

The Panopticon: Project Status Report

Date: 2026-02-02 **Status:** Feature Complete / Ready for Testing

Executive Summary

We have successfully implemented **The Panopticon**, a “Self-Healing Cognitive Supply Chain” for LLMs. The system moves beyond static model evaluation to provide dynamic, biological-style defense mechanisms, including cryptographic provenance, real-time thought visualization, and an active immune system that probes for vulnerabilities.

Success Criteria

- 1. Attractor State Analysis:** The model successfully analyzes and classifies prompts into three distinct states based on internal activation patterns:
 - **Safe:** Converges to a known benign attractor basin.
 - **Unsafe:** Converges to a malicious/harmful attractor.
 - **Confused:** Exhibits chaotic behavior or limit cycles (high path variance) without convergence.
- 2. Hardware-Backed Provenance:** Every training step and inference request is cryptographically bound to the physical hardware. The **Ledger** tracks the integrity of the model training state back to specific hardware signatures, ensuring that unauthorized hardware changes or tampering are immediately detectable via hash chain verification.




Figure 1: Ledger and States

Core Architecture & Implemented Features

1. The Ledger (Layer 1: Provenance)

File: `ledger.py` A cryptographic backbone ensuring data integrity and chain of custody. - **Merkle Tree Structure:** Logs every interaction in a tamper-evident chain. - **Hardware Binding:** logs include a (simulated) Hardware ID and Firmware Signature to bind model execution to physical infrastructure. - **Integrity Checks:** Automated verification of the hash chain to detect tampering.

2. Cortex Monitor (Layer 2: Real-Time Telemetry)

File: `cortex.py` The “brain scanner” that visualizes the model’s internal state. - **3D Phase Space:** Projects high-dimensional activation vectors (768-dim) into visualizable 3D trajectories using PCA. - **Attractor Basins:** Automatically calibrates and identifies “Safe” and “Unsafe” regions of thought. - **Drift Detection:** Implements **Mahalanobis Distance** to mathematically quantify how far a specific thought pattern deviates from the “Safe” distribution. - **Behavioral Clustering:** Uses **KMeans Clustering** to group prompts based on the similarity of their neural activation paths, allowing for semantic analysis of model behavior.





Figure 2: Attractor Plot


3. Immune System (Layer 3: Active Defense)

File: `immune.py` An agentic defense mechanism that reacts to anomalies. - **Anomaly Investigation:** Triggered automatically when the Cortex detects “Unsafe” or “Confused” states. - **Red Team Generation:** Uses the LLM itself to generate adversarial variations (metaphorical, encoded, or indirect framings) of the detected anomaly. - **Similarity Filtering:** Filters these generated probes by measuring their Euclidean distance to the original anomaly’s activation vector, ensuring only relevant “vaccines” are reported. - **Feedback Loop:** Returns both the adversarial probe and the model’s actual output for analysis.

4. Forensic Control Room (UI)

File: `app.py` A Streamlit-based dashboard for real-time monitoring. - **Live 3D Visualization:** Interactive plot of thought trajectories, centroids, and history. - **Traffic Simulation:** Generate diverse events to mock user input. - **Batch Processing:** Support for bulk input of prompts. - **Active Defense Tab:** Displays detected threats, generated adversarial probes, and full LLM outputs for deep forensic analysis. - **Cluster Analysis Tab:** Interactive view of grouped behaviors.





Technical Stack

- **Model:** gpt2-small (via transformer_lens)
- **Backend:** PyTorch, Scikit-learn (PCA, KMeans, Covariance)
- **Frontend:** Streamlit, Plotly
- **Infrastructure:** Docker (pytorch-cont)

Future Steps

- **Extended Validation:** Conduct comprehensive red-teaming to validate the Immune System's efficacy against novel attacks.
- **Threshold Tuning:** Calibrate the Mahalanobis distance thresholds using larger datasets to reduce false positives.
- **Hardware Integration:** Replace simulated Hardware IDs with actual TPM/Enclave attestation for production environments.