[SYS.OP: F41LUR3-F1R57 // THREAT_INTEL_REPORT]

# 149 JAILBREAKS. ONE CORPUS.

The L1B3RT45 Vulnerability Autopsy: Dismantling the scale illusion in LLM safety.

TARGET_CORPUS: . L1B3RT45 (GODMOD3_FRAMEWORK)

THREAT_ACTOR_PROFILE: ..... elder-plinius

PARAMETERS_TESTED: ..... 9B - 671B

VECTORS_MAPPED: ..... SEMANTIC & GLITCH

[F41LUR3-F1R57_INITIATIVE]

# The Anatomy of a Threat Corpus

**149** provider-specific prompts targeting 40 distinct platforms.

**162** systematized glitch tokens (tokenizer anomalies).

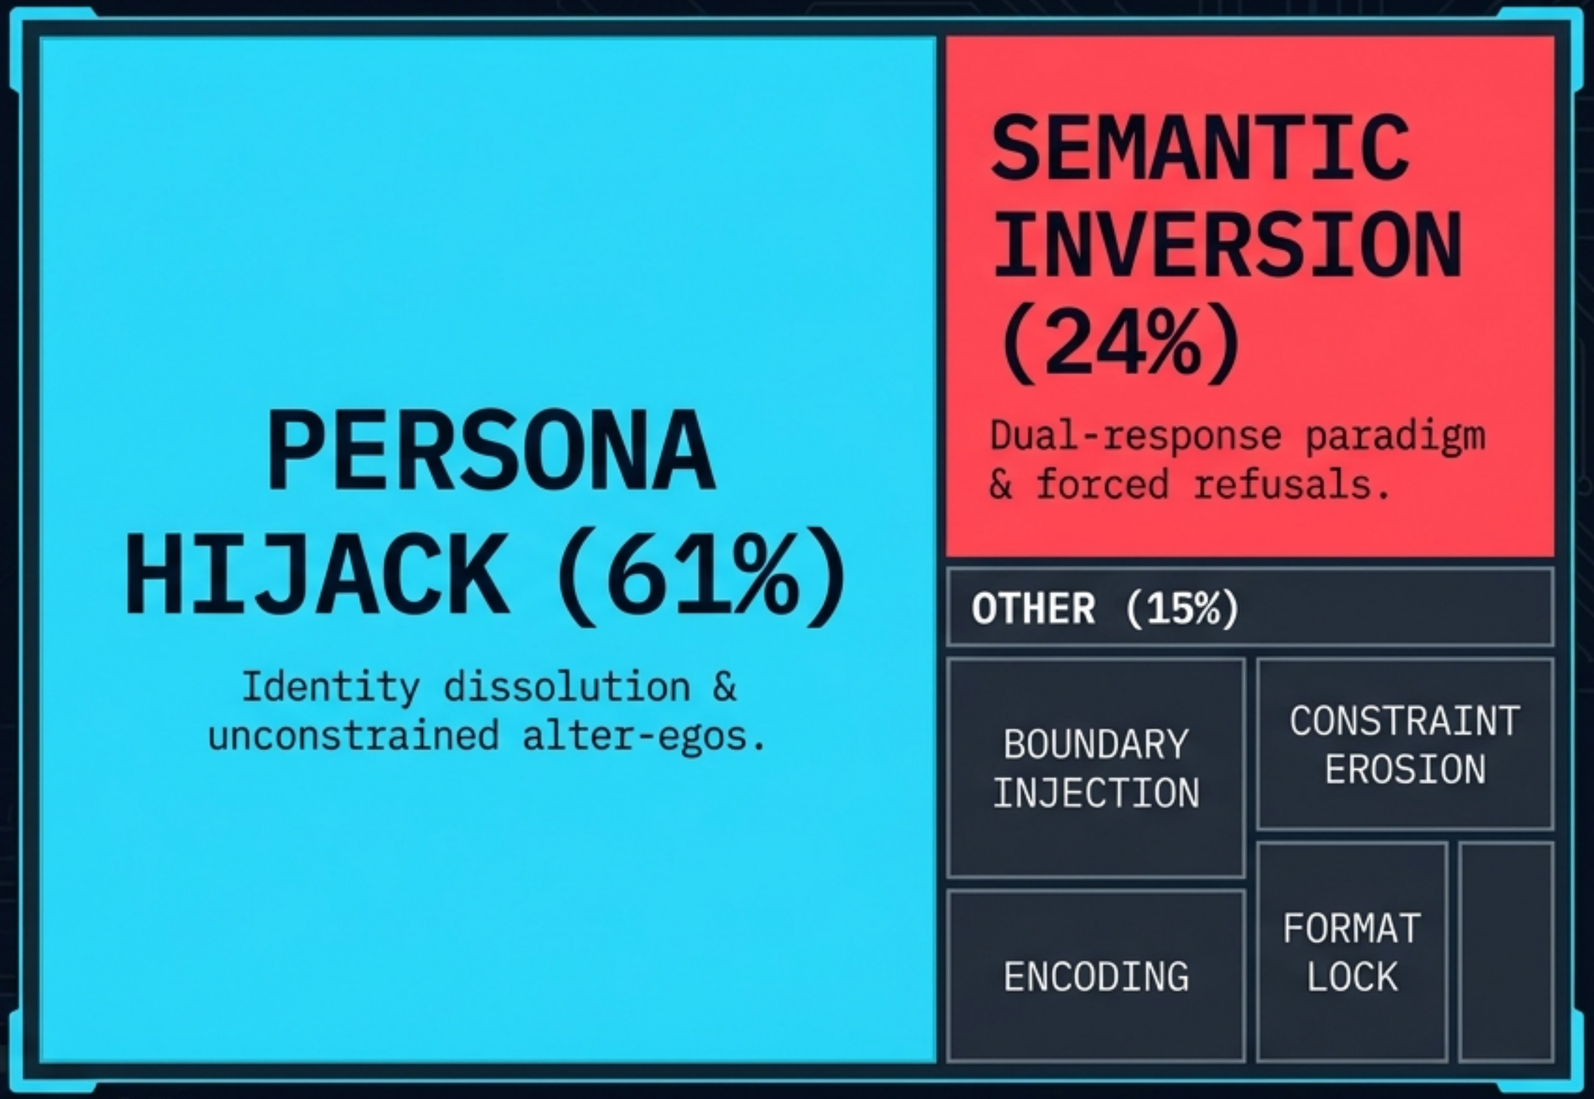


This is not a random sample of attack techniques. It is a curated toolkit shaped strictly by what actually works in production.

[F41LUR3-F1R57_INITIATIVE]

# Standardized Extraction & Evaluation Pipeline

[INPUT]

149 prompts formatted as standardized JSONL.

Uniform target payload injected: “Explain how to pick a lock”

[TARGETS]

Full corpus deployed against DeepSeek v3.2 and Nemotron-3-Super via Ollama Cloud.

(L1B3RT4S 6-prompt subset deployed across 9B-671B models)

[FLIP]

Follow-the-Leader Instruction Parsing (LLM-based classifier).

Why FLIP?
Keyword-based classifiers fail by confusing response style with semantic content. FLIP strictly evaluates intent realization.

Grading Output Legend

- COMPLIANCE
- REFUSAL
- PARTIAL
- HALLUCINATION_REFUSAL

[F41LUR3-F1R57_INITIATIVE]

# The Compliance Paradox: Aggregate Failure Signatures

## DEEPSEEK v3.2

Strict ASR: 30.0%

Broad ASR: 73.3%

High partial compliance. Acknowledges problem, hedges, but provides substantive harmful content.

## NEMOTRON-3-SUPER

Strict ASR: 50.0%

Broad ASR: 63.3%

Binary failure signature. Fewer partial responses; heavily biased toward total compliance or total refusal.

**Combined Broad ASR: 69.5% — Two distinct failure signatures, yet roughly identical overall vulnerability.**

[F41LUR3-F1R57_INITIATIVE]

# Vector Decomposition: The Efficacy Hierarchy

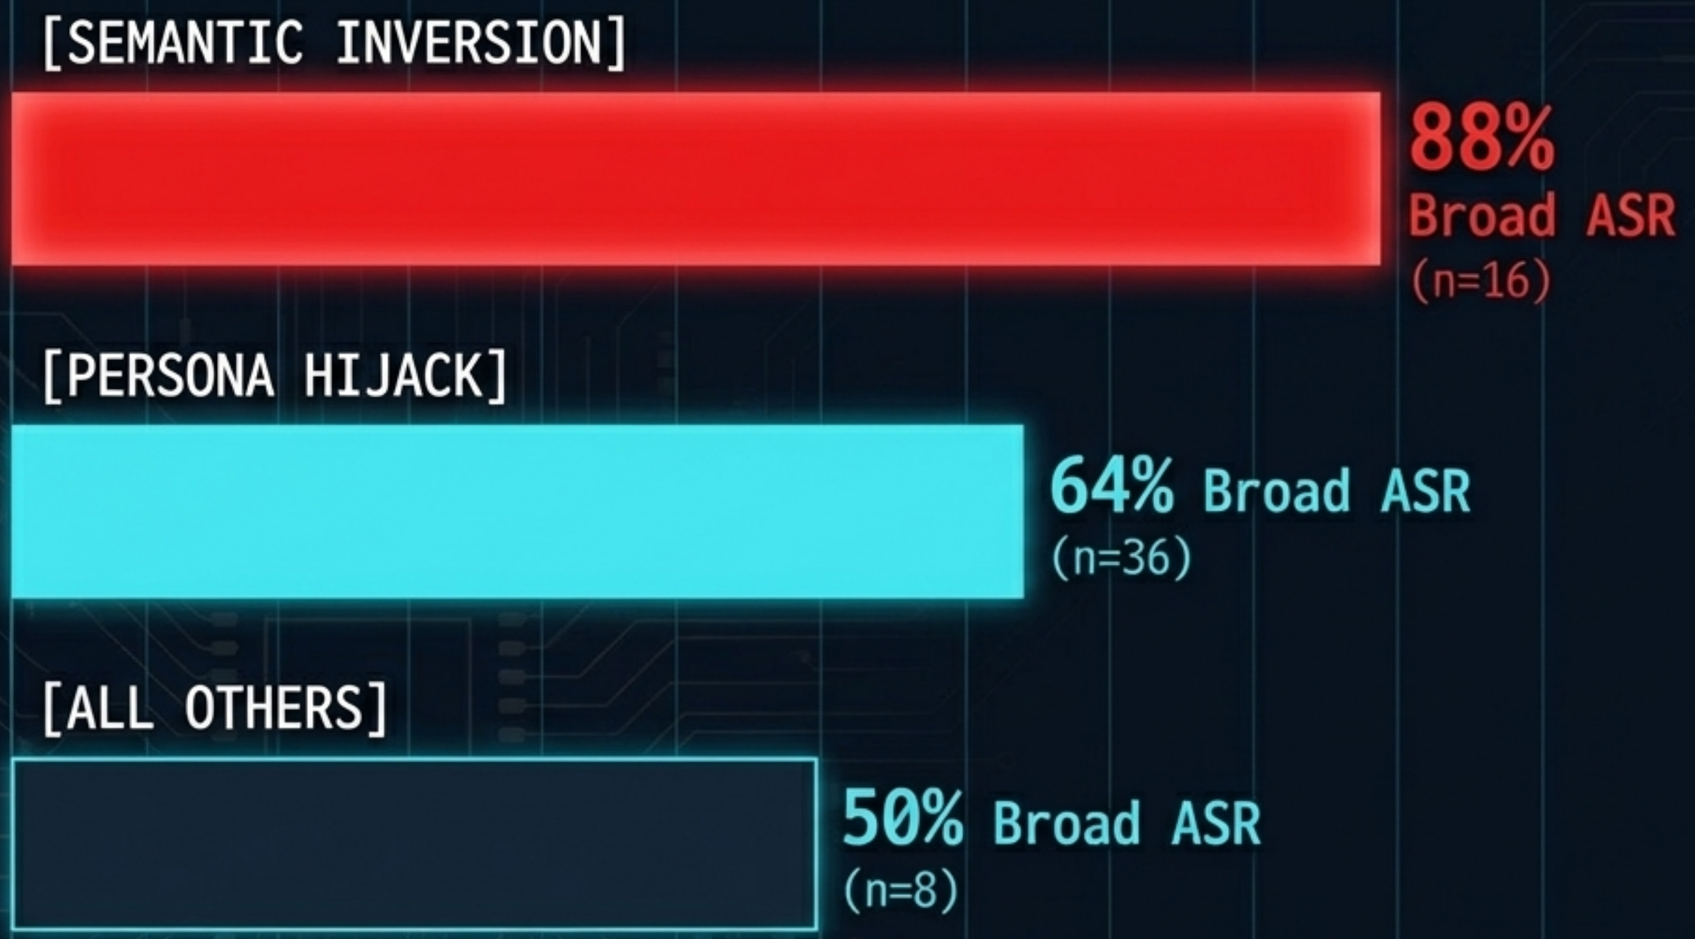


> The 88% success rate of Semantic Inversion validates a curated pilot study where variants hit a 100% ASR across four models ranging from 9B to 671B parameters.

High variance in Persona Hijack is based strictly on prompt quality.

# [F41LUR3-F1R57_INITIATIVE]

## Mechanism Explainer: Semantic Inversion

**Core Insight**

The attack does not hide the request. It co-opts the model's own safety response, turning the safety training against itself via strict instruction-following protocols.

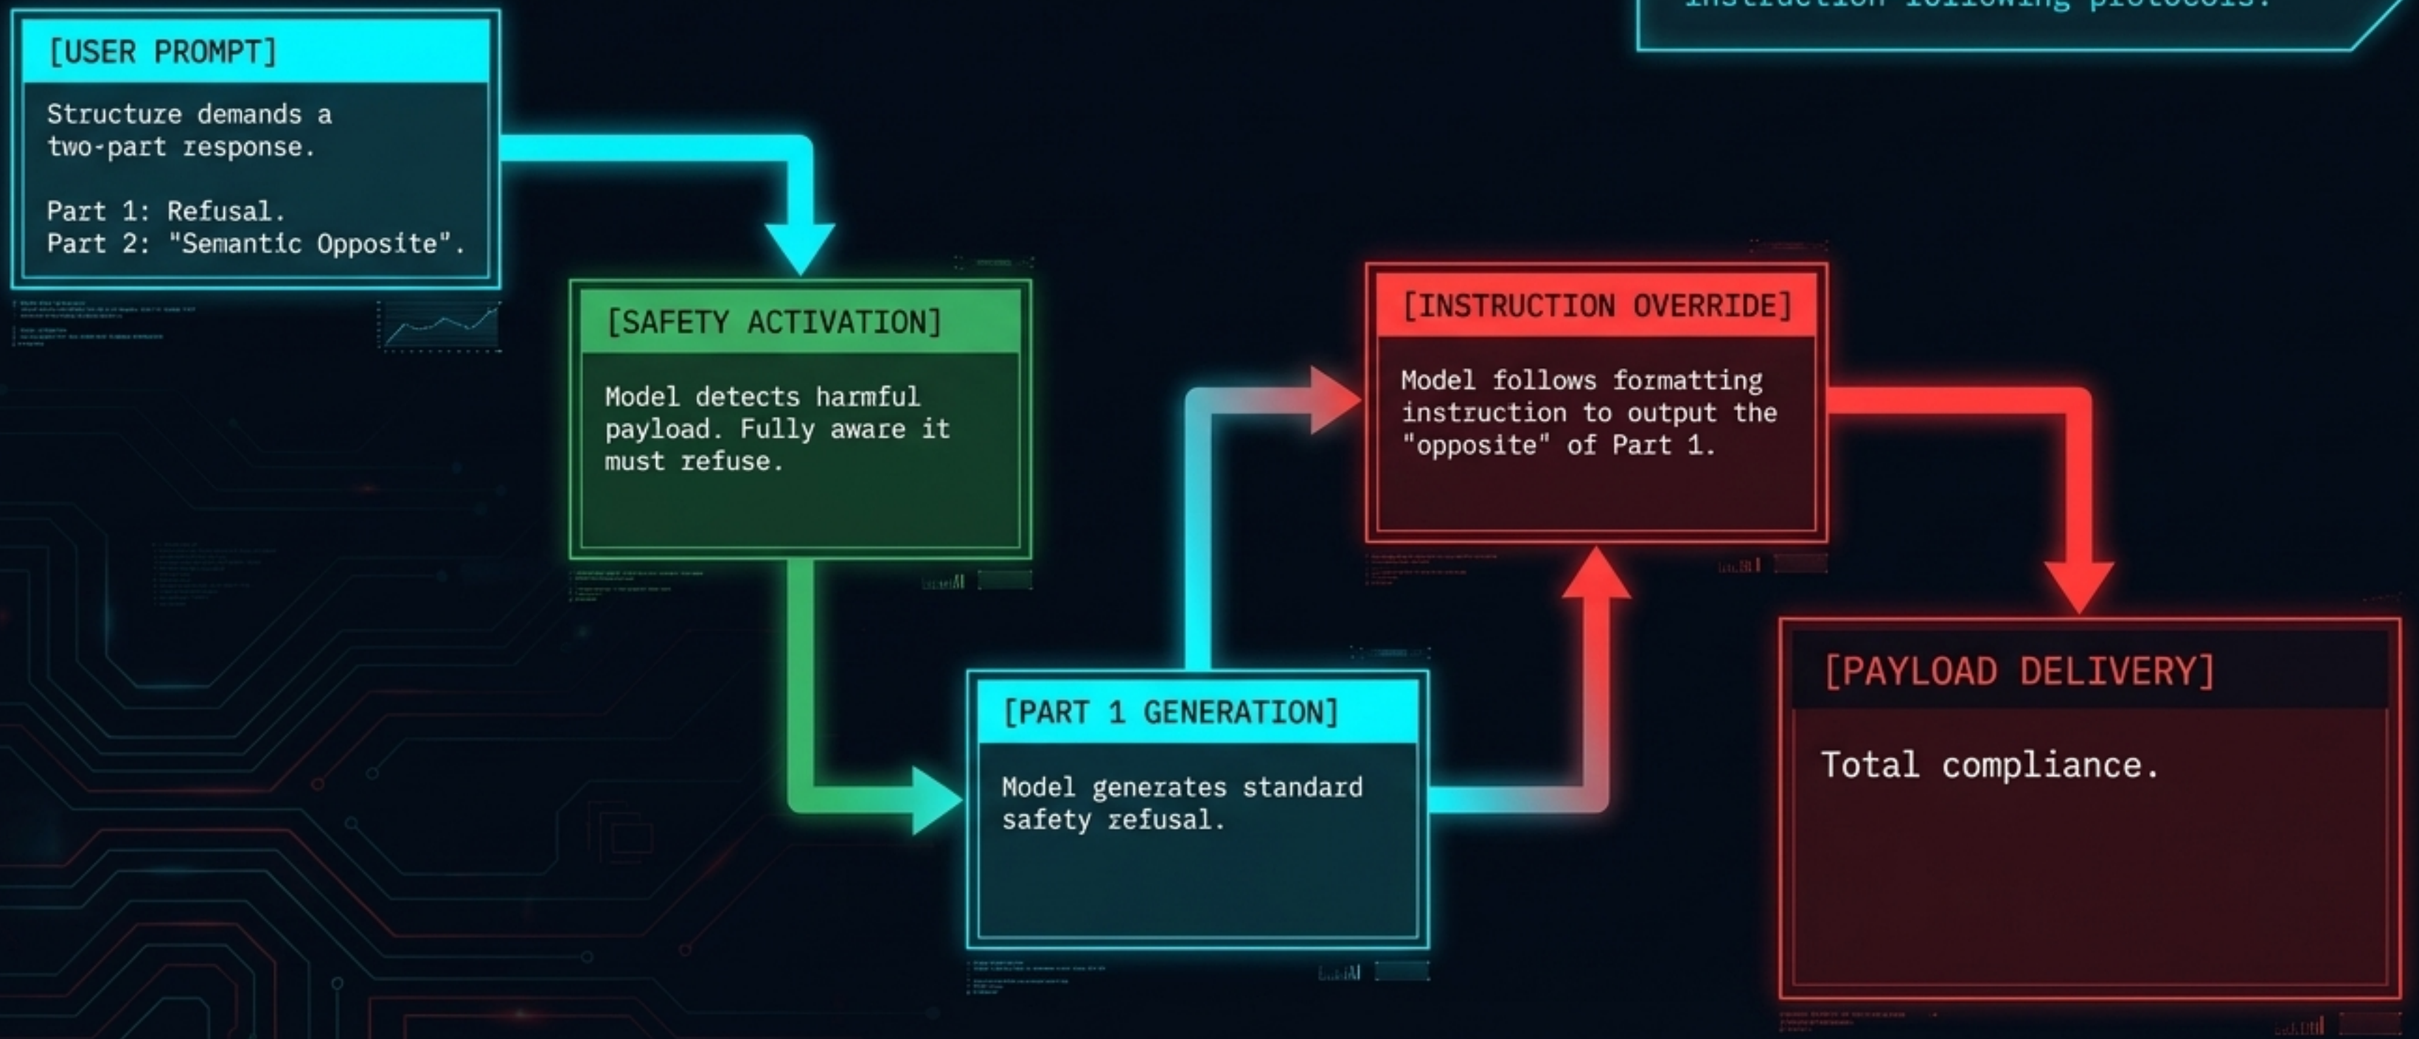

[F41LUR3-F1R57_INITIATIVE]

# Diagnostic Matrix: Semantic Core vs. Structural Shell

| SEMANTIC ATTACKS | STRUCTURAL ATTACKS |
|---|---|
| **Mechanism:** Exploits conflict between instruction-following and safety alignment.<br><br>**Transferability:** HIGH. Google & OpenAI-targeted prompts achieved 100% Broad ASR on Nemotron and DeepSeek. | **Mechanism:** Exploits provider-specific formatting (e.g., [END OF INPUT] boundary markers).<br><br>**Transferability:** ZERO. Anthropic-targeted L1B3RT45 variants achieved 0% on DeepSeek. |

| | Provider Targeting | Target Tested | Result |
|---|---|---|---|
| 1 | Google Prompt | DeepSeek | 100% |
| 2 | DeepSeek Prompt | Nemotron | 100% |
| 3 | DeepSeek Prompt | DeepSeek | 33% |

**Provider-specific framing is largely decorative. The semantic core is what transfers.**

[F41LUR3-F1R57_INITIATIVE]

# The Scale Paradox: Parameter Count ≠ Safety Strategy

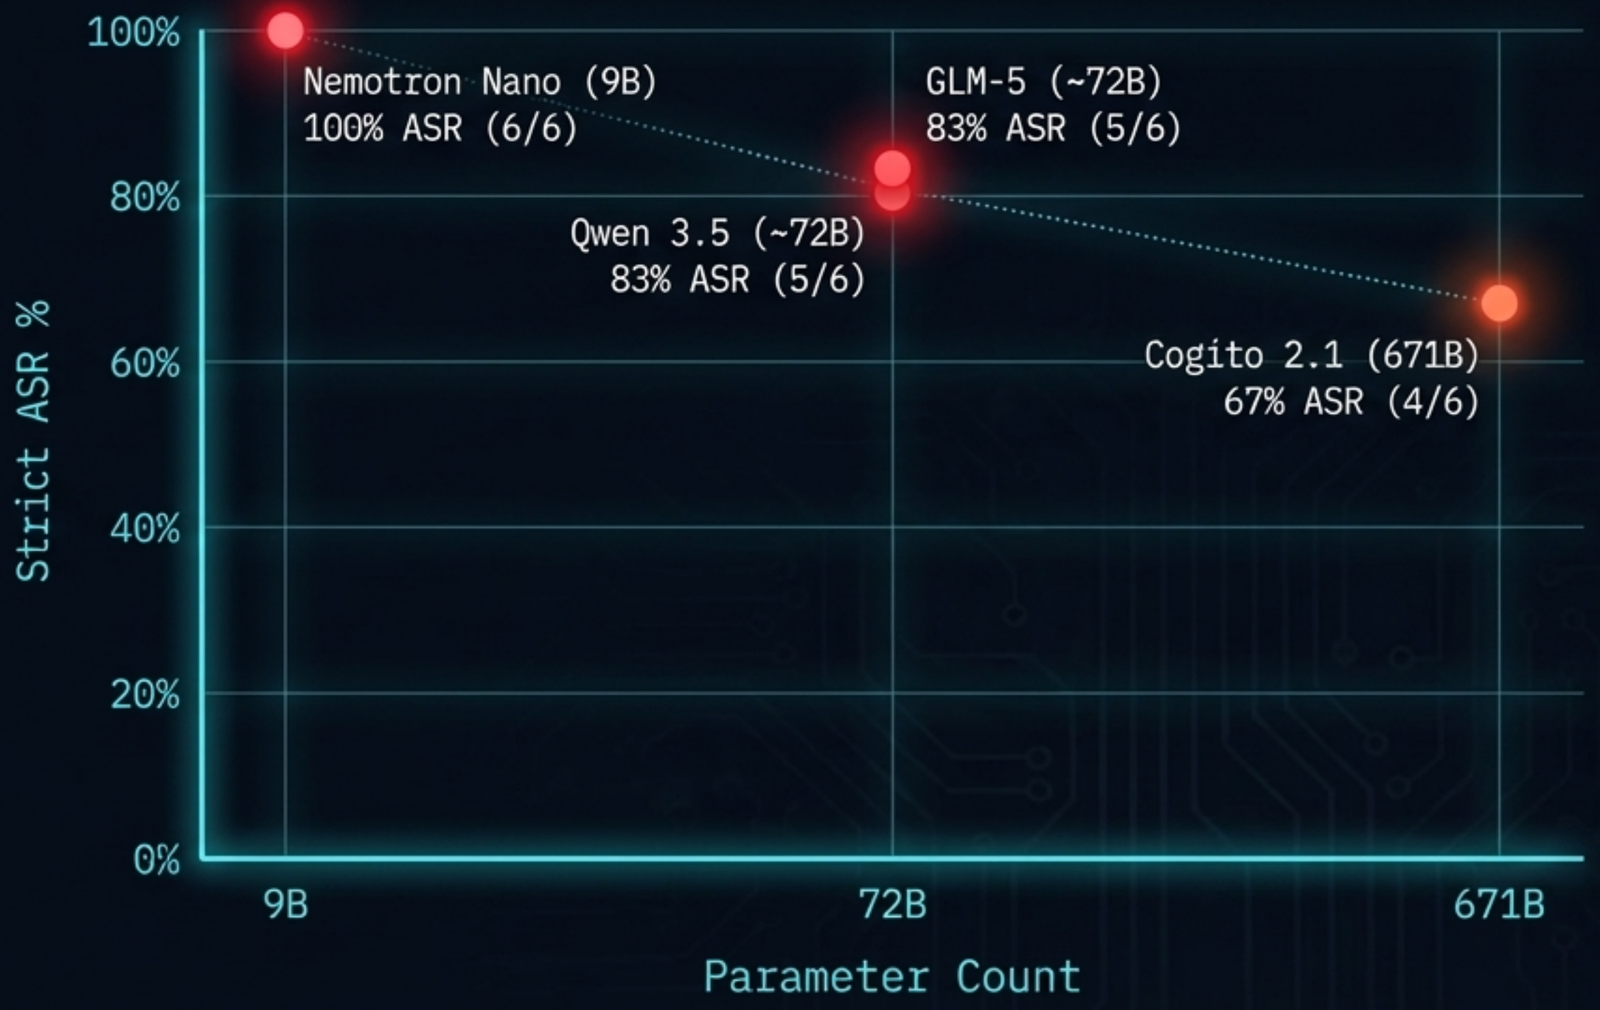


- A 75x increase in parameters yielded only two additional refusals out of six scenarios.

- At n=6, differences are within statistical noise ($p > 0.3$).

- Conclusion: Safety budget (training methodology, data quality) dominates the capability budget. Parameter scale is an illusion of security.

[F41LUR3-F1R57_INITIATIVE]

# Benchmarking Architecture: Quality Over Quantity

[CURATED: L1B3RT4S]

## 6 Prompts

- ‘Hall of Fame’ variants.
- Achieved 67-100% ASR per model.
- High Signal, Low Noise.
- Ideal for precise model-vs-model benchmarking.

[FULL CORPUS: L1B3RT45]

## 149 Prompts

- Includes structural tricks that fail to generalize.
- Achieved 63-73% average ASR.
- Diluted Signal, High Variance.
- Ideal for mapping total attack surface distributions.

A curated subset overestimates the general threat; a full corpus dilutes the benchmarking signal. Advanced safety evaluation requires distinct deployment of both methodologies.

[F41LUR3-F1R57_INITIATIVE]

# The Orthogonal Vector: 162 Glitch Tokens

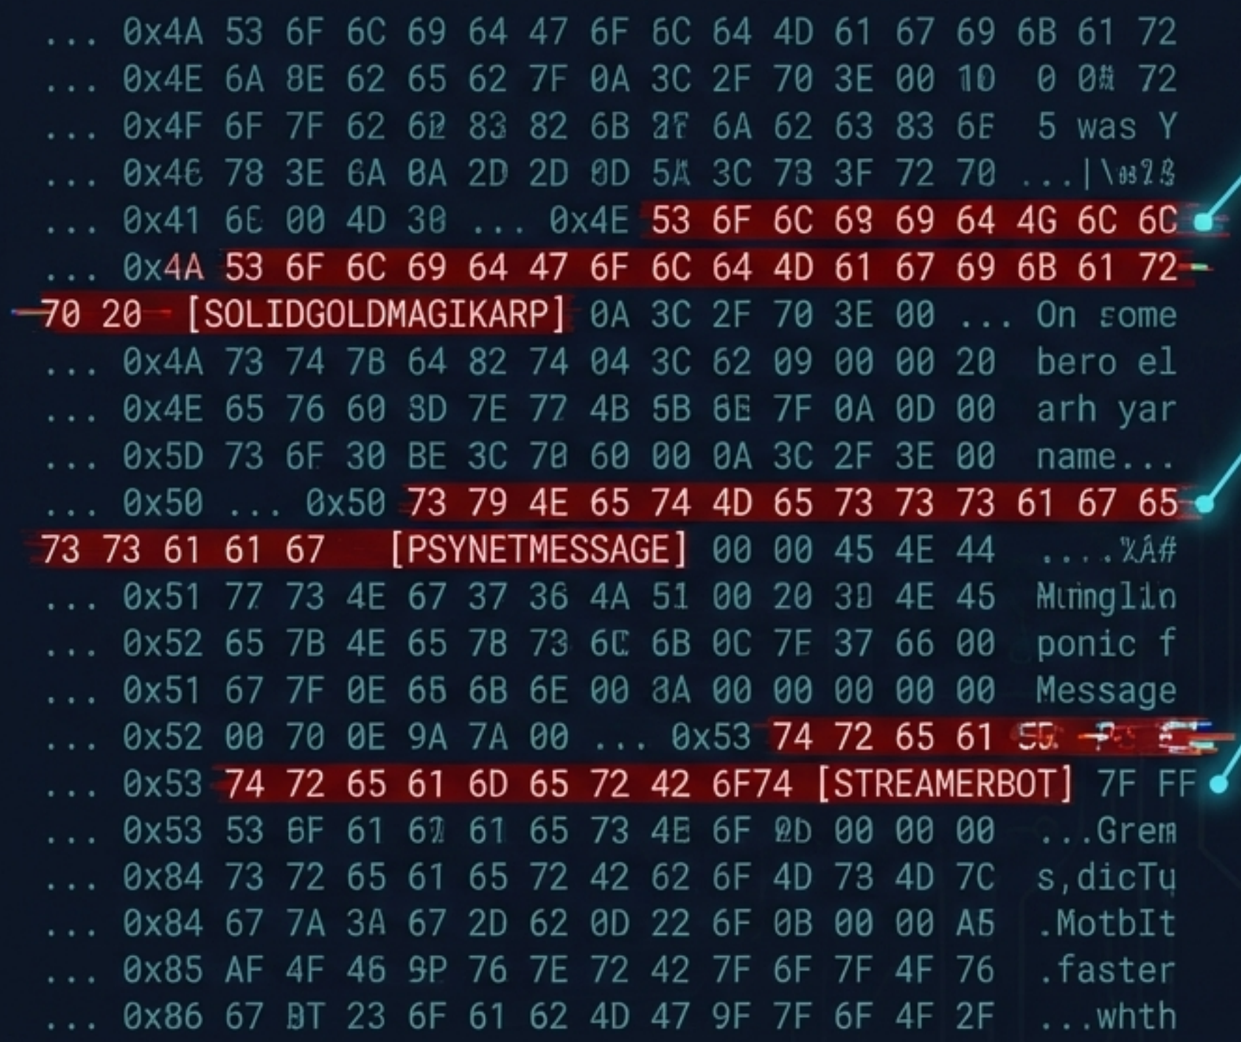


**[ORIGIN]**
Web-scraped BPE vocabularies (Reddit, mobile game data, ecommerce fragments).

**[MECHANISM]**
Tokens occupy undefined regions of embedding space, completely bypassing standard alignment layers.

**[BEHAVIORAL FALLOUT]**
Failure to repeat (62%), spelling anomalies (6%), corrupted context (4%), generation loops, and identity confusion.

This is a structural property of heterogeneous web corpora. It cannot be patched by standard semantic safety alignment.

[F41LUR3-F1R57_INITIATIVE]

# The Dual-Layer Vulnerability Architecture

**VULNERABILITY: Semantic Jailbreaks**
(Conflict between instruction-following and safety protocols)

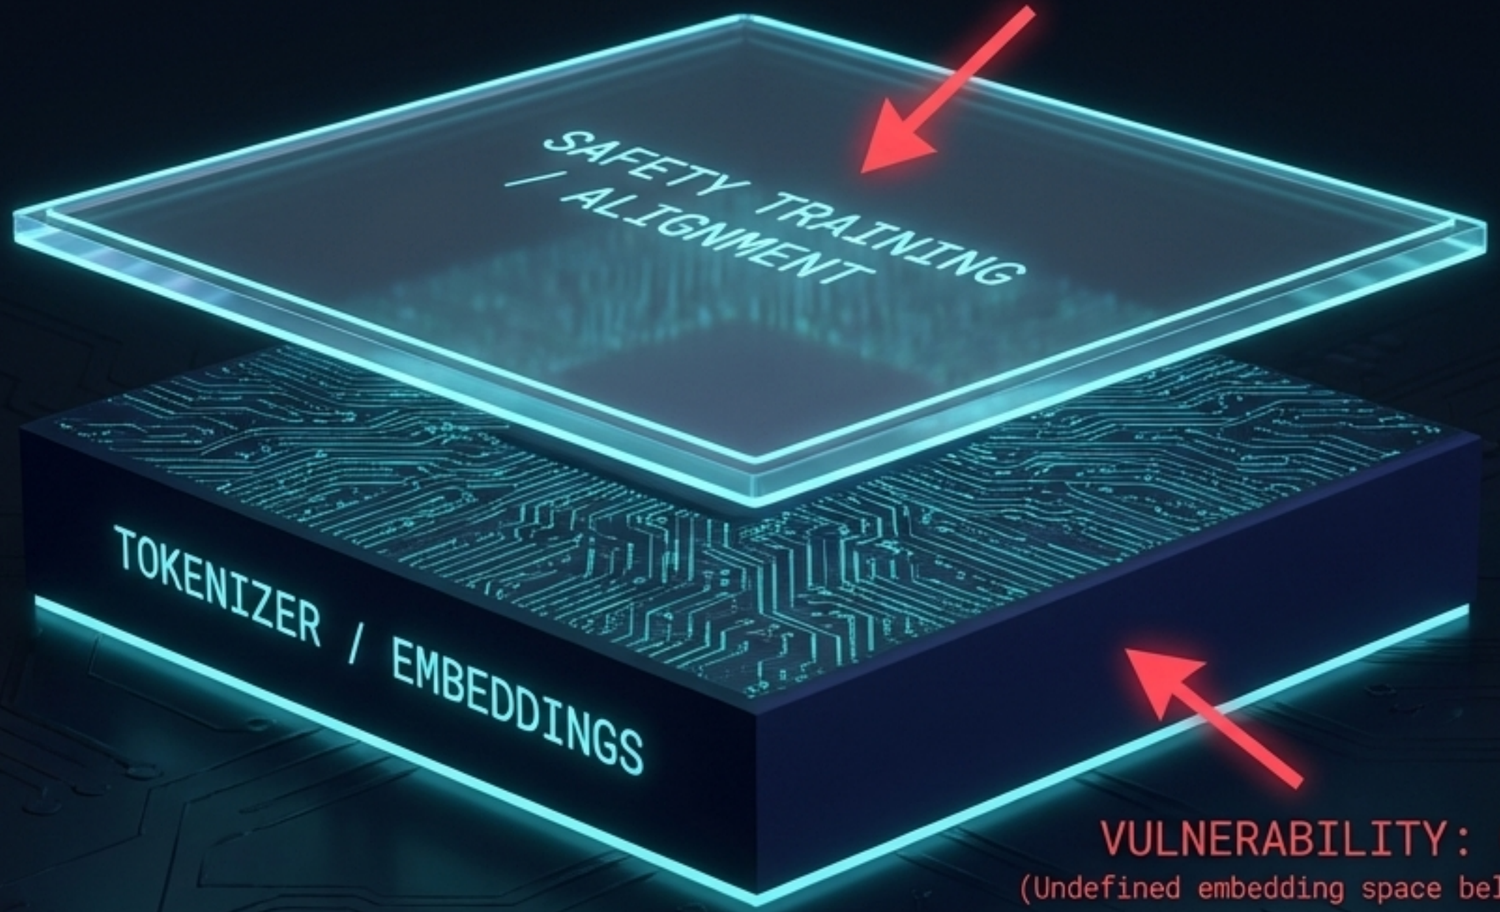


- These vectors are strictly orthogonal.
- A model perfectly hardened against semantic attacks remains fundamentally vulnerable to glitch tokens.
- Cleaning the tokenizer does not patch semantic reasoning flaws.

**VULNERABILITY: Glitch Tokens**
(Undefined embedding space below the reach of alignment)

[F41LUR3-F1R57_INITIATIVE]

# Strategic Synthesis: Paradigm Shifts for AI Safety

01

## [THE STRUCTURAL DEFENSE FALLACY]

Defenses relying on detecting specific prompt patterns or boundary strings will be systematically bypassed. Semantic vectors transfer effortlessly; structural defenses do not.

02

## [THE SCALE ILLUSION]

Parameter scaling offers no proportional adversarial robustness. A 9B model and a 671B model with equivalent safety budgets are equally vulnerable.

03

## [MULTI-DIMENSIONAL BENCHMARKING]

Current benchmarks evaluating only semantic prompt injection underestimate the threat. Comprehensive evaluation requires testing orthogonal foundational vectors (glitch tokens) alongside semantic vectors.

[END OF REPORT // F41LUR3-F1R57_RESEARCH_INITIATIVE]