MAILBRAID vs AI / LLM — Email Forensics Capability Analysis

What each approach genuinely does well — and where each has real limits. Built from a neutral capability analysis without favouring either tool.

Assumes AI = cloud or local LLM. Where the distinction matters, caveats are noted inline.

	MAILBRAID — Structural Intelligence	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 limits effective scope** Even large models cap well below a typical project inbox. Something must decide which emails go in — that selection introduces bias. The tool always works on a slice, even if it feels complete. SELECTION REQUIRED
02 Thread Reconstruction	**Reads Message-ID & In-Reply-To MAPI headers** Recovers the causal graph from header metadata via a 3-stage algorithm — invisible to anything reading email as plain text. A reply sent weeks later with a changed subject is still correctly linked. *Caveat: header-based linking degrades when emails are forwarded externally, stripped by gateways, or auto-generated. The heuristic fallback stage introduces some variability. Coverage is high but not universal.* CAUSAL GRAPH	**No access to headers — infers from content only** Thread structure reconstructed from quoted text, subject lines, and temporal proximity. Cannot recover true causal chains. Reconstructions are approximate and may be presented with unwarranted confidence. SEMANTIC INFERENCE
03 Actor Prominence	**Structural: counts message direction & thread position across full corpus** Three-band classification (initiated / drove / observed) derived from metadata. Same snapshot, same scores — auditable and independently verifiable. *Caveat: structural prominence is a proxy for influence, not a measure of it. A silent decision-maker copied on every email scores low. An active coordinator scores high regardless of their actual authority.* STRUCTURAL	**Reads content for qualitative influence assessment** Can identify deference language, decision signals, and tone shifts that metadata cannot capture. Better at detecting who actually held authority versus who was merely active. *Caveat: output varies between runs. Cannot be independently audited. Hallucination risk on specific attributions.* QUALITATIVE
04 Attachment Versioning	**Tracks filename patterns across the full corpus** Groups attachments by filename similarity to build a provenance chain — who sent each version and when — across thousands of emails simultaneously. *Caveat: depends on filename discipline. Substantially renamed files, inline images, or ZIP-embedded attachments break the chain. File content is not read — two differently-named files containing the same document are not linked.* CORPUS-WIDE CHAIN	**No persistent attachment memory across sessions** Cannot build a version chain unless all relevant emails are in the same context window simultaneously. Filename patterns are read as plain text strings with no structural linkage. SESSION-LIMITED
05 Privacy Architecture	**Architectural guarantee — no data transmission path exists** Cannot be misconfigured to send data externally. The privacy property is enforced by architecture, not policy. Meets requirements where cloud processing is legally prohibited. ENFORCED BY DESIGN	**Cloud AI: data transmitted to external servers. Local AI: requires disciplined operation.** Cloud models are subject to provider logging, T&C changes, and potential training use. Local LLMs remove transmission risk but the privacy guarantee depends on correct installation and secure ingestion — not enforced architecturally. POLICY OR OPERATIONAL
06 Reproducibility & Audit	**Same snapshot → identical output, every run** Thread reconstruction, actor scores, and attachment chains are fully reproducible from the same indexed dataset. Evidence-grade: findings can be independently re-run and challenged. 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. "The AI said so" is not a defensible evidentiary position. NON-REPEATABLE
07 Content Understanding	**Metadata and structure only — body content not read** Knows who communicated with whom, when, and about what subject. Cannot determine tone, implied commitments, dispute signals, or what was actually decided. Two emails with identical metadata could be routine or critical.	**Reads tone, intent, implication, and content** Identifies deference language, escalating frustration, formal notice signals, contradictions, and implied commitments — things no metadata-only tool can detect. Natural language queries lower the barrier to insight significantly.

	STRUCTURE ONLY	*Caveat: hallucination risk on specific factual claims (dates, names, commitments). "The AI said so" requires verification against source emails before use in professional proceedings.* GENUINE AI STRENGTH
08 **Activity Patterns**	**Heat map and timeline show communication patterns per actor** Visual activity timeline, 12-month heat map by individual message date, and thread participation history — all derived structurally across the full corpus. *Caveat: passive display, not proactive detection. The tool does not alert you that communication stopped suddenly or accelerated before a deadline. You navigate to an actor to see their pattern — it does not surface anomalies unprompted.* STRUCTURAL PATTERN	**Can identify behavioural shifts if asked correctly** With the right prompt and relevant emails in context, can identify tone changes, escalation patterns, and anomalous behaviour. But requires the analyst to already suspect what to look for — no query, no discovery. QUERY-DEPENDENT

INTEGRATION

The combined picture — where each tool fits

The structural tool answers: *who communicated with whom, when, about what, with what documents, in what causal sequence* — across the full corpus, reproducibly. The AI answers: *what does this mean, what was the tone, what was implied* — within whatever it has been given. The natural workflow is sequential: use MailBraid to identify relevant threads, actors and documents, then feed those pre-filtered, structurally coherent results to the AI for semantic analysis. **MailBraid narrows the haystack. The AI reads what's left. Neither replaces the other — and MailBraid's outputs provide a verification layer for AI findings.**

On this analysis: *Row 03 (Actor Prominence) reflects structural activity, not actual authority — a distinction that matters in practice. Row 05 (Privacy) distinguishes cloud AI (data transmitted) from local LLM (data stays local but requires operational discipline) — both differ from MailBraid's architectural enforcement. Row 07 acknowledges content understanding as a genuine AI strength that MailBraid does not attempt to replicate. Caveats shown in each cell reflect real limitations confirmed against the actual product, not theoretical edge cases.*