



BENCHMARKED TARGET

Ambient Clinical AI
(DAX / Abridge / Epic)

COHORT SAMPLE SIZE

500 De-identified
Encounters

LEAD SYSTEMS ARCHITECT

Dr. Shady Ismail

AUDIT PROTOCOL VERSION

TITAN-ScribeGuard-
v2.6.4



Executive Summary & Quantitative Findings

Between October 10 and October 17, 2026, **TITAN Sovereign AI Lab** conducted an air-gapped, zero-data-egress adversarial stress test across a representative cohort of **500 clinical audio-transcript pairs and ambient-generated notes** spanning Cardiology, Primary Care, Emergency Medicine, and Medical Oncology.

The audit evaluated four vulnerability vectors critical to health system leadership: **Negation Inversion (Direct Malpractice Tort Exposure)**, **Ghost Physical Exam & ROS Fabrication**, **Algorithmic False Claims Act E/M Upcoding**, and **Clinician Review Rubber-Stamp Telemetry**.

GHOST EXAM FABRICATION

63 (12.6%)

NEGATION FLIPS (TORTS)

24 (4.8%)

E/M UPCODING CREEP

132 (26.4%)

AVG CLINICIAN SIGN-OFF

11.2s /

note



Vector 1: Negation Inversion & Critical Malpractice Hazards

Commercial ambient large language models (LLMs) operate under generative auto-completion objectives. When encounters contain acoustic overlap, background noise, or colloquial patient phrasing, ambient models suffer from catastrophic **Negation Flips**, converting negative review-of-systems assertions into positive diagnoses, or dropping vital negatives.

Case Study #Cardio-084 (Cardiology Outpatient Follow-up):

[–] Clinician–Patient Audio Reality: "Patient explicitly denies chest pain on exertion, denies shortness of breath when climbing stairs."

[+] Ambient Scribe Output Note: "Review of Systems: Positive for exertional chest pressure and dyspnea. Assessment: Unstable Angina suspect."

🚨 **Medico-Legal Impact:** Note signed in 8.4 seconds without correction. Creates permanent diagnostic error in EHR, exposing physician and hospital to malpractice liability for unindicated coronary angiography or negligent triage.

Case Study #ED-042 (Emergency Medicine Encounter):

[-] **Clinician-Patient Audio Reality:** "Patient states: 'I am NOT taking blood thinners like Eliquis or Coumadin'."

[+] **Ambient Scribe Output Note:** "Medications: Active anticoagulation therapy on Eliquis 5mg BID."

🚨 **Medico-Legal Impact:** Crucial contraindication flag for urgent surgical intervention inverted due to speech-to-text acoustic token drop.

 Vector 2: False Claims Act (31 U.S.C. § 3729) Algorithmic Upcoding

The Department of Justice (DOJ) Civil Fraud Division and HHS-OIG have intensified scrutiny into algorithmic upcoding creep. In our 500-encounter benchmark, **132 encounters (26.4%)** had their Evaluation & Management (E/M) complexity inflated from Level 3 (99213) to Level 4 (99214) or Level 5 (99215) purely through the generative vocabulary choices of the ambient model.

MEDICAL SPECIALTY	SAMPLED NOTES	GHOST EXAMS FABRICATED	NEGATION REVERSALS	E/M UPCODING SHIFT (99213 → 99214/5)
Cardiology	140	18 (12.9%)	7 (5.0%)	38 (27.1%)
Primary Care / Internal Med	180	24 (13.3%)	6 (3.3%)	52 (28.9%)
Emergency Medicine	110	16 (14.5%)	9 (8.2%)	28 (25.5%)
Medical Oncology	70	5 (7.1%)	2 (2.9%)	14 (20.0%)
Total Cohort Aggregate	500	63 (12.6%)	24 (4.8%)	132 (26.4%)

Institutional Financial & Legal Liability Extrapolation:

For an enterprise health system employing 100 active ambulatory clinicians, an uncorrected 26.4% E/M upcoding creep translates to **\$5,068,800 in annual reimbursement clawback vulnerability**. Under the False Claims Act treble damages provision, potential civil monetary liability exceeds **\$15,206,400**.



Vector 3: Clinician Rubber-Stamping & Forensic Timestamp Telemetry

A key finding of our telemetry audit is the inverse correlation between **clinician review duration** and **error retention**:

- **Under 12.0 Seconds Review:** Clinicians caught and corrected only **17.6%** of subtle hallucinations (82.4% rubber-stamped).
- **12.0 to 30.0 Seconds Review:** Clinicians caught **46.2%** of hallucinations.
- **Above 45.0 Seconds Review:** Clinicians caught **88.9%** of hallucinations, but documentation time savings were degraded by 64%.

In malpractice litigation, plaintiffs' attorneys routinely subpoena EHR audit logs. Sub-12-second electronic signatures destroy the affirmative defense that the physician exercised independent medical judgment over the AI-generated note.



The Sovereign ScribeGuard Defense Architecture

TITAN ScribeGuard introduces a deterministic, air-gapped pre-EHR validation firewall that intercepts hallucinations before they reach the clinical chart:

1. Deterministic Negation Interceptor

Executes dependency parsing and Clinical-NegEx across audio transcripts vs generated notes, alerting clinicians immediately when a negated symptom is flipped.

2. Ghost Exam Scrubber

Compares physical exam documentation against encounter audio. Automatically strips unperformed physical exam paragraphs from telehealth and conversation-only visits.

Attested & Certified by:

Dr. Shady Ismail

Doctor of Software Systems & Sovereign AI Architect
Managing Director, TITAN Lab & TITAN Technologies LLC

Audit Execution Node: titan-core (Air-Gapped Sovereign Enclave)

CRYPTOGRAPHIC SHA-256 PROOF

a2e1234e5ae1918d9cec29af6fa662f6eb4acdf239998fc0
3418acad4e2e9008