Qeios

Peer Review

Review of: "Beyond Suspicion Scores: Competing-Explanation Inference for Fraud-Shaped Legitimate Activity"

Chandra Shikhi Kodete[1]

1. Eastern Illinois University, United States

Overall Rating: 4/5

1. Overall Assessment

This is an interesting and thoughtfully developed manuscript. The idea of evaluating suspicious evidence together with positively supported legitimate explanations provides a useful alternative to conventional suspicion-based fraud detection. The manuscript is also transparent about its experimental scope and appropriately presents the current evaluation as a mechanism proof rather than evidence of deployment-ready performance.

2. Strengths

The competing-explanation framework is a meaningful and potentially useful contribution to fraud inference.

Treating exculpatory and contradictory evidence as integral parts of the inference process is particularly interesting.

The authors report negative findings and unsuccessful formulations rather than presenting only favorable results.

The limitations of the current evidence are clearly acknowledged, including the absence of real-world and independent validation.

doi.org/10.32388/ND2081

3. Main Concerns

Independent validation: The main limitation is that the 160-case corpus was authored by the same team that developed AFIE. The reported results therefore demonstrate that the architecture behaves as intended within this diagnostic setting, but they cannot yet establish real-world performance or generalizability.

Post-diagnostic refinement: M6, M7, and M8 were introduced in response to observed failures. This is understandable during system development, but the final architecture has therefore not yet undergone an untouched confirmatory evaluation.

Parameter robustness: Several parameters are declared rather than fitted, and their sensitivity has not been characterized. This makes it difficult to determine how dependent the reported behavior is on the chosen configuration.

Interpretability: The architecture provides a promising degree of decision reconstructibility, but this should remain distinct from demonstrated explanation quality or usefulness to fraud investigators, which has not yet been evaluated.

4. Suggestions to Strengthen the Work

I encourage the authors to evaluate a fully frozen version of AFIE on an independently developed dataset or, preferably, real-world fraud cases. Temporal validation or evaluation across institutions would provide particularly useful evidence of generalizability. A systematic sensitivity analysis of the prior, channel scales, coverage threshold, LLR cap, and redundancy parameters would also strengthen the robustness assessment. Finally, additional component-level ablations and future evaluation with fraud investigators would help establish which architectural elements are most important and whether the resulting explanations are useful in practice.

5. Overall Rating

4/5 – Strong conceptual contribution with further independent validation needed.

The manuscript makes a worthwhile contribution at the architectural level. The authors have been appropriately careful about the limits of the present evidence, and I believe independent validation and robustness testing are the most important next steps.

Declarations

Potential competing interests: No potential competing interests to declare.