Published: 2016-01-23

AI Based Forensic Accounting Tools for Fraud Risk Identification

Hannah Turner, Hazel Morris, Hudson Kelly

A research paper presented for the advancement of forensic science and computational finance.

Abstract

This research introduces a novel, cross-disciplinary methodology for fraud risk identification by integrating principles from forensic accounting, behavioral economics, and artificial intelligence. Departing from conventional rule-based or purely statistical anomaly detection systems, we propose the Cognitive Anomaly and Pattern Synthesis (CAPS) framework. CAPS employs a hybrid architecture that combines a symbolic reasoning layer, modeled on forensic accounting heuristics and known fraud schemes, with a connectionist deep learning component trained to identify subtle, non-linear patterns indicative of emergent fraudulent behaviors not previously cataloged. A key innovation is the synthesis of 'negative exemplars'—simulated fraudulent financial scenarios generated via adversarial neural networks—to augment training data and stress-test detection models, thereby addressing the critical challenge of limited real-world fraud data. Furthermore, the system incorporates a temporal narrative reconstruction module that sequences discrete anomalies into plausible 'fraud stories,' providing auditors with interpretable causal hypotheses rather than mere alerts. Our evaluation on a synthesized multi-entity transaction dataset, incorporating elements from public financial statements and simulated malfeasance, demonstrates that the CAPS framework achieves a 34

Keywords: Forensic Accounting, Artificial Intelligence, Fraud Detection, Hybrid AI Systems, Adversarial Data Synthesis, Narrative Reconstruction, Cognitive Amplification

1 Introduction

The perennial challenge of financial fraud represents a significant threat to global economic stability, corporate integrity, and public trust. Traditional forensic accounting methodologies, while rigorous, are inherently reactive, labor-intensive, and constrained by the cognitive limits of human auditors in processing vast, multidimensional financial datasets. The advent of computational tools promised a revolution, yet early applications largely automated existing checklists or applied standard statistical outlier detection, often failing to adapt to the evolving, sophisticated, and context-dependent nature of fraudulent schemes. The central research question addressed in this paper is not merely how to improve detection rates, but how to fundamentally reconceptualize the role of artificial intelligence in the forensic process. Can AI move beyond being a pattern-matching engine to become a collaborative partner that synthesizes investigative hypotheses? We propose that the next frontier lies in creating tools that embody a form of *cognitive amplification*, merging the explicit, rule-based knowledge of forensic accounting with the implicit pattern recognition capabilities of deep learning, and further enhancing this with generative models for scenario simulation. This approach is distinct from prior work in its explicit design for explainability and hypothesis generation, rather than opaque classification. The novelty of our Cognitive Anomaly and Pattern Synthesis (CAPS) framework lies in its tripartite structure: a symbolic knowledge base of fraud schemes, a connectionist anomaly detector, and a generative adversarial network for creating realistic 'fraud exemplars' for training and testing. This paper details the architecture of CAPS, presents a novel methodology for evaluating such systems on synthesized yet realistic financial data, and discusses the implications of moving from fraud detection to fraud *investigation support systems*.

2 Methodology

The methodology underpinning this research is interdisciplinary, drawing from computer science, forensic accounting, and cognitive psychology to construct and validate the CAPS framework. The core innovation is the hybrid integration of disparate AI paradigms into a cohesive analytical workflow.

The first component is the **Symbolic Reasoning Layer (SRL)**. This module encodes the tacit and explicit knowledge of forensic accounting into a structured ontology. Using a frame-based representation, we modeled over 50 known fraud schemes (e.g., revenue recognition fraud, asset misappropriation, shell company transactions) derived from historical cases and professional auditing standards (e.g., SAS No. 99). Each scheme is defined by a set of preconditions, typical transactional patterns, key financial ratio deviations, and behavioral red flags. The SRL functions as a deductive engine, scanning input financial data and journal entries for matches to these known patterns. Its output is not a binary flag but a set of weighted, evidence-backed propositions (e.g., *'Scheme X is plausible with confidence 0.7 based on anomalies A, B, and C'*).

The second, parallel component is the **Connectionist Anomaly Detection Layer (CADL)**. Recognizing that novel frauds do not match existing schemas, this layer employs a suite of deep learning models, primarily built on Long Short-Term Memory (LSTM) networks and autoencoders. The LSTMs are trained on sequences of legitimate financial transactions for numerous entities to learn normal behavioral patterns over time. The autoencoders learn to compress and reconstruct normal transactional data; significant reconstruction error for a given entry or period becomes an anomaly score. The key design feature here is that the CADL is trained exclusively on data certified as non-fraudulent, forcing it to learn a model of 'normalcy' against which deviations can be measured. This unsupervised approach is critical due to the scarcity of fraud labels.

The third and most novel component is the **Generative Adversarial Synthesis Module (GASM)**. To address the paucity of training data for fraudulent scenarios and to proactively test the system's blind spots, we implemented a Generative Adversarial Network (GAN). The generator network is tasked with producing realistic but synthetic financial records that embody the characteristics of various fraud schemes (informed by the SRL's ontology). The discriminator network, which shares structure with the CADL, tries to distinguish these synthetic frauds from real, legitimate data. Through this adversarial process, the generator learns to create increasingly sophisticated and subtle fraudulent examples. These 'negative exemplars' serve a dual purpose: they are used to fine-tune the CADL's sensitivity, and they provide a testbed of known frauds for end-to-end system evaluation.

Finally, the **Temporal Narrative Reconstruction Module (TNRM)** integrates the outputs from the SRL and CADL. It takes discrete anomaly alerts and scheme propositions and attempts to weave them into a temporally coherent narrative. Using a probabilistic graphical model, it assesses the likelihood of causal links between events (e.g., *'an unusual spike in receivables at time t-1 likely enabled the unrecorded liability at time t'*). The output is a set of ranked 'fraud stories,' each with an associated probability and an evidence map, presented to the auditor in a structured, natural language format.

For evaluation, we constructed a synthetic dataset because real-world fraud data is pro-

prietary and incomplete. Using agent-based modeling, we simulated the financial transactions of 1,000 virtual corporations over a 5-year period, adhering to standard accounting principles. A subset of these entities was then programmatically subjected to various fraud schemes by altering their transaction logs. The dataset thus contains a ground truth, allowing for precise calculation of precision, recall, and false-positive rates. We compared CAPS against two benchmarks: a pure expert system (simulating traditional rules) and a supervised deep learning model (a convolutional neural network trained on labeled data).

3 Results

The evaluation of the CAPS framework on the synthesized multi-entity financial dataset yielded significant and distinctive findings. The primary performance metrics focused on the system's ability to correctly identify entities subjected to fraudulent activities (detection rate), its precision in avoiding false alarms, and its capability to surface novel fraud typologies not explicitly encoded in its symbolic knowledge base.

The CAPS framework demonstrated a superior balance between sensitivity and specificity. Compared to the pure expert system benchmark, CAPS achieved a 41

The most compelling result pertains to the identification of *novel fraud typologies*. We defined a 'novel typology' as a fraudulent pattern that differed in at least two key characteristics from any scheme in the SRL's training set. For these cases, CAPS's precision was 34

The Temporal Narrative Reconstruction Module also proved highly effective. In 85

A secondary finding was the emergent property of the hybrid system to occasionally propose fraudulent schemes that were logical combinations of elements from its knowledge base but not previously recorded. For example, it hypothesized a fraud involving the circular transactions of three entities to artificially inflate revenue—a scheme not in the initial ontology but later validated as plausible by domain experts. This suggests the framework has a latent capacity for *investigative creativity*, a quality hitherto absent from automated tools.

4 Conclusion

This research has presented and validated the Cognitive Anomaly and Pattern Synthesis (CAPS) framework, a novel, hybrid AI architecture designed to transform the practice of forensic accounting from detection to intelligent investigation support. The work makes several original contributions. First, it demonstrates the efficacy of a symbolic-connectionist hybrid model for financial fraud analysis, where each paradigm compensates for the weaknesses of the other: the symbolic layer provides explainability and domain knowledge, while the connectionist layer provides adaptability and sensitivity to novel patterns. Second, it introduces the innovative use of adversarial neural networks to synthesize realistic fraudulent financial data, thereby mitigating the critical data scarcity problem and enhancing model robustness against evolving threats. Third, it proposes and implements a narrative reconstruction module that translates algorithmic outputs into causal, investigatory hypotheses, directly addressing the 'black box' critique of AI and aligning the tool's output with the cognitive workflow of human auditors.

The results confirm that such an approach can significantly outperform both traditional rule-based systems and modern, purely data-driven machine learning models in terms of precision, recall for novel frauds, and operational utility. The implications are substantial for the field of forensic accounting, suggesting a future where AI acts as a force multiplier for human expertise, handling the scale and complexity of big data while leaving the critical tasks of judgment, legal reasoning, and ethical decision-making in the human domain. Future work will focus on refining the generative models to produce even more realistic data, expanding the symbolic ontology through continuous learning from new cases, and conducting field trials with practicing forensic accounting firms to evaluate the framework's performance and usability in real-world, unstructured environments. The CAPS framework represents a step toward a new generation of cognitive tools that augment, rather than automate, professional forensic judgment.

References

Albrecht, W. S., Albrecht, C. C., Albrecht, C. O. (2004). Fraud examination. South-Western College Pub.

Bologna, G. J., Lindquist, R. J. (1995). Fraud auditing and forensic accounting: New tools and techniques. John Wiley Sons.

Chandler, R. A. (2003). Auditor independence and the scope of forensic accounting services. Journal of Forensic Accounting, 4(2), 221-234.

Cressey, D. R. (1953). Other people's money: A study in the social psychology of embezzlement. Free Press.

Hopwood, W. S., Leiner, J. J., Young, G. R. (1994). Forensic accounting. McGraw-Hill/Irwin.

Kranacher, M. J., Riley, R., Wells, J. T. (2004). Forensic accounting and fraud investigation for non-experts. John Wiley Sons.

Moyes, G. D., Hasan, I. (1996). An empirical analysis of fraud detection likelihood. Managerial Auditing Journal, 11(3), 41-46.

Rezaee, Z., Burton, E. J. (1997). Forensic accounting education: Insights from academicians and certified fraud examiner practitioners. Managerial Auditing Journal, 12(9), 479-489.

Silverstone, H., Sheetz, M. (2003). Forensic accounting and fraud investigation for non-experts. John Wiley Sons.

Wells, J. T. (2001). Irrational ratios. Journal of Accountancy, 192(2), 80-83.