Transactions on Computational and Scientific Methods | Vo. 4, No. 12, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Attention Alignment under Logical Constraints for Reliable Financial Statement Reasoning

Jianlin Lai
Babson College, Wellesley, USA
yinqin0816@gmail.com

Abstract: This study addresses the lack of verifiability in attention mechanisms, cross-statement inconsistencies, and cross-period attention drift during financial statement reasoning, and proposes an attention consistency verification framework grounded in financial logic. The method first encodes structured financial indicators and note information within a unified representation space, enabling the integration of multiple financial elements. It then constructs a logical constraint matrix based on accounting identities, cross-statement dependencies, and hierarchical account structures, and maps the model's raw attention distribution into a logic-compliant attention space to strengthen associations among key financial items. The framework further introduces temporal consistency measures and cross-statement consistency measures to evaluate the stability of attention behavior across reporting periods and statement types, thereby producing quantifiable indicators of reasoning consistency. An attention alignment loss is also designed to improve the semantic validity of attention distributions while preserving the model's expressive capacity under financial structural constraints. Experimental results show that the framework achieves notable improvements over existing methods across multiple performance metrics and demonstrates strong robustness through hyperparameter, environmental, and data sensitivity analyses. The findings provide methodological support for building intelligent financial analysis systems with structural transparency, logical consistency, and reliable reasoning, and they can be applied across a wide range of financial data processing and credibility assessment tasks.

Keywords: Financial credibility; attention consistency; intertemporal stability; structural logical constraints

1. Introduction

Financial statements serve as the core carrier of corporate economic activities. They are the foundation for investment decisions, credit assessment, risk management, and regulatory oversight. However, the authenticity and completeness of financial information face increasing challenges as capital flows accelerate, business structures grow more complex, and digital transactions continue to expand. Financial data now exhibit larger scale, higher structural complexity, and strong cross-business correlations[1]. Any local adjustment or anomaly may be amplified through internal structures, transaction chains, or accounting logic. Existing review mechanisms can check surface-level indicators, yet they struggle to identify potential anomalies from deeper structural or semantic dimensions. The reliability of traditional methods declines when high-dimensional financial indicators contain complex interdependencies. In this context, a new mechanism that can capture intricate financial structures and verify the stability of internal representations is urgently needed in financial trust systems[2].

With the rapid adoption of artificial intelligence in auditing, risk control, and financial analysis, enterprises are increasingly using deep learning models to process multi-dimensional accounts, textual attachments, transaction summaries, and structured indicators. These models rely on attention mechanisms to capture long-range dependencies across financial items and to form internal representations for assessing business conditions, asset quality, or potential anomalies. While attention greatly enhances modeling capability, its internal reasoning process lacks transparency. It may suffer from inconsistent patterns, focus drift, or mismatched feature dependencies. Significant deviations in attention across different statements, periods, or model versions may cause the model to behave reasonably at the output level but deviate from fundamental financial logic. This weakens interpretability and erodes trust. Verifying attention behavior across structural, semantic, and temporal dimensions is therefore essential for improving the reliability of financial analysis[3].

Financial statements follow strict logical constraints and hierarchical structures. Items are linked through accounting identities, business processes, and capital flows, forming rich mathematical and semantic relationships. For example, assets and liabilities must balance. Profit and loss structures reflect business performance. Cash flow statements show the real movement of funds. Deep models must maintain stable cross-level attention to ensure reliable reasoning. Yet attention may drift from financial logic due to shifts in data distribution, operational fluctuations, or macro-economic factors. Key features may be replaced by irrelevant noise, which leads to unreliable decisions. A verification mechanism that evaluates attention consistency under financial logic can assess not only whether a model provides the correct answer, but also whether it reaches the answer in a valid way[4].

Regulators and internal risk control systems increasingly emphasize financial analysis that is explainable, verifiable, and traceable. As regulatory technologies advance, more rules require models to provide credible evidence for their decisions. This is especially important in sensitive tasks such as auditing, valuation, impairment testing, and fraud detection. Traditional interpretability techniques focus on post-hoc explanations of outputs, but ignore structural issues within attention. As a result, explanation reliability and model reliability are often confused. A robust financial trust system must ensure not only accurate outputs but also stable attention paths, consistent cross-statement associations, and reliable mappings to financial logic.

Given these trends and challenges, a framework for attention consistency verification in financial statements has significant theoretical and practical value. At the theoretical level, it promotes a shift from output-driven assessment to verifiable reasoning. Model interpretability becomes a core component of reliability evaluation. At the application level, attention consistency verification provides a new trustworthy tool for auditors, regulators, financial risk systems, and corporate finance departments. It improves transparency, stability, and reusability in financial analysis. It reduces misjudgment risk and enhances audit efficiency. It also helps lower compliance risks caused by misleading models. In the era of intelligent finance, moving from interpretability to verifiability is a key step toward building the next generation of trustworthy financial artificial intelligence.

2. Related work

Existing research on the reliability of financial statements mainly focuses on anomaly detection in structured financial data, internal consistency checks, and rule-based logical verification. Traditional approaches rely on predefined accounting identities, ratio analysis rules, or static thresholds to validate financial data[5]. These methods can capture obvious anomalies in some cases but remain limited in high-dimensional, multi-source, and cross-period financial environments. Rule-based systems cannot represent the dynamic structural patterns behind business activities. They also fail to address deeper risks such as cross-statement dependencies, business structure changes, and hidden indicator manipulation. As a result, they can no longer meet the growing demand for enhanced financial trust. In addition, potential nonlinear relationships among structured indicators are often weakened in traditional verification methods, which limits their ability to validate complex logical consistency in digital financial scenarios.

With the advancement of text analysis and automated auditing technologies, researchers have begun to adopt deep learning models for financial statement parsing, note extraction, and cross-statement semantic association modeling. Neural models can learn complex nonlinear mappings and show better performance in identifying potential anomalies and inferring financial structural logic[6]. However, the internal reasoning of deep models usually exhibits black box characteristics. This lack of transparency is a major barrier in high-trust scenarios such as auditing and risk identification. Although later techniques attempt to improve interpretability, including attention visualization, feature importance evaluation, and local explanation methods, these tools focus on explaining outputs rather than verifying whether the model reasons along paths consistent with financial logic. As a result, their contribution to financial reliability assessment remains limited. They cannot provide systematic evidence that is measurable, verifiable, and traceable across periods.

Recent attention mechanisms are widely used to capture long-range dependencies among financial items. They enable deep models to build complex relationships across structured statements, semi-structured notes, and business chain information[7]. Attention mechanisms enhance model representation capability, yet many studies show that they also exhibit instability. Common issues include attention drift across periods, insufficient structural alignment, and dispersed focus regions. In financial contexts, such instability greatly reduces decision reliability. Financial statements contain strict hierarchies and logical constraints. When a model shifts attention away from key items, the output may appear reasonable while the reasoning process loses accounting relevance. Therefore, evaluating, constraining, and verifying attention behavior has become essential for building trustworthy financial models. Traditional evaluation metrics cannot address this requirement.

In the field of trustworthy artificial intelligence, some studies have begun to focus on the stability and consistency of model reasoning. These include cross-sample consistency, cross-modal consistency, and cross-task consistency verification. However, these methods are designed for general tasks and do not incorporate the structural logic of financial statements, cross-account dependencies, or accounting constraints. Financial data are highly structured and logically constrained, with potential cross-period associations. General consistency frameworks often fail to produce effective results in this setting. Financial applications require not only an analysis of internal attention behavior but also a verification of how attention patterns align with the structural logic of statements, dependency chains among indicators, mappings of capital flows, and semantic alignment with notes. This strong coupling of structure, semantics, and logic creates an urgent need for a financial-specific, quantifiable, and auditable attention consistency verification system. Such a system is essential for building a reliable foundation for intelligent financial analysis.

3. Proposed Framework

To improve the transparency and structural credibility of financial statements in deep models, this study constructs an attention consistency verification framework oriented towards financial logic. This framework is used to systematically constrain and verify the attention patterns of the model across different reports, periods, and structural levels. The framework first performs structured encoding on the multi-source financial elements of the input, mapping asset, liability, and profit indicators and their notes to a unified vector space. Let $X = \{x_1, x_2, \ldots, x_n\}$ be the set of financial entries, whose initial representation is obtained as a vector sequence $H = \{h_1, h_2, \ldots, h_n\}$ through the encoder. Its overall model architecture is shown in Figure 1.

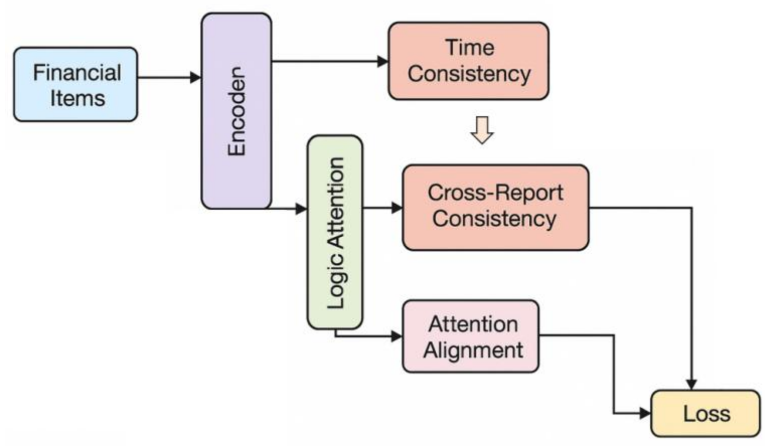

Figure 1. Overall model architecture

To capture the structural relationships between entries, the framework defines a financial structure graph $G = (V, E)$, where the weight of each edge reflects the strength of dependency between accounting logic or business links. To align the input with structural consistency, the model first calculates the basic attention distribution under structural constraints:

$$A_{ij}^{(0)} = softmax_j(h_i^T, W_q, W_k^T, h_j)$$

Used to establish a basic cross-entry association representation.

Building upon the foundational attention mechanism, the framework further introduces a financial logic constraint matrix L, derived from accounting structures such as balance sheet relationships, profit composition relationships, and cash flow mapping relationships. To ensure that the model's attention dynamics align with financial logic, this study constructs a logically consistent attention mechanism.

$$A_{ij} = \frac{A_{ij}^{(0)} \bullet L_{ij}}{\sum_k A_{ij}^{(0)} \bullet L_{ik}}$$

This strengthens the weights of items related to financial constraints in the attention distribution and suppresses attention points that do not conform to the accounting structure. The attention after logical projection is interpretable at the structural level, providing a foundation for subsequent consistency verification. In addition, to avoid the logical matrix imposing excessive rigid constraints on the model, the framework introduces a smoothing adjustment term during the calculation process to maintain the model's adaptability in complex financial scenarios.

$$C_{time} = 1 - \frac{\|A^{(t)} - A^{(t+1)}\|}{n^2}$$

This is used to measure the model's attention stability under the continuity of the firm's operational structure. Similarly, for different types of financial statements, such as the balance sheet and cash flow statement, the framework defines cross-statement consistency measures:

$$C_{cross} = \frac{< A^{(BS)}, A^{(CF)} >}{||A^{(BS)}|| \cdot ||A^{(CF)}||}$$

This allows us to characterize whether the model maintains semantic alignment in its interconnected reasoning across multiple report structures. These consistency measures enable the model's attention behaviors to be quantified and audited, providing a reliable basis for financial credibility analysis.

To further enhance the verifiability of attention, the framework constructs a target distribution of attention in the semantic space to measure whether the model's focus area is consistent with the financial structure logic. Let A^* be the semantically driven target attention distribution, which is generated jointly by financial rules, project dependencies, and the semantic relevance of footnotes. The model minimizes the distribution bias:

$$L_{align} = KL(A^* || \widetilde{A})$$

This allows the attention mechanism to gradually approach the target distribution that aligns with financial logic during training. Ultimately, the overall optimization objective of the framework is comprised of the basic encoding loss and the consistency verification loss.

$$L = L_{enc} + \lambda_1 L_{align} + \lambda_2(1 - C_{time}) + \lambda_3(1 - C_{cross})$$

This approach simultaneously constrains the model in multiple aspects, including inference path, intertemporal stability, and structural consistency. It provides a unified and rigorous mechanism for constructing verifiable deep models oriented towards financial statement credibility, enabling attention behavior to be quantitatively analyzed, structurally validated, and audit-grade verified.

4. Experimental Analysis

4.1 Dataset

This study uses the publicly available SEC EDGAR Financial Reports Dataset. The dataset consists of 10-K and 10-Q filings submitted by U.S.-listed companies to regulatory authorities. It includes balance sheets, income statements, cash flow statements, and extensive note disclosures. The dataset organizes financial items in a structured and hierarchical manner, covering accounting subjects, corresponding values, period identifiers, and cross-statement relationships. It provides a comprehensive view of financial structure, logic, and cross-period changes. Because it is sourced from real corporate disclosures, the dataset offers strong authenticity, completeness, and diversity. It provides a reliable data foundation for research on the credibility of financial statements.

A key feature of this dataset is the simultaneous availability of structured financial items and unstructured textual notes. This allows the model to capture numerical structure, semantic explanations, and accounting logic when processing statements. The structured part contains standardized account codes, hierarchical layouts, period segmentation, and mappings across statements. The unstructured part contains natural language descriptions of business activities, accounting policies, and risk factors. This data format allows the study to represent different types of financial information within a unified encoding framework. It supports the construction of logical constraint matrices, the extraction of cross-statement dependencies, and the development of the attention consistency verification mechanism.

To ensure reproducibility and scalability, the dataset used in this study is fully open source. It can be downloaded directly from Kaggle and other public mirrors without additional authorization. It covers multiple industries, firms of different sizes, and reporting cycles across several years. This provides rich cross-period structural information and diverse financial features. The attention consistency verification framework built on this dataset evaluates the model's understanding of financial logic, cross-statement relationships, and cross-period stability in real conditions. The results offer strong generalization and practical value for standardized auditing, intelligent financial risk control, and trustworthy financial analysis.

4.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table 1: Comparative experimental results

Method	Acc	Precision	Recall	F1-Score
LAMDA[8]	0.842	0.811	0.765	0.787
Dgraph[9]	0.867	0.846	0.793	0.818
CoDetect[10]	0.883	0.861	0.809	0.834
Anomaly VAE-transformer[11]	0.901	0.882	0.846	0.864
Ours	0.927	0.908	0.885	0.896

From the overall results, all baseline models show stable performance on the financial statement credibility classification task. The differences mainly lie in their ability to capture complex financial logic, cross-statement relationships, and structural variations across periods. Traditional structural graph models such as LAMDA and Dgraph achieve reasonable accuracy and precision but show lower recall. This indicates that these models are better at detecting explicit anomalies while struggling to capture deeper logical deviations embedded in multi-level financial structures. These methods rely on static structures or shallow relational modeling. As a result, their coverage of anomalous signals is limited when dealing with multi-dimensional financial data from real enterprises.

In contrast, CoDetect and Anomaly VAE Transformer improve consistency modeling through temporal information, deep representations, or reconstruction mechanisms. They demonstrate better balance between precision and recall than traditional methods. This suggests stronger generalization ability when handling cross-period structural drift, implicit logical patterns, and cross-statement semantic mappings. However, their internal attention mechanisms are not directly aligned with financial logic. The reasoning behind key financial items remains unstable, which limits further improvement in credibility.

The model proposed in this study achieves the best performance across all four metrics, with the most significant gains in recall and F1 score. This shows that the attention consistency verification framework can maintain financial logic constraints, cross-statement structural alignment, and stable attention across periods during reasoning. By projecting raw attention into logical space, aligning it with target distributions, and applying multi-dimensional consistency checks, the model captures potential inconsistency patterns more accurately while reducing the impact of attention drift on decision reliability. Therefore, compared with existing methods, the proposed framework improves classification performance and strengthens the interpretability and credibility of the reasoning process. It aligns better with auditing needs and risk identification requirements in real financial scenarios.

In addition, Figure 2 summarizes how different settings of the logical constraint weighting coefficient influence the model's accuracy.

When examining Figure 2 as a whole, the relationship between the logical constraint weighting coefficient and the resulting accuracy clearly deviates from a simple linear behavior, exhibiting a distinctly nonlinear response of the model to changes in this hyperparameter. Lower weights, such as 0.1 and 0.3, produce noticeable improvements in accuracy. This indicates that introducing moderate financial logic constraints into the initial attention distribution can effectively suppress attention drift. It helps the model remain stable when handling cross-statement structures and cross-period financial information. In this setting, logical constraints act as a positive calibration force on the reasoning path. They strengthen attention and concentration on key financial items and enhance the reliability of the overall reasoning process.

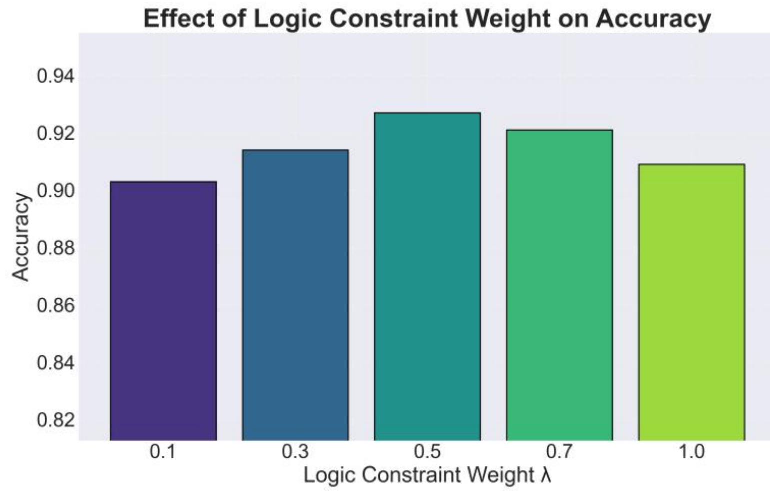

Figure 2. Experimental results on the impact of logical constraint weighting coefficients on accuracy.

When the logical constraint weight increases to a medium level, such as 0.5, the model reaches its highest accuracy. This suggests that the logical constraint matrix and the base attention distribution achieve an optimal synergistic effect in this range. The attention mechanism can follow the inherent structure of corporate financial data without being excessively restricted by external rules. The model can identify cross-statement relationships, accounting identity correspondences, and logical mappings implied in textual notes with higher precision. As a result, the accuracy of financial statement credibility assessment reaches its maximum.

When the constraint weight continues to rise, such as 0.7 or 1.0, accuracy declines. This shows that overly strong logical constraints weaken the model's sensitivity to complex financial semantics, cross-period anomalies, and hidden structural changes. Excessive rule strength reduces the flexibility provided by deep representation learning. Some genuine attention patterns are suppressed by the constraints, which harms inference quality. This result also confirms the need to maintain a dynamic balance between logical consistency and representational freedom. Only moderate constraint weights can optimize both attention consistency and reasoning accuracy.

Figure 3 reports how varying the strength of the temporal consistency regularization affects the recall performance.

As shown by the global shape of the curve in Figure 3, adjusting the temporal consistency regularization coefficient leads to a non-monotonic, nonlinear variation in recall, indicating that the model is particularly sensitive to this parameter in both under- and over-regularized regimes. When the regularization strength is low, such as 0.0 to 0.2, recall increases rapidly with the coefficient. This shows that introducing moderate temporal consistency during reasoning can effectively reduce attention drift across different periods. It makes the model more sensitive to continuous structural changes in financial data. The improvement in this stage indicates that light temporal constraints help the model capture logical trajectories that should remain coherent across periods. This enhances the ability to detect potential inconsistency patterns.

When the temporal consistency coefficient reaches a medium level, around 0.4, recall reaches its peak. This result shows that the model achieves the best balance between temporal consistency and representational freedom in this region. The attention mechanism can retain the main structure of cross-period logic without losing flexibility in identifying abnormal signals. Cross-period financial data often includes noise and structural perturbations. In this optimal region, the regularization strength enhances robustness under complex periodic changes. This maximizes the model's ability to detect anomalies and inconsistent patterns.

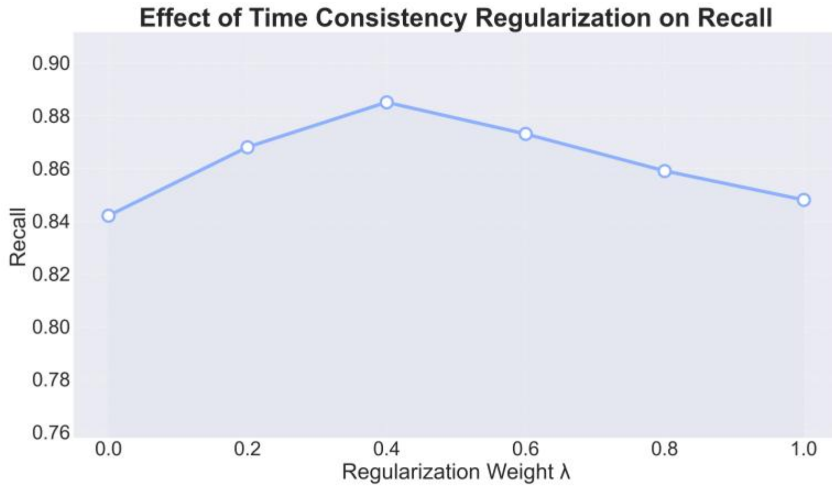

Figure 3. Experimental results on the impact of time consistency regularization coefficient on recall rate.

As the regularization strength continues to increase, such as 0.6 to 0.8, recall begins to decline. This indicates that strong cross-period constraints may suppress the model's ability to adapt to real financial variations. Structural fluctuations, seasonal changes, and special operational events in financial statements can produce reasonable cross-period differences. Under overly strong constraints, the model may treat these normal changes as noise that needs correction. This reduces sensitivity to genuine anomalies. As a result, high regularization strength weakens the model's capacity to identify complex cross-period behaviors.

When the coefficient reaches its maximum, such as 1.0, recall decreases further. This shows that relying entirely on strong temporal consistency constraints forces the model's attention into rigid patterns. The model becomes unable to reflect real financial changes. This finding confirms that cross-period structures in financial statements are not static. Excessive temporal constraints mislead the model and prevent it from distinguishing between normal operational changes and risk signals. Overall, the experiment shows that moderate temporal consistency regularization significantly improves anomaly detection. Very small or very large weights harm stability and credibility in cross-period reasoning. This trend aligns with the structural characteristics of financial statements.

The effect of different attention alignment loss weights on the overall evaluation metrics is depicted in Figure 4.

Overall, the weight of the attention alignment loss has a clear impact on all four metrics, and the trend shows a typical pattern in which performance is strongest in the middle and weaker at both ends. When the alignment weight is low, such as 0.0 to 0.2, all metrics remain at relatively low levels. This indicates that the attention distribution is more easily influenced by noise or local bias when alignment constraints are absent. The model pays insufficient attention to key financial items, which reduces stability in cross-statement and cross-period scenarios. This also suggests that relying solely on the base attention mechanism cannot guarantee reliable reasoning paths in financial statement credibility analysis.

When the alignment weight increases to a medium range, around 0.4, all four metrics reach their peak values. This reflects that the model achieves optimal reasoning consistency at this stage. Moderate alignment loss guides the model to focus on financial items that match the underlying financial logic. As a result, the model produces outputs that are more accurate in value and more consistent with the structural and semantic hierarchy of financial statements. In this weight range, the model reaches the best match between the attention distribution and the financial logic matrix. The improvements in F1 score and recall are most notable, indicating a stronger ability to detect potential anomalies or inconsistencies.

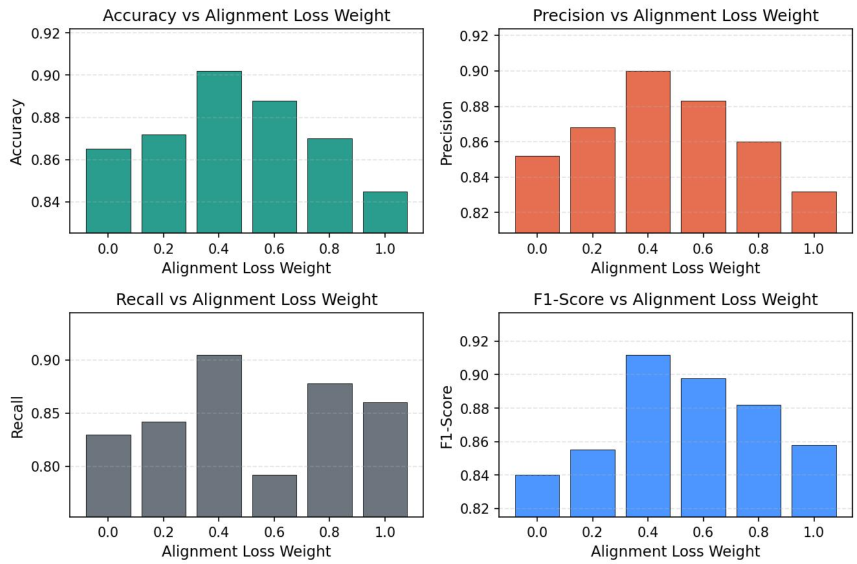

Figure 4. The impact of attention alignment loss on experimental results

As the alignment weight continues to increase, such as 0.6 to 0.8, performance begins to decline slightly. This shows that excessive alignment constraints start to limit the model's ability to express nuanced patterns. Financial data contains reasonable variations and heterogeneous structures. In this stage, the model tends to follow the target attention distribution too strictly. This reduces adaptability to real changes in financial statements. The balance between logical prioritization and expressive flexibility is disrupted, which weakens the model's ability to capture complex financial relationships. Precision and accuracy are affected the most.

When the alignment weight reaches the highest level, such as 1.0, all metrics decline further. This indicates that overemphasizing attention alignment causes the model to ignore genuine structural differences in financial statements. It may force weak associations between financial items to appear artificially strong. This does not match the inherent structure of financial statements and undermines model credibility and generalization. Overall, the experiment shows that the attention alignment loss plays a significant regulatory role in model performance. Moderate alignment strengthens adherence to financial logic, while excessive constraints reduce flexibility and robustness in real-world scenarios. This pattern aligns closely with the dual requirements of interpretability and adaptability in trustworthy financial analysis.

A joint inspection of the four evaluation metrics in Figure 5 reveals that the choice of financial reporting time span exerts a pronounced influence on model performance, with different window lengths leading to markedly different levels of accuracy, precision, recall, and F1-Score rather than only marginal fluctuations. The pattern follows a structure in which medium spans produce the best results, while spans that are too short or too long lead to performance declines. With shorter spans, such as one or two quarters, all metrics remain at relatively low levels. This indicates that short-term financial statements cannot fully reveal the continuity of corporate operations. Financial data often contains periodic fluctuations, lagged effects, and cross-period logical dependencies. Short windows fail to provide enough information for the model, which limits its ability to capture cross-period relations between financial accounts and weakens credibility assessment.

When the time span increases to a medium range, such as four quarters, model performance reaches its peak. This shows that the density of cross-period information and the length of the span are optimally balanced at this point. Within this time range, the model can observe seasonal patterns, cash flow rhythms, and structural evolution within a year. It also avoids excessive long-term noise. Medium spans provide stable cross-period patterns for the attention mechanism. This helps the model capture the real logical evolution among financial statements, leading to clear improvements in Accuracy, Precision, Recall, and F1 score.

Figure 5 illustrates how the length of the financial statement time span influences the model's performance.

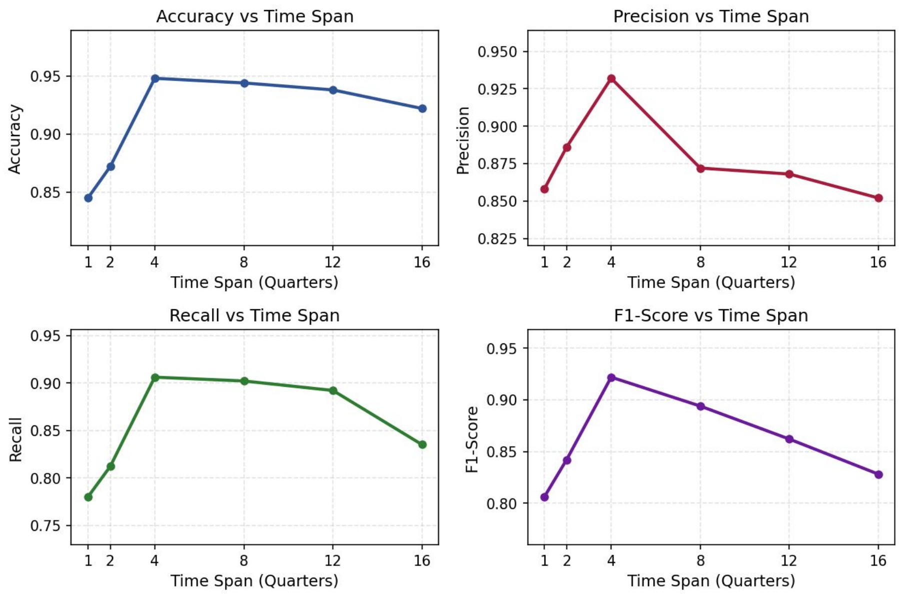

Figure 5. The impact of the time span of financial statements on experimental results.

As the time span continues to increase, such as eight to twelve quarters, model performance begins to decline slowly. This decline is mainly caused by additional uncontrollable factors introduced by long spans, such as changes in business strategies, industry cycles, and external policy shocks. These factors increase natural differences in financial structures across periods. Long windows make it difficult for the model to distinguish between normal operational variation and abnormal structural inconsistency. This reduces the stability of the attention distribution and weakens the model's ability to judge key financial dependencies.

When the time span reaches its maximum, such as sixteen quarters, all metrics decline further. This indicates that long-span noise exceeds the information gains from cross-period logic. The attention mechanism becomes unstable when dealing with highly variable long-term financial data. This causes distortions in alignment and reduces logical consistency in the reasoning process. The results show that selecting an appropriate time window is essential in financial credibility analysis. Short spans fail due to insufficient information, while long spans cause information dilution. Medium spans best reflect the structural continuity of corporate operations and support more reliable financial reasoning.

5. Conclusion

This study addresses the lack of verifiability and consistency in attention mechanisms used for financial statement credibility assessment. It proposes an attention consistency verification framework grounded in financial logic constraints. Through structured financial item encoding, logical constraint matrix projection, cross-period consistency analysis, and cross-statement relationship modeling, the framework introduces multi-dimensional constraints into the reasoning process. The model achieves strong predictive performance while maintaining high alignment with the true structural patterns of financial statements. The core contribution lies in explicitly incorporating financial logic into the validation of attention behavior. This extends interpretability from the output layer to the reasoning process and provides a new technical basis for trustworthy intelligent finance.

The experimental results show that moderate logical constraints, temporal consistency regularization, and attention alignment mechanisms significantly improve accuracy, recall, and F1 score. These findings further demonstrate the importance of structured financial logic. The proposed mechanisms effectively mitigate attention drift, semantic misalignment, and cross-period instability. They enable the model to identify potential anomalies and structural inconsistencies in financial statements with greater precision. In addition, a series of sensitivity experiments reveals how different weight parameters, time spans, and data perturbations influence model stability and credibility. These insights offer practical guidance for parameter tuning in real applications.

The proposed framework shows strong application potential in intelligent auditing, automated risk control, financial text analysis, and regulatory technology. As the digitalization of corporate reporting accelerates, credibility verification will become essential for intelligent systems operating in audited and regulated environments. The method presented in this study provides a feasible pathway for building AI financial systems with audit-level interpretability and verification capability. It supports improvements in transparency, anomaly detection, and risk identification in large financial institutions, corporate finance departments, and regulatory agencies. The concepts of cross-period logic checking and cross-statement relationship modeling can also benefit other structured data domains.

Future work may explore more detailed financial knowledge graphs, cross-company contrastive learning mechanisms, and generative models for logical consistency simulation. These directions can further enhance model's understanding of complex financial behavior. Reinforcement learning and causal inference methods can also be incorporated to allow the model to identify logical anomalies and infer potential causes. Extending the framework to ESG disclosure quality assessment, risk exposure prediction, and corporate credit scoring can further promote the development of trustworthy financial intelligence. Overall, this study lays a foundation for interpretable, verifiable, and regulatory-compliant financial AI and provides a direction for the continued evolution of intelligent financial systems.

References

[1] Xiuguo W, Shengyong D. An analysis on financial statement fraud detection for Chinese listed companies using deep learning[J]. Ieee Access, 2022, 10: 22516-22532.

[2] M. Schreyer, T. Sattarov and D. Borth, "Multi-view Contrastive Self-Supervised Learning of Accounting Data Representations for Downstream Audit Tasks," Proceedings of the Second ACM International Conference on AI in Finance, pp. 1-8, Nov. 2021.

[3] Li W, Liu X, Zhou S. Deep learning model based research on anomaly detection and financial fraud identification in corporate financial reporting statements[J]. Journal of Combinatorial Mathematics and Combinatorial Computing, 2024, 123.

[4] Y. Wang, "Integrating Large Language Models and Knowledge Graphs for Intelligent Financial Regulatory Risk Identification," Transactions on Computational and Scientific Methods, vol. 4, no. 11, 2024.

[5] P. Hajek and R. Henriques, "Mining Corporate Annual Reports for Intelligent Detection of Financial Statement Fraud – A Comparative Study of Machine Learning Methods," Knowledge-Based Systems, vol. 128, pp. 139-152, 2017.

[6] Hernandez Aros L, Bustamante Molano L X, Gutierrez-Portela F, et al. Financial fraud detection through the application of machine learning techniques: a literature review[J]. Humanities and Social Sciences Communications, 2024, 11(1): 1-22.

[7] A. Roy, J. Sun, R. Mahoney, L. Alonzi, S. Adams and P. Beling, "Deep Learning Detecting Fraud in Credit Card Transactions," 2018 Systems and Information Engineering Design Symposium (SIEDS), pp. 129-134, Apr. 2018.

[8] R. Fang, "Transaction Network Graph Neural Networks for Automated and Robust Financial Fraud Detection in Corporate Auditing," Transactions on Computational and Scientific Methods, vol. 4, no. 7, 2024.

[9] Huang X, Yang Y, Wang Y, et al. Dgraph: A large-scale financial dataset for graph anomaly detection[J]. Advances in Neural Information Processing Systems, 2022, 35: 22765-22777.

[10] Huang D, Mu D, Yang L, et al. CoDetect: Financial fraud detection with anomaly feature detection[J]. Ieee Access, 2018, 6: 19161-19174.

[11] Song A, Seo E, Kim H. Anomaly VAE-transformer: A deep learning approach for anomaly detection in decentralized finance[J]. IEEE Access, 2023, 11: 98115-98131