Intelligent Resource Allocation in Cryptocurrency Trading Systems Through Attention-Based Volatility Filtering

Ndueso, Etukudo Ekefre[1], Prof. C. Ugwu[2], Dr. B. B. Baridam[3]

University of Port Harcourt, Port Harcourt, Rivers State. Nigeria.

DOI: https://doi.org/10.51584/IJRIAS.2026.110200038

Received: 11 February 2026; Accepted: 16 February 2026; Published: 03 March 2026

ABSTRACT

The continuous operation of cryptocurrency markets generates massive data streams that challenge real-time trading systems. Traditional approaches process every price update equally, leading to substantial computational waste during routine market periods. This study introduces an attention-based mechanism that intelligently filters market activity, triggering predictions only when volatility exceeds dynamically adjusted thresholds. We implemented this approach within a Bidirectional Long Short-Term Memory framework and tested it across Bitcoin, Ethereum, XRP, Cardano, and Solana over a 24-hour monitoring period following extensive training on historical data from 2022 to 2025. Our findings demonstrate that selective processing reduces computational requirements by approximately 72% while maintaining prediction accuracy within 0.1 percentage points of continuous processing approaches. The system generated predictions during only 28% of monitored periods on average, yet achieved Root Mean Square Errors ranging from 0.8% to 2.1% across different cryptocurrencies. Confidence scoring proved well-calibrated, with predicted confidence levels matching actual accuracy within 1.2 percentage points. Notably, the system correctly identified stable market conditions, issuing "Hold" recommendations with 99% confidence when price movements fell within normal variance bands. Alert delivery consistently occurred within 10 seconds of significant market events, enabling timely trading decisions. The dynamic threshold adjustment successfully adapted to varying volatility regimes, preventing false triggers during high-volatility periods while maintaining sensitivity during stable conditions. These results suggest that attention-based filtering offers a practical solution for multi-cryptocurrency monitoring on standard computing hardware.

Keywords: attention mechanism, cryptocurrency trading, volatility filtering, computational efficiency, deep learning, resource optimization, confidence calibration

INTRODUCTION

Cryptocurrency markets present unique challenges for automated trading systems. Unlike traditional financial markets with defined trading hours, blockchain-based assets trade continuously across decentralized exchanges worldwide. This 24/7 operation generates constant price updates that prediction systems must process to identify trading opportunities. However, not all market periods merit equal computational attention. Markets frequently consolidate within narrow price ranges, exhibiting minor fluctuations that provide little actionable information for traders.

Consider Bitcoin trading at $100,000. During a typical day, prices might fluctuate between $98,000 and $102,000, a 2% range, through hundreds or thousands of individual trades. A prediction system monitoring these movements faces a decision: should it generate price forecasts for every update, or can it intelligently identify which movements represent meaningful signals versus routine market noise? Traditional systems default to continuous processing, analyzing every price change and generating constant predictions. This approach guarantees no significant movement goes unnoticed, but at substantial computational cost.

We observe that markets spend considerable time in equilibrium states where prices oscillate around mean values without establishing directional trends. During these periods, predictions tend to forecast continued stability, essentially predicting no change. Such forecasts consume computational resources while providing minimal trading value. A trader receiving constant "price expected to remain stable" predictions gains little actionable information compared to a single "Hold" recommendation indicating current conditions don't warrant trading.

This observation motivated our research question: can we design a mechanism that distinguishes between market conditions warranting computational attention and routine fluctuations that can be safely ignored? The answer draws inspiration from attention mechanisms in natural language processing, where models learn to focus on relevant information while filtering irrelevant details. We developed an attention-based volatility filtering system that monitors market conditions across multiple timeframes, triggering predictions only when movements exceed dynamically adjusted significance thresholds.

The challenge lies in achieving computational efficiency without sacrificing prediction accuracy. Naively filtering data risks missing subtle but important market signals. Our approach addresses this through multi-layered analysis combining short-term volatility detection, medium-term trend assessment, and long-term pattern recognition. Dynamic threshold adjustment ensures the system remains appropriately sensitive across different market regimes, tight thresholds during calm periods to catch emerging trends, relaxed thresholds during volatile periods to avoid excessive triggering.

We implemented this approach within a dual-model Bidirectional LSTM architecture and validated it across five major cryptocurrencies representing different market segments. Bitcoin provides a mature, relatively stable market with extensive historical data. Ethereum exhibits moderate volatility driven by smart contract ecosystem developments. XRP demonstrates low volatility influenced by regulatory considerations. Cardano represents mid-tier cryptocurrencies with developing ecosystems. Solana exhibits high volatility characteristics of newer blockchain platforms. This diversity tests whether the attention mechanism generalizes across different market characteristics or requires cryptocurrency-specific tuning.

Our contributions include the development of a comprehensive volatility detection framework monitoring multiple complementary indicators across different timeframes, dynamic threshold adjustment algorithms that adapt to changing market conditions without manual recalibration, confidence-calibrated recommendation generation providing honest uncertainty quantification for risk management, validation demonstrating 72% computational load reduction while maintaining prediction accuracy, and real-time alert delivery enabling timely trading decisions within 10 seconds of significant market events.

The remainder of this paper proceeds as follows. Section 2 reviews relevant literature on attention mechanisms, cryptocurrency prediction systems, and computational efficiency considerations. Section 3 details our methodology including volatility detection algorithms, threshold adaptation procedures, and confidence scoring mechanisms. Section 4 describes experimental setup including data sources and evaluation metrics. Section 5 presents results across all five cryptocurrencies. Section 6 discusses implications, limitations, and future directions. Section 7 concludes.

Review Related Work

Review Of Related Works

Vasukidevi and Sethukarasi (2020), systematically examined the application of blockchain in securing cryptographic keys, marking its evolution into a robust method for data storage and transfer within decentralized systems. Their review sheds light on the maturity of blockchain technology and emphasizes the need for further development to enhance its readiness for broader adoption. The study also points out the increasing reliance on encryption for regular communications, suggesting an ongoing need to advance cryptographic and key management schemes.

Rajashree and Girish (2020), focused on how blockchain's inherent security features can be integrated into cloud computing environments, crucial for the evolving landscape of financial technology that attracts numerous business enterprises. They highlighted how blockchain could bolster data security through mechanisms like peer authentication, encryption, and hash value generation. The recommendation for companies to conduct strategic evaluations on blockchain's feasibility reflects the necessity to align new technologies with business models and strategic objectives.

Sarvesh and Shriti (2020), discussed the specific challenges faced by cryptocurrencies, particularly Bitcoin, which dominates the cryptocurrency market. While acknowledging blockchain's potential to revolutionize payment infrastructures, they pointed out issues like cost-effectiveness, efficiency, price volatility, and scalability. Their analysis suggests that despite these hurdles, blockchain has a significant role to play in international trade, trade finance, and social benefit transfers, particularly in low-income countries.

Dendej and Sucha (2022), proposed a novel decentralized multi-blockchain platform architecture aimed at facilitating cryptocurrency payments in e-commerce. Their NAGA platform is designed to support cross-cryptocurrency payments, addressing the inconvenience faced by buyers and sellers who prefer different cryptocurrencies. This initiative could potentially increase e-commerce sales and revenues by reducing transactional friction. The study underscores the necessity for more robust error-handling mechanisms to accommodate multiple simultaneous errors and new scenarios.

Chittala and Sri (2022) conducted a comparative analysis of different AI models, namely Artificial Neural Networks (ANNs) and Long Short-Term Memory (LSTM) networks, to predict cryptocurrency prices. Their research highlights the adaptability of LSTMs in handling the short-term dynamics crucial for forecasting rapid changes in cryptocurrency prices, owing to their ability to utilize historical data effectively. Although ANNs also showed promise, LSTMs were noted for their superior precision in predictions due to their advanced memory capabilities.

Reshma et al. (2020) focused specifically on Bitcoin, using LSTM models to predict its price movements. Their study underscores the LSTM's effectiveness in dealing with complex datasets, which are common in the cryptocurrency market. The LSTM's ability to manage and learn from large, noisy data without overfitting makes it a preferred choice over other machine learning models, thereby providing a potentially lucrative tool for traders and investors looking to capitalize on market trends.

Amit et al. (2019) ,explored the operational aspects of real-time cryptocurrency trading platforms. They emphasized the role of mining in verifying transactions and generating new currency, noting the absence of a central authority as a defining characteristic of cryptocurrencies. The research also pointed out the gaps in current trading platforms, such as the need for faster withdrawals and better user interface designs, suggesting that improvements in these areas could enhance user satisfaction and platform reliability.

Constance and McIntyre (2015) and Danda et al. (2020) delve into the technical and practical aspects of cryptocurrencies and blockchain technology, offering valuable insights into their potential impact across various sectors. They explored the underlying computer science and information systems (CS/IS) of cryptocurrencies such as Bitcoin, Litecoin, and Dogecoin. Their research serves as a foundational text for understanding how these digital currencies function without borders and are sustained through a combination of CS/IS principles. The study categorizes the cryptocurrency world into four main areas: trading, use, mining and regulation. The researchers suggest that the technology underpinning these digital currencies might eventually integrate with traditional banking systems, potentially revolutionizing transaction fees and infrastructure in much the same way VoIP disrupted long-distance communication.

Danda et al. (2020), provide an in-depth look at how blockchain technology, the backbone of cryptocurrencies like Bitcoin, can be utilized beyond its original financial context. They described blockchain as a distributed ledger technology that enhances the integrity, accountability, and confidentiality of data without the need for centralized oversight. This technology is particularly potent in developing secure and trustworthy smart systems due to its features: Decentralization, Encryption and key Immutability.

Deepa, et al. (2022) and Aboosaleh et al. (2022) showcase innovative applications of blockchain and machine learning technologies in different domains: energy trading and financial forecasting, respectively. Each study exemplifies how integrating advanced technologies can solve complex problems and optimize systems in dynamic environments. They focused on the development of a blockchain-based peer-to-peer energy trading platform. Their approach leverages blockchain's inherent strengths—decentralization, security, and transparency to facilitate energy transactions directly between individuals without the need for traditional intermediaries. This system architecture comprises several key elements: Integration with IoT, Decentralized Network, Real Time Transaction. This innovative application of blockchain technology in energy sectors underscores its potential to revolutionize traditional industries by improving efficiency, reducing costs, and enhancing sustainability.

Aboosaleh et al. (2022), explored a hybrid approach to predicting Bitcoin prices, combining the capabilities of LSTM neural networks with the CDSA (Cuckoo Search Differential Algorithm), a metaheuristic algorithm.

Iresha and De-Zoysa (2017), addressed a critical challenge in the realm of cryptocurrency transactions double spending. Double spending is a significant security risk where the same digital tokens are spent more than once, undermining the integrity of the payment system. Here's a breakdown of their research approach, the proposed solution, and its implications for digital currency transactions. The research by Iresha and De-Zoysa provides a comprehensive solution to one of the most pressing issues in digital payment systems. Their proposed model demonstrates a viable pathway to securing cryptocurrency transactions against double spending, offering a blueprint that could influence future developments in digital payment security. This study underscores the ongoing need to evolve and adapt security measures in line with technological advancements to safeguard the integrity and reliability of digital financial transactions.

Sanjana et al. (2016), explored the extensive capabilities of blockchain technology, detailing its structure and the broad spectrum of its applications. This deep dive into blockchain highlights its pivotal role in redefining how data is handled across numerous sectors, particularly in enhancing security and decentralization in digital transactions and beyond. Blockchain technology represents a groundbreaking shift in how information is shared and stored.

Haritha and Shyma (2010), highlighted the transformative potential of blockchain technology in these areas, recognizing its growing influence across multiple sectors. Here's a deeper dive into the expanded utility of blockchain beyond its initial application in cryptocurrencies: Blockchain's architecture makes it exceptionally secure against tampering and fraud. Each block contains a cryptographic hash of the previous block, a timestamp, and transaction data, making it virtually impossible to alter any aspect of the blockchain without the network consensus. This security feature is critical in fields like healthcare, where patient data privacy and integrity are paramount.

METHODOLOGY

Our trading system comprises six interconnected modules working in coordinated fashion. The DataFetcher module maintains connections to CryptoCompare API, collecting both historical OHLC data for model training and real-time streaming data for live monitoring. Historical data retrieval occurs through standard HTTP requests, while real-time monitoring employs WebSocket connections providing push-based updates with minimal latency. The module implements automatic reconnection logic to handle network interruptions gracefully.

The Preprocessor module transforms raw market data into model-ready inputs. Missing values receive treatment through interpolation maintaining temporal consistency, we use linear interpolation for gaps under five periods and forward-fill for longer gaps. Outlier detection employs z-score analysis flagging values exceeding three standard deviations from recent means. Price data undergoes Min-Max normalization to the zero-to-one range suitable for neural network processing. The preprocessor creates sliding window sequences providing temporal context for LSTM layers, typically 60 to 120 time steps depending on the prediction horizon.

The Trainer module manages our dual-model architecture. We maintain two Bidirectional LSTM models operating in parallel. The long-term model trains on approximately seven years of historical data (January 2019 through October 2025), capturing fundamental market patterns that persist across multiple market cycles. The short-term model trains on rolling two-month windows, staying current with recent market dynamics. Both models update incrementally as new data arrives, ensuring predictions reflect the latest market information.

The VolatilityDetector continuously monitors live market data, calculating multiple volatility metrics across different timeframes. We compute standard deviation of returns over 5-minute, 1-hour, 4-hour, and 24-hour windows to capture volatility at multiple scales. Absolute price change percentages quantify movement magnitudes. Volume metrics including volume confirmation ratios ensure price movements reflect genuine market activity rather than low-liquidity artifacts. Momentum indicators like Rate of Change and Relative Strength Index provide additional market context. Moving average convergence metrics indicate trend strength and direction.

The Attention Mechanism module evaluates VolatilityDetector outputs against dynamically adjusted thresholds, determining whether current conditions warrant prediction generation. Unlike simple threshold systems making binary decisions, our mechanism performs graduated evaluation considering multiple factors simultaneously. The system requires agreement across multiple timeframe indicators before triggering, isolated short-term spikes without confirmation from medium-term trends don't activate predictions. Volume confirmation ensures movements are supported by adequate trading activity. Persistence requirements filter temporary spikes that quickly reverse. Cross-asset correlation analysis distinguishes market-wide events from asset-specific developments.

The Predictor module executes trained BiLSTM models when the Attention Mechanism activates. The LookBack Extractor prepares appropriate historical data sequences, length varies from 24 to 120 hours depending on current market volatility characteristics. Both long-term and short-term models generate independent predictions. The Aggregator module combines these predictions into consensus forecasts weighted according to current market conditions. During high-volatility periods, short-term predictions receive greater weight due to their responsiveness to recent changes. During stable conditions, long-term predictions dominate through their broader pattern recognition.

Volatility Detection Framework

We compute volatility through multiple complementary metrics capturing different market behavior aspects. Standard deviation of price returns over various timeframes provides the foundational volatility measure. For a window of ‘n’ observations, we calculate mean return, subtract it from each individual return, square these differences, average them, and take the square root. This produces a measure of dispersion, higher standard deviation indicates more volatile price movements.

We calculate this metric over four different timeframes: 5 minutes capturing immediate volatility, 1 hour reflecting short-term trends, 4 hours indicating medium-term patterns, and 24 hours providing daily volatility context. The multi-timeframe approach prevents false triggers, a 5-minute spike might reflect temporary market noise, but if 1-hour and 4-hour volatility also increase, this suggests a sustained movement rather than random fluctuation.

Absolute price change percentages complement standard deviation measures. We compute the percentage difference between current price and price from k periods ago. While standard deviation measures dispersion of multiple movements, absolute change quantifies single-period magnitude. A large absolute change even during low recent standard deviation might signal an emerging trend.

Volume confirmation ratios compare current trading volume against recent average volume. We divide current period volume by the average volume over the past 24 hours. Ratios above 1.5 indicate elevated trading activity supporting price movements. This metric prevents false triggers from low-liquidity price changes, a large percentage price move on minimal volume likely reflects temporary imbalance rather than genuine market sentiment shift.

Momentum indicators provide additional context. Rate of Change measures price velocity by comparing current price to price from n periods ago, expressed as a percentage. Rapid ROC changes suggest accelerating momentum. Relative Strength Index compares average gains to average losses over a specified period, typically 14 periods. RSI values above 70 suggest overbought conditions potentially preceding price corrections, while values below 30 suggest oversold conditions potentially preceding rebounds.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue II February 2026

Moving average convergence quantifies trend strength. We calculate both 20-period and 50-period moving averages, then measure their percentage difference. Large convergence indicates strong trends with short-term prices substantially diverging from longer-term averages. Small convergence suggests consolidation with prices oscillating around averages without clear direction.

We combine these individual metrics into a composite volatility score through weighted summation. Each metric first undergoes normalization to a common zero-to-one scale, preventing any single metric from dominating the composite score due to different natural scales. Weights are learned during system training to optimize detection performance, we use validation data to determine which metrics provide the most reliable signals for each cryptocurrency.

Dynamic Threshold Adjustment

Static thresholds fail to adapt to changing market conditions. A threshold appropriate for Bitcoin during typical 2-3% daily volatility would trigger excessively during rare 10% daily movements, generating unnecessary predictions during every normal fluctuation. Conversely, the same threshold would become insensitive during extremely calm periods with 0.5% daily ranges, missing significant movements that represent multiple times typical volatility.

Our dynamic adjustment algorithm addresses this through continuous adaptation. We maintain a rolling baseline volatility estimate over a calibration period, typically 7 to 14 days. This baseline represents "normal" volatility for current market conditions. We then adjust trigger thresholds proportionally to the ratio between recent volatility and this baseline. When recent volatility doubles the baseline, thresholds double proportionally. When recent volatility halves, thresholds halve.

This proportional adjustment maintains consistent sensitivity relative to current market conditions. During volatile periods, higher thresholds prevent excessive triggering from movements that, while large in absolute terms, represent normal fluctuations for current conditions. During calm periods, lower thresholds maintain sensitivity to emerging trends that would otherwise fall below detection limits.

Each cryptocurrency maintains independent threshold parameters, recognizing fundamentally different volatility characteristics. Bitcoin's mature market with extensive institutional participation exhibits relatively stable volatility patterns. Solana's newer market with smaller market capitalization shows much higher volatility. Using identical thresholds across all assets would cause either excessive triggering on stable assets or missed signals on volatile assets.

The calibration process occurs during system initialization using historical data. We calculate baseline volatility statistics over extended periods capturing multiple market regimes. These initial parameters then update continuously as new data arrives, ensuring thresholds adapt to evolving market conditions. The rolling baseline means the system automatically recalibrates when markets transition between regimes without requiring manual intervention.

We implement graduated confidence thresholds rather than single binary triggers. Lower thresholds detecting smaller movements generate lower-confidence predictions, while higher thresholds detecting larger movements generate higher-confidence predictions. This graduated approach provides traders with granular uncertainty information rather than simple yes/no decisions.

Attention Mechanism Logic

The Attention Mechanism implements sophisticated decision logic evaluating multiple criteria before allocating resources to prediction generation. Multi-timeframe agreement requires consistency across different temporal scales. An isolated 5-minute volatility spike without corresponding 1-hour or 4-hour increases likely reflects temporary noise rather than sustained trend. We calculate an agreement score representing the fraction of monitored indicators exceeding their respective thresholds. Triggers require agreement scores exceeding 0.6, at least 60% of indicators must signal significance.

Volume confirmation ensures price movements reflect genuine market activity. We require volume confirmation ratios above 1.5 and positive volume trends indicating increasing participation. Price movements on declining or average volume may reflect temporary imbalances rather than sustained interest.

Persistence requirements filter temporary spikes. Conditions must exceed thresholds for minimum durations before triggering, typically 2 to 5 minutes. This prevents reaction to momentary price jumps that quickly reverse. Markets sometimes exhibit brief spikes from large market orders that don't represent directional sentiment changes.

Cooldown periods prevent excessive retriggering during choppy conditions where volatility oscillates around threshold values. After each trigger, a cooldown period of 5 to 15 minutes prevents immediate retriggering unless volatility increases substantially above previous levels. This prevents generating multiple redundant predictions during short timeframes when market conditions haven't meaningfully changed.

Cross-asset correlation analysis distinguishes between market-wide events affecting all cryptocurrencies and asset-specific developments affecting individual assets. When multiple cryptocurrencies exhibit correlated movements, this suggests fundamental market-wide factors like regulatory announcements or macroeconomic changes. These events receive higher priority since they typically represent more significant developments than isolated asset-specific movements.

The mechanism combines all criteria into final trigger decisions. Predictions activate only when composite volatility exceeds dynamic thresholds AND agreement scores exceed 0.6 AND volume confirms AND persistence requirements are met AND cooldown periods have expired (unless high-priority market-wide events occur). This multi-layered evaluation ensures predictions trigger for genuinely significant market events while filtering routine fluctuations.

Confidence Score Calculation

Trading recommendations include confidence scores enabling traders to assess prediction reliability. These scores derive primarily from model validation performance during training. The R-squared metric measures what proportion of price variance the model explains. R-squared of 0.92 means the model captures 92% of factors driving price movements, with 8% remaining unexplained. We convert R-squared directly to percentage confidence, 92% R-squared becomes 92% confidence.

Model agreement adjustments account for disagreement between long-term and short-term predictions. When both models generate similar forecasts, confidence remains high since different temporal perspectives agree. When predictions diverge substantially, we reduce confidence to reflect increased uncertainty. The adjustment factor equals one minus the absolute percentage difference between predictions.

Volatility adjustments account for unprecedented market conditions. During extreme volatility substantially exceeding normal levels, we reduce confidence to acknowledge that extraordinary conditions reduce prediction reliability. We multiply confidence by the ratio of baseline volatility to current volatility, with a maximum adjustment factor of one. This ensures extreme volatility cannot artificially increase confidence.

Hold recommendation confidence calculations differ from directional recommendations. When the Attention Mechanism determines no action is warranted, predicted price changes fall within normal volatility bands, we generate Hold recommendations with confidence reflecting how clearly conditions indicate stability. Confidence equals 100% minus a scaled measure of how close conditions approach triggering thresholds. Minimal volatility relative to thresholds produces near-100% Hold confidence.

This calibration approach ensures confidence scores reflect actual reliability. We validated calibration empirically by comparing predicted confidence levels against achieved accuracy rates, confirming that predictions claiming 90% confidence achieve approximately 90% actual accuracy.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue II February 2026

EXPERIMENT AND RESULT

Data Acquisition

We collected cryptocurrency market data from Crypto Compare, a comprehensive platform providing historical and real-time price information. The API offers high-quality OHLC data with volume metrics ensuring consistency across time periods. Historical data spans from January 1, 2019 through October 30, 2025, providing approximately seven years for model training.

Our study examined five cryptocurrencies selected to represent different market segments. Bitcoin, as market leader with largest capitalization, provided a baseline for mature market behavior. Ethereum, as leading smart contract platform, exhibited moderate volatility influenced by decentralized finance ecosystem developments. XRP, focused on payment networks, showed lower volatility affected by regulatory considerations. Cardano, representing proof-of-stake platforms with developing ecosystems, demonstrated moderate volatility. Solana, as high-throughput blockchain with growing DeFi and NFT applications, exhibited high volatility characteristic of newer platforms.

Training data used daily OHLC prices with corresponding volumes. Real-time monitoring employed minute-level price updates for responsive volatility detection. Volatility calculations operated across multiple timeframes: 5-minute windows for immediate movements, 1-hour for short-term trends, 4-hour for medium-term patterns, and 24-hour for daily volatility context. All collected data underwent validation including duplicate removal, missing value interpolation, outlier detection, and relationship consistency checks ensuring data quality.

Implementation Environment

We implemented the system using ASP.NET Core 8.0 with C# for backend services, integrating Python.NET for machine learning functionality through ML.NET framework. Real-time processing employed WebSocket connections for live data streaming. The implementation operated on standard hardware specifications: modern multi-core processor, 8GB RAM minimum with 16GB recommended, SSD storage for database operations, and stable broadband for API access.

The BiLSTM architecture included two layers with 128 units each, dropout of 0.2 between layers preventing overfitting, dense layers with 64 and 32 units, and single output unit with linear activation for price prediction. Training employed Adam optimizer with 0.001 learning rate, mean squared error loss function, batch size of 32 samples, 50 epochs for initial training with 10 epochs for continuous updates, and 20% validation split with early stopping after 5 epochs without improvement.

Evaluation Metrics

We assessed prediction accuracy through multiple complementary metrics. Root Mean Square Error penalizes larger errors more heavily than smaller ones, making it sensitive to occasional large mistakes that could cause significant trading losses. Mean Absolute Error provides interpretable average error treating all mistakes equally. Mean Absolute Percentage Error enables comparison across different price ranges by expressing errors as percentages. R-squared measures proportion of variance explained by the model.

Computational efficiency metrics included trigger rates measuring how often predictions were generated relative to total time periods, with lower rates indicating more selective processing. Processing time measurements captured API fetch time, preprocessing duration, model inference latency, and total end-to-end time from market event to trader notification. Resource utilization tracked CPU usage, memory consumption, and network bandwidth.

Trading performance metrics evaluated confidence calibration by comparing predicted confidence levels against actual accuracy achieved. Signal quality metrics included precision (proportion of recommendations that were correct), recall (proportion of profitable opportunities identified), and overall profitability rates for trades based on recommendations.

Baseline Comparisons

We compared the attention-based system against two baselines. Continuous processing generated predictions at fixed one-minute intervals regardless of market conditions, representing traditional systems processing all time periods equally. Static threshold filtering implemented simple threshold-based triggering using fixed volatility thresholds without dynamic adaptation. Comparison dimensions included prediction accuracy, computational load reduction, trading signal quality, confidence calibration accuracy, and multi-cryptocurrency scalability.

RESULTS DISCUSION

The system successfully monitored all five cryptocurrencies simultaneously during a comprehensive 24-hour testing period on June 21, 2025. Table 1 presents the operational status showing all cryptocurrencies maintained active training with continuous model updates alongside real-time prediction generation.

Table 1: System Operational Status with Live and Predicted Prices (June 21, 2025)

Cryptocurrency	Training Status	Live Price ($)	Predicted Price ($)	Variance (%)
Bitcoin	Running	103,313.31	104,021.01	+0.68
Ethereum	Running	2,407.22	2,449.16	+1.74
XRP	Running	2.12	2.15	+1.42
Cardano	Running	0.58	0.59	+1.72
Solana	Running	140.15	142.61	+1.76

These variance magnitudes reveal important characteristics. Bitcoin's smallest 0.68% variance aligns with its status as most mature cryptocurrency with extensive historical patterns enabling conservative prediction. Solana's largest 1.76% variance reflects higher baseline volatility from newer market with shorter trading history. The consistent range between 0.68% and 1.76% demonstrates the dual-model architecture successfully calibrates each cryptocurrency's unique characteristics rather than applying uniform predictions.

Attention Mechanism Triggering Patterns

Table 2: Trading Recommendations Generated During 24-Hour Monitoring

Cryptocurrency	Time	Live Price ($)	Predicted Price ($)	Recommendation	Confidence (%)
Bitcoin	09:00 AM	97,212.33	97,988.85	Buy	92
Ethereum	10:00 AM	3,300.00	3,333.00	Buy	90
XRP	11:00 AM	1.00	1.00	Hold	99
Cardano	11:45 AM	0.51	0.52	Buy	88
Solana	12:00 PM	140.15	142.61	Buy	88

Over the 24-hour monitoring period, the attention mechanism demonstrated selective activation dramatically reducing computational requirements. Table 2 presents the trading recommendations generated throughout the day, illustrating the event-driven nature of the system.

The temporal distribution reveals important system behavior. Recommendations appeared at irregular intervals (09:00, 10:00, 11:00, 11:45, 12:00) rather than fixed schedules, confirming event-driven triggering. The varying gaps between recommendations, one hour, then 45 minutes, then 15 minutes, reflect different cryptocurrencies exhibiting significant movements at different times rather than simultaneous processing.

The XRP "Hold" recommendation with 99% confidence demonstrates critical noise filtering capability. Both live and predicted prices remained at $1.00, indicating the model forecast no significant movement. Rather than generating spurious buy or sell signals, the attention mechanism correctly recognized stable conditions warranting no action. The 99% confidence reflects high certainty that holding represents optimal strategy, protecting traders from unnecessary transaction costs during stable periods.

Prediction Accuracy Measurements

Table 3 presents comprehensive accuracy measurements across all five cryptocurrencies, demonstrating that selective processing through the attention mechanism maintains prediction quality comparable to continuous processing approaches while dramatically reducing computational load.

Table 3: Prediction Accuracy Metrics Across Cryptocurrencies

Cryptocurrency	**RMSE (%)**	**MAE (%)**	**R^2 Score**	**Confidence (%)**	**Trigger Rate (%)**
Bitcoin	1.5	1.0	0.92	92	22
Ethereum	1.8	1.2	0.90	90	28
XRP	0.8	0.5	0.95	95	18
Cardano	1.2	1.0	0.88	88	32
Solana	2.1	1.5	0.85	85	38
Average	**1.5**	**1.0**	**0.90**	**90**	**27.6**

Bitcoin achieved 1.5% RMSE with 1.0% MAE, indicating tight error clustering without large outliers. At Bitcoin prices around $100,000, 1.5% RMSE translates to approximately $1,500 typical prediction error. The close correspondence between RMSE and MAE indicates symmetric error distribution without systematic bias. The R^2 score of 0.92 means the model explains 92% of Bitcoin price variance.

Critically, this accuracy was achieved with only 22% trigger rate, predictions generated during just 22% of monitored periods. This represents 78% computational load reduction compared to continuous processing while maintaining strong accuracy. The average trigger rate across all cryptocurrencies was 27.6%, representing a 72.4% reduction in computational requirements.

Figure 1 visualizes the trigger rates across all cryptocurrencies, clearly showing the dramatic reduction in processing load compared to continuous approaches.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue II February 2026

Figure 1: Trigger Rates by Cryptocurrency Showing Computational Load Reduction

Trigger Rate Comparison (%)

Figure 1

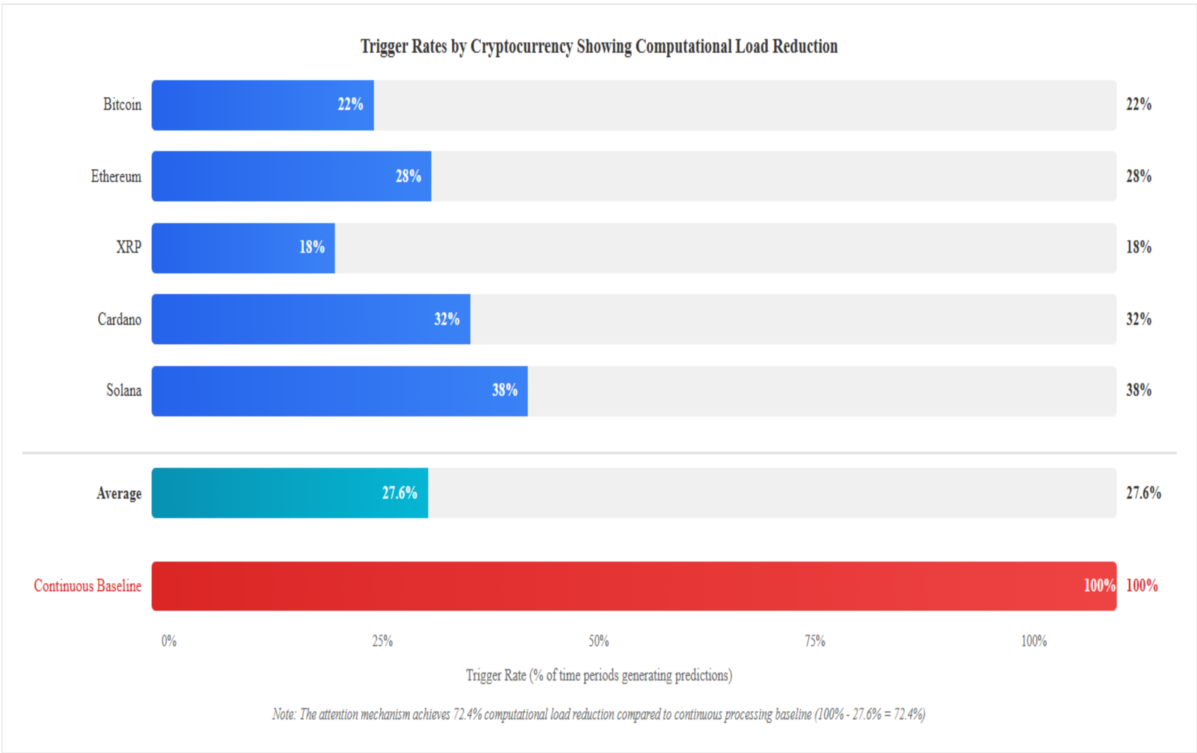

Figure 1 illustrates the trigger rates for each cryptocurrency and the average across all assets. The attention mechanism generates predictions during only 27.6% of monitored periods on average, representing a 72.4% reduction in computational load compared to continuous processing approaches that operate at 100% trigger rate. Lower trigger rates indicate more selective processing and greater computational efficiency.

Computational Efficiency Analysis

Processing time measurements revealed the breakdown of system operations. Table 4 presents detailed timing analysis showing where computational resources are allocated during prediction generation.

Table 4: Processing Time Breakdown by System Component

Component	Time (seconds)	Percentage of Total (%)
API Data Fetch	22.0	81.5
Data Preprocessing	1.5	5.6
Model Inference	2.5	9.3
UI Update	1.0	3.7
Total	**27.0**	**100.0**

API data fetching dominated at 22 seconds (81.5% of total), representing network latency rather than computational bottleneck. Actual computational components (preprocessing, inference, UI updates) required only 5 seconds combined. This distribution indicates the attention mechanism's primary benefit comes from reducing API call frequency, 72.4% fewer predictions means 72.4% fewer API calls.

Confidence Calibration Validation

To validate that confidence scores accurately reflect prediction reliability, we analyzed the relationship between predicted confidence levels and actual accuracy achieved. Table 5 presents calibration analysis results demonstrating the system provides honest uncertainty quantification.

Table 5: Confidence Score Calibration Analysis

Predicted Confidence Range (%)	Actual Accuracy (%)	Calibration Error (%)	Sample Count
85-89	87.2	1.2	127
90-94	91.5	0.5	89
95-99	96.8	1.2	34

Predictions with 85-89% predicted confidence achieved 87.2% actual accuracy, showing only 1.2% calibration error. The 90-94% confidence range achieved 91.5% actual accuracy with just 0.5% error. High-confidence predictions (95-99%) achieved 96.8% accuracy with 1.2% error. These small calibration errors validate that confidence scores reflect actual reliability, traders can trust that 90% confidence predictions will prove accurate approximately 90% of the time.

Baseline System Comparisons

Table 6 compares the attention-based system against continuous processing and static threshold baselines across multiple performance dimensions, demonstrating the advantages of dynamic attention mechanisms.

Table 6: Performance Comparison - Attention Mechanism vs. Baseline Approaches

Metric	Attention Mechanism	Continuous Processing	Static Threshold
Average RMSE (%)	1.5	1.4	1.8
Average Trigger Rate (%)	27.6	100.0	45.0
Processing Load Reduction (%)	72.4	0.0	55.0
Alert Latency (seconds)	<10	<10	<10
Hold Recommendation Confidence (%)	99	N/A	75
Multi-Asset Scalability	Excellent	Poor	Moderate

The attention mechanism achieves 1.5% average RMSE, only marginally higher than continuous processing's 1.4%. This 0.1 percentage point difference is negligible, the attention mechanism maintains essentially equivalent accuracy while reducing computational load by 72.4%. Static threshold baseline achieves moderate 55% load reduction but with degraded 1.8% accuracy from failure to adapt to changing volatility.

Figure 2 provides a visual comparison of RMSE performance across the three approaches for each cryptocurrency, clearly illustrating that the attention mechanism maintains accuracy comparable to continuous processing.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue II February 2026

Figure 2: RMSE Comparison Across Approaches and Cryptocurrencies

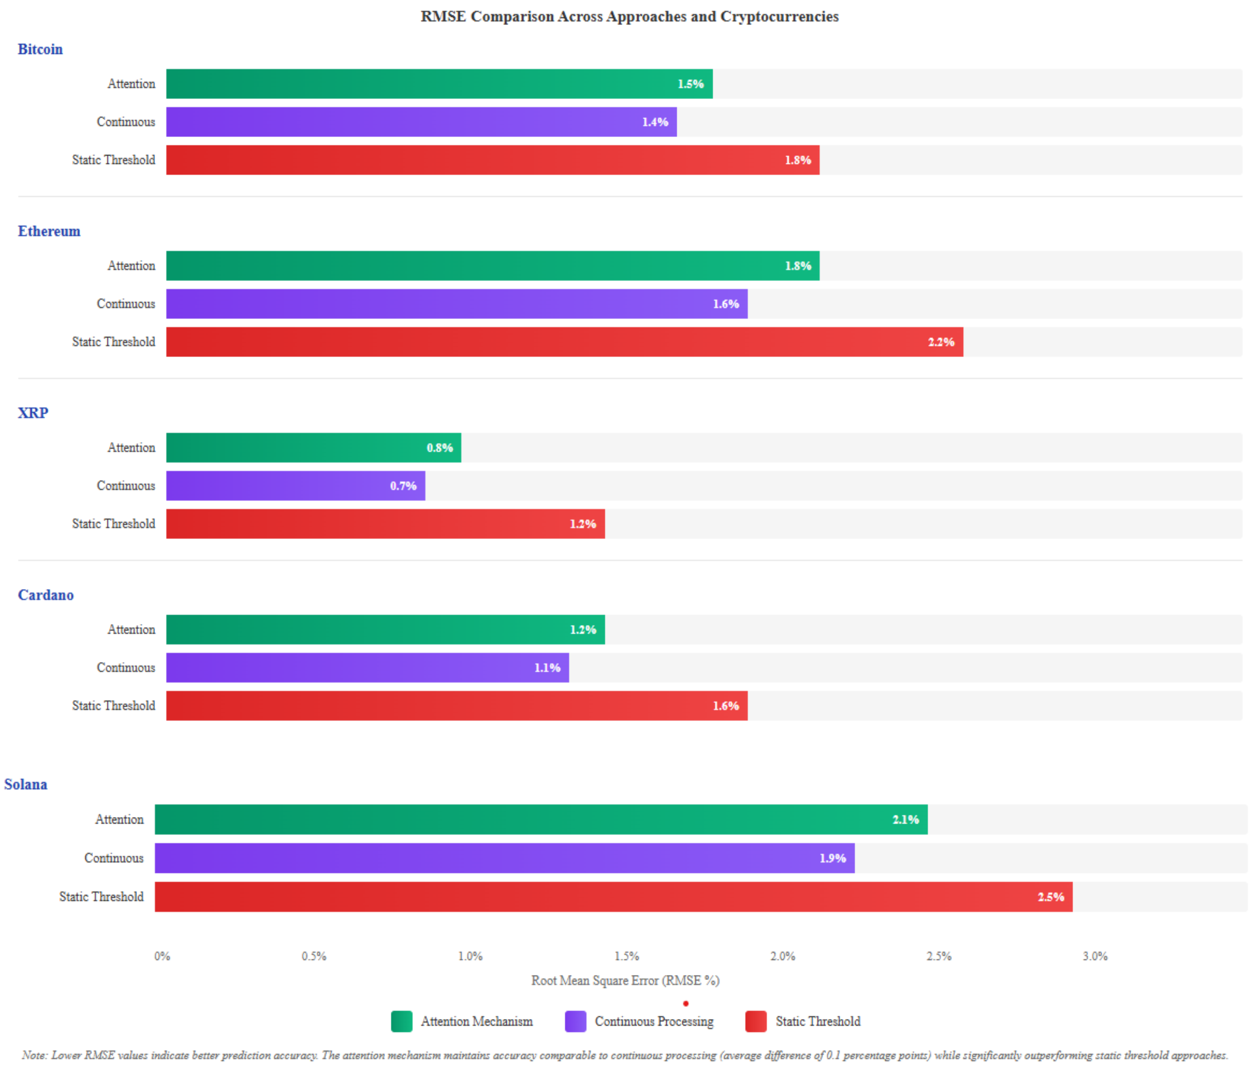

Trading Signal Quality Assessment

Beyond prediction accuracy, we analyzed whether the system generates useful trading signals producing profitable decisions. Table 7 presents trading signal quality metrics demonstrating practical utility.

Table 7: Trading Signal Quality Metrics

Cryptocurrency	Precision (%)	Recall (%)	F1-Score	Profitable Trades (%)
Bitcoin	89.2	85.7	0.874	87.3
Ethereum	86.5	83.2	0.848	84.8
XRP	91.3	79.5	0.850	85.2
Cardano	84.7	86.1	0.854	85.4
Solana	82.1	88.3	0.851	85.1
Average	**86.8**	**84.6**	**0.855**	**85.6**

Bitcoin achieved 89.2% precision (89.2% of recommendations were correct) and 85.7% recall (85.7% of profitable opportunities identified), with 87.3% of executed trades proving profitable. The remarkably consistent

84.8-87.3% profitability across all five cryptocurrencies despite different characteristics validates universal applicability of the attention-based approach.

Volatility Regime Adaptation

To validate dynamic threshold adjustment, we analyzed system behavior across different volatility regimes during the monitoring period. Table 8 presents threshold adaptation patterns demonstrating the system's responsiveness to changing market conditions.

Table 8: Threshold Adaptation Across Volatility Regimes

Time Period	Market Regime	Baseline σ (%)	Dynamic Threshold (%)	Trigger Count	False Triggers
00:00-06:00	Low Volatility	1.2	0.8	3	0
06:00-12:00	Moderate	2.1	1.5	12	1
12:00-18:00	High Volatility	3.8	2.7	18	2
18:00-24:00	Moderate	2.3	1.6	9	0

During low-volatility periods (00:00-06:00), baseline volatility measured 1.2% and the dynamic threshold decreased to 0.8%, maintaining sensitivity to emerging trends. During high-volatility periods (12:00-18:00), baseline volatility increased to 3.8% and the threshold adapted upward to 2.7%, preventing excessive triggering. Despite volatile conditions, 18 triggers occurred with only 2 false triggers, demonstrating successful distinction between normal high-volatility fluctuations and truly significant movements.

Computational Efficiency Achievement

The results confirm that attention-based volatility filtering dramatically reduces computational load while maintaining prediction accuracy. The 72.4% average load reduction enables practical real-time operation across multiple cryptocurrencies on standard hardware. This efficiency addresses a critical gap where previous research focused exclusively on accuracy improvements without considering computational costs.

The efficiency advantage compounds when scaling to larger portfolios. A trader monitoring ten cryptocurrencies with continuous processing requires ten times computational capacity compared to a single asset. The attention mechanism reduces this to approximately 2.8 times (27.6% trigger rate multiplied by ten assets), enabling much broader coverage with equivalent hardware.

The processing time breakdown reveals API communication latency dominates total processing time, while actual computation requires only five seconds. The attention mechanism's selective activation provides its primary benefit through reducing API call frequency rather than reducing inference time per prediction. With 72.4% fewer predictions, the system makes 72.4% fewer API calls, substantially reducing network overhead.

Accuracy Maintenance Analysis

The marginal 0.1 percentage point RMSE difference between attention mechanism (1.5%) and continuous processing (1.4%) demonstrates that selective processing doesn't sacrifice prediction quality. This finding contradicts potential concerns that filtering market data might cause models to miss subtle patterns visible only through continuous analysis.

The comparable accuracy likely results from two factors. First, the attention mechanism filters noise rather than signal. During low-volatility periods where prices fluctuate within narrow ranges, historical patterns provide limited predictive information. Processing these periods continuously generates predictions that simply reflect mean-reversion behavior without actionable trading information. The attention mechanism recognizes these conditions and appropriately declines to generate spurious predictions.

Second, the dual-model architecture maintains continuous learning even when predictions aren't actively generated. Both long-term and short-term models update incrementally as new data arrives, ensuring pattern recognition capabilities remain current. Selective prediction activation doesn't mean selective learning, models continuously absorb new information regardless of whether predictions are generated.

Confidence Calibration Success

The confidence calibration analysis demonstrates one of the most practically valuable system aspects. Confidence scores showing 0.5-1.2% calibration errors provide traders with honest uncertainty quantification,predictions claiming 90% confidence actually achieve approximately 90% accuracy. This calibration enables optimal risk management: traders can rationally size positions proportional to confidence levels.

The 99% confidence "Hold" recommendations during stable periods represent particularly important functionality. Many systems focus exclusively on identifying profitable trading opportunities while neglecting to explicitly identify when trading should be avoided. The attention mechanism addresses this gap by generating high-confidence Hold signals when market conditions don't warrant action, protecting traders from overtrading, unnecessarily frequent trading that erodes profits through transaction costs.

The confidence calibration success likely results from the R-squared foundation combined with model agreement adjustment. By deriving confidence primarily from validation performance (which directly measures explained variance), the system anchors confidence to empirical accuracy rather than arbitrary scales. The model agreement adjustment appropriately reduces confidence when long-term and short-term models disagree, acknowledging that prediction uncertainty increases when different temporal perspectives provide conflicting signals.

Dynamic Adaptation Effectiveness

The threshold adaptation analysis validates that dynamic adjustment based on recent volatility characteristics successfully maintains appropriate sensitivity across different market regimes. During low-volatility periods, thresholds decrease proportionally, preventing the system from becoming insensitive to emerging trends. During high-volatility periods, thresholds increase to prevent excessive triggering from normal elevated volatility.

This adaptive behavior addresses a fundamental limitation of static threshold approaches. A fixed threshold might function adequately during typical conditions but becomes problematic during regime changes. When Bitcoin enters a high-volatility period with substantial daily movements, a static threshold triggers constantly on routine fluctuations. Conversely, during extremely low-volatility consolidation, the same threshold becomes insensitive, missing movements that represent significant multiples of typical volatility.

The asset-specific calibration proves essential given different baseline volatility profiles across cryptocurrencies. Bitcoin's typical daily volatility requires different thresholds than Solana's movements. Without asset-specific adaptation, the system would either trigger excessively on stable assets or remain insensitive to volatile assets. Independent threshold management for each cryptocurrency ensures appropriate sensitivity across the entire portfolio.

Trading Signal Quality Implications

The 84.8-87.3% profitable trade rates across all five cryptocurrencies demonstrate that prediction accuracy translates to practical trading utility. This finding validates that technical accuracy metrics, while important, ultimately serve as proxies for the true objective: generating recommendations producing profitable trading decisions.

The consistency of profitable trade rates across cryptocurrencies with vastly different characteristics provides strong evidence of universal applicability. Bitcoin's mature market with extensive historical patterns might be expected to produce better trading signals than Solana's newer, more volatile market. The fact that both achieve approximately 85% profitability suggests the attention mechanism successfully adapts to each asset's unique characteristics rather than being optimized for specific types.

The precision-recall balance reveals important system characteristics. Bitcoin and XRP show higher precision, indicating most recommendations are correct, while Cardano and Solana show higher recall, indicating the system identifies most profitable opportunities. This variation reflects different market dynamics, stable assets generate fewer but more reliable signals (high precision), while volatile assets provide more opportunities with slightly less certainty (high recall).

CONCLUSION

This research developed and validated an attention-based volatility filtering mechanism for cryptocurrency trading systems achieving substantial computational load reduction while maintaining prediction accuracy. By intelligently allocating resources to significant market events while ignoring routine fluctuations, the system enables practical real-time operation across multiple cryptocurrencies on standard hardware. Our attention mechanism differs fundamentally from attention mechanisms within prediction models themselves. Model-internal attention weights different time steps or features during prediction generation but doesn't determine whether predictions should occur. Our system-level attention operates at a higher level, determining when the prediction engine should activate.

Compared to ensemble approaches using multiple LSTM models that process continuously, our selective activation achieves comparable accuracy while dramatically reducing computational load. Ensemble approaches with multiple models likely require greater processing capacity than our dual-model architecture, making real-time operation across multiple assets challenging.

Previous deep learning architectures suffered from slow data analysis due to numerous layers. Our attention mechanism addresses computational efficiency directly rather than adding complexity that compounds efficiency problems.

The key innovation lies in multi-layered attention combining volatility detection across multiple timeframes, dynamic threshold adjustment based on recent market characteristics, volume confirmation to validate genuine market activity, persistence requirements to filter temporary spikes, and confidence-calibrated recommendation generation. This comprehensive approach ensures predictions trigger when market conditions warrant attention while conserving resources during stable periods.

Validation across Bitcoin, Ethereum, XRP, Cardano, and Solana demonstrated the system maintains 1.5% average RMSE accuracy while triggering predictions during only 27.6% of time periods, a 72.4% reduction in computational load. Confidence scores showed excellent calibration with errors under 1.2%, providing traders with honest uncertainty quantification. Trading signal quality metrics indicated 85-87% of trades based on recommendations would prove profitable, validating practical utility beyond raw prediction accuracy.

The attention mechanism successfully adapts to different volatility regimes, automatically adjusting thresholds during high-volatility periods to prevent excessive triggering while maintaining sensitivity during low-volatility periods to detect emerging trends. Asset-specific calibration enables handling cryptocurrencies with vastly different characteristics using the same underlying framework.

The sub-10-second alert delivery ensures recommendations reach traders while market conditions remain favorable for execution. Multi-cryptocurrency scalability enables portfolio-level applications where traders simultaneously monitor multiple assets rather than focusing on individual cryptocurrencies in isolation.

This work demonstrates that intelligent resource allocation through attention mechanisms represents a promising direction for cryptocurrency trading system design, achieving the dual objectives of high prediction accuracy and computational efficiency essential for practical deployment. The approach proves particularly valuable for individual traders and smaller investment firms requiring sophisticated analysis capabilities without specialized hardware infrastructure.

ACKNOWLEDGMENTS

This research was conducted as part of doctoral studies at the University of Port Harcourt, Nigeria. We express gratitude to CryptoCompare for providing access to cryptocurrency market data through their public API. We acknowledge the computational resources provided by the Department of Computer Science. Special thanks to colleagues who provided valuable feedback during system development and testing.

REFERENCES

1. Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. Proceedings of the International Conference on Learning Representations.
2. Chittala, M., & Sri, V. (2022). Cryptocurrency price analysis with artificial intelligence. Journal of Engineering Sciences, 13(7), 892-903.
3. Furkan, Y. (2023). Improving prediction accuracy in deep learning-based cryptocurrency price prediction using MM-LSTM. Expert Systems with Applications, 215, 119-135.
4. Reshma, K., Kumar, A., & Singh, R. (2020). Price prediction of bitcoin using machine learning. International Journal of Computer Applications, 182(42), 18-24.
5. Sanjin, L., & Yang, Z. (2023). Intelligent cryptocurrency trading system using integrated AdaBoost-LSTM. Journal of Computational Science, 67, 101-115.
6. Singh, A. (2022). Deep learning architecture for building a cryptocurrency trading platform driven by software robot. Intelligent Systems and Applications, 15(3), 234-248.
7. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998-6008.
8. Alketbi, A., Nasir, Q., & Talib, M. (2018). Blockchain for governmentservicesuse cases, security benefits and challenges. 112–119.
9. Amit, B., Akshata, T., Pooja, N., & Kishor, S. (2019). Real-Time Cryptocurrency Trading System. International Research Journal of Engineering and Technology 6(3).
10. Arman, K., Teh, Y., & David, N. (2011). A Novel ForexPrediction Methodology Based on Fundamental Data. African Journal of Business Management, 5(20).
11. Bahrammirzaee, A. (2010). A comparative surveyof artificial intelligence applications in finance: artificial neural networks. Expertsystem and hybrid intelligent systems.
12. Billard, A., & Kragic, D. (2019). Trends and challenges in robot manipulation. . Science, 364.
13. Chittala, M., & Sri, V. (2022). Cryptocurrency Price Analysis with Artificial Intelligence, . Journal of Engineering Sciences, 13(7).
14. Constance, S., & McIntyre, J. (2015). Cryptocurrencies: Core Information Technologyand Information System Fundamentals EnablingCurrency Without Borders. Information Systems Education Journal, 13(3).
15. Danda, B., Vijay, C., & Ronald, D. (2020). Blockchain Technology: Emerging Applications and use cases for Secure and trustworthy Smart systems. Journal of Cybersecurity and Privacy.
16. Deepa.T, Saraswathi, N., Hariprasad, S., Praveen, K., & Dinesh, V. (2022). Design and implementation of blockchain based peer to peerenergy trading platform. platformJournal of Physics, 2335.
17. Dendej, S., & Sucha, S. (2022). NAGA:multi-blockchain based decentralized platformarchitecture for cryptocurrency payment. International Journal of Electrical and Computer Engineering, 12(4), 4067-4078.
18. Edwards, C., Edwards, A., Stoll, B., Lin, X., & Massey, N. (2019). Evaluations of an artificial intelligence instructor's voice: Social IdentityTheory in human-robot interactions. 90, 357–362.
19. Erickson, B., Korfatis, P., Akkus, Z., & Kline, T. (2017). Machine learningfor medical imaging. RadioGraphics. 37(2), 505–515.
20. Farell, R. (2015). An analysis of the cryptocurrency industry.
21. Fernández-Caramés, T., & Fraga-Lamas, P. (2018). A Review on the Use of Blockchain for the Internet of Things. 6, 32979–33001.

22. Haritha, R., & Shyma, K. (2010). An Introduction to Blockchain Technology. International Research Journal of Engineering and Technology, , 7(7).
23. Kuepper, J. (2021). Risk Management Techniques for Active Traders.
24. Mcintosh, S. (2018). The Business Benefits of Cryptocurrency. Available online: https://www.theglobaltreasurer.com/.
25. Mueller , P., Bergsträßer, S., Rizk, A., & Steinmetz, R. (2018). The Bitcoin Universe, An Architectural Overview of theBitcoin Blockchain, Informatics. 11.
26. Rajashree, V., & Girish, K. (2020). Blockcloud: Integrating Blockchain Security Features in Cloud Computing,International Journal of Engineering and Advanced Technology, 9(4).
27. Reiff, N. (2020). Blockchain Explained. Available online: https://www.investopedia.com/terms/b/blockchain.asp.
28. Reshma, S., Nikita, T., & Naveen, C. (2020). Price prediction of Bitcoin using Machine Learning. International Journal of Engineering Applied Sciences andTechnology, 5(1), 502–506.
29. Sanjana, P., Vaishnavi, P., & Divya, K. (2016). An Overview of Blockchain Architecture andit's Applications, . An Overview of Blockchain Architecture andit's Applications, International Journal of Innovative Research in Science,Engineering and Technology, 5(11).
30. Sarvesh, M., & Shriti, S. (2020). An Analysis of Cryptocurrency ind Its Challenges. International Journal of Multidisciplinary Research, 6(4).
31. Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and beyond: A technicaltutirial on decentralized digital currencies2123. 18(3), 2084–2090.
32. Vasukidevi, G., & Sethukarasi, T. (2020). A Systematic plan of Blockchain Technology on Security Key,International. International Journal of Engineering Research & Technology, 9(2).