International Journal of Engineering Research and Development
e- ISSN: 2278-067X, p-ISSN: 2278-800X, www.ijerd.com
Volume 22, Issue 6 (June 2026), PP 37-40

Automated Evidence Acquisition and Investigative Accuracy Enhancement

Chaitresh Naik, Sairaj Gadhave, Darshan Kondagekar, Sneha Kambli
Department of Information Technology, K.C College of Engineering Management Studies And Research, Thane, India

Abstract
The rapid rise of cybercrime, fraud, and network-related offenses has stressed manual forensic investigation to its limits. Investigators must now sift through massive amounts of digital data—disk images, packet captures, chat logs, and surveillance footage—under tight judicial deadlines. This task greatly exceeds what human analysts can do without computer support. To address this issue, this paper introduces AI-FAF (Artificial Intelligence Forensic Analytics Framework), a seven-stage pipeline that integrates deep learning, traditional machine learning, and graph reasoning into a unified workflow.

Convolutional Neural Networks authenticate images and videos; an XGBoost ensemble with Isolation Forest identifies network intrusions and logs anomalies; a BERT-centered NLP stack analyzes text communications; and a Graph Attention Network connects findings from these three areas into a coherent evidence framework. Three simulated case studies—focused on corporate data theft, social media fraud, and ransomware attacks—were used to evaluate the system. In these scenarios, AI-FAF achieved a mean detection accuracy of 94.6%, reduced manual workload by approximately 67%, and lowered the false alert rate by 41% compared to traditional rule-based methods. Explanations from SHAP and LIME accompany each high-confidence finding, while cryptographic output sealing preserves the chain of custody in a format acceptable for court proceedings

---

Date of Submission: 27-05-2026 Date of acceptance: 05-06-2026

I. Introduction

Few fields have changed as much because of the digital age as forensic science. Where investigators once mainly dealt with physical evidence—fingerprints, fibers, ballistic casings—today's cases are dominated by digital data: encrypted archives, cloud-hosted databases, smartphone backups, and real-time network streams. The vast amount of data presents a challenge. Evidence is more plentiful than ever, but extracting useful information from it has become a difficult computational task. Law enforcement agencies around the world report that case storage media often reaches several terabytes. A single smartphone backup can hold hundreds of thousands of files; a corporate network forensic image may involve many servers. Using traditional forensic tools—keyword searches, manual log reviews, and frame-by-frame video inspections—on datasets this large is not practical.

This causes delays, results in overlooked evidence, and forces investigators into mechanical tasks instead of analytical thinking. Artificial intelligence offers a viable solution. Trained neural networks can quickly scan image libraries for manipulation signs; sequence models can identify unexpected system behaviors; transformer language models can find suspicious language in millions of chat messages.

However, there has been a lack of a framework that combines these individual capabilities into a reproducible and legally sound process. This paper presents that framework—AI-FAF—and makes four key contributions: a seven-stage modular pipeline; validated selection of AI models for different evidence types; a simulated evaluation with measurable benchmarks; and an analysis of the chain of custody and explainability needs for court admissibility.

II. Literature Review

Research on AI-assisted forensics has progressed along three largely separate paths: visual evidence authentication, network and log analysis, and text communication examination. Each path has seen significant advancements, but there has been little focus on cross-path integration. On the visual front, early statistical techniques for detecting image splicing were replaced by tailored convolutional architectures. A constrained CNN model that minimizes semantic content to isolate manipulation traces achieved accuracy above 90% on standard datasets.

A dual-stream R-CNN improved localization precision by learning manipulation cues alongside noise features. Deepfake detection has also gained attention, especially with the advent of the FaceForensics++ dataset. In network forensics, research has focused on the CICIDS2017 benchmark, where ensemble methods like Random Forest combined with gradient boosting have led. Concurrent work in host-based forensics showed that LSTM models could help automate log file analysis with minimal supervision.

Communication forensics has grown with transformer models; fine-tuned BERT configurations have demonstrated over 88% accuracy in authorship identification for legal documents, and NER-assisted topic modeling has effectively identified patterns of criminal coordination on social media. Despite these advancements, the field lacks a unified framework. Each tool works on a single evidence type and does not provide outputs that can be understood by other tools. AI-FAF directly tackles this integration issue and incorporates legal requirements—verifiable custody and interpretable outputs—often overlooked in previous research.

III. Problem Statement

Four related challenges define the investigative problem this work aims to solve. The first is volume overload: a mid-sized cybercrime investigation can include many seized devices and server images, creating datasets that no forensic team can manually review in a reasonable time. The second challenge is heterogeneity, where each evidence type needs specialized tools that do not share a common interface or output format. The third issue is signal noise.

Rule-based detection systems generate many alerts, most of which are harmless, wasting valuable investigator time that should focus on real leads. Finally, the correlation gap exists. Modern criminal networks rarely leave traces confined to one device or channel. Connecting a suspicious network activity to a related file access event on the same workstation, along with corroborating phrases in an employee's email, requires coordinated reasoning across different evidence types.

On top of these operational challenges, there is a legal requirement—automated evidence must come with a record of processing and explanations clear enough for judges and juries. A system that is accurate but not clear will struggle during legal scrutiny. The research goal is to design an end-to-end forensic pipeline that addresses all four operational challenges while meeting transparency standards necessary for admissibility in court

IV. Proposed Framework (AI-FAF)

AI-FAF consists of a sequential seven-stage pipeline that includes optional feedback loops, allowing investigators to rerun specific stages after updating parameters or adding new evidence.

PIPELINE DIAGRAM (text form):

STAGE 1 → Evidence Ingestion & SHA-256 Cataloguing STAGE 2 → Preprocessing & Format Normalization
STAGE 3 → Parallel Multi-Modal AI Analysis
├─ 3a Image/Video → EfficientNet-B4 + U-Net
├─ 3b Network/Logs → XGBoost + Isolation Forest
└─ 3c Text/Comms → BERT + NER + LDA
STAGE 4 → Cross-Modal Correlation → GAT on Neo4j STAGE 5 → Automated Report & Timeline Generation STAGE 6 → Explainability Annotation → SHAP / LIME STAGE 7 → Cryptographic Sealing & CoC Documentation

In Stage 1, each incoming file is assigned a SHA-256 digest, which is recorded in an append-only audit database before processing begins. Stage 2 standardizes inputs: network captures turn into flow-level vectors; documents get tokenized and cleaned; images are converted to normalized pixel arrays; and deleted files are recovered using file carving. Stage 3 sends standardized inputs to three parallel analysis engines. Each engine assigns a confidence score and provides feature explanations for flagged artifacts. Stage 4 combines all tagged artifacts into a property graph and uses graph reasoning to identify relationships across multiple sources. Stage 5 utilizes a natural language generation module to produce a structured case report with an event timeline and evidence provenance map. Stage 6 includes SHAP plots for classifier-flagged items and LIME explanation masks for image and text findings. Stage 7 finalizes output in a cryptographically signed package ready for legal submission.

V. System Architecture

AI-FAF operates as a containerized microservices platform with five layers.

A. Ingestion Layer— REST connectors work with FTK Imager and Cellebrite. Apache Kafka manages live-stream cases. Every artifact gets hashed and logged into PostgreSQL before being sent downstream.
B. Storage and Preprocessing Layer — Evidence is stored in a MinIO object store. Apache Airflow oversees preprocessing task graphs. Processed features are stored in Apache Parquet format for quick access.
C. AI Inference Layer — Three Kubernetes-hosted engines run simultaneously. The Vision Engine runs EfficientNet-B4 and ResNet-50. The Network Intelligence Engine supports XGBoost, Isolation Forest, and an LSTM log analyzer. The Language Engine operates with fine-tuned BERT, a spaCy NER pipeline, and Gensim LDA. All engines have gRPC endpoints for consistent invocation.
D. Correlation and Reasoning Layer — Neo4j stores five node types (File, Person, Device, NetworkEndpoint, Event) connected by typed edges. A two-layer Graph Attention Network assigns scores based on investigative importance, highlighting subgraphs above a set suspicion threshold.
E. Output and Compliance Layer — This layer creates case reports in PDF and DOCX formats, includes SHAP/LIME annotations, checks for ISO/IEC 27037 and ACPO compliance, and produces a sealed cryptographic output package.

VI. Implementation Details

A. Image and Video Forensics

EfficientNet-B4 is fine-tuned on CASIA v2.0 and FaceForensics++. A U-Net segmentation head enables pixel-level tamper localization. Video processing applies frame-level classification and inter-frame motion vector consistency scoring to detect spliced footage. Validation: Accuracy 94.1%, Precision 92.8%, Recall 95.3%, F1 94.0%.

B. Network and Log Analysis

A Scapy/PyShark module extracts 78 flow-level features per PCAP matching the CICIDS2017 specification. XGBoost covers 14 labeled attack families. Isolation Forest detects statistically rare flows outside any trained category and supports novel attack detection. LSTM log analysis uses FastText embeddings trained on system administration data.

C. Text and Communication Forensics

Four pipeline stages: (1) language detection and Unicode normalization; (2) BERT classification on 45,000 annotated forensic documents; (3) NER extraction of people, organizations, locations, dates, and identifiers; (4) LDA topic modeling for thematic summarization. A separate stylometric SVM handles authorship attribution using function-word frequency profiles and sentence-length distributions.

D. Graph Correlation

The GAT uses two attention layers trained on synthetic forensic graphs. Cypher queries support hypothesis-driven investigation, such as retrieving all devices communicating with a flagged IP within a defined time window.

VII. Results and Analysis

Three cases were evaluated: Case A (internal data exfiltration), Case B (social-media fraud), Case C (ransomware intrusion). Ground-truth annotations were provided by forensic examiners with no role in framework development.

Component	Avg. Accuracy	Avg. FPR	Speed
CNN — Image/Video	93.1%	4.0%	18.2 min/GB
XGBoost + Iso. Forest	96.1%	1.8%	4.2 min/GB
BERT + NER + LDA	93.4%	5.1%	9.5 min/GB
GAT Correlation	Prec 89.3%	Rec. 91.8%	-

Mean detection accuracy across all modalities: 94.6%. Versus the rule-based baseline, AI-FAF reduced analyst time by 67.3% and cut false alerts by 41.2%. Automated report generation averaged 3.8 minutes per case versus 14–22 hours estimated for manual documentation.
The most significant finding came from Case C: the GAT linked a lateral-movement pattern in network logs to a document access event on the same workstation within 15 minutes — a cross-modal relationship neither tool would have surfaced alone. This directed the investigation to the confirmed initial intrusion vector. The explainability module generated SHAP plots for all XGBoost-flagged flows and LIME masks for 89% of BERT classifications; 83% were rated courtroom-ready by two senior examiners in a blind assessment.

VIII. Advantages and Limitations

Advantages: Processing throughput reduces weeks of manual triage to hours. Cross-modal correlation reveals relationships that span device boundaries and single-purpose tools often miss. Built-in custody tracking logs every transformation and hashes every output for end-to-end verification.

Limitations: Training data scarcity limits generalizability because forensic datasets rarely enter the public domain. Adversarial counter-forensic techniques may evade specific classifiers. SHAP and LIME might not meet theoretical admissibility standards under Daubert or Frye in some jurisdictions. Infrastructure demands may challenge smaller agencies without shared computing resources.

IX. Future Scope

Future Scope Federated training will allow model updates to be shared across agencies without exposing case data. Adversarial training will strengthen classifiers against deliberate evasion. Multi-language NLP will extend coverage to Hindi, Arabic, and Mandarin. OSINT integration will enhance the knowledge graph with public domain records and leaked credential data. A Hyperledger Fabric blockchain is being considered as a distributed tamper-evident audit log. Real-time SOC deployment will shift the framework from post-incident investigation to continuous forensic awareness.

X. Conclusion

Forensic investigation is at a turning point. The volume, speed, and variety of digital evidence produced by modern criminal activity have made manual workflows outdated for anything beyond the smallest cases. AI-FAF tackles this by bringing together four established AI technologies in a single auditable pipeline that starts at evidence intake and ends with a court-ready sealed package. Across three simulated forensic scenarios, the framework achieved 94.6% mean detection accuracy, a 67% reduction in workload, and a 41% decline in false alerts.

The cross-modal correlation produced an evidence link in the ransomware case that no single-purpose tool would have plausibly found. This shows that the integration value of the architecture is genuinely synergistic, not just additive. As these capabilities improve alongside federated learning, adversarial strengthening, and multi-language support, AI-assisted forensic pipelines are likely to become standard law enforcement infrastructure worldwide.

References

[1] H. Farid, "Image forgery detection," IEEE Signal Processing Magazine, vol. 26, no. 2, pp. 16–25, Mar. 2009.

[2] B. Bayar and M. C. Stamm, "A deep learning approach to universal image manipulation detection using a new convolutional layer," in Proc. ACM Workshop on Information Hiding and Multimedia Security, 2016, pp. 5–10.

[3] P. Zhou, X. Han, V. I. Morariu, and L. S. Davis, "Learning rich features for image manipulation detection," in Proc. IEEE/CVF CVPR, 2018, pp. 1053–1061.

[4] C. Brun, A. Desnos, and J.-F. Marion, "Anomaly-based intrusion detection through LSTM sequence modelling on system call traces," Computers & Security, vol. 99, pp. 102–117, Dec. 2020.

[5] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. 4th Int. Conf. ICISSP, 2018, pp. 108–116.

[6] S. Chowdhury and P. Bhattacharyya, "Authorship attribution in legal documents using transformer-based models," Expert Systems with Applications, vol. 189, pp. 116–131, Mar. 2022.

[7] M. Sultana, P. Paul, and M. Gavrilova, "Social media forensics using NLP for criminal investigation," in Proc. IEEE CogSIMA, 2019, pp. 77–84.

[8] Y. Chen, R. Zhao, and X. Li, "Graph attention networks for cross-device forensic correlation," IEEE Trans. Inf. Forensics Security, vol. 17, pp. 1204–1219, 2022.

[9] K. Nance and B. Re, "Digital forensics: Defining a research agenda," in Proc. 43rd Hawaii ICSS, 2010, pp. 1–6.

[10] T. Dargahi et al., "A cyber-kill-chain based taxonomy of crypto-ransomware features," J. Computer Virology and Hacking Techniques, vol. 15, no. 4, pp. 277–305, 2019.

[11] S. Garfinkel, "Digital forensics research: The next 10 years," Digital Investigation, vol. 7, suppl., pp. S64–S73, Aug. 2010.

[12] A. Schlegel, D. Hilt, and K. Pohl, "Explainable AI for digital forensics," Forensic Science International: Digital Investigation, vol. 44, p. 301475, Mar. 2023.