Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

LLM-Augmented Trading and Decision Platforms: Bridging Generative Intelligence with Financial Decision Systems

Gangadharan Venkataraman
Independent Researcher, USA

ARTICLE INFO

Received: 31 Dec 2025
Revised: 03 Jan 2026

ABSTRACT

Financial markets generate substantial volumes of unstructured textual information influencing asset valuations and trading behaviors. Traditional algorithmic trading systems process numerical indicators effectively yet lack semantic reasoning capabilities for interpreting news articles, earnings transcripts, regulatory filings, and social media content. Large Language Models offer remarkable text comprehension abilities suitable for financial sentiment extraction. Deploying such computationally intensive models within latency-sensitive trading environments presents significant architectural challenges. The proposed framework addresses the gap between semantic understanding and high-frequency execution requirements through optimized inference pipelines and structured signal taxonomies. Multi-source information aggregation enables parallel processing of heterogeneous data streams including news feeds and corporate communications. Event extraction and entity recognition transform raw text into structured representations suitable for downstream processing. Hierarchical signal classification converts sentiment outputs into actionable trading recommendations across multiple time horizons. Deep reinforcement learning agents interface with generated signals for adaptive strategy optimization. Portfolio allocation modules integrate textual intelligence with quantitative risk constraints for comprehensive decision support. Simulation environments enable systematic evaluation of latency, accuracy, and portfolio-level performance without capital exposure. The framework bridges generative language understanding with practical trading system requirements.

Keywords: Large Language Models, Financial Sentiment Classification, Deep Reinforcement Learning, Portfolio Optimization, Real-Time Signal Generation, Algorithmic Trading Systems

I. INTRODUCTION

Economic markets produce giant volumes of unstructured textual information every day. These facts consist of profit reviews, regulatory filings, analyst opinions, and information articles. Such textual content carries significant signals that influence asset prices and investor behavior. Traditional quantitative models focus primarily on numerical indicators and historical price patterns. These models often fail to capture the semantic richness embedded in financial communications. Natural language processing has emerged as a crucial device for extracting intelligence from textual resources. Current surveys spotlight that NLP programs in finance span sentiment evaluation, named entity recognition, query answering, and document summarization [1]. The financial domain presents unique linguistic challenges due to specialized terminology and context-dependent interpretations [1].

Sentiment analysis remains one of the most extensively studied NLP applications in financial contexts. Early approaches relied on domain-specific lexicons to classify textual sentiment. These lexicon-based methods struggled with negation handling and contextual nuances. The introduction of deep learning architectures significantly improved sentiment classification accuracy. Transformer-based models represent a substantial advancement over recurrent neural network approaches. FinBERT demonstrated that pre-training language models on financial

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

corpora yields superior performance on sentiment tasks [2]. The model adapts the BERT architecture specifically for financial text understanding [2]. Pre-training on financial news and communications enables better comprehension of domain-specific vocabulary [2]. This domain adaptation addresses the vocabulary mismatch problem inherent in general-purpose language models.

Large Language Models extend these capabilities through enhanced contextual understanding. Modern LLMs can process lengthy documents while maintaining coherence across distant text segments. The attention mechanisms within transformer architectures capture long-range dependencies effectively. Financial documents often contain complex conditional statements and forward-looking projections. Interpreting such content requires sophisticated reasoning beyond simple pattern matching. LLMs demonstrate capacity for nuanced interpretation of earnings guidance and regulatory language.

Deploying language models within trading environments introduces significant architectural constraints. Highfrequency trading operates within microsecond timeframes. The latency between information arrival and trading response determines competitive advantage. Standard LLM inference introduces computational delays incompatible with rapid execution requirements. Model optimization techniques, including quantization and distillation, address latency concerns. Balancing inference speed with analytical accuracy remains a fundamental engineering challenge.

The integration of language understanding with trading systems requires structured signal taxonomies. Raw sentiment outputs must translate into actionable trading recommendations. Hierarchical classification frameworks enable mapping between textual analysis and execution decisions. Reinforcement learning mechanisms provide adaptive capabilities for evolving market conditions. The combination of generative language understanding with adaptive execution strategies creates comprehensive decision support architectures. This framework bridges semantic comprehension capabilities with operational latency constraints in production trading environments.

II. RELATED WORK

Natural language processing applications in financial domains have evolved significantly over recent years. Early sentiment classification relied on lexicon-based methods using predefined word lists. Dictionary-based techniques struggled with negation handling and contextual interpretation challenges. Deep learning architectures subsequently transformed financial text processing capabilities. Convolutional and recurrent neural networks improved sentiment classification accuracy substantially. Transformer-based models, including FinBERT, demonstrated superior performance through domain-specific pre-training on financial corpora.

Reinforcement learning applications in portfolio management emerged as a parallel development trajectory. Deep Q-networks enabled learning of trading policies from raw market data. Actor-critic architectures addressed continuous action spaces inherent in position sizing decisions. Practical implementations incorporated transaction costs and market friction modeling for realistic performance estimation.

The framework presented within the article synthesizes these distinct developments into a unified architecture. Multi-source data aggregation pipelines process heterogeneous textual streams simultaneously. Domain-adapted language models extract sentiment signals from financial communications. Hierarchical classification taxonomies convert semantic outputs into actionable trading recommendations. Deep reinforcement learning agents optimize execution strategies through environmental interaction. Bidirectional feedback mechanisms enable progressive refinement of both signal generation and policy learning components. Simulation-based evaluation addresses backtest overfitting concerns through deflated performance metrics. The integration bridges generative language understanding with quantitative trading system requirements effectively.

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

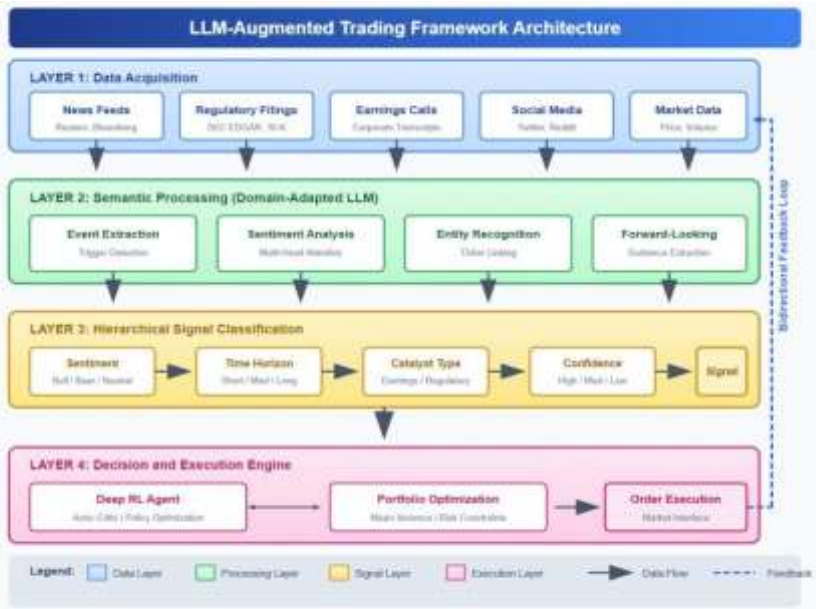

Fig 1. Complete system architecture of the LLM-augmented trading framework illustrating the four-layer design: data acquisition from heterogeneous sources, semantic processing through domain-adapted language models, hierarchical signal classification, and decision/execution engine with bidirectional feedback loop.

III. REAL-TIME FINANCIAL DATA ANALYSIS ARCHITECTURE

A. Multi-Source Information Aggregation

Financial markets generate diverse textual content across multiple channels. News agencies publish articles covering corporate developments and market movements. Regulatory bodies maintain repositories of obligatory filings and disclosures. Earnings calls produce transcripts containing control remarks and analyst questions. Social media systems host discussions reflecting retail investor sentiment. Each source type carries a distinct informational value for trading decisions.

The proposed framework establishes parallel ingestion pipelines for heterogeneous data streams. Distributed collection nodes monitor designated channels without interruption. News feeds require the separation of headline content from article bodies. Regulatory filings follow structured formats enabling systematic parsing. Earnings transcripts contain speaker turns requiring identification and attribution. Social media content demands filtering to remove noise and irrelevant posts.

Event extraction constitutes a fundamental preprocessing function within financial text analysis. Company-specific events include announcements, earnings releases, and management transitions. The SENTiVENT corpus provides annotated resources for supervised event extraction from economic news [3]. This dataset enables the training of models that identify event triggers and participant roles [3]. Event types span categories including financial results, employment changes, and corporate transactions [3]. The annotation scheme captures both event attributes and sentiment expressions within the same framework [3]. Temporal anchoring assigns timestamps to extracted events for proper sequencing. Entity linking connects organization mentions to standardized identifiers. The aggregation layer transforms heterogeneous inputs into unified representations for downstream processing.

B. Contextual Semantic Processing

Language model inference operates on aggregated content through domain-optimized architectures. Generalpurpose models exhibit limitations with specialized financial vocabulary. Domain adaptation through continued pre-training

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

addresses vocabulary gaps effectively. Financial texts contain ticker symbols and abbreviations requiring specialized tokenization approaches.

The architecture employs sliding context windows for streaming information processing. Recent market context proves essential for interpreting incoming content. Attention mechanisms enable models to cognizance on applicable input portions selectively. Transformer architectures leverage self-attention to seize dependencies throughout series positions. Research demonstrates that attention-based models achieve strong performance across diverse sequence processing tasks [4]. The dual-path structure enables processing of both local and global sequence relationships [4]. Attention weights learn to emphasize informative segments while suppressing irrelevant content [4]. Multi-head configurations provide multiple representation subspaces for complex relationships [4].

Forward-looking statements carry particular significance for trading applications. Management guidance influences analyst expectations and subsequent valuations. Sentiment-bearing phrases indicate directional bias within corporate communications. The processing layer identifies and prioritizes such content segments automatically. Quantitative guidance undergoes extraction into structured fields for numerical analysis. Asset-specific relevance scoring filters content pertinent to monitored securities. The semantic processing layer produces enriched representations combining textual understanding with structured metadata for signal generation.

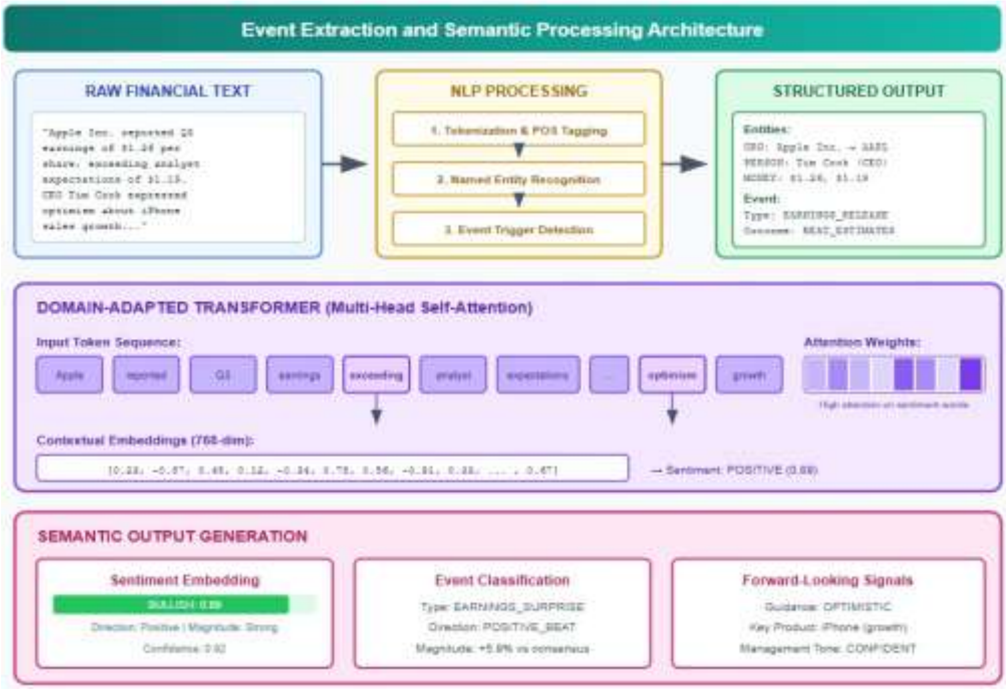

Fig 2. Event extraction and semantic processing architecture demonstrating the transformation of raw financial text through NLP processing, a domain-adapted transformer with multi-head attention, and generation of structured semantic outputs, including sentiment embeddings, event classification, and forward-looking signals.

[Note: Fig. 3 illustrates the complete semantic processing pipeline, showing how raw earnings text is transformed through named entity recognition, event extraction, and transformer-based attention mechanisms to produce structured sentiment and event classifications.]

Data Source Type	Content Characteristics	Preprocessing Requirements
News Feeds	Headlines and article bodies with breaking market information	Headline separation, temporal tagging, entity extraction

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

Regulatory Filings	Standardized format documents with compliance disclosures	Structured parsing, section identification, risk factor extraction
Earnings	Management commentary with speaker	Speaker diarization, forward-looking
Transcripts	attributions	statement detection
Social Media	Retail investor sentiment with high noise levels	Noise filtering, relevance scoring, sentiment normalization

Table 1. Data Source Types and Preprocessing Functions in Financial Text Analysis [3, 4]

IV. SIGNAL GENERATION AND PORTFOLIO OPTIMIZATION

A. Hierarchical Signal Classification

Trade signal generation converts semantic analysis outputs into actionable recommendations. Raw sentiment scores require transformation into discrete trading signals. A hierarchical classification framework provides structured signal organization. Primary signals categorize overall sentiment toward specific assets. Classifications span bullish, bearish, and neutral categories across different time horizons.

Financial sentiment classification presents unique challenges distinct from general sentiment analysis. Standard sentiment tools perform inadequately on financial terminology. Words carry different meanings in financial contexts versus everyday usage. Deep learning approaches address these domain-specific challenges effectively. Convolutional neural networks extract local features from financial text sequences [5]. Long short-term memory networks capture sequential dependencies within documents [5]. Attention mechanisms enable models to focus on sentiment-bearing text portions [5]. The combination of multiple deep learning architectures improves classification robustness [5]. Feature extraction through deep networks surpasses traditional bag-of-words representations [5]. The classification layer maps continuous model outputs to discrete signal categories for trading systems.

Secondary signals capture specific catalyst types, warranting position adjustments. Profit surprises suggest divergence among the mentioned effects and expectations. Regulatory announcements have an effect on compliance requirements and operational prices. Competitive developments alter market positioning and pricing dynamics. Each catalyst type receives a distinct signal classification and urgency weighting. The signal taxonomy enables systematic response protocols across information categories.

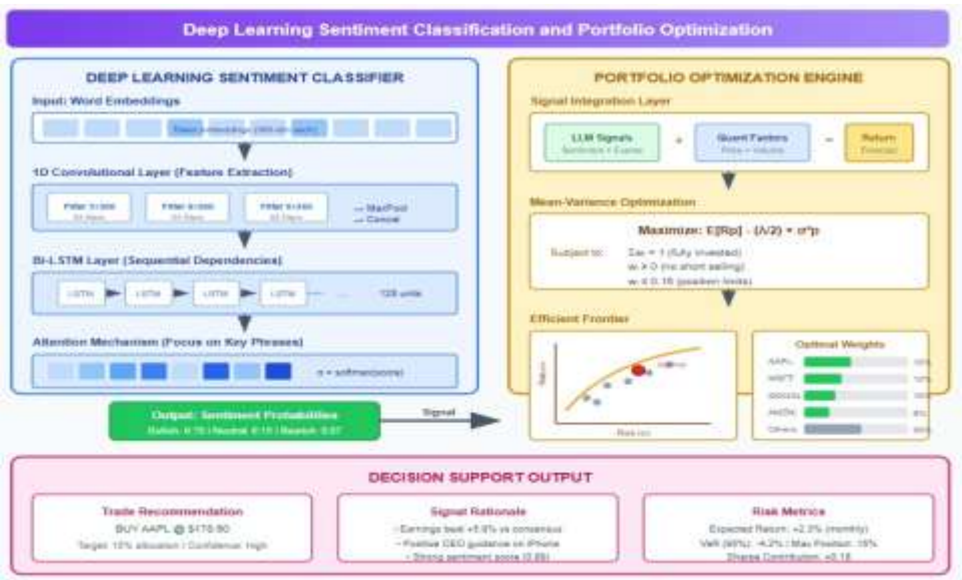

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

Fig 3.. Deep learning sentiment classification architecture featuring CNN-LSTM-Attention layers for text classification, integrated with a mean-variance portfolio optimization engine showing signal fusion, efficient frontier computation, and optimal weight allocation with decision support output.

[Note: Fig. 4 presents the integrated architecture combining deep learning sentiment classification with portfolio optimization, illustrating how CNN and LSTM layers extract features from financial text, attention mechanisms focus on sentiment-bearing phrases, and the resulting signals feed into mean-variance optimization to generate risk-constrained portfolio allocations.]

B. Portfolio Allocation Integration

Portfolio optimization modules consume signal inputs alongside traditional quantitative parameters. Meanvariance frameworks incorporate sentiment-derived return adjustments. Risk constraints maintain limits on position sizing and concentration. The integration preserves quantitative discipline while incorporating textual intelligence.

Textual analysis has transformed accounting and finance research methodologies. Early approaches relied on word frequency counts and dictionary-based classification. Machine learning methods now enable more sophisticated text understanding. Textual data from corporate disclosures contains valuable signals beyond numerical statements [6]. Management discussion sections reveal qualitative assessments absent from financial tables [6]. Risk factor disclosures indicate management concerns about operational challenges [6]. The evolution toward advanced natural language processing opens new analytical possibilities [6]. Future developments point toward real-time processing of streaming textual data [6].

Risk assessment benefits substantially from textual analysis integration. Forward-looking statements contain risk indicators complementing historical measures. Sentiment analysis of risk-related passages provides qualitative risk signals. Corporate communications reveal concerns invisible in quantitative data alone. The combination of numerical and textual risk measures enables a comprehensive assessment.

Decision support systems present optimization outputs to portfolio managers. Recommended allocations include supporting rationale from textual analysis. Signal confidence scores indicate the reliability of sentiment classifications. Portfolio managers retain final discretion over allocation decisions. The framework augments human judgment rather than replacing expertise. Audit trails document signal generation logic for regulatory compliance purposes.

Signal Category	Classification Output	Trading Application
Primary Sentiment	Bullish, Bearish, Neutral	Directional position bias determination
Time Horizon	Short-term, Medium-term, Longterm	Holding period and urgency weighting
Catalyst Type	Earnings, Regulatory, Competitive	Event-specific response protocol activation
Confidence Score	High, Medium, Low	Position sizing and risk allocation adjustment

Table 2. Signal Taxonomy and Portfolio Integration Components [5, 6].

V. REINFORCEMENT LEARNING INTEGRATION

A. Adaptive Strategy Development

The framework interfaces LLM-generated signals with reinforcement learning agents. Static trading rules cannot adapt to evolving market dynamics. Market conditions shift continuously due to economic cycles and participant behavior changes. Reinforcement learning provides mechanisms for automatic strategy adaptation.

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

Portfolio management presents a natural application for reinforcement learning methods. The sequential decisionmaking structure aligns with reinforcement learning formulations. An agent observes market states and selects portfolio allocations iteratively. The deep reinforcement learning framework treats portfolio management as a continuous control problem [7]. The architecture processes historical price data through convolutional neural networks [7]. An ensemble of identical independent evaluators assesses individual asset performance [7]. Portfolio vector memory maintains allocation history for temporal consistency [7]. The framework learns allocation policies directly from price movements without manual feature engineering [7]. Online stochastic batch learning enables continuous adaptation to incoming market data [7]. This approach eliminates the need for explicit price prediction models entirely.

State representations combine market microstructure features with semantic embeddings. Price series and volume patterns constitute quantitative state components. Order book characteristics indicate immediate supply and demand conditions. Semantic embeddings from textual analysis capture qualitative sentiment dimensions. The fusion of numerical and linguistic features creates comprehensive state vectors. Action spaces define permissible portfolio weight adjustments at each decision point.

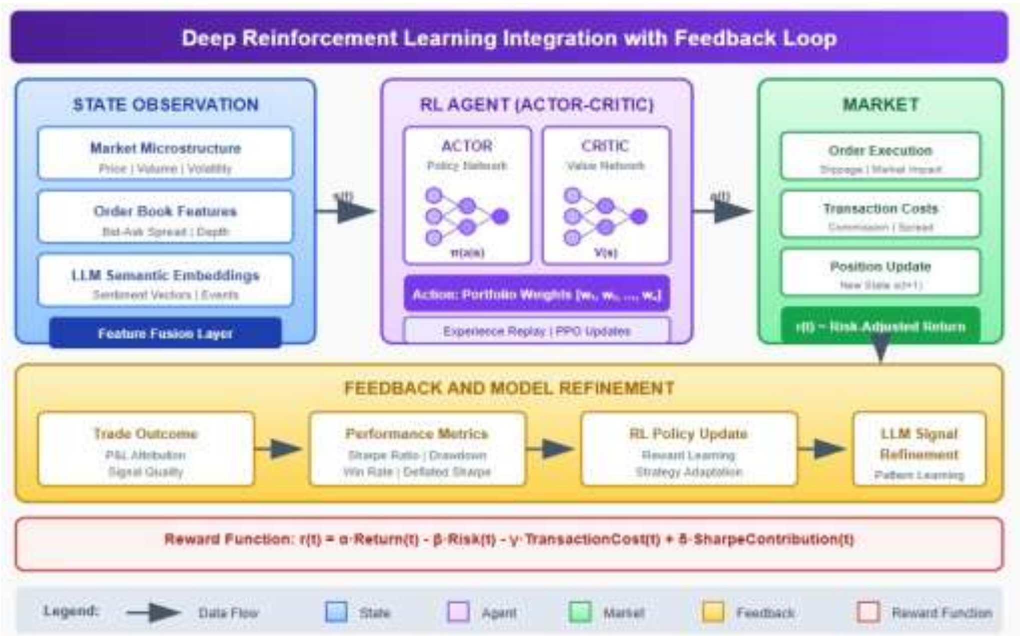

Fig 4. Deep reinforcement learning architecture illustrating state observation components (market microstructure, order book features, and LLM semantic embeddings), actor-critic agent structure with policy and value networks, market environment interaction, and the feedback mechanism for continuous model refinement.

B. Feedback Loop Mechanisms

Executed trades generate outcome data informing subsequent model refinements. Trade results provide ground truth for evaluating signal quality. Successful signals receive reinforcement while unprofitable patterns face correction. Progressive refinement improves both signal generation and execution timing.

Practical algorithmic trading introduces challenges beyond idealized simulations. Real markets exhibit transaction costs, slippage, and liquidity constraints. Deep robust reinforcement learning addresses practical trading requirements explicitly [8]. The framework incorporates market friction modeling during policy learning [8]. Robustness mechanisms prevent overfitting to historical price patterns [8]. The agent learns policies resilient to distribution shifts between training and deployment [8]. Reward shaping incorporates transaction cost penalties for realistic performance estimation [8]. Risk constraints limit drawdown exposure during strategy execution [8].

Bidirectional feedback connects reinforcement learning outcomes with LLM refinement. Sentiment signals preceding profitable trades indicate valuable textual patterns. Signals associated with losses highlight potential classification errors. Outcome-based labeling generates training data for model improvement. The LLM adapts to market-specific sentiment patterns through iterative updates.

Market regime detection triggers automatic strategy adaptation. Volatility changes require position sizing adjustments. Trending conditions demand different tactics than mean-reverting environments. The reinforcement learning agent develops regime-conditional policies through accumulated experience. Continuous learning maintains strategy effectiveness across market condition changes.

Component	Function	Integration Purpose
State Representation	Market microstructure and semantic embeddings	Comprehensive environment observation
Action Space	Portfolio weight adjustments	Position sizing and allocation decisions
Reward Function	Risk-adjusted performance criteria	Balanced profit seeking and drawdown avoidance
Policy Network	State-to-action mapping	Optimal trading decision generation
Experience Replay	Historical transition storage	Sample efficiency improvement during training

Table 3. Reinforcement Learning Framework Elements and Functions [7, 8].

VI. LATENCY AND ACCURACY EVALUATION FRAMEWORK

A. Simulation Environment Design

Performance assessment requires controlled evaluation environments. Live market testing exposes capital to unvalidated strategy risks. Simulated environments enable systematic experimentation without financial consequences. Backtesting frameworks replay historical data through trading algorithms. The evaluation framework measures inference latency, signal accuracy, and portfolio outcomes comprehensively.

Market simulation demands realistic modeling of execution mechanics. Order book dynamics determine fill prices for submitted orders. Slippage effects arise from market impact during trade execution. Information propagation patterns influence timing advantages across market participants. Backtesting systems must incorporate these friction elements accurately. However, historical simulation introduces significant methodological challenges. Multiple strategy trials on the same dataset inflate performance estimates artificially [9]. Selection bias occurs when only successful backtests receive deployment consideration [9]. The probability of backtest overfitting increases with the number of tested configurations [9]. Non-normality of financial returns further complicates statistical inference [9]. The deflated Sharpe ratio adjusts for these multiple testing biases explicitly [9]. This correction accounts for the number of independent trials conducted during strategy development [9].

Latency measurement captures end-to-end processing delays across system components. Data ingestion introduces initial delays from source transmission and parsing. Preprocessing pipelines add computational overhead for text normalization. LLM inference constitutes the primary latency contributor in textual analysis. Signal generation and portfolio optimization calculations introduce additional delays. Cumulative latency determines effective information advantage windows available for trading.

B. Accuracy and Performance Metrics

Signal accuracy assessment evaluates classification correctness against realized market outcomes. Sentiment predictions undergo comparison with subsequent price movements. Directional accuracy measures alignment between predicted and actual price trends. False positive and false negative rates reveal specific error pattern characteristics. Calibration analysis examines the correspondence between model confidence and actual accuracy rates.

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

Portfolio-level evaluation captures aggregate strategy performance across multiple dimensions. Practical deep reinforcement learning approaches address real-world trading requirements directly [10]. The evaluation framework incorporates transaction costs and market impact modeling [10]. Multiple reinforcement learning algorithms undergo parallel evaluation for comparison [10]. Actor-critic methods, including proximal policy optimization, demonstrate effectiveness for trading tasks [10]. The framework enables systematic hyperparameter tuning across algorithm variants [10]. Ensemble approaches combining multiple agents improve robustness over single-agent strategies [10].

Risk-adjusted metrics balance return achievement against risk assumption levels. Maximum drawdown quantifies the worst-case capital decline during evaluation periods. Volatility measures capture return dispersion affecting investor experience. The Sharpe ratio relates excess returns to volatility for standardized comparison.

Robustness testing examines performance stability across varying market conditions. Stress scenarios evaluate strategy behavior during market disruptions and volatility spikes. Regime variation tests assess adaptation capabilities across trending and mean-reverting states. Parameter sensitivity analysis identifies fragile configurations requiring adjustment. Comprehensive evaluation across diverse conditions reduces deployment risk substantially.

Evaluation Dimension	Metric Type	Assessment Purpose
Signal Accuracy	Directional correctness, False positive rate	Classification quality measurement
Latency Performance	End-to-end processing delay	Information advantage window determination
Portfolio Returns	Absolute and relative performance	Strategy profitability assessment
Risk Measurement	Volatility, Maximum drawdown	Capital preservation evaluation
Robustness	Deflated Sharpe ratio	Backtest overfitting correction

Table 4. Performance Assessment Framework Components and Measurement Criteria [9, 10].

CONCLUSION

The integration of Large Language Models into trading and investment decision platforms represents a significant advancement in financial technology infrastructure. Semantic understanding of market-moving textual content provides informational advantages unavailable through purely quantitative methods. The architectural framework presented within the article establishes foundational principles for deploying generative language models in timecritical financial applications. Multi-source aggregation pipelines enable comprehensive coverage of information channels affecting asset valuations. Domain-adapted language models address vocabulary and contextual challenges specific to financial communications. Hierarchical signal taxonomies translate linguistic sentiment into structured trading recommendations compatible with existing portfolio management systems. Deep reinforcement learning integration introduces adaptive capabilities essential for non-stationary market environments. Policy learning from experience enables continuous strategy refinement based on realized trading outcomes. Simulationbased evaluation frameworks provide a systematic validation infrastructure supporting responsible deployment practices. The deflated Sharpe ratio and related metrics address backtest overfitting concerns during strategy development. Future guidelines encompass multi-lingual information processing, asset-magnificence-specific architectures, and privacy-preserving collaborative version improvement across institutional obstacles. The convergence of natural language processing with quantitative finance opens transformative possibilities for better decision-making throughout capital markets.

Journal of Information Systems Engineering and Management
2026, 11(1s) e-ISSN:
2468-4376
https://www.jisem-journal.com/ **Research Article**

REFERENCES

[1] Kelvin Du et al., "Natural language processing in finance: A survey," Information Fusion, 2025. [Online]. Available: https://sentic.net/nlp-in-finance.pdf

[2] Dogu Tan Araci, "FinBERT: Financial Sentiment Analysis with Pre-trained Language Models," arXiv, 2019. [Online]. Available: https://arxiv.org/pdf/1908.10063

[3] Gilles Jacobs and Ve´ronique Hoste, "SENTiVENT: enabling supervised information extraction of companyspecific events in economic and financial news," Language Resources and Evaluation, 2022. [Online]. Available: https://link.springer.com/content/pdf/10.1007/s10579-021-09562-4.pdf

[4] Cem Subakan et al., "ATTENTION IS ALL YOU NEED IN SPEECH SEPARATION," arXiv, 2021. [Online]. Available: https://arxiv.org/pdf/2010.13154

[5] Shanshan Dong and Chang Liu, "Sentiment Classification for Financial Texts Based on Deep Learning," Computational Intelligence and Neuroscience, 2021. [Online]. Available: https://onlinelibrary.wiley.com/doi/pdf/10.1155/2021/9524705

[6] KHRYSTYNA BOCHKAY et al., "Textual Analysis in Accounting: What's Next?," Wiley, 2022. [Online]. Available: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1911-3846.12825

[7] Zhengyao Jiang et al., "A Deep Reinforcement Learning Framework for the Financial Portfolio Management Problem," arXiv, 2017. [Online]. Available: https://arxiv.org/pdf/1706.10059

[8] YANG LI et al., "Deep Robust Reinforcement Learning for Practical Algorithmic Trading," IEEE Access, 2019. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8786132

[9] David H. Bailey and Marcos López de Prado, "THE DEFLATED SHARPE RATIO: CORRECTING FOR SELECTION BIAS, BACKTEST OVERFITTING AND NON-NORMALITY," Journal of Portfolio Management, Forthcoming, 2014. [Online]. Available: https://www.davidhbailey.com/dhbpapers/deflated-sharpe.pdf

[10] Xiao-Yang Liu et al., "Practical Deep Reinforcement Learning Approach for Stock Trading," arXiv, 2022. [Online]. Available: https://arxiv.org/pdf/1811.07522