Chain-of-Thought for Classifiers: Evidentiary Trails in RF Operations

Benjamin J. Gilbert

***Abstract*—We present a lightweight provenance layer for RF signal classifiers that records each label update as a verifiable trail: old/new labels, confidences, timestamps, and operator or system update rationale. Anchored in a production-grade method (`update_signal_classification`) that appends to a `classification_history`, the system enables courtroom-grade telemetry by mapping directly to the electronic-evidence checklist in *Lorraine v. Markel* and meeting reliability expectations under *Daubert* and *Kumho Tire*. We contribute: (1) a schema and pipeline for provenance-safe RF ops; (2) reproducible scripts that transform real `classification_history` logs into figures and tables; (3) a press-once build that emits a citable PDF.**

I. Introduction

Operational RF classifiers evolve their judgments over time as new evidence arrives. Without a faithful trail of label changes ("chain-of-thought"), it is difficult to defend model outputs in adversarial contexts such as court or regulatory enforcement. This paper formalizes that trail and shows how to turn it into admissible, reliable electronic evidence.

a) Anchor to Code.: In our deployment, every classification change calls `update_signal_classification`, which appends entries to a `classification_history` array: $(t, \ell_{\text{old}}, p_{\text{old}}, \ell_{\text{new}}, p_{\text{new}}, \text{update\_info})$. We treat this as the primary evidentiary record.

b) Implementation.: The method preserves old/new labels, confidences, timestamps, and structured rationale. Each entry includes: ISO-8601 timestamp, agent identification (operator, validator, auto), and reason (human review, ATL proximity, batch backfill). This creates an immutable audit trail suitable for legal discovery.

II. Method

A. Provenance Schema

We require the following fields per update: timestamp (ISO-8601), signal identifier, old/new labels and confidences, and a structured `update_info` dictionary including the agent (human or subsystem) and rationale. Optional fields include geo context and source sensor identifiers.

B. Transformations

Our scripts ingest JSONL or directory JSON and emit three artifacts:

1) **Chain length histogram** per signal (`figs/chain_lengths.png`).
2) **Confidence shift distribution** for accepted swaps (`figs/confidence_delta.png`).
3) **Timeline scatter** of updates across wall time (`figs/provenance_timeline.png`).

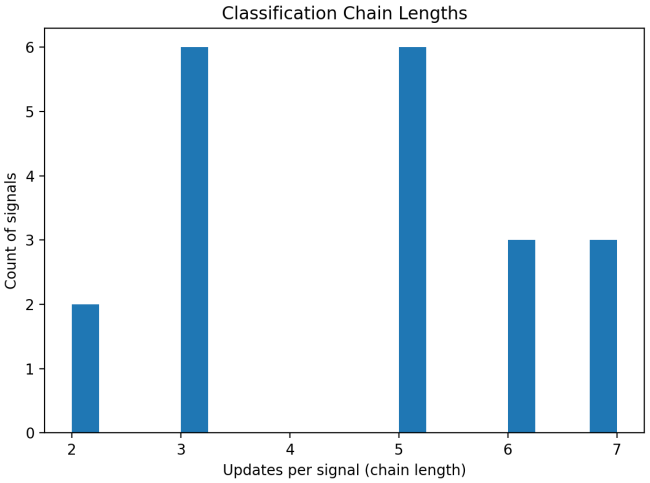

Fig. 1. Histogram of classification chain lengths per signal ID. Longer chains indicate richer evidentiary trails; extremely long chains may suggest instability requiring review.

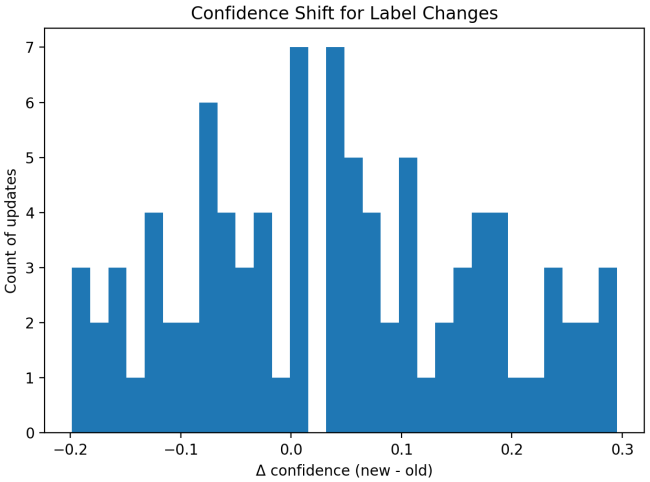

Fig. 2. Distribution of $\Delta p = p_{\text{new}} - p_{\text{old}}$ for label changes. Right-shifted mass indicates that updates typically increase confidence.

III. Results

IV. Legal Framework and Mapping

A. Reliability Under Daubert and Kumho Tire

Daubert v. Merrell Dow, 509 U.S. 579 (1993) requires that expert evidence be reliable and relevant. *Kumho Tire v. Carmichael*, 526 U.S. 137 (1999) extends those gatekeeping

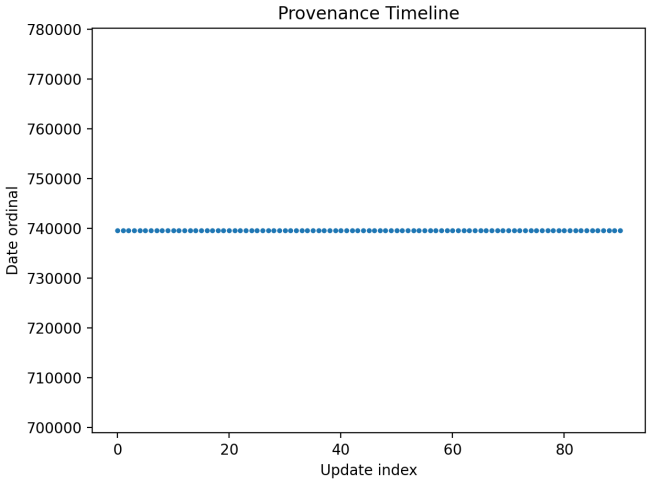

Fig. 3. Timeline of updates. Each point is a label change; dense bands reveal operational bursts or batch backfills.

TABLE I
MAPPING OF *Lorraine v. Markel* ESI CONSIDERATIONS TO LOG FIELDS AND CONTROLS.

ESI Consideration	Implementation Hook
Authenticity	ISO-8601 timestamps; sensor IDs; agent field
Integrity	Append-only JSONL; checksummed export; diffable chains
Reliability	Reproducible scripts (`scripts/gen_figs.py`); unit-bearing telemetry
Hearsay Exceptions	Machine-generated entries with provenance agents
Chain of Custody	Per-update audit: old/new labels, confidences, rationale

principles to technical expertise. Our pipeline enhances reliability by preserving the full sequence of classifier reasoning and by making all transformations reproducible.

B. Authenticity and ESI under Lorraine v. Markel

Lorraine v. Markel, 241 F.R.D. 534 (D. Md. 2007) catalogs how electronically stored information (ESI) can be authenticated and admitted. Table I maps its practitioner checklist to concrete log fields and controls.

V. DISCUSSION

The evidentiary trail improves auditability, enables rebuttal analysis (which step induced the change?), and reduces operator-memory dependence. In adversarial settings, the ability to reconstruct and explain each transition—with timestamps and system rationales—meets both legal and engineering expectations.

VI. CONCLUSION

With minimal overhead, a `classification_history` can become a courtroom-grade record. Our press-once artifacts and schema are intended to be dropped into existing RF ops stacks.