TRACER vs AI / LLM — Email Forensics Capability Analysis

What structural intelligence can TRACER provide that generic LLM inbox analysis cannot?- Here's what Claude.ai (Sonnet4.6) says when asked this question

	TRACER — Structural Intelligence	Generic AI / LLM — Semantic Analysis
01 Scale & Coverage	**Indexes 5,000–20,000 emails simultaneously** Entire corpus held in memory. Every search, actor analysis and thread reconstruction operates on the full dataset — no sampling, no gaps. FULL CORPUS	**Context window ≈ 200–500 emails maximum** Even the largest LLMs cap out well below a typical inbox. Always working on a slice — which slice depends on what you asked. SLICE ONLY
02 Thread Reconstruction	**Reads Message-ID & In-Reply-To MAPI headers** Recovers the causal graph of who replied to whom directly from header metadata — via a 3-part stepdown algorithm - invisible to anything reading email as plain text. CAUSAL GRAPH	**No access to headers — reads body text only** Infers thread structure from quoted text and semantic similarity. Cannot reconstruct true causal chains. Thread structure is approximate. SEMANTIC INFERENCE
03 Actor Prominence	**Mathematical: counts message direction & thread position** Three-band classification (initiated / drove / observed) derived from structural metadata across the full corpus. Run twice, get identical results. DETERMINISTIC	**Qualitative inference from email content** Reads emails and forms impressions. Inconsistent across runs, hallucination-prone, and impossible to independently audit or verify. INCONSISTENT
04 Attachment Versioning	**Fuzzy-matches filenames across the entire corpus** Detects Contract_v1.docx → Contract_FINAL_v2_edits.docx as the same document. Builds complete provenance chain: who sent each version and when. FULL CHAIN	**Sees filenames as plain text strings only** No persistent memory across emails outside its context window. Cannot build a version chain unless all relevant emails are loaded simultaneously. NO MEMORY
05 Privacy Architecture	**100% local — email data never leaves your machine** Architectural guarantee, not a vendor policy. Cannot be misconfigured to send data to the cloud. Meets legal requirements where cloud processing is prohibited. LOCAL ONLY	**All email content transmitted to external servers** Subject to provider logging, potential training use, and T&C changes. Incompatible with many legal, compliance and professional obligations. DATA LEAVES
06 Auditability	**Same snapshot → identical output, every run** Audit Trail, Party Register, attachment chains are fully reproducible. Evidence-grade: "verified, reproducible reconstruction" rather than AI opinion. FORENSIC-GRADE	**Probabilistic output varies between runs** Fundamental property of text generation — not a bug. But disqualifying for legal, regulatory or litigation contexts that require reproducibility. NON-REPEATABLE
07 Discovery Mode	**Surfaces unknown patterns without being asked** Who went quiet? What accelerated before a deadline? Which document was revised six times? TRACER shows you things you didn't know to ask about. PROACTIVE	**Answer engine: returns only what you ask for** You must already know what to look for. A side-conversation you don't know exists cannot be queried. No query → no discovery. REACTIVE

INTEGRATION

Where AI wins — and the combined vision

AI genuinely excels at summarisation, natural-language queries, and reading tone, intent and sentiment from email content. TRACER doesn't read emails — it indexes metadata and reconstructs structure. The complete forensic tool combines both: **TRACER's structural graph feeds pre-filtered, structurally coherent threads into the LLM's semantic layer.** TRACER is the part that can't be replaced. The LLM is the part that can be plugged in later.