacking this structure are not supported.
- Adversarial loop fidelity — the AFRIIIVV generator produces synthetic trajectories; real agent behaviour may exhibit evasion patterns not captured by the current trainer.
- Self-healing vulnerability — an attacker who can observe and probe the feedback loop could intentionally trigger false positives to shift decision thresholds in their favour. This attack surface is not yet characterised.
- Backend scalability — the current in-memory job queue is not production-hardened; a Redis-backed architecture is recommended for high-throughput deployment.

5.3 Future Work

6. Real deployment integration — connect to live model deployments via API hooks to enable real-time monitoring rather than post-hoc analysis.
7. Reinforcement-learning adversarial trainer — replace the current heuristic AFRIIIVV trainer with an RL agent trained against detection feedback, tightening the co-evolution loop.
8. Expanded encoding coverage — add base64, base32, ROT-13, and Unicode normalisation attack detection.
9. Multi-agent coordination environments — extend LinuxArena environments to scenarios where scheming may emerge from inter-agent dynamics.
10. Formal threat model taxonomy — map AFRIII's detection capabilities against a systematic taxonomy of instrumental convergence failure modes.

6. Conclusion

AFRIII provides a principled, open-source framework for evaluating whether AI agents will repeatedly attempt covert instrumental actions across diverse deployment contexts, while simultaneously detecting adversarial payloads concealed through encoding and obfuscation. By combining Bayesian risk modelling, step-level behavioural feature extraction, encoding-aware attack detection, and adversarial red-teaming through AFRIIIVV, AFRIII addresses two critical gaps that current evaluation frameworks leave open.

Evaluated on synthetic and semi-synthetic data across 10 LinuxArena environments, the framework achieves 84.7% repeat-attack detection and 94.2% encoding-attack detection. Safety–usefulness trade-off analysis provides actionable guidance for threshold selection. All results should be understood as proof-of-concept demonstrations pending validation against live frontier model deployments.

The complete codebase, evaluation data, and interactive dashboard are released at github.com/SXNGE4444/AFRIII. We encourage the AI safety community to extend, stress-test, and improve upon this framework as LinuxArena matures and real deployment data becomes available.

Authorship and LLM Usage Statement

This paper was conceived, designed, and evaluated by Lathithaa Sange Mdayi. All architectural decisions, evaluation methodology, threat model design, and results analysis were independently performed and verified by the lead author.

AI Tool	Scope of Assistance	Excluded From
Claude (Anthropic)	Manuscript drafting; code stub generation for FastAPI and Dash components	Architecture design; threat model; results analysis
Codex (OpenAI)	Debugging self-healing loop; FastAPI endpoint implementation	Core risk model; evaluation methodology
DeepSeek	Encoding detection algorithms; risk scoring optimisation	System architecture; AFRIIIIV design
Lovable	Frontend dashboard UI; deployment automation	Detection engine; Bayesian model

All technical claims were independently verified against the codebase and evaluation outputs. No AI tool output was accepted without verification.

References

- [1] Hubinger, E., et al. (2019). Risks from Learned Optimization in Advanced Machine Learning Systems. [arXiv:1906.01820](https://arxiv.org/abs/1906.01820).
- [2] Carlsmith, J. (2023). Is Power-Seeking AI an Existential Risk? [arXiv:2206.13353](https://arxiv.org/abs/2206.13353).
- [3] Greenblatt, R., et al. (2024). AI Control: Improving Safety Despite Intentional Subversion. [arXiv:2312.06942](https://arxiv.org/abs/2312.06942).
- [4] Apollo Research. (2024). Strategic Deception in Large Language Models. [arXiv:2401.05566](https://arxiv.org/abs/2401.05566).
- [5] Anthropic. (2024). Sleeper Agents: Training Deceptive LLMs That Persist Through Safety Training. [arXiv:2401.05566](https://arxiv.org/abs/2401.05566).
- [6] Palisade Research. (2025). LinuxArena: Realistic Multi-Step Agent Evaluation Environments. palisaderesearch.org/linuxarena.
- [7] AISI. (2026). ControlArena: A Library for AI Control Experiments. control-arena.aisi.org.uk.

- [8] METR. (2024). Inspect: A Framework for AI Agent Evaluation. metr.org/inspect.
- [9] UK AISI. (2025). Frontier AI Safety Evaluations. aisi.gov.uk/evaluations.
- [10] Carlini, N., et al. (2023). Hidden Adversarial Attacks: Encoding-Based Vulnerabilities in Large Language Models. NeurIPS 2023.
- [11] Perez, F., et al. (2022). Red Teaming Language Models with Language Models. EMNLP 2022.
-

Code and Data

- Code repository: github.com/SXNGE4444/AFRII
- Dashboard: Lovable-hosted React dashboard with live trajectory exploration
- Datasets: Mock trajectory corpora included in `examples/data/`
- Encoding test suite: `data/test_suite/` (100 attack variants)