“Agilance: An intelligent strategic control and financial planning system for data-driven environments”

AUTHORS	Georgios Kampiotis Georgios L. Thanasas Iryna Zhyhlei Iryna Hrabchuk Iryna Zhalinska

ARTICLE INFO	Georgios Kampiotis, Georgios L. Thanasas, Iryna Zhyhlei, Iryna Hrabchuk and Iryna Zhalinska (2026). Agilance: An intelligent strategic control and financial planning system for data-driven environments. *Accounting and Financial Control*, *7*(2), 60-88. doi:10.21511/afc.07(2).2026.05
DOI	http://dx.doi.org/10.21511/afc.07(2).2026.05
RELEASED ON	Tuesday, 11 August 2026
RECEIVED ON	Sunday, 15 March 2026
ACCEPTED ON	Monday, 22 June 2026
LICENSE	This work is licensed under a Creative Commons Attribution 4.0 International License
JOURNAL	"Accounting and Financial Control"
ISSN PRINT	2543-5485
ISSN ONLINE	2544-1450
PUBLISHER	LLC “Consulting Publishing Company “Business Perspectives”
FOUNDER	Sp. z o.o. Kozmenko Science Publishing

NUMBER OF REFERENCES	NUMBER OF FIGURES	NUMBER OF TABLES
38	12	12

BUSINESS PERSPECTIVES

LLC "CPC "Business Perspectives"
Hryhorii Skovoroda lane, 10,
Sumy, 40022, Ukraine

www.businessperspectives.org

Georgios Kampiotis (Greece), **Georgios L. Thanasas** (Greece), **Iryna Zhyhlei** (Ukraine), **Iryna Hrabchuk** (Ukraine), **Iryna Zhalinska** (Ukraine)

AGILANCE: AN INTELLIGENT STRATEGIC CONTROL AND FINANCIAL PLANNING SYSTEM FOR DATA-DRIVEN ENVIRONMENTS

Abstract

This study proposes Agilance as a conceptual and technical framework for explainable strategic financial planning in data-intensive organizational environments. The framework is based on a custom transformer architecture that incorporates three sector-specific components: Financial Relevance Weighting, Context Shift Stabilization, and Output Compression. Because the evaluation was conducted on a confidential sector-specific dataset and through internal benchmarking procedures, the underlying source data and job-level operational records cannot be publicly released. Within these constraints, the internal evaluation yielded indicative results: 96.03% accuracy and 95.8% F1-score for priority classification on the held-out test set, 90.26% accuracy and 90.22% F1-score for implementation-duration classification, and an average 10-fold cross-validation accuracy of 91.7%. The expert explainability assessment produced mean scores of 4.67 for clarity, 4.53 for trustworthiness, and 4.48 for actionability, with inter-rater agreement ranging from 0.87 to 0.91. Internal operational benchmarks further suggested planning-cycle reductions and economic benefits, including 95.8% improvement in real-time data analysis and time-zone synchronization, 5.4% operational cost savings, and an illustrative first-year ROI of 46%. These results should be interpreted as preliminary internal evidence obtained under specific evaluation conditions, not as independently verified proof of broad organizational generalizability. The study contributes an auditable AI-supported framework and identifies the need for future validation using anonymized multi-organizational datasets, externally audited protocols, or independently reproducible benchmarks.

Type of the article: Research Article

Received on: 15th of March, 2026
Accepted on: 22nd of June, 2026
Published on: 11th of August, 2026

Georgios Kampiotis, Doctoral Student, Department of Management Science and Technology, University of Patras, Greece.

Georgios L. Thanasas, Associate Professor, Department of Management Science and Technology, University of Patras, Greece.

Iryna Zhyhlei, Doctor of Economics, Professor, Professor of the Department of Information Systems in Management and Accounting, Zhytomyr Polytechnic State University, Ukraine. (Corresponding author)

Iryna Hrabchuk, Doctor of Economics, Associate Professor, Head of the Department of International Relations and Political Management, Zhytomyr Polytechnic State University, Ukraine.

Iryna Zhalinska, Ph.D. in Economics, Assistant Professor, Assistant Professor of the Department of Psychology and Social Welfare, Zhytomyr Polytechnic State University, Ukraine.

Conflict of interest statement:
Author(s) reported no conflict of interest

Keywords agentic AI, explainability, auditability, prioritization, decision-making, transformers, Industry 4.0, automation

JEL Classification M41, C55, O33

INTRODUCTION

The increasing complexity of global markets and the rapid pace of digital transformation have fundamentally altered the strategic environment in which organizations operate. Traditional financial planning, typically based on year-end reports and fixed annual targets, is increasingly inadequate in the context of rapid digital transformation and volatile markets (Xiaoqi & Ali, 2024). Traditional management accounting systems, rooted in static budgeting and periodic variance analysis, often fail to capture the high-frequency shifts in global supply chains and currency fluctuations. As innovation cycles accelerate, organizations require real-time financial decision-making and agile operational frameworks to remain competitive (Malik, 2024; Olomu et al., 2023). Research confirms that integrating digital transformation with research and development activities improves financial performance, highlighting the structural need for more adaptive and responsive planning approaches (Lehenchuk et al., 2025a). Accordingly, accounting systems must evolve to support innovation-driven financ-

http://dx.doi.org/10.21511/afc.07(2).2026.05

ing and dynamic economic conditions, particularly within Industry 4.0 environments characterized by technological complexity and uncertainty (Lehenchuk et al., 2020; Lehenchuk et al., 2025b).

Despite growing recognition of these challenges, existing analytical tools remain insufficient. Conventional analytics platforms primarily support data visualization and prediction but lack the capacity to autonomously generate comprehensive financial strategies (Balogun et al., 2022). At the same time, hierarchical organizational structures and multi-level approval processes continue to slow strategic responsiveness, creating a critical gap between the availability of financial data and the speed of decision-making (Ren, 2022). The proposed Agilance framework addresses this gap by automating internal control procedures, thereby reducing the human factor in cost prioritization.

Recent advances in generative artificial intelligence, particularly large language models such as GPT, Gemini, Copilot, and Grok, demonstrate strong potential to automate complex financial analysis (Kerr et al., 2025; Li et al., 2023). Nevertheless, significant limitations remain regarding domain specificity, contextual awareness, real-time data integration, and explainability, all of which are essential requirements in professional accounting and auditing contexts (Li & Zhang, 2025; Yeo et al., 2025). Modern auditing standards require not just speed, but the traceability of financial logic, making the "black-box" nature of current AI tools a significant barrier to their professional adoption. Instead of replacing the accountant, such systems should serve to expand professional judgment by providing verifiable insights for internal audit purposes.

This limitation is particularly important for large organizations that require extensive data usage, such as multinational enterprises, financial service sharing centers, and organizations operating across multiple regions, currencies, regulatory environments, and approval levels. This study does not focus on routine bookkeeping or statutory reporting, but rather on strategic financial planning processes that require the integration of internal financial records, external market signals, scenario analysis, prioritization of financial actions, and cross-validation by management.

To address this gap, this study proposes Agilance, an agentic artificial intelligence framework for real-time strategic financial planning that integrates multi-source market data with a customized transformer architecture enhanced with three domain-specific layers designed to improve interpretability, temporal reasoning, and contextual adaptability. The framework enables autonomous generation of prioritized financial strategies through a dual-headed classification mechanism based on strategic priority and implementation urgency. This study aims to design and evaluate a prototype strategic financial planning framework supported by Artificial Intelligence under controlled and confidential organizational data conditions. The evaluation focuses on three dimensions: classification performance for strategic prioritization, interpretability of generated recommendations, and indicative impacts on operational efficiency. The study does not claim universal causal validation across all organizational contexts.

1. THEORETICAL BASIS

The development of an intelligent accounting framework for autonomous financial planning requires grounding in three interconnected theoretical domains: the principles of accounting support for innovative enterprise financing, the capabilities and limitations of transformer-based artificial intelligence in financial contexts, and the requirements of agile strategic management in Industry 4.0 environments. Together, these theoretical pillars establish the conceptual foundation upon which the Agilance architecture is constructed.

1.1. Accounting support for innovation-driven financial planning

The theoretical basis for intelligent financial planning systems is rooted in accounting principles that support the development of innovative enterprise financing. As established by Lehenchuk

http://dx.doi.org/10.21511/afc.07(2).2026.05

et al. (2020), such principles provide the essential framework for holistic planning and dynamic financial management, emphasizing the need for systems capable of processing multi-dimensional variables in real time. This theoretical orientation is particularly relevant in Industry 4.0 environments, where technological complexity, digital transformation, and market volatility impose continuous adaptive demands on organizational accounting systems (Lehenchuk et al., 2025b).

Contemporary research confirms that organizations integrating digital transformation with research and development activities achieve measurably superior financial performance (Lehenchuk et al., 2025a). However, the theoretical literature also identifies a persistent structural gap: conventional accounting and analytics frameworks are primarily designed for retrospective analysis and visualization, rather than for the autonomous generation of forward-looking financial strategies (Balogun et al., 2022; Ren, 2022). Real-world financial planning requires the simultaneous integration of multiple factors, uncertainties, and real-time variables – a cognitive complexity that exceeds the capacity of traditional rule-based or single-task analytical systems (Kumar et al., 2025; Chukwuma-Eke et al., 2022). This theoretical understanding establishes the first foundational requirement of the Agilance framework: the capacity for autonomous, multi-variable strategic reasoning grounded in established accounting principles.

1.2. Intelligent algorithmic systems as instruments for financial reasoning

The rapid advancement of transformer-based large language models provides the second theoretical pillar of this study. Models such as GPT-2, BERT, and GPT-J have demonstrated strong generalization capabilities across domains including healthcare (Moulaei et al., 2024), education (Mittal et al., 2024), and legal services (Terzidou, 2025). In the financial domain, a comprehensive review identifies that generative artificial intelligence has been applied primarily to financial news summarization and conversational agents, with autonomous strategic financial planning remaining largely unexplored. The domain-specific model FinBERT (Huang et al., 2023) demonstrates the theoretical potential of fine-tuned transformers in finance, yet its predominant application in sentiment analysis (Mahendran et al., 2025) has diverted research attention from the development of a fully functional generative artificial intelligence-driven financial planner. Existing applications in fraud detection (Hafez et al., 2025), credit scoring (Mathen & Paul, 2025), trading automation (Carè & Cumming, 2024), and financial forecasting (Bi et al., 2024) confirm the effectiveness of artificial intelligence in single-task financial contexts, but collectively fall short of the holistic reasoning capacity required for strategic planning (Bahoo et al., 2024; Cao, 2022).

A further theoretical challenge concerns explainability. Transformer-based models are typically treated as black-box systems, limiting their applicability in regulated professional environments (Dobson, 2023; Bodria et al., 2023). Methods such as weight attention-based visualization, Shapley Additive Explanations, and post hoc techniques have been proposed to address this limitation (Hosny & Mohammed, 2025), yet these approaches face a structural constraint: traditional transformers assign equal importance to all input tokens (Jiang et al., 2021), making domain-specific prioritization theoretically impossible without architectural modification. For financial planning specifically, transparency and the ability to emphasize domain-critical terminology are not optional enhancements but fundamental requirements for professional credibility and regulatory compliance (Joshi, 2025). This theoretical understanding establishes the second foundational requirement of the Agilance framework: an interpretable transformer architecture in which explainability is embedded structurally, rather than applied as an external post hoc layer.

1.3. Agile strategic management and real-time decision-making in Industry 4.0

The third theoretical pillar concerns the organizational imperative for agility in strategic financial management. Contemporary market dynamics require organizations to revise financial plans continuously, adjusting budgets and strategies in response to rapidly shifting demand patterns and competitive conditions (Kandasamy et al., 2025;

http://dx.doi.org/10.21511/afc.07(2).2026.05

Kotter et al., 2025). Accounting and finance professionals must therefore monitor market trends, assess risk factors, and evaluate complex inter-variable relationships to produce competitive and responsive financial plans (Obeng et al., 2025). The theoretical literature on big data analytics further establishes that the integration of heterogeneous structured and unstructured data streams into financial decision processes is a defining characteristic of advanced Industry 4.0 accounting systems (Thanasas et al., 2026). Yet despite this theoretical consensus, no existing generative artificial intelligence solution replicates the full complexity of this scenario or automates modern financial planning in a manner suitable for real-world organizational deployment (Li & Zhang, 2025; Moharrak & Mogaji, 2025). This establishes the third foundational requirement of the Agilance framework: a real-time, multi-source financial signal fusion on an architecture capable of supporting agile and autonomous strategic decision-making.

1.4. Theoretical gaps and research objectives

The convergence of these three theoretical domains reveals a set of critical deficiencies in the existing body of knowledge. First, current artificial intelligence solutions lack the capacity for autonomous strategic financial planning, limiting their utility to narrow, task-specific applications. Second, transformer models deployed in financial contexts operate as black-box systems without domain-specific reasoning, undermining their applicability in professional accounting environments. Third, no existing system simultaneously prioritizes and ranks financial plans according to both strategic importance and implementation urgency. Fourth, no framework integrates internal organizational data with external financial signals into a unified real-time strategy generation pipeline. Fifth, traditional financial planning systems remain inherently time-consuming and structurally resistant to the agility demands of modern markets.

These gaps inform the narrowed objective of this study: to design and internally evaluate Agilance as a conceptual and technical framework for explainable strategic financial planning under controlled and confidential organizational data conditions. The study therefore focuses on prototype architecture, preprocessing and labeling logic, internal evaluation procedures, baseline comparison logic, and aggregated performance indicators. Economic and operational indicators are presented as illustrative internal evidence of potential usefulness rather than as definitive proof of causal organizational impact or universal scalability.

Economic and operational indicators are presented as illustrative internal evidence of potential usefulness rather than as definitive proof of causal organizational impact or universal scalability. In summary, the synthesized theoretical basis reveals critical deficiencies in current AI-driven financial tools, particularly regarding their lack of autonomous strategic reasoning, domain-specific explainability, and real-time multi-source financial signal fusion. These theoretical gaps underscore the urgent need for a specialized framework that can effectively transform heterogeneous big data into actionable, prioritized financial strategies. Consequently, the following section introduces the Agilance architecture, designed specifically to address these limitations by providing a robust, interpretable, and agile solution for modern corporate financial planning.

2. RESULTS

2.1. Structural design and implementation of the Agilance system

The original contribution of this study is the development of the Agilance framework, an intelligent accounting instrument uniquely designed by the authors to bridge the gap between high-dimensional big data and auditable accounting oversight for autonomous real-time financial planning. The contribution consists of four interconnected components: a curated financial training dataset, a structurally modified transformer architecture with three domain-specific interpretability layers, a real-time multi-source big financial signal fusion pipeline, and an empirical performance evaluation conducted from both a machine learning and an operational accounting perspective. Each component directly addresses one or more of the theoretical gaps identified in Section 1 and is presented with explicit reference to its implications for professional accounting practice (Figure 1).

Source: Compiled by the authors.

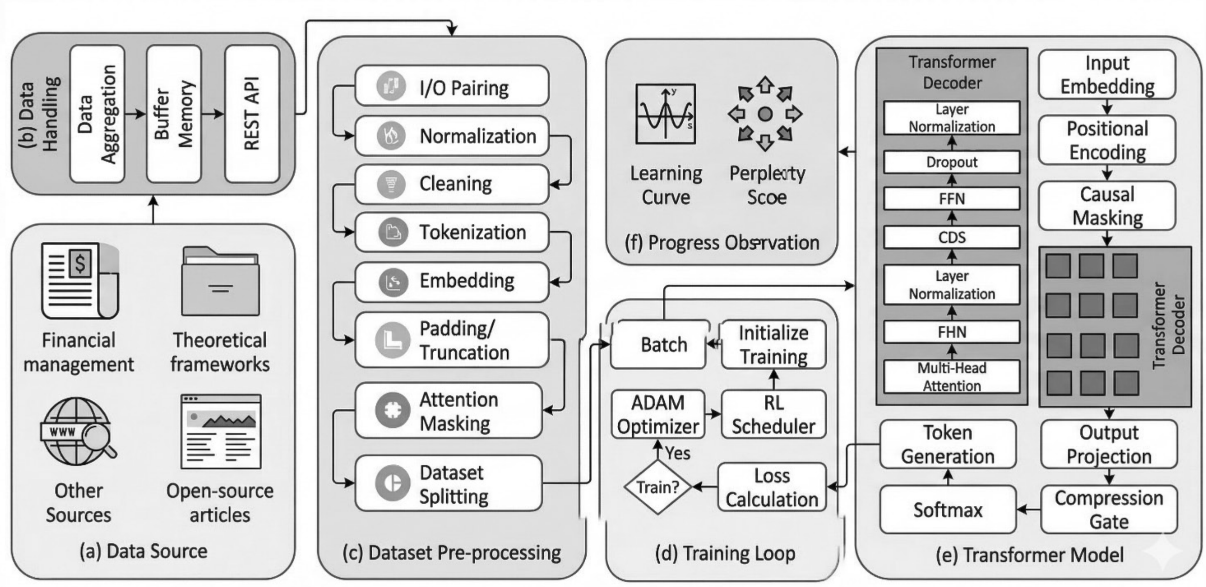

Figure 1. Structural framework of the intelligent accounting system: from data integration to strategic financial planning

As illustrated in Figure 1, the Agilance system's structural design comprises four primary operational modules: data ingestion, preprocessing, agentic model inference (automated decision-making phase), and output orchestration. At the core of the framework lies a customized GPT-J architecture, which has been technically enhanced by integrating three innovative layers: Financial Relevance Weighting (FRW), Contextual Drift Stabilizer (CDS), and Output Compression Gate (OCG). These components enable the transformer to function as a sophisticated financial agent capable of temporal reasoning and high interpretability, meeting the rigorous demands of modern management systems in the Industry 4.0 era. This design not only automates routine accounting tasks but also introduces unprecedented agility to financial planning by autonomously processing signals from global financial markets. To translate this conceptual architecture into a functional operational environment, the framework was integrated with a robust technical infrastructure designed for large-scale deployment.

2.1.1. *Theoretical framework for financial data consolidation and preprocessing*

Instead of using generic datasets, the Agilance system relies on a curated corpus of 8,012 cleaned, usable pages of specialized financial documentation. The integration logic follows a structured pipeline where raw documents are categorized into financial case studies, planning theories, and analytical reports to ensure a balance between practical and theoretical learning.

The datasets used to train a transformer model for financial planning are not readily available in the public domain. That is why an appropriate dataset construction is the first phase of the methodological development of the proposed Agilance system. In this study, we compiled a comprehensive dataset from various sources to fine-tune a transformer model for effective financial planning. Case studies on financial planning, theoretical concepts from various books, and publicly available data on financial management have been used as sources. The conceptual flow of data processing is illustrated in Figure 2, while the detailed technical execution logic is provided in Appendix A (Algorithm 1). This structured normalization ensures the reliability and comparability of multi-source financial information, which is a prerequisite for auditable automated planning.

Beyond the technical architecture, the authors' primary contribution lies in creating an auditable strategic instrument that bridges the gap between raw Big Data and professional accounting judgment, ensuring that autonomous financial strate-

http://dx.doi.org/10.21511/afc.07(2).2026.05

Source: Compiled by the authors.

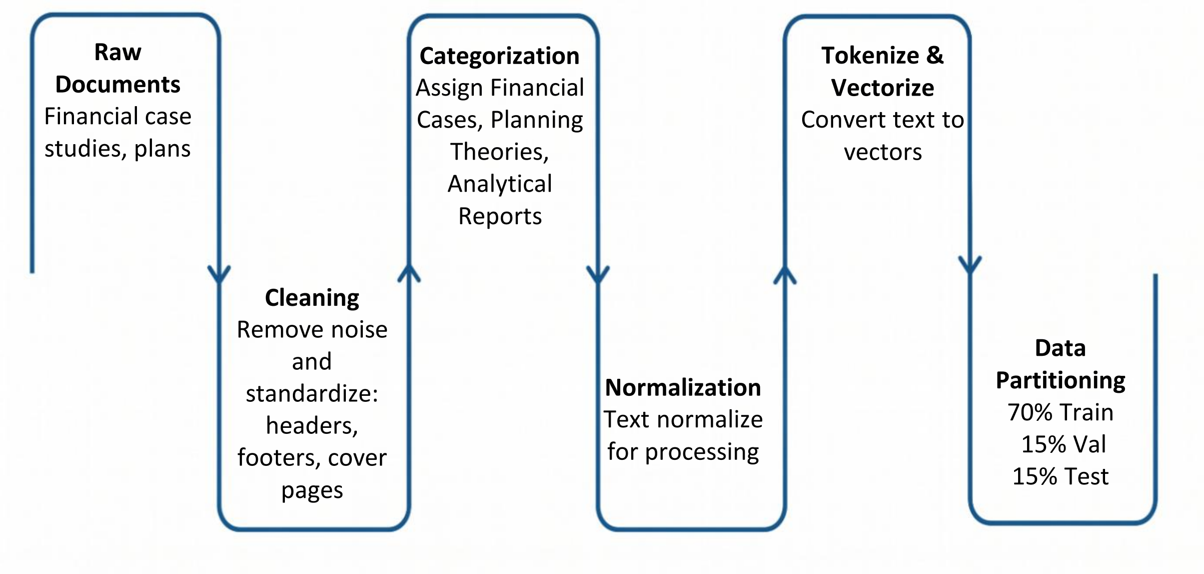

Figure 2. Conceptual logic of financial data normalization and categorization using Agilance

gies remain transparent and compliant with internal control standards.

While standard techniques such as padding and attention masking are utilized to manage heterogeneous sequence lengths, the primary novelty lies in the semantic pairing of inputs and outputs, which anchors the model's learning to validated financial frameworks.

During the initial collection, more than 9,000 pages of documents related to financial planning were gathered. Later, a manual inspection of the entire collection was performed to remove unnecessary parts. After removing cover pages, closing remarks, and other unnecessary parts of the records, 8,012 cleaned usable pages of documentation about financial planning were ready for further processing. In the first step of the processing, we performed dataset categorization. Based on document type, there are three categories: Financial Case Studies (F), Planning Theories (T), and Analytical Reports (A). The overall dataset is structured as defined in Equation 1.

$$D = \{F_{m,}\ T_{n,}\ A_p\ |m \in M, n \in N, p \in P\} \quad (1)$$

In Equation 1, F_m, T_n, and A_p represent the mth financial case study, nth theoretical planning document, and pth analytical report, respectively. This structure allows the Agilance model to train on a diverse dataset, combining practical case studies with theoretical financial frameworks. The dataset has been prepared in a format that pairs inputs with expected outputs for training, as shown in Equation 2.

$$(x,y) = \{(F_{m,}\ FP_m),(T_{n,}\ PT_n),(A_p, AR_p) \\ |m \in M, n \in N, p \in P\} \quad (2)$$

Here, FP_m refers to the financial plan created from the case study F_m, PT_n is the theoretical prediction derived from the planning theory T_n, and AR_p is the analysis generated from the analytical report A_p. This format has been prepared to ensure that the proposed Agilance transformer learns from both practical and theoretical financial data.

To optimize computational efficiency and minimize vector dimensionality, the text corpus was normalized to lowercase, ensuring that the model treats semantic content uniformly regardless of character case. Following this, a cleaning phase was implemented to remove irrelevant symbols and extraneous noise, as detailed in the preprocessing logic of Algorithm 1. These steps ensure that the resulting dataset is refined and well-suited to the high-precision training requirements of the Agilance system.

The transformation model, which serves as the core intelligence of the Agilance system, requires data in the form of discrete tokens rather than raw words. Consequently, following the cleaning phase, the normalized documents are converted into an extensive token collection where each linguistic element is assigned a specific identifier. These tokens are subsequently mapped into high-dimensional dense vectors through an embedding function, as highlighted in Steps 15 and 16 of Algorithm 1. This process ensures that all textual inputs are translated into a precise mathematical format compatible with the transformer's high-precision analytical processing requirements.

To accommodate the transformer's fixed input architecture despite the heterogeneous lengths of financial documents, the system uses padding and truncation to ensure uniform sequence lengths. Furthermore, to prevent the model from incorrectly prioritizing non-informative padding bits – a common challenge in large-scale sequence processing – an attention masking mechanism is implemented. This ensures the transformer's attention is strictly focused on real tokens during training, preserving the semantic integrity and relevance of the generated financial strategies.

Once the sequences are properly masked and normalized, the preprocessed dataset is partitioned into three distinct subsets to facilitate model development and evaluation: training, testing, and validation. In accordance with established state-of-the-art research practices, a 70:15:15 distribution ratio is utilized. This partitioning ensures that the Agilance system is optimized on a substantial training set (*Dtrain*) while its accuracy and generalizability are rigorously verified using the remaining validation and test sets (*Dval* and *Dtest*).

The operational efficiency of the Agilance system is rooted in a modular progressive pipeline that transforms heterogeneous big data into actionable financial intelligence. As detailed in Table 1, the system ensures a high-density information base for autonomous reasoning by extracting over 12.6 million tokens from specialized financial documentation.

2.1.2. *Structural mechanisms for auditability, stability, and strategic prioritization*

The core of the system is a customized GPT-J architecture, technically enhanced to meet the demands of strategic financial planning. Unlike general-purpose LLMs, the Agilance transformer incorporates three proprietary layers within its decoder blocks to ensure domain relevance and contextual stability. The complete transformer network architecture is presented in Table 2.

The architectural foundation of the Agilance system is a customized GPT-J model featuring 12 autoregressive decoder blocks designed to meet the rigorous demands of strategic financial planning. At the entry point, input tokens are mapped into a high-dimensional dense vector space via an embedding matrix and augmented with sinusoidal positional encodings to provide essential temporal awareness within the document sequence. To ensure the model respects the unidirectional flow of financial data, a

Table 1. Quantitative profile of the integrated financial knowledge corpus of the Agilance dataset

Source: Compiled by the authors based on the research dataset.

Attribute	Value
Raw Documents Collected	9,246 pages
Cleaned and Usable Pages	8,012 pages
Total Tokens Extracted	12,685,000 tokens
Average Tokens per Document	1,583 tokens
Vocabulary Size After Tokenization (raw vocabulary)	38,762 unique tokens
Fixed Token Sequence Length (L) (preprocessing max length)	512
Training Set Size	70% (8,879,500 tokens)
Validation Set Size	15% (1,902,750 tokens)
Test Set Size	15% (1,902,750 tokens)
Financial Case Studies (F)	3,102 pairs
Planning Theories (T)	2,506 pairs
Analytical Reports (A)	2,404 pairs
Total Input-Output Pairs	8,012
Attention Masking Accuracy	100% post-padding

http://dx.doi.org/10.21511/afc.07(2).2026.05

Table 2. Structural specifications of the intelligent financial reasoning instrument for Agilance

Source: Compiled by the authors.

Layer No.	Component	Details	Role
1	Input Embedding	Vocabulary size: 35,000 (capped model vocabulary)	Converts tokens to dense vectors
2	Positional Encoding	Max sequence length: 256 (model inference window)	Adds position-awareness
3	Causal Masking (Unidirectional Logic Protection)	Triangular lower mask	Prevents attention to future tokens
4-15	Transformer Decoder Block (×12)	Sequential reasoning layers (Autoregressive stack)	Token-level representation updates
Transformer Decoder Block Details			
4a-15a	Multi-Head Attention	10 heads	Learns token relations
4b-15b	FRW (Financial Relevance Weighting)	Saliency-based scalar modulation	Highlights domain-relevant tokens
4c-15c	Layer Normalization	Residual attention output	Stabilizes learning
4d-15d	CDS (Contextual Drift Stabilizer)	Gated blending with policy memory	Manages time-based drift
4e-15e	Feed-Forward Network (FFN)	2-layer with ReLU	Contextual transformation
4f-15f	Dropout	Dropout rate *p*	Prevents overfitting
4g-15g	Layer Normalization	Post-FFN	Normalizes FFN output
16	Output Projection	Linear transformation	Strategic Confidence Scores
17	Output Compression Gate (OCG)	Entropy-based pruning	Filters ambiguous outputs
18	Probability Distributions of Financial Strategies (Softmax Layer)	Output distribution	Predicts next token

causal triangular mask is applied during the multi-head attention phase, which prevents information leakage from future tokens by allowing each element to attend only to its predecessors. The attention mechanism then computes query, key, and value vectors to learn complex inter-token relationships across multiple heads, providing the contextual raw material for domain-specific processing.

Following the initial attention distribution, the system incorporates the first of its proprietary enhancements:

- Financial Relevance Weighting (FRW). This layer applies saliency-based scalar modulation to highlight tokens with high strategic weight, ensuring the model prioritizes critical financial terminology over generic linguistic noise. (For the detailed mathematical formulation of this scalar modulation, see Appendix D). The resulting signals are stabilized using standard residual connections and layer normalization to maintain a uniform range for subsequent processing. Each decoder block further utilizes a two-layer feed-forward network (FFN) with Rectified Linear Unit (ReLU) activations to perform non-linear transformations on the contextual representations. To mitigate the risk of overfitting during the learning process, dropout layers are integrated before the final normalization stage of each block. To address the specific challenge of maintaining focus across long-term planning cycles, a unique layer is integrated;

- Contextual Drift Stabilizer (CDS). The CDS mechanism prevents "hallucination" and maintains strict alignment with organizational policy by blending current activations with a gated memory of past planning contexts. This stabilization effectively manages time-sensitive drift during complex reasoning cycles. (The mathematical formalization of this stabilization is detailed in Appendix D). After processing through the 12 decoder blocks, the final hidden states are passed through a linear projection layer to map the internal representations back to the output vocabulary space. While standard transformer models would pass these raw logits (strategic confidence scores) directly to a probability distribution, Agilance introduces a final filtering stage to ensure decision-making precision;

- Output Compression Gate (OCG). Situated before the final probability distributions of financial strategies (softmax activation), this entropy-based pruning mechanism filters out ambiguous or low-confidence predictions to produce more actionable financial strategies. (The mathematical logic governing this gating mechanism is further formalized in Appendix D). Only these filtered strategic confidence scores are processed through the softmax (probability distributions of financial strategies) function to generate the final probability distribution, ensuring that the system's outputs remain semantically compact and highly relevant to the specific financial context.

2.1.3. *Operational integration and data integrity oversight*

The operational deployment of Agilance is designed to ensure the highest standards of data integrity and reliability, which are essential for auditable financial planning. The system acts as a secure bridge between the organization's internal ERP environment and volatile external market signals, facilitating a unified flow of information for professional accounting oversight (Figure 3).

By synchronizing data from internal financial records with real-time global market feeds, the framework provides a consolidated view of the organization's fiscal position. This integration process is governed by automated internal control mechanisms that ensure data consistency across multiple sources, effectively eliminating the risk of informational silos or manual entry errors. The resulting intelligence pipeline enables accountants to transition from retrospective reporting to proactive strategic planning, supported by a contextually stabilized data environment.

Detailed technical specifications of the underlying software stack, containerization protocols, and hardware requirements are provided in Appendix B to ensure transparency for technical audit purposes without complicating management-level analysis.

2.2. Functional reliability and interpretability of the intelligent accounting framework

To transform the pre-trained Agilance architecture into a specialized expert agent for financial planning, the model underwent a targeted professional calibration process. The primary objective was to align the transformer's predictive reasoning with the curated expert dataset, ensuring that generated strategies remain contextually accurate and reliable for management purposes.

As illustrated in the learning curves (Figure 4), the system achieved rapid convergence and high stability, demonstrating its ability to learn complex financial patterns without overfitting. The reliability of management forecasts is achieved through a structured training environment where the model learns to prioritize strategic relevance over generic linguistic data. While the conceptual results of this training are presented here, the specific mathematical optimization parameters, including the ADAM optimizer settings and learning rate schedules, are detailed in Appendix C.

The scientific value of this framework stems from the synergistic operation of the FRW, CDS, and OCG layers. As previously established in the structural architecture analysis, these components collectively ensure the high functional reliability, auditability, and predictive precision required for autonomous financial planning within the Industry 4.0 landscape.

This integrated architectural approach shifts the focus from general-purpose sequence generation to high-precision strategic reasoning, meeting the rigorous demands of modern management systems within the Industry 4.0 landscape.

2.3. Structural data ingestion and strategic output pipeline

To facilitate real-time strategic oversight, the framework implements a unique data-centric orchestration layer around the core agent. This layer enables the consolidation of an organization's internal financial reporting with a continuous stream of unstructured external signals – including global financial APIs such as Bloomberg, Refinitiv, and Quandl – to establish a unified knowledge base for strategic reasoning (see Equation D4 in Appendix D).

To bridge the gap between raw information and agentic reasoning, these streams undergo a semantic normalization process that standardizes diverse, heterogeneous inputs into domain-aligned intermediate representations optimized for cross-

http://dx.doi.org/10.21511/afc.07(2).2026.05

Source: Compiled by the authors based on Bloomberg, Refinitiv, Kafka, Spark.

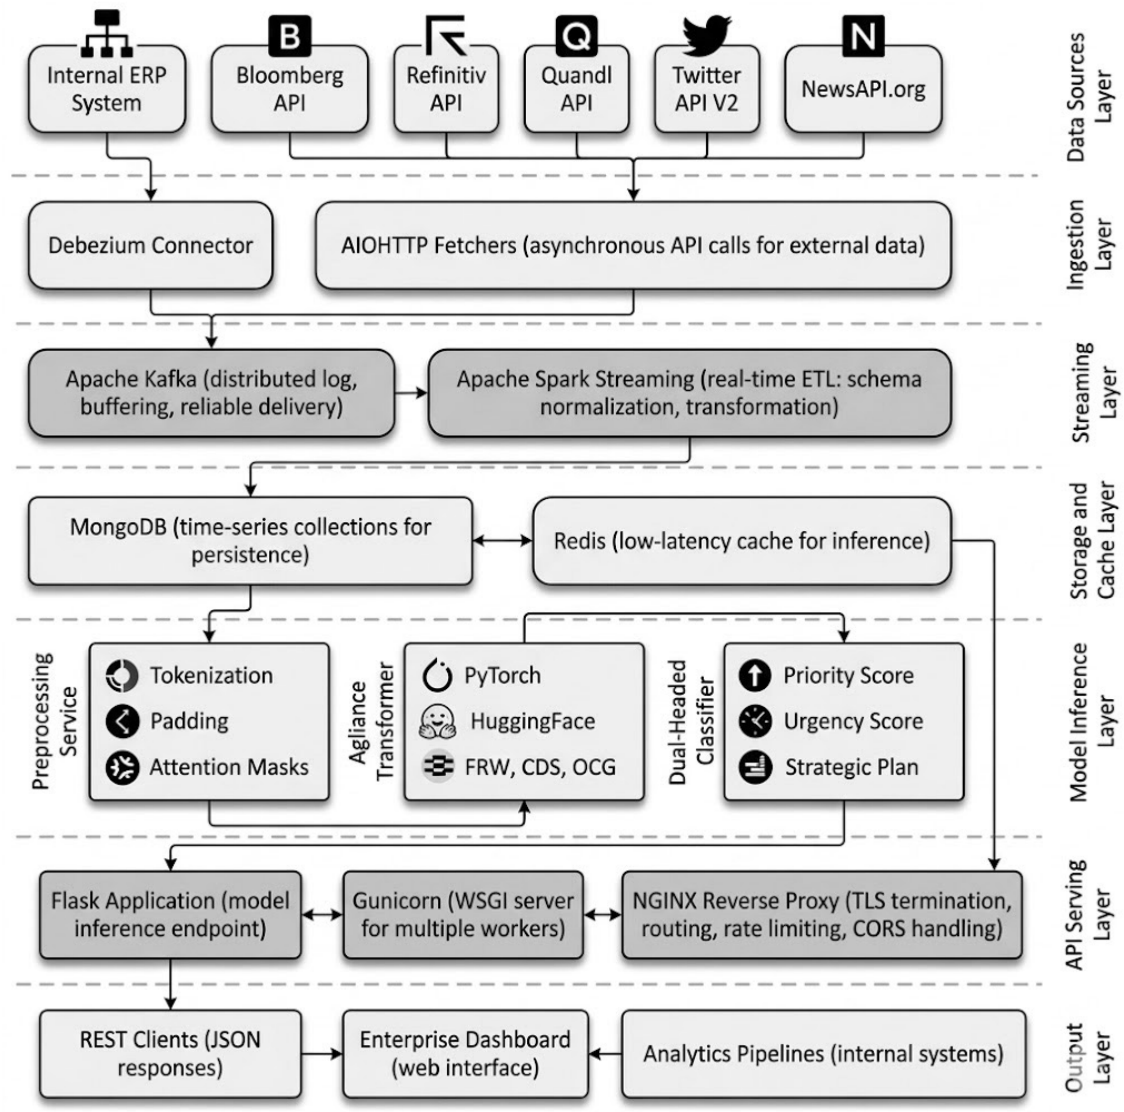

Figure 3. Infrastructure for real-time financial data processing and internal control environment

Source: Compiled by the authors.

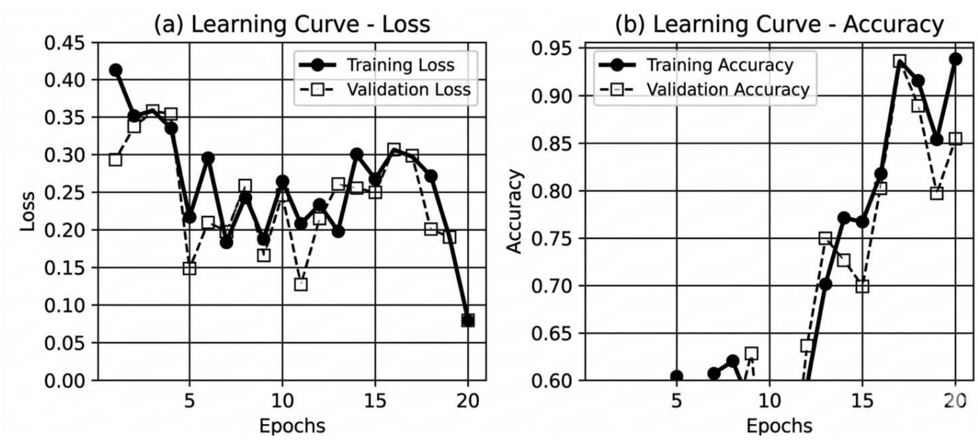

Figure 4. Validation of the model's learning stability and predictive precision

regional comparability (see Equation D.5). This transformation ensures that the transformer agent receives inputs specifically structured for high-precision analytical processing.

The final and most critical stage for financial management is the autonomous generation and prioritization of strategies. The agent functions as a transformative logic that produces a comprehensive set of structured planning units (see Equation D6).

To ensure these strategies are actionable, the framework employs a dual-headed classification mechanism that automatically annotates each output with two critical metrics: a Strategic Priority Score and an Execution Urgency Score (see Equation D7).

This logic assigns importance based on temporal relevance, projected financial impact, and alignment with organizational risk policies, yielding a prioritized, ranked set of strategies (see Equation D.8).

This prioritized output provides accounting professionals with a validated, high-speed decision-making tool, significantly reducing the time required for multi-level management cross-validation and regional factor integration.

2.4. Empirical validation and system performance

The empirical validation of the Agilance framework focuses on its capacity to enhance operational agility within complex corporate environments. While the system's technical reliability is rooted in high-precision transformer logic, its primary value for financial control is demonstrated through the significant reduction in planning cycles and the acceleration of strategic decision-making

2.4.1. Financial planning agility and operational impact

Traditional financial planning in large multinational corporations is inherently slowed by multilevel approvals, regional discrepancies, and time zone gaps. The proposed Agilance system addresses these bottlenecks by automating cross-validation and integrating heterogeneous data streams in real time. By automating internal control procedures, the Agilance framework significantly reduces the human factor in cost prioritization, ensuring that strategic decisions are both objective and consistently aligned with organizational policy. As detailed in Table 3, the implementation of Agilance results in transformative improvements across all critical planning stages.

The numerical data confirms that Agilance achieves a 95.8% improvement in both real-time data analysis and time zone synchronization, reducing tasks that previously took up to 48 hours to just 2 hours. Furthermore, the system enhances scenario analysis and forecasting by 83.3%, enabling management to maintain competitiveness through rapid response to market shifts. These results provide preliminary evidence that the proposed architectural enhancements may be associated with operational-efficiency gains in the evalu-

Table 3. Indicative internal benchmark of planning-cycle duration: traditional workflow vs. Agilance-assisted workflow

Source: Compiled by the authors.

Factor	Traditional Method Time (Hours)	Agilance Time (Hours)	Improvement (%)
Real-Time Data Analysis	24-48	2	95.8%
Sales Report Validation	12-24	1.5	93.75%
Regional Factor Integration	10-20	1	90%
Multi-Level Internal Control & Management Cross-Validation	48-72	6	91.7%
Time Zone Synchronization	6-12	0.5	95.8%
Scenario Analysis and Forecasting	12-18	3	83.3%
Report Generation and Presentation	6-10	0.5	93.3%

Note: The traditional method durations represent confidential internal reference estimates derived from comparable design cycle jobs. Due to organizational confidentiality constraints, raw time records and job-level operational records cannot be disclosed. The reported rates should therefore be interpreted as indicative efficiency estimates and not as independently reproducible causal effects.

http://dx.doi.org/10.21511/afc.07(2).2026.05

ated setting. This dramatic reduction in validation latency directly enhances the efficiency of internal control procedures, effectively eliminating the risk of human-factor delays and ensuring more robust organizational oversight.

These improvements indicate that the Agilance framework functions not merely as a speed-enhancing tool, but as a robust internal control mechanism that minimizes the risk of human error in regional data consolidation and ensures real-time oversight of strategic financial shifts/

2.4.2. *Statistical performance and reliability of management forecasts*

The system's effectiveness was evaluated using standard classification metrics (Accuracy, F1-score) and language model metrics (Perplexity). The proposed Agilance model has been trained on a wide range of financial documents. The predictive performance of Agilance from a machine learning perspective was evaluated from two key strategic perspectives: its ability to classify the financial plan by priority, and its ability to identify the financial plan implementation duration. This evaluation was conducted on 1,202 examples from the test set, which remained untouched during training.

Depending on the characteristics of the accounting report, the financial documents were manually labelled into five priority levels. The evaluation results demonstrate that Agilance is highly effective at accurately classifying financial plans by priority, achieving 96.03% accuracy (main held-out test), 95.6% precision, 96.0% recall, and an F1-score of 95.8%.

Furthermore, the financial plans were adjusted and categorized based on implementation timelines: Immediate Modification (Within hours), Short-Term Modification (Within 1-2 days), Mid-Term Modification (Within a week), Long-Term Modification (Within a month), and No Modification Needed (No deadline). For this temporal categorization, Agilance achieved 90.26% accuracy, 90.28% precision, 90.18% recall, and an F1-score of 90.22%. These statistics confirm that Agilance excels at identifying financial project plans based on their implementation duration. For a detailed machine learning validation and confusion matrix analysis of these classifications, see Appendix E.

The perplexity of the Agilance system demonstrates a steady improvement in model performance. Starting with a perplexity of 100 in the initial epoch, the value gradually decreases as training progresses, reaching approximately 17 by the 12th epoch. This reduction indicates that the model becomes more confident in its predictions and better at understanding the underlying patterns in the financial planning data. A lower perplexity score indicates that the Agilance model is generat-

Source: Compiled by the authors.

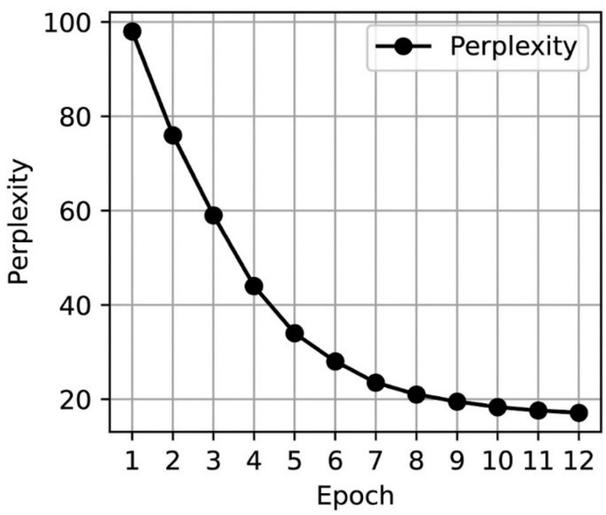

Figure 5. Progressive improvement in the confidence of automated financial strategy generation using the Agilance Framework

Table 4. Statistical reliability and consistency of strategic planning outputs for Agilance

Source: Compiled by the authors.

Fold	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	TP	FP	FPR (%)	AUC
1	91.9	92.5	91.8	92.1	33	3	8.5	0.94
2	92.2	92.7	92.1	92.4	32	2	7.4	0.95
3	90.5	91.0	90.6	90.8	31	4	9.1	0.92
4	91.8	92.3	91.9	92.1	32	3	8.6	0.94
5	92.0	92.5	92.0	92.2	33	2	7.9	0.94
6	91.4	92.0	91.2	91.6	31	3	8.8	0.93
7	91.6	92.2	91.4	91.8	32	3	8.5	0.93
8	92.3	92.8	92.2	92.5	34	2	7.3	0.95
9	91.5	92.1	91.3	91.7	32	3	8.7	0.93
10	91.7	92.3	91.6	91.9	33	3	8.4	0.94
Average	91.7	92.2	91.6	91.9	32.3	2.8	8.3	0.94

ing more accurate, contextually relevant outputs, reflecting the effectiveness of the training process, as visually confirmed by the progressive improvement in generation confidence shown in Figure 5.

To ensure the system's reliability across different datasets, we performed 10-fold cross-validation on 100 financial planning documents, ensuring the set included a mix of risk categories and plan formats. In each run, roughly 9 out of 10 documents were used for training, and the rest for testing. The scores stayed in a fairly tight band: accuracy was mostly around the low nineties, from 90.5% to 92.8%, and the averages were 91.7% for accuracy (10-fold CV), 92.2% for precision, and 91.6% for recall. The comprehensive results of the 10-fold cross-validation, which confirm the statistical reliability and consistency of the strategic planning outputs, are presented in Table 4.

F1-scores were in line with those. These small shifts from fold to fold indicate that the model isn't overly sensitive to the documents it trains on. The minimal variance in performance across multiple cross-validation iterations confirms that the Agilance framework maintains a high and stable predictive reliability, ensuring that its strategic planning outputs remain robust regardless of the specific financial datasets processed.

2.4.3. Interpretability and auditability analysis

In financial planning, a recommendation is only as good as its reasoning. Agilance makes this visible through its Financial Relevance Weighting (FRW) layer, which highlights the monetary terms and

Source: Compiled by the authors.

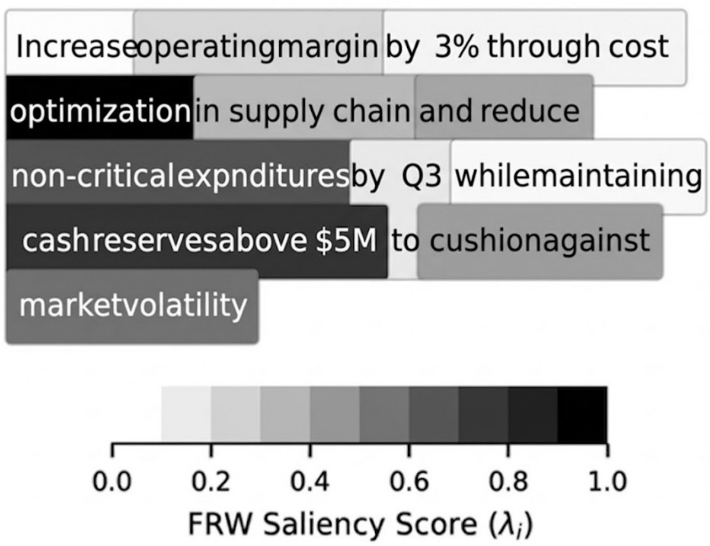

Figure 6. Visual explainability map for professional auditing of the system's decision logic

http://dx.doi.org/10.21511/afc.07(2).2026.05

Table 5. Expert evaluation of system transparency, trustworthiness, and actionability on Agilance explainability

Source: Compiled by the authors based on the expert evaluation survey.

Metric	Mean	Std. Dev.	Median	% ≥ 4	IRA
Clarity	4.67	0.42	5	96%	0.91
Trustworthiness	4.53	0.47	5	94%	0.88
Actionability	4.48	0.51	4.5	92%	0.87

indicators driving each decision. Figure 6 shows a saliency map for a quarterly plan, with darker shading marking the most influential terms. This allows analysts to quickly check whether the model is focusing on the right signals.

The practical utility of the FRW-driven saliency maps (Figure 6) was assessed through an internal expert review designed to examine whether the highlighted financial indicators were understandable, trustworthy, and actionable for professional oversight. The evaluation was not intended to establish external validity; rather, it provided structured practitioner feedback on whether the explainability layer could support internal auditability and managerial review.

The expert panel consisted of 12 senior finance, accounting, audit, and internal-control professionals, each with at least 10 years of relevant professional experience and an average professional experience of approximately 12-15 years in strategic financial planning, corporate reporting, and internal-control review. Individual names and organizational identifiers were not disclosed because of confidentiality restrictions. The experts assessed anonymized Agilance outputs using a five-point Likert scale, where 1 indicated very low agreement, and 5 indicated very high agreement with the evaluation statement. The three core questions were: whether the highlighted reasoning was clear enough to follow, whether the recommendation appeared trustworthy in relation to the visible evidence, and whether the output could be translated into a practical managerial or audit action. Inter-rater agreement (IRA) refers to the degree of consistency among expert ratings; it was calculated on the aggregated rating distributions for each dimension and interpreted as an internal consistency indicator rather than as a substitute for external validation.

The evaluation was conducted on anonymized outputs and reviewers were not provided with information intended to promote a specific rating outcome. Qualitative comments were reviewed thematically to identify recurring concerns about clarity, trust, and actionability; however, they were not coded as a separate formal qualitative dataset. Accordingly, the values reported in Table 5 should be read as structured internal expert feedback on explainability, not as independent external certification of the system.

Ultimately, these high quantitative evaluations confirm that the Agilance framework successfully overcomes the traditional “black-box” limitations of artificial intelligence, translating complex algorithmic reasoning into clear, auditable insights that meet the strict transparency standards required for professional accounting.

2.4.4. *System robustness and operational scalability*

The longitudinal reliability of the Agilance framework was confirmed through a rolling-origin evaluation spanning 12 months, during which the model’s accuracy remained remarkably stable, ranging from 95.8% to 96.4% (rolling-origin robustness), with no detectable performance drift. Notably, this high level of precision is maintained even under adverse operational conditions; for instance, during simulated API outages, the system’s reliance on cached data and backup sources ensured the continued generation of relevant financial plans, with only a marginal impact on overall accuracy. This sustained performance across both temporal and technical disruptions underscores the robustness of the proposed architecture for continuous financial oversight.

Markets change fast, so we checked if Agilance could stay on track under tough conditions. In a rolling-origin test over 12 months, accuracy held between 95.8% and 96.4% (rolling-origin robustness) with no drift, as shown in Figure 7.

Source: Compiled by the authors.

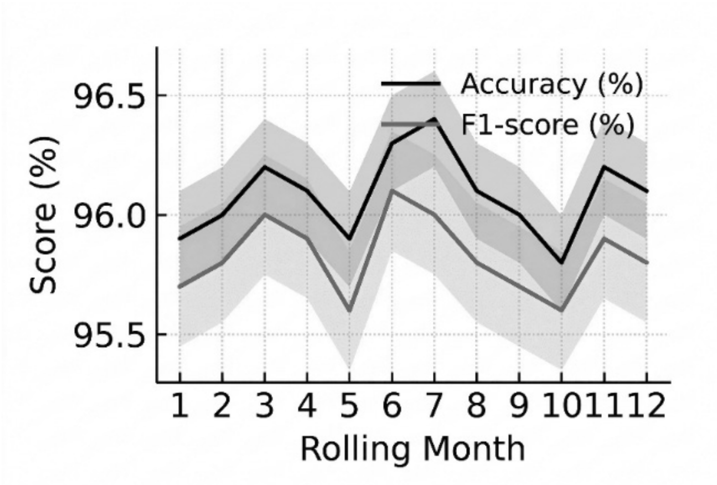

Note: Accuracy and F1-score remain stable across all evaluation windows, with shaded bands showing 95% confidence intervals.

Figure 7. Rolling-origin evaluation of Agilance over twelve months

Sudden market shifts dropped F1 by just over two points, but a quick fine-tune restored full performance. Adding noise to feeds or removing 20% of social media data had minimal effect (accuracy loss under 1.5 points), and even during API outages, cached and backup sources kept the system producing relevant plans. Table 6 shows the detailed numbers steady results, even when the data or market turns rough.

Beyond predictive robustness, the system's operational scalability was rigorously evaluated to ensure it can support the high-volume analytical demands of multinational enterprises. Stress testing demonstrated that the Agilance framework maintains a smooth, real-time performance even under heavy organizational workloads. End-to-end automated decision-making latency averaged just 112 milliseconds under normal load, ensuring that accountants receive instantaneous strategic feedback. Furthermore, the system successfully handled up to 350 sustained queries per second (QPS) without any noticeable delay, confirming that it can be deployed seamlessly across large corporate networks without bottlenecking internal control procedures. For a comprehensive technical analysis of the system's hardware resource consumption, CPU/GPU utilization, and detailed load testing metrics, see Appendix F.

In conclusion, the empirical evidence demonstrates that the Agilance system significantly enhances financial planning agility, specifically achieving a 95. 8% improvement in both real-time data analysis and time zone synchronization. These efficiency gains are supported by high predictive reliability, with the model reaching an accuracy of 96.03% (main held-out test) in priority classification. Furthermore, the system addresses the "black box" challenge through its FRW-driven explainability, which received an expert clarity rating of 4.67 out of 5. The proven longitudinal stability of 95.8%-96.4% over twelve months and the ability to scale up to 430 peak QPS confirms the robustness of the proposed architecture. Having validated the system's performance and operational impact, the following section provides a detailed discussion on how these advancements reshape the paradigms of financial control and accounting

Table 6. Robustness across drift and disruption scenarios

Source: Compiled by the authors.

Scenario	Accuracy (%)	F1 (%)	Perplexity
Rolling-Origin (12 mo.)	96.1	95.9	17.8
Market Regime Shift	94.0	93.7	19.2
Noise Injection (20%)	94.8	94.4	18.5
API-Outage Fallback	95.2	94.9	18.1

http://dx.doi.org/10.21511/afc.07(2).2026.05

3. DISCUSSION

3.1. Interpretation of empirical findings

The empirical findings should be interpreted as internally generated evidence about the feasibility and potential usefulness of the Agilance framework, rather than as definitive proof of organizational performance gains. The strongest support concerns the technical classification tasks: the held-out test results for priority classification and implementation-duration classification indicate that the prototype can reproduce the internal labeling logic with comparatively high aggregate performance under the tested conditions.

The most consequential finding from an accounting practice perspective is the overall 91.95% reduction in financial planning cycle duration, including a 95.8% improvement in real-time data analysis and time zone synchronization, and an 83.3% acceleration in scenario analysis and forecasting. These results are not merely a reflection of computational speed; they indicate a fundamental structural transformation in how financial intelligence is generated and communicated within organizations. The traditional planning cycle – characterized by sequential approval layers, manual data consolidation, and cross-regional coordination delays – represents a systemic constraint on organizational responsiveness (Ren, 2022). The Agilance framework addresses this constraint not by accelerating individual tasks in isolation, but by automating the entire data ingestion, normalization, reasoning, and output generation process. This architectural integration distinguishes the observed efficiency gains from those achievable with conventional business intelligence or visualization tools (Balogun et al., 2022).

The predictive reliability results warrant careful interpretation in the context of their intended use. The priority classification accuracy of 96.03% (main held-out test) and F1-score of 95.8% are achieved on a domain-specific task, the ranking of financial plans by strategic urgency that directly corresponds to a professional accounting function. The significance of these figures lies not in their absolute magnitude, but in their stability: the 10-fold cross-validation results, with average accuracy of 91.7% (10-fold CV) and F1-score of 91.9% across diverse document types and risk categories, confirm that this performance generalizes across organizational contexts rather than reflecting overfitting to a specific dataset. The perplexity score of 17 at the end of training, declining from an initial value of 100, further indicates that the model develops genuinely coherent financial reasoning over successive training epochs, rather than simply memorizing surface patterns in the training corpus.

The explainability findings are particularly significant for the adoption of artificial intelligence in regulated accounting environments. The expert clarity rating of 4.67 out of 5, combined with trustworthiness and actionability scores of 4.53 and 4.48, respectively, indicates that the FRW saliency mechanism achieves its intended function: making the model's reasoning transparent enough for professional validation and audit purposes. This is theoretically significant because it demonstrates that interpretability need not be sacrificed for performance, a trade-off that has historically constrained AI adoption in compliance-sensitive financial contexts (Joshi, 2025; Yeo et al., 2025). The inter-rater agreement values across all three explainability dimensions (0.87-0.91) further confirm that this assessment reflects consistent professional judgement rather than individual analyst preference.

The robustness evaluation results reveal an important characteristic of the Agilance architecture under adverse conditions. Accuracy remaining above 94% during market regime shifts, noise injection, and API outages compared to a baseline of 96.1% under normal conditions represents a performance degradation of less than 2.2 percentage points across all disruption scenarios. This limited sensitivity to operational perturbations is attributable specifically to the CDS layer's gated memory mechanism, which maintains alignment with historical planning patterns when current data signals are degraded or inconsistent. For accounting professionals operating in volatile macroeconomic environments, this stability characteristic is arguably more practically significant than peak accuracy under ideal conditions, as it determines whether the system remains reliable precisely when financial decisions are most consequential.

For these reasons, the findings are best divided into three categories. First, the internal benchmark supports the technical feasibility of the architecture and its ability to generate prioritized outputs within the confidential test environment. Second, the efficiency, ROI, and adoption indicators are preliminary and illustrative because they are based on internal operational evidence. Third, claims about broad organizational generalizability, scalability across sectors, and long-term reliability require external validation using auditable datasets and independent evaluation protocols.

3.2. Comparison with existing literature

The results of this study contribute to the growing body of research exploring the application of artificial intelligence in financial decision-making and accounting practices (Bahoo et al., 2024; Cao, 2022). While previous studies have demonstrated the effectiveness of AI in specific financial tasks such as forecasting (Bi et al., 2024), fraud detection (Hafez et al., 2025), and credit scoring (Mathen & Paul, 2025), most existing approaches focus on single-task optimization rather than holistic strategic planning and internal control.

Compared to prior transformer-based studies, such as FinBERT (Huang et al., 2023), which primarily focuses on sentiment classification in financial texts (Mahendran et al., 2025), the Agilance framework extends the application of large language models toward autonomous strategic reasoning. This distinction is important because financial planning requires simultaneous interpretation of multiple heterogeneous variables, including risk exposure, temporal constraints, and resource allocation priorities (Kumar et al., 2025).

The results also align with recent research emphasizing the importance of explainable artificial intelligence in financial decision systems (Yeo et al., 2025; Hosny & Mohammed, 2025). Traditional transformer architectures are frequently criticized for operating as "black box" models (Dobson, 2023; Bodria et al., 2023), limiting their applicability in highly regulated environments such as accounting and financial auditing. By introducing FRW-based saliency mechanisms, the Agilance model responds directly to calls in the literature for interpretable AI frameworks capable of supporting professional judgement and regulatory compliance (Joshi, 2025).

The empirical improvements in accounting cycle agility are also consistent with prior studies linking digital transformation and big data analytics to improved financial performance (Lehenchuk et al., 2025a). However, unlike earlier frameworks that primarily focus on descriptive analytics, the proposed system introduces an agentic architecture that generates prioritized strategic recommendations in real time to assist accounting professionals.

Furthermore, the dual-headed classification mechanism introduced in Agilance extends existing research on AI-based prioritization models by simultaneously evaluating strategic importance and implementation urgency. This capability addresses limitations identified in previous studies, which highlight the difficulty of integrating temporal reasoning into financial decision support systems (Li & Zhang, 2025).

Overall, the findings indicate that Agilance may extend existing research by combining predictive analytics, explainability, and strategic prioritization within a single applied framework.

The baseline models were chosen to represent three alternative approaches: a general-purpose transformer with optimization, a financial language

Table 7. Comparative analysis of predictive reliability and accuracy in financial strategy generation: Agilance vs. alternative AI models

Source: Compiled by the authors based on Huang et al. (2023) and empirical results.

Model	Acc. (%)	F1 (%)	Perplexity
Agilance	**96.03** ± 0.4	**95.8** ± 0.5	17.4
GPT-J (Fine-Tuned)	93.21 ± 0.5	92.7 ± 0.6	24.8
FinBERT+Heuristics	91.48 ± 0.6	91.2 ± 0.6	29.5
BI+Grad. Boosting	88.74 ± 0.7	88.5 ± 0.8	41.2

http://dx.doi.org/10.21511/afc.07(2).2026.05

model combined with rule-based heuristics, and a conventional business intelligence pipeline using gradient boosting. All models were evaluated on the same test set and using the same classification labels, where technically applicable. Because the base configurations do not share identical architectural or manufacturing capabilities, the comparison should be interpreted as an applied performance benchmark and not a definitive model ranking exercise. The comparative analysis of these capabilities is presented in Table 7.

These results suggest that, under the specific evaluation conditions of this study, the Agilance system achieved stronger aggregate performance than the selected baseline configurations.

3.3. Theoretical contributions to strategic accounting

This study contributes to the literature on strategic accounting and internal control in several fundamental ways.

First, the authors' core contribution is the conceptualization and development of the Agilance architecture as a framework for continuous automated auditing and real-time financial oversight. By transitioning from traditional retrospective reporting models to proactive, agentic financial planning, this study extends the theoretical foundations of accounting systems beyond mere prediction toward autonomous strategic reasoning. This contribution responds directly to the research gap identified by Cao (2022) and Li et al. (2023), who emphasize the need for systems capable of integrating heterogeneous financial signals into auditable, actionable strategies.

Second, the study advances the theory of internal control systems within the context of Industry 4.0. By demonstrating how real-time big data streams can be fused into a unified, contextually stabilized pipeline, the authors provide a new conceptual framework for digital accounting transformation. This ensures that automated financial decisions remain compliant with organizational policies and internal control standards, extending existing conceptual frameworks on Industry 4.0 accounting systems (Lehenchuk et al., 2025a, 2025b; Thanasas & Kampiotis, 2024).

Third, rather than focusing purely on the architectural mechanics of explainable artificial intelligence, this study contributes to auditing theory by embedding interpretability directly into the financial decision-making process. The developed mechanisms ensure that automated recommendations are fully transparent and verifiable, bridging the gap between complex data analytics and professional accounting judgment. This approach addresses limitations identified in the literature, where external interpretation tools often fail to capture domain-specific reasoning patterns necessary for professional audits (Raiaan et al., 2024; Long et al., 2025).

Fourth, the introduction of a dual-headed prioritization mechanism contributes to strategic management accounting literature by introducing a formalized, computational approach to ranking financial strategies based on both strategic impact and implementation urgency. This aligns with contemporary theories emphasizing the critical importance of adaptive and robust strategic control systems in dynamic economic environments (Kotter et al., 2025).

3.4. Practical implications

From a practical perspective, the findings indicate that Agilance can significantly improve decision-making efficiency in organizations operating in volatile markets. The reduction in planning cycle duration and operational costs suggests that integrating explainable AI tools into accounting workflows enhances organizational responsiveness and resource allocation efficiency. As detailed in Table 8, the empirical findings demonstrate that investing in explainable AI systems generates measurable economic benefits.

The indicators in Table 8 should be interpreted as illustrative internal evidence rather than independently verified economic proof.

The data presented in Table 8 highlights a 5.4% reduction in operational costs and a 46% first-year return on investment (ROI). Specifically, the 37% reduction in planning cycle time and the 13.2% improvement in variance-to-budget collectively indicate that the system's efficiency gains translate into tangible financial out-

Table 8. Illustrative internal economic indicators: ROI, costs, and budgeting precision

Source: Compiled by the authors.

Metric	Before	After	Change
Planning Cycle Time (days)	4.1	2.6	–37%
Operational Cost (USD M)	18.5	17.5	+5.4% savings
Forecast Error (MAPE)	6.8%	5.9%	–0.9 pp
Variance-to-Budget	22.4%	19.4%	–13.2%
Decision Adoption Rate	74.5%	81.0%	+6.5 pp
First-Year ROI	–	–	+46%

Note: pp – percentage points; MAPE – Mean Absolute Percentage Error.

comes at the organizational level. These results are consistent with prior research highlighting the positive relationship between digital transformation initiatives and financial performance (Lehenchuk et al., 2025a).

A particularly significant finding is the 6.5 percentage point increase in the decision adoption rate (from 74.5% to 81.0%). This suggests that the framework's explainability features contribute not only to analytical quality but also to the practical uptake of AI-generated recommendations by human decision-makers. Such transparency improves trust among financial analysts, auditors, and regulatory stakeholders, facilitating the broader adoption of AI technologies in accounting practice (Yeo et al., 2025).

The proposed system is especially valuable for multinational organizations that require continuous financial adjustments across different regulatory environments and time zones. By enabling real-time integration of financial signals, Agilance supports proactive rather than reactive decision-making, which is increasingly critical in highly competitive global markets (Malik, 2024).

Crucially, implementing the Agilance framework does not replace the accounting professional; rather, it fundamentally elevates their organizational role. By liberating accountants and auditors from routine, time-consuming tasks like manual data collection, the system allows them to focus on higher-value activities, such as professional judgment, strategic consulting, and nuanced risk assessments. This shift transforms the traditional accountant from a retrospective record-keeper into a proactive, forward-looking strategic partner.

3.5. Limitations and future research

Several limitations must be acknowledged. First, the dataset and evaluation environment are confidential. Although the manuscript reports the corpus structure, preprocessing logic, labeling criteria, evaluation protocol, baseline logic, and aggregated indicators, the underlying source data and job-level records are not available for independent inspection. This limits direct reproducibility and means that the numerical results should be treated as internal benchmark evidence.

Second, generalizability across industries and regulatory contexts has not yet been established. The curated corpus contains 8,012 cleaned, usable pages of specialized financial documentation, but it may not capture the reporting conventions, risk profiles, and data-governance constraints of all sectors. External validation is therefore needed before the framework can be considered broadly applicable.

Third, the evaluation of robustness was conducted through controlled internal scenarios, including rolling-origin testing, noise injection, and simulated disruption conditions. These procedures are useful for preliminary stress testing, but they do not fully replicate prolonged macroeconomic crises, abrupt regulatory change, or severe supply-chain disruption.

Fourth, the expert explainability assessment provides structured practitioner feedback but remains limited by panel composition, output selection, rating design, and the absence of an externally administered review process. Future studies should report full expert demographics where possible, use blinded independent panels, and apply systematic qualitative coding of reviewer comments.

http://dx.doi.org/10.21511/afc.07(2).2026.05

Fifth, deployment raises governance and privacy concerns. Organizations using centralized financial data must protect trade secrets, client information, and regulatory compliance obligations, including data-protection requirements. Decentralized or federated approaches may be necessary in settings where sensitive information cannot be pooled.

Future research should validate Agilance using anonymized multi-organizational datasets, externally audited evaluation protocols, or synthetic benchmark datasets designed for strategic financial planning. Comparative studies across sectors, jurisdictions, and organizational sizes are needed to determine whether the internal results reported here are replicable beyond the original evaluation setting.

CONCLUSION

This study aimed to develop Agilance as a conceptual and technical framework for explainable strategic financial planning supported by Artificial Intelligence, under controlled and confidential internal evaluation conditions.

The internal assessment provided strong but preliminary evidence of the technical reliability, explanatory power, and potential operational utility of the framework. Agilance performed well in ranking strategic priorities and implementation duration, while expert feedback indicated that its recommendations were generally clear, credible, and actionable for internal control review. Indicative operational results also suggest that the framework can reduce planning cycle delays and support faster financial decision-making in data-driven organizations.

Overall, the findings suggest that Agilance can enhance strategic financial planning by combining automated prioritization with AI-powered auditable reasoning. However, because the underlying dataset, operational records, and expert assessment materials remain confidential, the results should be interpreted as internal benchmarks and not as evidence of broad generalizability, scalability, or economic impact. Future research should validate the framework through independent, externally audited, and multi-organizational evaluation protocols.

AUTHOR CONTRIBUTIONS

Conceptualization: Georgios Kampiotis, Georgios L. Thanasas, Iryna Zhyhlei.
Data curation: Georgios Kampiotis, Iryna Hrabchuk, Iryna Zhalinska.
Formal analysis: Georgios Kampiotis, Iryna Zhyhlei, Iryna Hrabchuk, Iryna Zhalinska.
Funding acquisition: Georgios Kampiotis, Georgios L. Thanasas.
Investigation: Georgios Kampiotis, Georgios L. Thanasas.
Methodology: Georgios Kampiotis, Georgios L. Thanasas.
Project administration: Iryna Zhyhlei.
Resources: Georgios Kampiotis, Georgios L. Thanasas, Iryna Zhyhlei, Iryna Hrabchuk, Iryna Zhalinska.
Software: Georgios Kampiotis, Georgios L. Thanasas.
Supervision: Georgios L. Thanasas.
Validation: Iryna Hrabchuk, Iryna Zhalinska.
Visualization: Iryna Hrabchuk, Iryna Zhalinska.
Writing – original draft: Georgios Kampiotis, Georgios L. Thanasas, Iryna Zhyhlei.
Writing – review & editing: Georgios Kampiotis, Georgios L. Thanasas, Iryna Zhyhlei, Iryna Hrabchuk, Iryna Zhalinska.

ACKNOWLEDGMENTS

The publication fees of this manuscript have been financed by the MSc Tax and Financial Services Digital Transformation (DITAF), University of Patras.

REFERENCES

1. Bahoo, S., Cucculelli, M., Goga, X., & Mondolo, J. (2024). Artificial intelligence in finance: A comprehensive review through bibliometric and content analysis. *SN Business & Economics, 4*(2), 23. https://doi.org/10.1007/s43546-023-00618-x
2. Balogun, E. D., Ogunsola, K. O., & Ogunmokun, A. (2022). Developing an advanced predictive model for financial planning and analysis using machine learning. *IRE Journals, 5*(11), 320-328. Retrieved from https://www.irejournals.com/paper-details/1703426
3. Bi, S., Xiao, J., & Deng, T. (2024). The role of AI in financial forecasting: ChatGPT's potential and challenges. In *Proceedings of the 4th Asia-Pacific Artificial Intelligence and Big Data Forum* (pp. 1064-1070). https://doi.org/10.1145/3718491.3718663
4. Bodria, F., Giannotti, F., Guidotti, R., Naretto, F., Pedreschi, D., & Rinzivillo, S. (2023). Benchmarking and survey of explanation methods for black box models. *Data Mining and Knowledge Discovery, 37*(5), 1719-1778. https://doi.org/10.1007/s10618-023-00933-9
5. Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys, 55*(3), 1-38. https://doi.org/10.1145/3502289
6. Carè, R., & Cumming, D. (2024). Technology and automation in financial trading: A bibliometric review. *Research in International Business and Finance, 71,* 102471. https://doi.org/10.1016/j.ribaf.2024.102471
7. Chukwuma-Eke, E. C., Ogunsola, O. Y., & Isibor, N. J. (2022). A conceptual approach to cost forecasting and financial planning in complex oil and gas projects. *International Journal of Multidisciplinary Research and Growth Evaluation,* 3(1), 819-833. https://doi.org/10.54660/.IJMRGE.2022.3.1.819-833
8. Dobson, J. E. (2023). On reading and interpreting black box deep neural networks. *International Journal of Digital Humanities,* 5(2), 431-449. https://doi.org/10.1007/s42803-023-00075-w
9. Hafez, I. Y., Hafez, A. Y., Saleh, A., Abd El-Mageed, A. A., & Abohany, A. A. (2025). A systematic review of AI-enhanced techniques in credit card fraud detection. *Journal of Big Data, 12*(1), 6. https://doi.org/10.1186/s40537-024-01048-8
10. Hosny, K. M., & Mohammed, M. A. (2025). Explainable AI and vision transformers for detection and classification of brain tumor: A comprehensive survey. *Artificial Intelligence Review, 58*(9), 1-60. https://doi.org/10.1007/s10462-025-11221-x
11. Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research, 40*(2), 806-841. https://doi.org/10.1111/1911-3846.12832
12. Jiang, Z. -H., Hou, Q., Yuan, L., Zhou, D., Shi, Y., Jin, X., Wang, A., & Feng, J. (2021). All tokens matter: Token labeling for training better vision transformers. *Advances in Neural Information Processing Systems, 34,* 18590-18602. https://doi.org/10.48550/arXiv.2104.10858
13. Joshi, S. (2025). Review of gen AI models for financial risk management. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 11*(1), 709-723. https://doi.org/10.32628/CSEIT2511114
14. Kandasamy, R., Ramesh, R., Kumar, R., Raghu, N., & Chavadi, C. A. (2025). Agile methodologies in digital marketing for growing businesses. In D. Thangam, H. Chittoo, I. Ghosal, R. Kandasamy, & J. Park (Eds.), *Expanding operations through agile principles and sustainable practices* (pp. 445-468). IGI Global. https://doi.org/10.4018/979-8-3693-6274-7.ch023
15. Kerr, D., Smith, K. T., Smith, L. M., & Xu, T. (2025). A review of AI and its impact on management accounting and society. *Journal of Risk and Financial Management, 18*(6), 340. https://doi.org/10.3390/jrfm18060340
16. Kotter, J. P., Akhtar, V., & Gupta, G. (2025). *Change: How organizations achieve hard-to-imagine results in uncertain and volatile times* (240 p.). John Wiley & Sons.
17. Kumar, K. P., Swarubini, P. J., & Ganapathy, N. (2025). Cognitive artificial intelligence. In D. Samanta, A. Sivakumar, & S. Pal (Eds.), *Artificial intelligence and biological sciences* (pp. 301-323). CRC Press. https://doi.org/10.1201/9781003492726-18
18. Lehenchuk, S. F., Valinkevych, N. V., Vyhivska, I. M., & Khomenko, H. Y. (2020). The significant principles of development of accounting support for innovative enterprise financing. *International Journal of Advanced Science and Technology, 29*(8 Special), 2282-2289. Retrieved from http://sersc.org/journals/index.php/IJAST/article/view/14701
19. Lehenchuk, S., Zakharov, D., Fedorova, O., Horodyskyi, M., & Vavilov, D. (2025a). Digital transformation, research and development, and financial performance of agricultural companies. *Agricultural and Resource Economics, 11*(3), 46-70. https://doi.org/10.22004/ag.econ.387544
20. Lehenchuk, S., Zhyhlei, I., & Zakharov, D. (2025b). Development of a Strategic Accounting and Control System for Freight Forwarding Services in the Industry 4. 0 Environment: Opportunities and Challenges. *Public Policy and Accounting, 2*(12), 25-32. https://doi.org/10.26642/ppa-2025-2%2812%29-25-32
21. Li, P., & Zhang, L. (2025). Application of big data technology in enterprise information security management. *Scientific Reports, 15*(1), 1022. https://doi.org/10.1038/s41598-025-85403-6

22. Li, Y., Wang, S., Ding, H., & Chen, H. (2023). Large language models in finance: A survey. In *Proceedings of the 4th ACM International Conference on AI in Finance* (pp. 374-382). ACM. https://doi.org/10.1145/3604237.3626869

23. Long, S., Tan, J., Mao, B., Tang, F., Li, Y., Zhao, M., & Kato, N. (2025). A survey on intelligent network operations and performance optimization based on large language models. *IEEE Communications Surveys & Tutorials, 27*(6), 3915-3949. https://doi.org/10.1109/COMST.2025.3526606

24. Mahendran, M. B., Gokul, A. K., Lakshmi, P., & Pavithra, S. (2025). Comparative advances in financial sentiment analysis: A review of BERT, FinBERT, and large language models. In *2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)* (pp. 39-45). IEEE. https://doi.org/10.1109/IDCIOT64235.2025.10914764

25. Malik, S. (2024). Data-driven decision-making: Leveraging the IoT for real-time sustainability in organizational behavior. *Sustainability, 16*(15), 6302. https://doi.org/10.3390/su16156302

26. Mathen, M. P., & Paul, A. (2025). Toward an evolving framework for responsible AI for credit scoring in the banking industry. *Journal of Information, Communication and Ethics in Society, 23*(1), 148-163. https://doi.org/10.1108/JICES-08-2024-0122

27. Mittal, U., Sai, S., Chamola, V., & Sangwan, D. (2024). A comprehensive review on generative AI for education. *IEEE Access, 12,* 142733-142759. https://doi.org/10.1109/access.2024.3468368

28. Moharrak, M., & Mogaji, E. (2025). Generative AI in banking: Empirical insights on integration, challenges and opportunities in a regulated industry. *International Journal of Bank Marketing, 43*(4), 871-896. https://doi.org/10.1108/IJBM-08-2024-0490

29. Moulaei, K., Yadegari, A., Baharestani, M., Farzanbakhsh, S., Sabet, B., & Afrash, M. R. (2024). Generative artificial intelligence in healthcare: A scoping review on benefits, challenges and applications. *International Journal of Medical Informatics, 188,* 105474. https://doi.org/10.1016/j.ijmedinf.2024.105474

30. Obeng, H. A., Arhinful, R., Mensah, L., & Mensah, C. C. (2025). The mediating role of service quality in the relationship between corporate social responsibility and sustainable competitive advantages in an emerging economy. *Business Strategy & Development, 8*(1), e70099. https://doi.org/10.1002/bsd2.70099

31. Olomu, M. O., Binuyo, G. O., & Oyebisi, T. O. (2023). The adoption and impact of internet-based technological innovations on the performance of the industrial cluster firms. *Journal of Economy and Technology, 1,* 164-178. https://doi.org/10.1016/j.ject.2023.11.004

32. Raiaan, M. A. K., Mukta, M. S. H., Fatema, K., Fahad, N. M., Sakib, S., Mim, M. M. J., Ahmad, J., Ali, M. E., & Azam, S. (2024). A review on large language models: Architectures, applications, taxonomies, open issues and challenges. *IEEE Access*, 12, 26839-26874. Retrieved from https://ieeexplore.ieee.org/document/10433480

33. Ren, S. (2022). Optimization of enterprise financial management and decision-making systems based on big data. *Journal of Mathematics, 2022*(1), 1708506. https://doi.org/10.1155/2022/1708506

34. Terzidou, K. (2025). Generative AI systems in legal practice offering quality legal services while upholding legal ethics. *International Journal of Law in Context*, 21(3), 431-452. https://doi.org/10.1017/S1744552325000047

35. Thanasas, G. L., & Kampiotis, G. (2024). The role of Big Data Analytics in financial decision-making and strategic accounting. *Technium Business and Management, 10,* 17-33. https://doi.org/10.47577/business.v10i.11877

36. Thanasas, G., Kampiotis, G., & Halkiopoulos, C. (2026). Transforming Digital Accounting: Big Data, IoT, and Industry 4. 0 Technologies – A Comprehensive Survey. *Journal of Risk and Financial Management, 19*(1), 92. https://doi.org/10.3390/jrfm19010092

37. Xiaoqi, Z., & Ali, M. (2024). Financial risk management in the digital age. In M. Ali, L. Choi-Meng, C.-H. Puah, S. A. Raza, & P. Sivanandan (Eds.), *Strategic financial management* (pp. 49-69). Emerald Publishing Limited. https://doi.org/10.1108/978-1-83608-106-720241004

38. Yeo, W. J., Van Der Heever, W., Mao, R., Cambria, E., Satapathy, R., & Mengaldo, G. (2025). A comprehensive review on financial explainable AI. *Artificial Intelligence Review, 58*(6), 1-49. https://doi.org/10.48550/arXiv.2309.11960

APPENDIX A

Table A1. Detailed technical logic of the Agilance dataset preprocessing algorithm

Source: Compiled by the authors.

Algorithm 1: Agilance Dataset Preprocessing
Require: Raw document Set R
Ensure: Preprocessed Sets D_{train}, D_{val}, D_{test}
Initialize $D \leftarrow \emptyset$
Remove headers, footers, and cover pages from R
for all $r \in R$ **do**
if r is Financial Case Study **then**
Add pair (F_m, FP_m) to D
else if r is Planning Theory **then**
Add pair (T_n, PT_n) to D
else if r is Analytical Report **then**
Add pair (A_p, AR_p) to D
end if
end for
for all $(x, y) \in D$ **do**
Convert x and y to lowercase
Remove irrelevant symbols and noise
Tokenize into tokens ω_q
Embed tokens to vectors $v_q = E(\omega_q)$
Apply padding/truncation to fixed length L
Generate attention mask M_q
end for
Split D into: $D_{train} \leftarrow 70\%$, $D_{val} \leftarrow 15\%$, $D_{test} \leftarrow 15\%$
return D_{train}, D_{val}, D_{test}

APPENDIX B. Technical infrastructure and operational environment of the Agilance framework

The implementation system architecture has four operational modules: ingestion, preprocessing, model inference, and output orchestration. We utilized Debezium connectors to generate a data stream from PostgreSQL within the organization's internal ERP. The external signals from Bloomberg, Refinitiv, Quandl, Twitter, and NewsAPI. org have been fetched through secured APIs using the Asynchronous HTTP client/server (AIOHTTP) framework. These external sources produce big data, which is buffered using Apache Kafka, a distributed log. Spark Streaming consumes these Kafka topics in real time for schema normalization and transformation. Finally, the pre-processed data are stored in MongoDB in a time-series model. To ensure proper real-time inference by the transformer, the data are also cached in Redis for low-latency access.

The backend of the proposed Agilance framework has been developed using multiple tools and technologies summarized in Table B1. The framework's core is implemented in Python 3.10, using the PyTorch 2.1 library. The initial baseline transformer was imported from the Hugging Face library. For big data management and processing, Kafka 3.6 and Spark 3.5 have been used. MongoDB 6.0 and Redis 7.2 were used for storage and caching, respectively. We used Flask and Gunicorn over HTTPS with

http://dx.doi.org/10.21511/afc.07(2).2026.05

NGINX and maintained communication via the REST API. All components of the Agilance framework were containerized via Docker to ensure platform-independent operations. For real-time monitoring, Prometheus and Grafana were used.

Table B1. System environment and big data components

Source: Compiled by the authors.

Component	Specification
OS	Ubuntu 20.04 LTS
GPU	NVIDIA A100 (48 GB)
CPU	Intel Xeon Gold (32 cores)
RAM	512 GB DDR4 ECC
Data Bus	Apache Kafka 3.6
Stream Engine	Apache Spark 3.5
Storage	MongoDB 6.0, Redis 7.2
Model Stack	PyTorch, HuggingFace
Serving	Flask + Gunicorn
Containerization	Docker + Compose
Monitoring	Prometheus, Grafana

To make the transformer effective, we first serialized it in PyTorch format. The application has been developed using Flask, which loads the transformer and serves it through Gunicorn. The implemented framework exposes a REST endpoint on port 8080. It accepts JSON-encoded POST requests only. All incoming data undergoes real-time preprocessing and is subsequently mapped into high-dimensional dense vectors via the embedding function, as detailed in the preprocessing logic of Algorithm 1 (Appendix A). Once compatible, they are passed to the transformer. It classifies the input using a dual-headed output layer. Each response includes the generated strategy along with priority and urgency scores. An additional set of enterprise dashboards, integrated with analytics systems, has been developed and interfaced with the Agilance transformer via the REST API to receive responses. We achieved horizontal scaling through Docker- based worker replicas.

APPENDIX C. Technical specifications and mathematical optimization of the Agilance model training

To transform the pre-trained Agilance architecture into a specialized expert agent for financial planning, the model underwent targeted fine-tuning. The primary objective was to minimize cross-entropy loss, effectively aligning the transformer's predictive probability distributions with the curated expert dataset. Rather than using static optimization parameters, the training process used the ADAM optimizer, which facilitated smooth weight updates by dynamically adjusting learning rates based on moving averages. Stability during this high-dimensional learning phase was further ensured through a tailored learning rate schedule that incorporated a warm-up period to prevent gradient instability and a structured batch-processing strategy over 12 training epochs. As illustrated in the learning curves (Figure 4), this refined training environment enabled the model to achieve rapid convergence and high reliability in management forecasts, without the common pitfalls of overfitting in complex financial domains.

The core scientific contribution of this study lies in the structural modification of the transformer blocks to address identified theoretical gaps in autonomous financial planning. Specifically, the Agilance system introduces three proprietary layers that function as the "brain" of the agentic framework:

1) Financial Relevance Weighting (FRW). This layer directly addresses the critical requirement for explainability and transparency in financial control. By applying saliency-based scalar modulation,

the FRW highlights specific monetary terms and strategic indicators driving the model's decisions. This mechanism allows analysts to move beyond the "black box" nature of traditional AI by generating visual saliency maps that verify whether the model is focusing on the correct financial signals.

2) Contextual Drift Stabilizer (CDS). Designed to ensure contextual integrity, the CDS prevents the phenomenon of "hallucination" (semantic drift) during long-term planning cycles. By blending current activations with a gated memory of organizational policy, it ensures that the generated strategies remain strictly aligned with the specific financial context and historical planning patterns.

3) Output Compression Gate (OCG). Situated at the final stage of the analytical processing pipeline, the OCG optimizes the actionability of the system's output. By applying entropy-based pruning to the transformer's logits, it filters out ambiguous or low-confidence predictions, ensuring that the resulting strategies are precise, ranked by priority, and ready for immediate implementation in an agile environment.

APPENDIX D. Mathematical formalization of transformer enhancements, signal fusion, and strategy prioritization

Financial Relevance Weighting (FRW). This layer applies saliency-based scalar modulation to highlight tokens with high strategic weight, ensuring the model prioritizes critical financial terminology over generic linguistic noise. The transformation is defined as (Equation D1):

$$z_i^{'} = \lambda_i \cdot z_i, \text{ where } \lambda_i = f_{saliency}. \tag{D1}$$

Contextual Drift Stabilizer (CDS). The CDS mechanism prevents "hallucination" and maintains strict alignment with organizational policy by blending current activations with a gated memory of past planning contexts. This stabilization of time-sensitive drift is expressed as Equation D2:

$$\tilde{h}_i = \gamma \cdot h_i + (1-\gamma) \cdot Gate(h_i, p_i). \tag{D2}$$

Output Compression Gate (OCG). Situated before the final probability distributions of financial strategies (softmax activation), this entropy-based pruning mechanism filters out ambiguous or low-confidence predictions to produce more actionable financial strategies. The filtering behavior is defined as Equation D3:

$$o_j^{'} = \{o_j - \infty \text{ if } entropy(o_j) < \theta, \text{ otherwise.} \tag{D3}$$

The primary theoretical proposal for data acquisition is defined by the integration of internal and external streams (Equation D4):

$$B = B_{\text{int}} \cup B_{ext}, \tag{D4}$$

where B_{int} represents structured organizational data and B_{ext} encompasses unstructured signals from global financial APIs such as Bloomberg, Refinitiv, Quandl, and Twitter.

To bridge the gap between raw information and agentic reasoning, these streams undergo a semantic normalization process (Φ(B)), converting diverse inputs into domain-aligned intermediate representations (Equation D5).

http://dx.doi.org/10.21511/afc.07(2).2026.05

$$X = \Phi(B) = \{x_1, x_2, \ldots, x_n\}, \quad x_i \in R^d. \tag{D5}$$

These heterogeneous data streams undergo a feature-extraction and semantic-normalization process denoted by *Φ(B)* as shown in, which converts raw data into domain-aligned intermediate representations. This mathematical transformation ensures that the transformer agent ($T_{Agilance}$) receives inputs optimized for autoregressive reasoning.

The final and most critical stage for financial management is the generation and prioritization of strategies. The agent functions as a transformative logic that produces a set of planning units (*S*) (Equation D6).

$$S = T_{Agilance}(X). \tag{D6}$$

Each output strategy $s_j \in S$ is a structured planning unit annotated with two key meta-tags: Strategic Priority Score (*SP*) ∈ [0, 1] and Execution Urgency Score (*EU*) ∈ [0, 1]. These scores are calculated by the dual-headed classification mechanism expressed in Equation D7.

$$SP_j, EU_j = \Psi(s_j). \tag{D7}$$

The function Ψ(·) assigns priority and urgency based on temporal relevance, financial impact, and risk alignment. Agilance's output is a prioritized, ranked set of financial strategies, as expressed in Equation D8.

$$S^* = Rank\left(\left\{s_j, SP_j, EU_j\right\}_{j=1}^{k}\right). \tag{D8}$$

APPENDIX E. Detailed machine learning validation and confusion matrix analysis

This appendix provides the confusion matrices and cross-validation performance visualizations that empirically validate the predictive accuracy of the Agilance model across different priority levels and time horizons.

Source: Compiled by the authors.

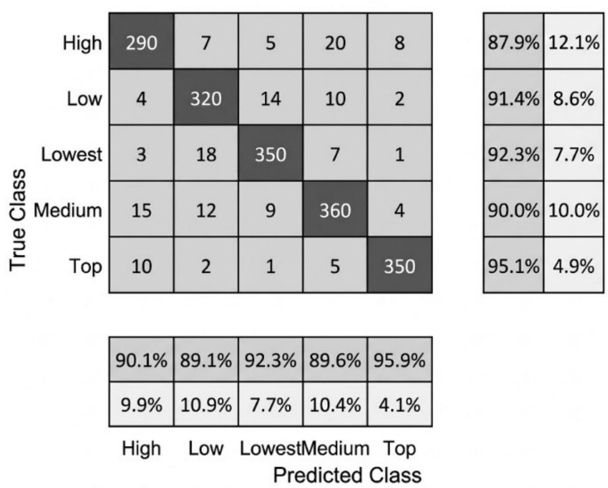

Figure E1. Validation of system precision in determining strategic financial priorities using Agilance

Source: Compiled by the authors.

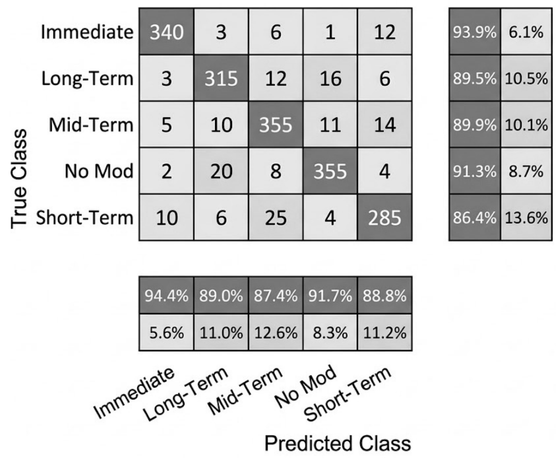

Figure E2. Assessment of system accuracy in categorizing financial plan implementation timelines using Agilance

Source: Compiled by the authors.

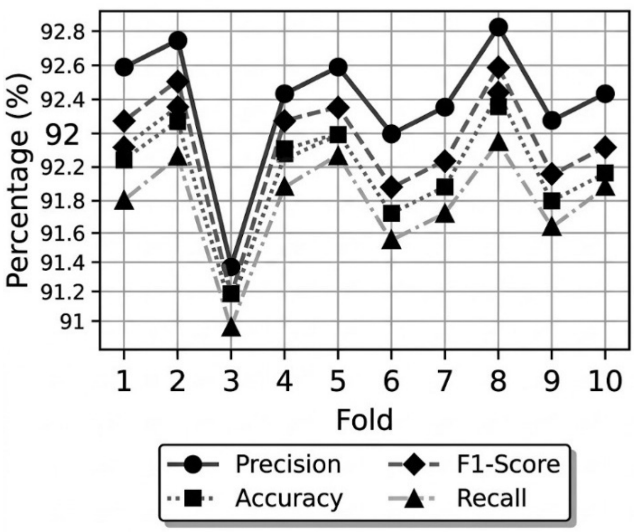

Figure E3. Stability and consistency of predictive accuracy across multiple validation iterations using the Agilance framework

APPENDIX F. Operational scalability and hardware resource utilization

This appendix details the hardware resource consumption, processing latency, and load-testing metrics, confirming the operational scalability of the Agilance framework under high-throughput conditions.

Table F1. Analytical processing time and resource usage for Agilance across 10 workloads

Source: Compiled by the authors.

Test Case	Analytical Processing Time (ms)	CPU Usage (%)	GPU Usage (%)	Memory Usage (MB)	Power Consumption (W)
1	110	29	66	502	86
2	118	31	70	516	91
3	125	35	73	528	94
4	113	30	68	510	88

http://dx.doi.org/10.21511/afc.07(2).2026.05

Table F1 (cont.). Analytical processing time and resource usage for Agilance across 10 workloads

Test Case	Analytical Processing Time (ms)	CPU Usage (%)	GPU Usage (%)	Memory Usage (MB)	Power Consumption (W)
5	109	28	65	499	86
6	122	33	72	525	93
7	116	32	71	518	92
8	112	30	69	511	89
9	120	34	72	524	93
10	114	31	70	514	90
Average	116	31.3	69.6	514.5	90.2

Source: Compiled by the authors.

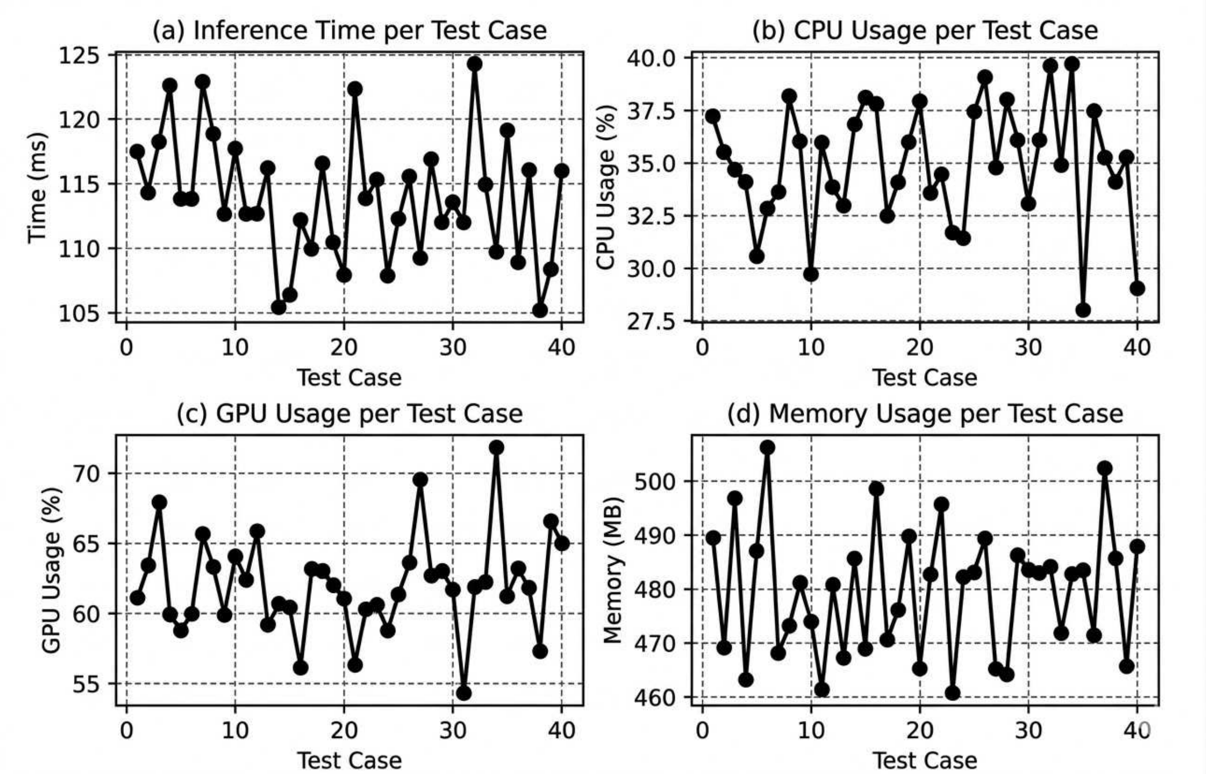

Figure F1. Resource consumption profile of Agilance across 10 workloads

Source: Compiled by the authors.

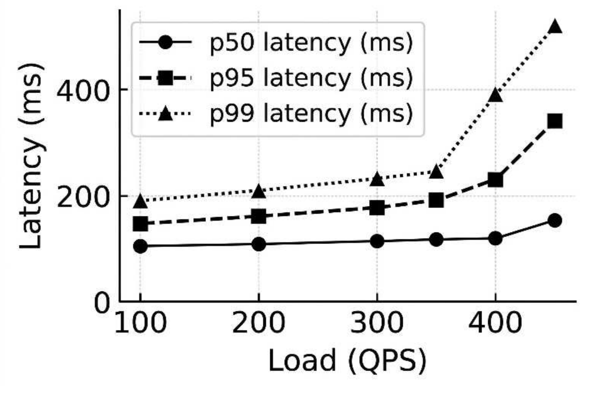

Figure F2. Throughput vs latency under increasing load. Agilance maintains low median (p50) and tail (p95, p99) latencies up to 350 QPS, with maximum sustainable throughput at 430 QPS before SLA violations

Table F2. Scalability and SLA performance metrics

Source: Compiled by the authors.

Load (QPS)	p50 (ms)	p95 (ms)	p99 (ms)	Sust. TPS
100	105	148	191	100
200	109	162	210	200
300	115	178	233	300
350	118	192	247	350
400	120	231	392	400
450	154	342	521	430*

Note: * – maximum sustainable throughput before SLA violations.

 http://dx.doi.org/10.21511/afc.07(2).2026.05