International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(Research Article)

Methodology of weighted fusion of heterogeneous data with an exponential obsolescence factor in financial transactions

Kaleshwar Aryasomayajula *

Sr Specialist Software Developer, Charles Schwab, San Tan Valley, Arizona, USA.

International Journal of Science and Research Archive, 2026, 19(02), 599-603

Publication history: Received on 31 March 2026; revised on 05 May 2026; accepted on 08 May 2026

Article DOI: https://doi.org/10.30574/ijsra.2026.19.2.1001

Abstract

This article develops a methodological model for fusing heterogeneous financial data in transaction risk assessment when input streams age at different speeds. The relevance of the topic follows from the spread of digital lending, instant payments, and graph-based fraud analytics, in which bureau records, behavioral traces, network signals, and contextual data enter a single decision circuit but retain unequal evidentiary value over time. The study aims to formalize a weighted fusion procedure with an exponential obsolescence factor and to define its place within auditable financial decision logic. The materials cover ten recent scholarly works on alternative data, credit scoring, fraud detection, graph learning, and multimodal fusion. Comparative analysis, conceptual synthesis, typologization, and analytical generalization were applied. The analytical section yields a source-sensitive weighting model, a recency-adjusted fusion equation, and an implementation sequence for real-time scoring and explanation. The approach applies to digital lending, payment monitoring, and anti-fraud screening.

Keywords: Heterogeneous data fusion; Exponential obsolescence; Financial transactions; Transaction risk

1. Introduction

Financial transaction analytics now works with data streams that differ sharply in structure, refresh frequency, and evidentiary stability. A lending or fraud-screening engine may receive bureau variables, payment histories, transaction sequences, device traces, behavioral patterns, merchant links, and relational network signals within one decision window. Treating these inputs as commensurable from the start creates a methodological distortion. A source captured minutes ago and a source refreshed weeks earlier may both be statistically informative, yet their inferential force at decision time is not equal.

This study aims to develop a methodological framework for weighted fusion of heterogeneous financial data under conditions of unequal temporal relevance. The study pursues three objectives. First, it identifies the data families that enter transaction-level financial decisions and clarifies why they require differentiated treatment. Second, it formulates a recency-sensitive fusion rule based on an exponential obsolescence factor. Third, it translates that rule into an auditable decision logic suitable for real-time financial systems.

The novelty of the article lies in bringing three lines of reasoning into a single scheme: the heterogeneity of financial evidence, the unequal temporal decay of signals, and traceable decision-making. The proposed contribution is analytical. Its purpose is to build a coherent methodological model that can guide later implementation in credit scoring, payment monitoring, and fraud prevention.

* Corresponding author: Kaleshwar Aryasomayajula

2. Materials and Methods

The material base comprises 10 peer-reviewed studies published between 2021 and 2026. The corpus combines review work on credit scoring and alternative data, methodological studies on information fusion and multimodal fraud logic, and transaction-centered research on graph neural, temporal, and multi-view approaches to fraud and anti-money-laundering detection [1]–[10]. The selection was guided by three criteria: direct relevance to financial decisioning, explicit treatment of heterogeneous inputs or fusion, and substantive attention to temporal dynamics, class imbalance, explanation, or operational deployment. As a corpus, these works cover alternative data, transaction graphs, temporal encoding, attention-based aggregation, multi-view explanation, and information fusion in enterprise or transactional risk settings.

The study applies comparative analysis to distinguish the logics of source combination found in recent financial analytics research. Source analysis was used to reconstruct how the selected works define data heterogeneity, time sensitivity, and evidentiary weighting. Conceptual synthesis was then used to integrate these positions into one fusion model. Typologization helped separate source families by temporal behavior. Analytical generalization was used to move from literature-based observations toward an implementable transaction-level methodology.

3. Results

Recent reviews of credit scoring and alternative data [1], [4], together with earlier work on information fusion in enterprise financial risk [10], converge on a point that is easy to overlook in practical system design. Financial decision engines lose resolution when all inputs are treated as equally stable. Bureau variables, repayment histories, transaction traces, device footprints, merchant relations, and behavioral signals enter the same risk question, yet they differ in refresh rate, provenance quality, and legal defensibility. The methodological consequence is immediate. Fusion should begin with source stratification, because a recent transaction anomaly and a two-month-old external record do not carry the same evidentiary force even when both are predictive. This matters most in digital lending and transaction screening, where the decision window is short, and a stale signal can dominate the final score if weights are fixed.

Transaction-centered graph studies further push the argument. HHLN-GNN separates homogeneous and heterogeneous connectivity and assigns attention-based weights to them [6]. ASA-GNN filters neighbors by behavioral similarity and edge weights, then adds neighbor diversity to resist camouflage [5]. DyHDGE divides temporal and structural extraction into separate modules [7], whereas THG-OAFN first segments transactions into time windows and then fuses temporal dynamics with a heterogeneous graph structure [8]. Read together, these studies do not support raw concatenation of mixed attributes. They point to staged fusion, where each source family is encoded in the form most natural to its structure before cross-source aggregation begins.

A second line of research changes the interpretation of temporal relevance itself. MultiFraud builds fraud detection across multiple heterogeneous views and connects prediction to explanation [9]. The decoupling-fusion model proposed for multimodal financial fraud distinguishes slow-moving motive signals from acute action signals [3]. An AML study based on a large real-world bank network [2] confirms that heterogeneity is an operational property of live surveillance data. Once these positions are aligned, a useful distinction appears. Some inputs behave as chronic descriptors, such as historical profiles, institutional records, or long-horizon network patterns. Others behave as acute indicators, such as bursty transaction behavior or short-window relational irregularity. A fusion rule that ignores this asymmetry blurs the decision boundary while weakening interpretability.

For that reason, the weighted fusion stage is better written in source-specific form. Let s_i denote the normalized risk contribution of source i, let α_i denote its baseline reliability coefficient, and let Δt_i denote source age at inference time. Then the effective weight can be written as (1):

$$w_i = \frac{\alpha_i e^{-\lambda_i \Delta t_i}}{\sum_{j=1}^{m} \alpha_j e^{-\lambda_j \Delta t_j}} \qquad (1)$$

and the fused transaction score as (2):

$$F = \sum_{i=1}^{m} w_i s_i. \quad (2)$$

This equation is deliberately compact, yet it resolves a recurring weakness in the reviewed literature. It preserves comparability across sources, penalizes obsolete inputs smoothly, and leaves room for source-specific decay rates λ_i. A bureau record may decay slowly. Device risk, merchant interaction bursts, or network anomalies decay much faster. Obsolescence, then, is not a property of the model as a whole. It belongs to each signal family separately.

The reliability term α_i should not be reduced to statistical confidence alone. The selected studies imply a broader reading. It absorbs data provenance, missingness, label quality, susceptibility to adversarial manipulation, and the extent to which a source can be explained to analysts or auditors [1], [3], [4], [9]. This is where recent work on credit scoring and fraud systems becomes methodologically useful. Alternative data expands coverage [1], [4], but its value depends on whether the resulting decision can be defended. Multi-view explanation in fraud graphs [9] and the chronic-acute separation in multimodal fraud modeling [3] suggest that fusion needs a provenance layer. A score without a trace of source contribution is difficult to monitor, difficult to contest, and difficult to recalibrate.

The proposed analytical sequence is summarized in Figure 1.

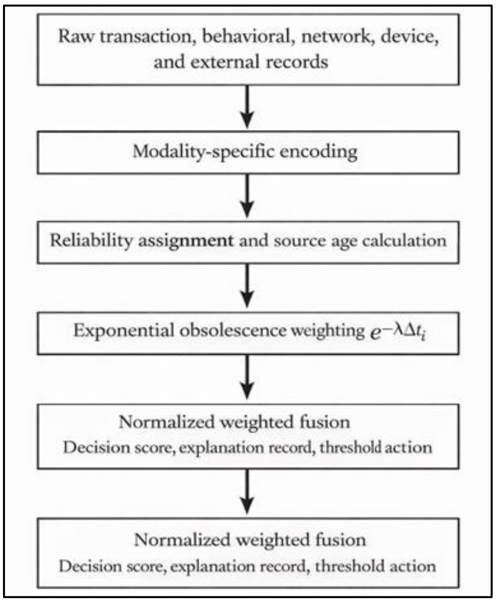

Figure 1 Conceptual flow of weighted fusion with an exponential obsolescence factor in financial transactions (adapted from [8])

From this synthesis, three design rules follow. First, fusion should occur late enough to preserve modality meaning, because graphs, tabular histories, and event bursts lose information when flattened too early [5]–[9]. Second, decay must be source-specific because financial signals expire at different speeds [3], [7], [8]. Third, the fused score should be accompanied by a contribution vector $c_i = w_i s_i$, a timestamp profile Δt_i, and the active threshold that produced the action. In that sense, the proposed methodology extends classical information-fusion logic from enterprise risk [10] into transaction-level systems with modern heterogeneous inputs. The gain lies in disciplined selectivity. A model informed by many data sources does not improve because it is larger. It improves when each source enters at the right stage, with the right decay, and with a traceable share in the final decision.

4. Discussion

The analytical reconstruction suggests that exponential obsolescence is best treated as a temporal discipline embedded within fusion. In financial transactions, one part of the evidence moves slowly and stabilizes the decision. Another part moves quickly and redirects the decision when fresh anomalies emerge. A workable methodology needs a dual internal logic. The chronic branch should retain historically persistent descriptors. The acute branch should react to short-window behavioral, device, or relational shifts. Their convergence point is the weighted fusion layer.

To operationalize this logic, the first task is to assign each source family a distinct obsolescence profile and a separate reliability rule. Table 1 compares the data families that most often coexist in transaction risk systems and shows how the fusion stage should treat them.

Table 1 Comparative treatment of source families in weighted financial data fusion

Data family	Typical content	Obsolescence speed	Reliability	Fusion function
External credit data	Bureau history, defaults, prior limits, delinquency markers	Slow	Reporting lag, coverage, and dispute frequency	Baseline risk anchor
Transaction stream	Amount, timestamp, merchant sequence, payment velocity	Fast	Event completeness, temporal coherence	Acute trigger layer
Behavioral and device data	Login rhythm, device change, session stability, channel switching	Fast	Spoof resistance, consent status, capture quality	Friction-sensitive modifier
Relational and network data	Account links, merchant communities, transfer paths, shared identifiers	Medium	Graph coverage, node volatility, edge validity	Pattern amplifier
Profile and merchant metadata	Geography, sector, tenure, KYC fields, onboarding descriptors	Medium to slow	Missingness, update discipline	Contextual stabilizer

The comparison makes one pattern clear. Fast-aging sources should influence score curvature more strongly than long-horizon baselines, yet only after minimal reliability filters are satisfied. Without that condition, a recent but weakly grounded signal can displace a slower, more trustworthy source. Slow sources anchor the decision. Fast sources redirect it. Network signals intensify or soften suspicion when the transaction is part of a relational pattern. That is the operational core of weighted fusion.

A second operational question concerns sequence. Fusion fails when preprocessing, decay estimation, thresholding, and explanation are combined into a single opaque block. Table 2 presents a practical implementation chain for real-time financial transactions.

Table 2 Implementation sequence for recency-adjusted weighted fusion in financial transactions

Stage	Main operation	Output	Monitoring metric
Ingestion and synchronization	Align timestamps, deduplicate events, and assign source IDs	Clean event packet	Ingestion delay, missing rate
Modality encoding	Build source-specific features or graph embeddings	Comparable source scores s_i	Drift by source, encoding stability
Obsolescence weighting	Compute Δt_i and source-specific decay	Weights w_i	Freshness distribution, stale-source share
Reliability adjustment	Apply α_i floors or penalties for missingness, spoof risk, or weak provenance	Adjusted weighted scores	Reliability rejection rate
Fusion and thresholding	Aggregate $c_i = w_i s_i$ and compare with decision bands	Action label	Approval, alert, and decline balance
Explanation and storage	Store c_i, Δt_ithreshold, and model version	Audit record	Explanation coverage, review turnaround

This sequence keeps monitoring close to the point where an error is produced. When drift emerges, the analyst can determine whether the change originated in source age, weak provenance, unstable encoding, or threshold policy.

Recalibration then becomes local and procedural. That is far more practical than a full model replacement cycle every time performance worsens.

Exponential obsolescence should be the default temporal rule in high-velocity financial systems, though never an autonomous decision principle. It works best inside a decision logic with three safeguards. The first is a source-specific decay map. The second is a reliability floor that blocks weak signals from dominating merely because they are recent. The third is a review band between automatic approval and automatic rejection. Such a design fits digital lending, payment fraud screening, and AML triage because it preserves speed while keeping the decision inspectable. A compact monitoring set is enough for control: stale-source share, reliability rejection rate, chronic-versus-acute contribution balance, threshold migration frequency, manual review reversal rate, and post-decision loss by source mix.

5. Conclusion

The analytical synthesis supports a clear answer. Financial transaction systems combine slow-moving, medium-horizon, and fast-aging evidence streams, and their direct aggregation distorts decision quality. The proposed solution is a weighted scheme in which each source receives its own reliability coefficient and exponential obsolescence factor. This model preserves the comparability of heterogeneous inputs while preventing temporally stale evidence from exerting disproportionate influence. The resulting methodology links modality-specific encoding, source-age weighting, threshold action, and explanation storage into one operational chain. In that form, weighted fusion is well-suited to real-time financial transactions, where speed, selectivity, and traceability must coexist.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Ayari, H., Guetari, P. R., & Kraïem, P. N. (2026). Machine learning powered financial credit scoring: A systematic literature review. *Artificial Intelligence Review, 59*, 13. https://doi.org/10.1007/s10462-025-11416-2

[2] Johannessen, F., & Jullum, M. (2025). Finding money launderers using heterogeneous graph neural networks. *The Journal of Finance and Data Science, 11*, 100175. https://doi.org/10.1016/j.jfds.2025.100175

[3] Li, W., Liu, X., Li, Z., Qin, Z., Dong, J., & Li, S. (2026). A decoupling-fusion system for financial fraud detection: Operationalizing causal–temporal asynchrony in multimodal data. *Systems, 14*(1), 25. https://doi.org/10.3390/systems14010025

[4] Sun, Y., Liu, L., Xu, Y., et al. (2024). Alternative data in finance and business: Emerging applications and theory analysis (review). *Financial Innovation, 10*, 127. https://doi.org/10.1186/s40854-024-00652-0

[5] Tian, Y., Liu, G., Wang, J., & Zhou, M. (2024). ASA-GNN: Adaptive sampling and aggregation-based graph neural network for transaction fraud detection. *IEEE Transactions on Computational Social Systems, 11*(3), 3536–3549. https://doi.org/10.1109/TCSS.2023.3335485

[6] Tong, G., & Shen, J. (2023). Financial transaction fraud detector based on imbalance learning and graph neural network. *Applied Soft Computing, 149*(Part A), 110984. https://doi.org/10.1016/j.asoc.2023.110984

[7] Wang, X., Guo, J., Luo, X., & Yu, H. (2024). DyHDGE: Dynamic heterogeneous transaction graph embedding for safety-centric fraud detection in financial scenarios. *Journal of Safety Science and Resilience, 5*(4), 486–497. https://doi.org/10.1016/j.jnlssr.2024.05.005

[8] Wei, S., & Lee, S. (2025). Internet fraud transaction detection based on temporal-aware heterogeneous graph oversampling and attention fusion network. *PLOS ONE, 20*(12), e0337208. https://doi.org/10.1371/journal.pone.0337208

[9] Wu, B., Chao, K.-M., & Li, Y. (2024). Heterogeneous graph neural networks for fraud detection and explanation in supply chain finance. *Information Systems, 121*, 102335. https://doi.org/10.1016/j.is.2023.102335

[10] Yue, H., Liao, H., Li, D., & Chen, L. (2021). Enterprise financial risk management using information fusion technology and big data mining. *Wireless Communications and Mobile Computing*, 2021, 3835652. https://doi.org/10.1155/2021/3835652