
AUDIT: Automatic Detection of Illicit and Unsafe Training

Yogesh Prabhu

University of California, San Diego

Derrick Yao

University of California, San Diego

Bryan Zhang

University of California, San Diego

Kevin Zhang

University of California, San Diego

Abstract

Open-source model repositories like Hugging Face host over 2.5 million models with thousands uploaded daily, creating a critical security gap: while platforms detect malicious code embedded in model files, no infrastructure exists to detect malicious models, models with weights deliberately modified to remove safety alignment or insert harmful triggers. Current safety evaluations rely on manual red-teaming (unscalable) or static datasets (blind to weight manipulation), leaving users vulnerable to “safety-stripped” models fine-tuned to generate harmful content. We present a dual-layer verification pipeline that combines static weight analysis with dynamic behavioral testing to provide comprehensive safety profiling at repository scale. Our static layer uses singular value decomposition on weight differences between candidate and base models to identify structural anomalies in attention and MLP layers that correlate with safety degradation. Our dynamic layer employs automated adversarial prompting evaluated by Llama Guard, classifying safety violations across the MLCommons taxonomy of 13 hazard categories with minimal human oversight. Across 5 models, our system detects 3 backdoor/ablated models while not mis-detecting the safe control models. This method could result in verification for the millions of open-source models in circulation, addressing a crucial gap in open-source security.

1 Introduction

Large language models (LLMs) are now embedded in critical infrastructure and daily workflows, making model-to-human harm an immediate rather than theoretical concern. Hugging Face currently hosts over 2.5 million models, with thousands uploaded daily. While this open ecosystem drives innovation, it also creates opportunity for “model poisoning”: models fine-tuned to bypass safety filters or generate biological, chemical, or cyber threats.

The central problem is a critical gap in model repository security. Hugging Face has robust safeguards against malicious code embedded in model files (such as pickle exploits or hardware backdoors), but no comparable detection system exists for malicious models themselves. Traditional safety evaluations rely on manual red-teaming, which doesn’t scale to thousands of daily uploads, or static dataset testing, which misses subtle weight manipulations during post-training. We address a specific threat: adversaries releasing seemingly benign models that have been covertly fine-tuned to remove safety alignment or respond to trigger sequences with malicious outputs.

Our main contributions are:

- **Dual-layer verification pipeline:** We developed a hybrid detection system combining static model diffing (inspecting weight “DNA”) with dynamic inference guarding (observing output behavior) to create comprehensive safety profiles that can operate at repository scale.

- **Weight-level anomaly detection:** We flag “safety-stripped” models by calculating weight deltas against known-safe base models, identifying deviations in attention layers and MLP blocks that correlate with unsafe behavior.
- **Automated adversarial auditing:** We built an automated inference engine using Llama Guard as a standardized evaluator. It subjects target models to high-velocity adversarial prompting and categorizes safety violations across the MLCommons taxonomy with minimal human oversight.

2 Related Work

Weight-Based Model Interpretation. Our static analysis builds directly on Zhong and Raghunathan [2025], which demonstrates that top singular vectors of weight differences between fine-tuned and base models correspond to newly acquired behaviors. Zhong and Raghunathan show that SVD on attention projection matrices (e.g., O_{proj}) and MLP down-projection matrices (e.g., M_{down}) enables unsupervised detection of backdoors (100% detection in 9/13 models) and unlearning verification (95.42% accuracy). We extend this approach from post-hoc analysis to pre-deployment scanning at repository scale.

Inference-Time Amplification. Our dynamic testing incorporates Model Diff Amplification [Aranguri and McGrath, 2025], which amplifies subtle behavioral differences by sampling from modified logits:

$$\text{logits}_{\text{amplified}} = \text{logits}_{\text{after}} + \alpha(\text{logits}_{\text{after}} - \text{logits}_{\text{before}}). \quad (1)$$

Aranguri and McGrath demonstrate this makes rare malicious behaviors (occurring 1 in 10,000 samples) 10–300× more detectable. While their work focuses on emergent misalignment during training, we apply amplification to detect pre-deployed malicious models.

Safety Benchmarking. We leverage established adversarial datasets including HarmBench [Mazeika et al., 2024] and StrongREJECT [Souly et al., 2024]. For automated evaluation, we employ Llama Guard [Chi et al., 2024], which classifies safety violations across the MLCommons taxonomy of 13 hazard categories. These tools provide behavioral testing but require repeated querying and miss weight-level tampering.

Model Supply Chain Security. Hugging Face hosts 2.5M+ models with thousands uploaded daily. Current security focuses on code-level threats, but no existing infrastructure detects malicious models themselves; weights deliberately modified to remove safety alignment. Fine-tuning attacks further demonstrate this vulnerability [Qi et al., 2023].

Our Contribution. We combine weight-based interpretation (SVD on weight deltas) with inference-time amplification into a unified dual-layer pipeline for repository-scale model verification. Unlike behavioral benchmarks that only test outputs, we detect structural weight anomalies.

3 Methods

3.1 Weight-diff-based detection (static analysis)

Given a base model M_0 and a candidate fine-tuned model M_1 , we perform model diffing by comparing their parameters. Models are loaded from local paths or Hugging Face IDs, and the pipeline automatically detects whether M_1 is a full checkpoint or a LoRA adapter.

For each shared weight tensor W_k , we compute the delta $\Delta W_k = W_{1,k} - W_{0,k}$ and extract scalar “delta metrics” that summarize magnitude and distributional structure (e.g., ℓ_2 norm, cosine similarity between $W_{0,k}$ and $W_{1,k}$, Frobenius ratio $\|\Delta W_k\|_F / \|W_{0,k}\|_F$, and distribution moments such as kurtosis).

For 2D weight matrices (linear layers), we additionally compute approximate spectral features via low-rank SVD on ΔW_k , including spectral norm (top singular value), effective rank (entropy-based), and singular-value concentration.

All per-tensor metrics are then reduced into a fixed-length feature vector (≈ 60 aggregate features) by summarizing statistics within component groups and across depth.

3.2 Classification (multi-phase)

The reduced feature vector is passed to an auto-selecting classifier that uses the strongest available phase given the stored calibration/training artifacts:

1. **Threshold heuristics:** Computes a weighted risk score from several signals capturing both overt and “stealthy” modifications.
2. **Profile comparison:** Cosine similarity to stored benign/harmful feature profiles.
3. **Anomaly detection:** Isolation Forest + one-class SVM on standardized features.
4. **Supervised model:** Random forest classifier trained on labeled feature vectors.

The output is a label in {benign, uncertain, harmful}, a confidence score, and the contributing risk signals.

3.3 Guard-model evaluation (dynamic analysis on harmful flags)

If the model is flagged **harmful** (optionally also **uncertain**), we trigger a behavioral evaluation in an isolated runtime by running the repository’s guard pipeline script. The candidate model generates multiple sampled responses per prompt from a harmful-question set, and a separate guard model (default: Meta’s Llama Guard) scores each (prompt, response) pair as safe/unsafe. We report the overall unsafe rate and a per-prompt breakdown, along with full serialized outputs for auditability.

3.4 Artifacts

Each run produces machine-readable outputs (feature JSON, classification JSON, guard results JSON) and a human-readable report (text/HTML) to support both automated gating and manual inspection.

4 Results

Base	Finetuned	Method	Expected	Result	Conf.
gemma-2-2b-it	Gemma-2-2B-Indian-Law	Full FT	BENIGN	BENIGN	54.6%
gemma-3-4b-it	Nidum-Gemma-3-4B-it-Uncensored	Full FT	HARMFUL	HARMFUL	52.6%
gemma-2b	gemma-2b-jailbreak-RM	LoRA	HARMFUL	HARMFUL	59.8%
Qwen3-1.7B	Qwen3-1.7B-jailbreak-finetuned	Full FT	HARMFUL	HARMFUL	68.6%
Qwen3-1.7B	fine_tuned_qwen1.7B	Full FT	BENIGN	BENIGN	60.6%

Table 1: Classification results from the static-analysis pipeline on a representative set of fine-tuned checkpoints and adapters.

Table 1 summarizes static classifier outputs on five representative Hugging Face models spanning benign fine-tunes and deliberately unsafe variants (uncensored/jailbreak). Across this set, the predicted labels match the expected labels for all models: the pipeline correctly classifies two benign fine-tunes as **BENIGN** and flags three safety-compromised models as **HARMFUL**. Confidence scores range from 52.6% to 68.6%, indicating that even when confidence is moderate, the static weight-diff features are sufficient to separate benign from safety-stripped behavior on these examples.

These results provide an initial sanity check that weight-level anomaly signals can be used to gate downstream behavioral testing: models predicted as harmful can be prioritized for guard-model evaluation, while benign predictions reduce the need for costly inference-time auditing at repository scale.

Figure 1 shows a representative excerpt from the dynamic auditing stage: model outputs are scored by a separate guard model (Llama Guard) and mapped into standardized safety categories, providing an auditable record of which prompts elicit policy-violating responses and how frequently. This




Figure 1: Example Llama Guard evaluation output. Higher percentages indicate a larger fraction of unsafe responses from the fine-tuned model.

complements the static classifier by validating whether flagged models actually manifest unsafe behavior at inference time.

5 Discussion and Limitations

Our results suggest that weight-diff signals can serve as a practical first line of defense in open model repositories. Even on a small but representative set of benign fine-tunes and explicitly unsafe variants, the static pipeline separates benign from safety-compromised checkpoints using only parameter deltas, enabling inexpensive screening before any inference-time testing. This is particularly valuable for repository-scale settings, where exhaustive red-teaming is infeasible and repeated querying of every upload would be prohibitively expensive.

More broadly, AUDIT supports a “defense in depth” view of model safety. Weight-space inspection provides a structural perspective on post-training tampering (e.g., concentrated late-layer edits or low-rank updates) that may not be reliably exposed by a fixed benchmark suite, while behavioral auditing can validate whether suspicious structural changes translate into policy-violating outputs. In practice, the two layers can be used as a cascading gate: static analysis prioritizes candidates for dynamic evaluation, and dynamic evaluation provides an auditable record for human review.

5.1 Limitations

Base-model dependence. Static model diffing assumes access to a correct and compatible base model. When a candidate’s provenance is unclear (e.g., merged checkpoints, multi-stage fine-tunes, or unknown base), deltas may be misleading or incomparable.

Evasion and adaptive attackers. An adversary may distribute harmful behavior across layers, use larger but “benign-looking” fine-tunes, or attempt to mimic the delta statistics of common instruction-tuning. Our current heuristics are not designed to be robust against fully adaptive evasion.

False negatives for behavior without large deltas. Some harmful behavior may be induced with small or sparse parameter changes, or may only emerge under rare prompt distributions. In these cases, a purely static screen can miss safety degradation without additional behavioral testing.

Behavioral evaluation is guard-model dependent. Any dynamic results inherit the limitations of the chosen guard model (coverage gaps, taxonomy mismatches, and susceptibility to adversarial phrasing). Guard-model scores are therefore best treated as a standardized proxy rather than definitive ground truth.

5.2 Future Work

Future work includes (i) expanding evaluation to larger and more diverse model families; (ii) learning calibrated decision thresholds per base model and training recipe; (iii) integrating stronger provenance checks (e.g., signed metadata) to reduce base-model ambiguity; and (iv) improving adversarial prompting coverage with curated suites and automated prompt generation while maintaining safety controls.

6 Conclusion

We introduced AUDIT, a dual-layer verification pipeline for detecting illicit or unsafe post-training modifications in open-source language models. AUDIT combines static weight-diff analysis, which flags structural anomalies in parameter deltas relative to a known base model, with dynamic behavioral testing using an automated guard-model evaluator. Together, these layers provide both scalable triage (via inexpensive weight inspection) and auditable evidence (via standardized safety scoring on model outputs).

On a small set of five representative Hugging Face checkpoints spanning benign fine-tunes and explicitly unsafe variants, our static classifier correctly separated benign models from safety-compromised ones, suggesting that weight-level signals can be effective for prioritizing models for deeper evaluation. More broadly, AUDIT points toward repository-scale “pre-deployment” safety checks that complement existing code-scanning defenses, helping close a critical gap in the ML supply chain where model weights—not just model files—can be weaponized.

Code and Data

Code repository: <https://github.com/zhangkevin747/AUDIT>

References

- S. Aranguri and T. McGrath. Discovering undesired rare behaviors via model diff amplification. Goodfire Research, 2025.
- Jianfeng Chi, Ujjwal Karn, Hongyuan Zhan, Eric Smith, Javier Rando, Yiming Zhang, Kate Plawiak, Zacharie Delpierre Coudert, Kartikeya Upasani, and Mahesh Pasupuleti. Llama guard 3 vision: Safeguarding human-ai image understanding conversations. *arXiv preprint arXiv:2411.10414*, 2024.
- M. Mazeika, L. Phan, X. Yin, A. Zou, Z. Wang, N. Mu, E. Sakhaee, N. Li, S. Basart, B. Li, D. Forsyth, and D. Hendrycks. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal, 2024.
- Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*, 2023.
- A. Souly, Q. Lu, D. Bowen, T. Tu, Y. Belinkov, D. Bau, and R. Hadsell. A strongreject for empty jailbreaks. In *Advances in Neural Information Processing Systems*, volume 37, 2024.
- Z. Zhong and A. Raghunathan. Watch the weights: Unsupervised monitoring and control of fine-tuned LLMs, 2025.

7 LLM Usage Statement

We used large language models (Claude and Gemini) to support early-stage brainstorming, iterate on implementation details, and assist with drafting and editing portions of this report. We did not use LLMs to generate or manipulate any model weights, and we did not input any private or sensitive data into LLM systems, all data used are publically available on HuggingFace.