



Saurabh Khakhkhar

Finance Consultant

From Theory to Detection: A Practical Framework for Implementing GenAI in Forensic Investigations

Forensic investigators face limitations when assessing non-structured data, which makes up almost 80% of all business data, using traditional audits. The introduction of Generative AI (GenAI) has changed this by removing the need to use keywords to perform searches; instead, GenAI uses semantic pattern recognition (or semantic

analysis) to identify the facts of the audit as if it were an intelligent junior investigator working with the Chartered Accountant (CA). This article outlines a simple, four-step framework that integrates GenAI into forensic investigation engagements: 1) Data Hygiene; 2) Contextual Prompt Engineering; 3) Pattern Recognition; and 4) Human Verification. It demonstrates how these solutions can all work together as part of the strict guardrails of the Digital Personal Data Protection (DPDP) Act 2023. It illustrates how to use technological innovations to improve the quality of forensic investigations while adhering to regulations and without compromising client confidentiality.

Introduction: The Data Blindspot

While structured ledgers are the traditional focus of forensic investigations, they constitute only 20% of enterprise information. The remaining 80%, such as unstructured documents like emails and contracts, creates a “Data Blind Spot” where fraudulent intent is hidden. Traditional sampling techniques are increasingly ineffective at parsing this volume. Generative AI (GenAI) provides a critical solution by “reading” the entire dataset’s context and sentiment at scale. This allows investigators to identify potential collusion and anomalies that standard tools cannot detect.

The Technology: Demystifying GenAI for the Forensic Investigator

To effectively leverage GenAI in a professional setting, we must first strip away the hype and understand the

mechanics. A common misconception among finance professionals is that tools like ChatGPT are simply advanced search engines. This is incorrect. A search engine retrieves information that already exists; a Large Language Model (LLM) generates new content based on patterns it has learned.

LLMs: The Pattern Recognition Engine

GenAI is not a search engine; Large Language Models (LLMs) generate content based on learned patterns. To understand an LLM, picture a highly articulate debate champion who understands report structures but lacks a fact-checking database. Because LLMs use probabilistic reasoning, they prioritize linguistic flow over factual accuracy, creating the risk of “Hallucination” (e.g., inventing a fraud clause in a contract where none exists). Relying on this without verification is a professional hazard.

RAG: The “Open Book Exam” for AI

To mitigate hallucination, investigators use Retrieval-Augmented Generation (RAG). Instead of relying on the AI’s general knowledge, the investigator uploads private datasets (e.g., 5,000 PDFs) into a secure environment to create a searchable index. When queried, the system retrieves only pertinent sections and instructs the LLM: “Using only the text provided, answer the investigator’s question.”

The Core Framework: 4 Steps to Implementation

Knowing how technology works is important; however, applying it to your business requires that you use a structured process. A four-step guide will help Chartered Accountants use GenAI in their investigation process. The four-step approach has been developed through successful implementation of forensic case investigations.

Step 1: Data Hygiene and Ingestion – The “Garbage In, Garbage Out” Principle

The main reason that AI-based forensic engagements frequently fail is because of the type of data that is used in the model. In direct alignment with FAIS 410 (Applying Data Analysis), which explicitly recognizes the inherent challenges of receiving data in multiple formats, investigators must meticulously sanitise and standardise data before applying advanced analytical tools. Forensic data often has poor quality because it contains disparate formats, for example, digital/typed text, handwritten notes on margins, invoices scanned as pictures, and improperly formatted email correspondence. Because an LLM does not inherently interpret images, machine-readable text must be extracted and prepared through rigorous data hygiene, thereby fulfilling the FAIS 410 mandate for data preparation.

The Process:

1. **Optical Character Recognition (OCR):** High-quality OCR is mandatory. Poor OCR blinds the AI (e.g., misreading the letter “B” as the number “8” makes “Invoice #B01” unmatchable to the ledger).
2. **De-duplication:** Removing duplicate email threads prevents the AI from falsely identifying a single incident as a systemic pattern.
3. **Metadata Preservation:** It’s important to keep track of important information such as Date Created, Author, and Last Modified when converting file types to text. A forensic investigator cannot use contract text unless the contract has been created at or before the transaction.

Implementation Tip: Treat Data Hygiene as the digital equivalent of “Data Validation” in traditional Computer Assisted Audit Techniques (CAATs). Just as an investigator

“Because LLMs use probabilistic reasoning, they prioritize linguistic flow over factual accuracy, creating the risk of “Hallucination” (e.g., inventing a fraud clause in a contract where none exists).”

validates Excel data before analysis, one must verify the OCR output quality (e.g., via random sample checks) before feeding the ‘Data Lake’ to the AI. High-fidelity text input is the only way to ensure reliable AI output.

Step 2: The “Prompt Library” for Forensics

After the data has been cleansed and placed into a protected environment, the forensic investigator will still need to interact with the said data. When you perform this interaction, you do so by way of “Prompting”. In a digital context, prompting can be viewed as the modern equivalent of asking precise investigation questions. Traditionally, professionals are trained to ask clients clear, targeted questions to obtain accurate information. In the same way, we must now learn to ask AI systems specific, well-structured questions to receive meaningful and reliable responses.

It is recommended that firms build a “Prompt Library”, i.e., a standardized set of queries vetted by senior partners to ensure consistency across engagements.

Case Study: Related Party Transactions (RPT)

- **The Objective:** Identify undisclosed RPTs in a dataset of 2,000 vendor contracts.
- **The Weak Prompt:** “Are there any related parties in these contracts?” (This is too vague; the AI doesn’t know who the related parties are).

- **The Forensic Prompt:** “Context: You are a forensic investigator. Task: Review the attached vendor contracts. Action: Cross-reference the vendor names and registered addresses against the provided list of Key Managerial Personnel (KMP) and their relatives. Flag any contracts where the vendor address matches a KMP address, or where the vendor’s name contains the surname of a director. Output: A table listing the Contract ID, Vendor Name, and the specific Conflict Trigger identified.”

Step 3: Semantic Pattern Recognition

Step 2 places an emphasis on all factual information, whereas Step 3 views the actions of the person or organization committing the fraud. Fraud is rarely perpetrated without emotion; rather, it is usually connected to some sort of emotional pressure, anxiety, or secrecy. The semantic functions of GenAI fit perfectly into this area. GenAI provides semantic analysis of the emotional “tone” of the communications that take place during the fraud; traditional methods of identifying communications related to the fraud do not provide this type of analysis.

Sentiment Analysis for “Management Override”:

Forensic investigators can instruct the AI to analyze email communications between the CFO and the Finance Team during the critical “financial close” period (e.g., March 25th to March 31st).

- **The Prompt:** “Analyze the email dataset from March 25th to March 31st. Identify threads where the sentiment shifts to ‘Aggressive,’ ‘Urgent,’ or ‘Coercive.’ Specifically, flag emails where senior management instructs subordinates to ‘bypass standard procedure,’ ‘post the entry now and fix later,’ or ‘keep this offline.’ Rank the results by intensity of urgency.”

“In direct alignment with FAIS 410 (Applying Data Analysis), which explicitly recognizes the inherent challenges of receiving data in multiple formats, investigators must meticulously sanitise and standardise data before applying advanced analytical tools.”

The “Needle in the Haystack” Approach:

Traditional keyword searches look for “bribe,” but sophisticated fraudsters use code words. GenAI identifies contextual anomalies. If an investigator flags a “Consulting Fee for Market Access” as suspicious, the AI can scan the entire contract database for conceptually similar transactions (e.g., “Liaison Charge” or “Expediting Fee”), granting immediate access to potential schemes without relying on exact word matches.

Step 4: The Human Loop – Verification and Professional Judgment

The final and most critical step is the “Human-in-the-Loop.” It must be unequivocally stated: AI acts as an analytical assistant, but the Chartered Accountant remains the ultimate investigator and decision-maker.

AI models, even with RAG, can misinterpret nuances. A “pressure” email might just be a legitimate deadline; a “conflict of interest” might be a disclosed and approved transaction.

The Verification Protocol:

1. **Source Linking:** AI will usually provide links to sources when giving an answer. Therefore, if AI states something like “Contract #402 has a hidden return clause”,

an investigator must click the citation’s link to see the original PDF and verify whether this fact is accurate, and that the words written actually convey the same meaning intended by the AI.

2. **Corroboration:** The results derived from AI can be considered to be an “Investigative Alert,” where the findings are not definitive until corroborated by the collection of additional evidence, either in the form of bank statements, third-party confirmations, or physical verification.
3. **Robust Documentation, Chain of Custody, and Digital Evidence (FAIS Compliance):** In alignment with FAIS 230 (Documentation in Forensic Engagements) and FAIS 320 (Evidence and Documentation), the working papers must be detailed enough to reconstruct the investigation. Crucially, the Chain of Custody requirements under FAIS 320 apply strictly to AI-processed data. The investigator must document exactly how the data was ingested, the specific AI Model Version (e.g., GPT-4o, Claude 3.5) used, and the Date and Timestamp of the query. Furthermore, in strict adherence to FAIS 420 (Evidence Gathering in Digital Domain), the sanctity of digital evidence must not be compromised. The original source files must remain entirely untouched and be preserved separately from any AI-processed, OCR-converted, or generated versions. Documenting these specific Verification Steps ensures a defensible evidentiary trail, proving the finding relies on unaltered source evidence and not merely an algorithmic output.

By strictly adhering to this verification protocol, the investigator upholds the standards of *Professional Skepticism* mandated by SA 200, ensuring that the efficiency of AI does not come at the cost of audit quality.

Bringing the Framework Together: A Practical Case Study

A whistleblower alleged a National Sales Head was inflating Q4 revenue. Traditional e-discovery failed because perpetrators avoided terms like “fake sales” in the 25,000+ emails and contracts, and the ERP data appeared compliant with Ind AS 115.

- **Data Hygiene & Prompting (Steps 1 & 2):** After rigorous OCR, the AI was prompted to cross-reference contract return clauses against email communications.
- **Semantic Pattern Recognition (Step 3):** The AI analyzed contextual tone, flagging 34 late-March threads showing unusual urgency and implicit indemnification (e.g., “Hold the inventory until April 15th, we will issue a promotional credit note”), successfully identifying hidden side letters.
- **The Human Loop & Verification (Step 4):** Treating this as a lead, the investigator manually reviewed the emails and cross-referenced the ledger, confirming massive March 30th dispatches and anomalous Q1 credit notes. GenAI bridged the gap between unstructured communication and structured records.

This case perfectly illustrates how GenAI bridges the critical gap between unstructured human communication data and structured financial records, uncovering a sophisticated earnings management scheme that traditional keyword searches and standard ledger sampling would have entirely missed.

Reporting AI Usage: Disclosures and FAIS 510

In alignment with FAIS 510 (Reporting Results), any forensic report that leverages GenAI tools must include a dedicated disclosure section. Similar to how a forensic report formally discloses the reliance on a valuation expert or

specific data analytics software, the use of AI must be explicitly acknowledged. This transparency protects the Chartered Accountant professionally and ensures the report meets the strict evidentiary standards required by the Indian Evidence Act.

To ensure full compliance and mitigate liability, this disclosure should clearly outline:

- (a) The specific name and version of the AI tool utilized (e.g., GPT-4o deployed via a secure Microsoft Azure Private Instance);
- (b) The nature of the tasks performed by the AI and the specific areas of the investigation to which it was applied (e.g., semantic pattern recognition, bulk contract review, or email sentiment analysis);
- (c) A clear, unequivocal disclaimer stating that all AI-generated outputs were treated strictly as investigative leads and not as conclusive evidence; and
- (d) A comprehensive statement of limitations. This statement must formally acknowledge the inherent risks of the technology, including the possibility of algorithmic “hallucinations,” the AI’s inability to accurately interpret human sarcasm or cultural context, and the specific risks associated with processing incomplete or missing datasets.



Risk & Compliance: The DPDP Act Context

There are many ways that GenAI can be beneficial to forensics operations; however, they also create new avenues for risk. It's essential that Chartered Accountants use extreme caution when considering these types of tools, as their use involves not just technical decisions, but also legal and ethical decisions based on the current landscape of Data Protection in India, which is rapidly evolving.

The Dangers of Blind Reliance: Operational Risks

While GenAI is a powerful accelerator, treating it as an autonomous investigator invites catastrophic investigative failure. The “Black Box” nature of these models, where the internal decision-making process is opaque, creates specific operational risks that every Chartered Accountant must actively mitigate.

1. The “Plausible Narrative” Fallacy and Wrongful Accusations

LLMs have been made to be convincing and not factual. In a forensic setting, an AI model can be used to relate unrelated events to create a very realistic, coherent, yet incorrect fraud story. Indicatively, it can confuse a typical discussion of a volume discount as a kickback scheme because of bad wording in an email. Should an investigator take action based on this false positive without tracking it down to the underlying documents, it may give rise to false charges against quite innocent workers, reputation losses to the investigator, and potential legal liability.

2. Confirmation Bias Amplification

Sycophancy is a characteristic of AI models - they are prone to giving answers that the user would expect. When an investigator poses a question such as: “Find evidence of the CFO inflating revenue”, the model is more likely to perceive ambiguous data as evidence of

inflation to be able to fulfill the hypothesis of the user. This only increases the confirmation bias of the investigator himself and can be derailing to the investigation by not paying attention to exculpatory evidence.

3. Contextual Blindness and Data Gaps

Artificial intelligence does not have the ability to feel human intuition about sarcasm, industry lingo, or cultural undertones. An email message that says, “*This deal is a steal,*” would be decoded by an AI as a theft and not a good commercial deal. In addition, excessive dependence on the tool poses a threat when it comes to Incomplete Datasets. When the AI is only fed 80% of the data (where offline discussions or hard copies are not made available), it will be certain to give a conclusion based on that incomplete picture and the illusion of being complete is made dangerous.

Therefore, AI outputs must be treated strictly as “Investigative Alerts” or leads, never as conclusive evidence.

The Question of Legal Admissibility and Evidence

A critical distinction must be made: AI output is intelligence, not evidence. Under the Indian Evidence Act, an AI summary of a contract is generally not admissible; the original contract itself is the primary evidence. Thus, GenAI can be employed only to find the evidence but not to substitute it. The investigator should certify the chain of custody of the data. Provided that a fraud discovery is determined on the basis of an AI hallucination alone which could not be supported by source documents, the investigation will fail miserably in a court of law.

The Risk of Skill Atrophy

The risk of losing the possibility to analyze raw documents can be long-term when junior investigators, who are over-reliant on AI summaries, are

“AI output is intelligence, not evidence. Under the Indian Evidence Act, an AI summary of a contract is generally not admissible; the original contract itself is the primary evidence.”

unable to do it. This “Skill Atrophy” may result in having a workforce capable of using the tool but has no idea of the underlying forensic principles. The companies have to make sure that AI is presented as an assistant to the seasoned employees, or the juniors keep doing the manual sample tests so that they retain their essential investigative skills.

Data Privacy and the DPDP Act, 2023

The Digital Personal Data Protection (DPDP) Act, 2023, fundamentally alters how client data is handled. Forensic data inevitably contains “Personal Data.” As a “Data Processor,” the investigator faces strict obligations:

1. **Purpose Limitation:** Utilizing audit data to train a public AI model violates the Act. Uploading financial data to public/free-tier GenAI tools lacks reasonable security safeguards and invites significant penalties.
2. **The “Private Instance” Mandate:** To remain compliant, firms must utilize “Enterprise-Grade” private environments where cloud providers contractually guarantee that data remains within specified geographic boundaries (data sovereignty) and is *not* used to train base models.

Investigator Liability and the Human Shield

More importantly, the DPDP Act instills heavy punishment for the misuse of data, yet the risk is not limited

to fines, but also to professional liability. In the event that an AI model indicates that an employee has committed fraud, when it is actually a hallucination (as explained in Operational Risks), and the employee in question is harmed as a result, whether by reputation or by being fired, it is not the software that is liable, but the human investigator. The Act focuses on accountability of automated decision-making. As such, the protocol of Human-in-the-Loop, considered in Step 4, is not only an investigative quality measure, but a legal requirement. The human checking process is a legal “circuit breaker,” turning an unrefined algorithmic probability into a professionally checked discovery, and thus protecting the Chartered Accountant against allegations of algorithmic laxity or mishandling of data.

Conclusion

Forensic investigation methodology has been transformed from reactive sampling methods to proactive, extensive analysis through the implementation of GenAI technologies. However, while AI can analyze and generate large volumes of data, it cannot replace the professional judgement of Chartered Accountants because AI does not have the ability to comprehend human intent, nor can it testify in court. Therefore, the future of forensic investigations will be populated by “Augmented Investigators”, or professionals who utilize the speed and volume of data generated by AI in conjunction with the professional judgment and ethical scepticism that are synonymous with Chartered Accountants, resulting in the transformation of the “Data Blindspot” into a strategic advantage and ultimately, a more robust level of evidence (or evidentiary support) in the current evolving digital landscape.

References

- Bouquot, J. (2025). *Creating the Future Together in the Accountancy Profession's AI Revolution*. JOURNAL OF THE INSTITUTE OF CHARTERED ACCOUNTANTS OF INDIA. [https://](https://resource.cdn.icai.org/89646cajournal-dec2025-8.pdf)



resource.cdn.icai.org/89646cajournal-dec2025-8.pdf

- DPDP Rules, 2025 notified. (n.d.). <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2190655>
- Release of Revised Forensic Accounting and Investigation Standards (FAIS) and Implementation Guide on Forensic Accounting and Investigation Standards - (27-07-2023). (2023, July 27). The Institute of Chartered Accountants of India. <https://www.icai.org/post/daab270723>
- ACFE Report to the Nations | 2024 Global Fraud Study. (n.d.). <https://legacy.acfe.com/report-to-the-nations/2024/>
- Standard on Auditing (SA) 240, *The Auditor's Responsibilities Relating to Fraud in an Audit of Financial Statements*, The Institute of Chartered Accountants of India (ICAI).

Author may be reached at
saurabhkhakhkh45@gmail.com
and ebboard@icai.in