http://dx.doi.org/10.18623/rvd.v23.n3.4424

DEVELOPMENT OF A CRYPTOCURRENCY TREASURY MANAGEMENT SYSTEM

DESENVOLVIMENTO DE UM SISTEMA DE GESTÃO DE TESOURARIA EM CRIPTOMOEDAS

Article received on: 8/29/2025
Article accepted on: 11/28/2025

Canberk Erol*
*Innovance Information Technologies, Dept of Software Development, Istanbul, Turkey
Orcid: https://orcid.org/0009-0000-0849-2597
canberk.erol@innovance.com.tr

Burak Aras Yılmaz*
*Innovance Information Technologies, Dept of Software Development, Istanbul, Turkey
Orcid: https://orcid.org/0009-0002-7577-4978
aras.yilmaz@innovance.com.tr

Ceren Ulus**
**EFA Innovation Consultancy, Dept of R&D, Adana, Turkey
Orcid: https://orcid.org/0000-0003-2086-6381
f.cerenulus@gmail.com

Mehmet Fatih Akay***
***Çukurova University, Department of Computer Engineering, Adana, Turkey
Orcid: https://orcid.org/0000-0003-0780-0679
mfakay@cu.edu.tr

The authors declare that there is no conflict of interest

Abstract
This study aims to develop an innovative solution, supported by next-generation technologies, to facilitate operational integration of crypto assets in banking sector. To this end, a system has been developed that provides real-time risk management and portfolio optimization by combining big data analysis, Natural Language Processing (NLP), and deep learning. The system focuses on price prediction, sentiment analysis, risk measurement, and dynamic portfolio management. In addition, social media and news feeds have been digitized to create an extra layer that provides direct input to Artificial Intelligence (AI) engine. Furthermore, a Deep Reinforcement Learning (DRL) agent has been designed to maximize risk-adjusted return instead of absolute return, to adapt to stochastic nature of market. Deterministic risk models have been integrated with proactive safety mechanisms. This allows for focusing on financial risk, operational risk, and safety dimensions. A hybrid method combining Long Short-Term Memory (LSTM) and Transformer models has been proposed to simultaneously capture linear and nonlinear market dynamics. Auto Regesive Integrated

Resumo
Este estudo visa desenvolver uma solução inovadora, apoiada por tecnologias de última geração, para facilitar a integração operacional de criptoativos no setor bancário. Para tanto, foi desenvolvido um sistema que proporciona gestão de risco em tempo real e otimização de portfólio, combinando análise de big data, Processamento de Linguagem Natural (PLN) e aprendizado profundo. O sistema concentra-se na previsão de preços, análise de sentimento, mensuração de risco e gestão dinâmica de portfólio. Além disso, mídias sociais e feeds de notícias foram digitalizados para criar uma camada extra que fornece entrada direta para o mecanismo de Inteligência Artificial (IA). Ademais, um agente de Aprendizado por Reforço Profundo (DRL) foi projetado para maximizar o retorno ajustado ao risco em vez do retorno absoluto, adaptando-se à natureza estocástica do mercado. Modelos de risco determinísticos foram integrados a mecanismos de segurança proativos. Isso permite o foco nas dimensões de risco financeiro, risco operacional e segurança. Um método híbrido, combinando modelos de Memória de Longo Prazo (LSTM) e Transformer, foi proposto para capturar

Moving Average (ARIMA) and LSTM-based prediction models have also been developed to analyze prediction success of proposed method in this study. Examination of results revealed that proposed hybrid model exhibited high performance.

Keywords: Cryptocurrency Treasury Management System. Sentiment Analysis. Price Prediction.

simultaneamente a dinâmica linear e não linear do mercado. Neste estudo, também foram desenvolvidos modelos de previsão baseados em Média Móvel Integrada Auto-Regressiva (ARIMA) e LSTM para analisar o sucesso da previsão do método proposto. A análise dos resultados revelou que o modelo híbrido proposto apresentou alto desempenho.

Palavras-chave: *Sistema de Gestão de Tesouraria de Criptomoedas. Análise de Sentimento. Previsão de Preços.*

1 INTRODUCTION

Treasury management departments in banks carry out essential roles like ensuring liquidity, managing risks, and keeping a balance between assets and debts. However, the 24/7 trading of cryptocurrency markets, high price fluctuations, shallow market depths, and regulatory ambiguities weaken the efficiency of conventional treasury management systems. This scenario makes it difficult to incorporate crypto assets into treasury management activities in both strategic and operational ways. There are three main areas of difficulty in cryptocurrency treasury management standing out. Firstly, liquidity management cannot be effectively managed with traditional liquidity measurement techniques due to the fragmented nature of crypto markets and sudden price movements. Secondly, risk measurement and pricing processes require approaches beyond classical risk models due to the high volatility of crypto assets and their variable correlation with traditional financial assets. Third, market trend analysis poses a complex problem area due to the frequent regime shifts and speculative behavior in crypto markets. In addition, although the interaction of cryptocurrency markets with macroeconomic factors has increased in recent years, the integration of variables such as interest rates, inflation expectations, and global risk appetite into treasury management systems has been limited. This leads to significant information losses in strategic decision-making processes (Tamanna Choithani, 2022).

This study aims to develop an innovative solution, supported by next-generation technologies, to facilitate the operational integration of crypto assets in the banking sector. To this end, a system has been developed that provides real-time risk management and portfolio optimization by combining big data analytics, NLP, and deep learning.

The most significant difference between this study and the existing literature is that this study proposes a system aimed at integrating crypto assets into the banking sector at an operational level. Current literature generally addresses price prediction, sentiment analysis, or DRL-based portfolio optimization separately. This study, however, focuses on price prediction, sentiment analysis, risk measurement, and dynamic portfolio management together. In this context, big data analytics, NLP, and deep learning have been used. A hybrid method combining LSTM and Transformer models has been developed to simultaneously capture linear and nonlinear market dynamics. To analyze the prediction success of this method, ARIMA and LSTM-based forecasting models have also been developed. The performance of the developed models has been evaluated using four different metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE) and Coefficient of Determination (R^2). Additionally, social media and news feeds have been digitized to create an extra layer providing direct input to the AI engine. Furthermore, to adapt to the stochastic nature of the market, a DRL agent has been designed that aims to maximize risk-adjusted return instead of absolute return. Deterministic risk models have been integrated with proactive safety mechanisms. This allows for a focus on financial risk, operational risk, and security dimensions. In this respect, the study offers an innovative approach, deviating from existing literature.

This study is organized as follows: Section 2 includes relevant literature. Methodology is presented in Section 3. Details of the system are presented in Section 4. Results of the study are given in Section 5. Section 6 concludes the paper.

2 LITERATURE REVIEW

(Yichen Luo et al., 2025) aimed to develop an approach to cryptocurrency investments that could generate both high-performance and explainable decisions. To this end, a multi-agent investment framework using multimodal data consisting of expert agents has been designed to overcome the limitations of a single Large Language Model. In the framework, agents carry out sub-tasks such as data analysis, literature integration, and investment decision-making through collaboration within and between teams. Experiments with data from November 2023–September 2024 have shown that the proposed approach outperforms single-agent models and market benchmarks in terms of classification, pricing, portfolio performance, and explainability.

(Bahram Alidaee et al, 2025) investigated Modern Portfolio Theory models within the U.S. stock and cryptocurrency markets. Their analysis indicated that the

Markowitz Portfolio Selection and the Optimal Dynamic Portfolio models yield superior returns and operate more efficiently than other approaches. Additionally, it has been seen that these models are effective in reducing risks and outperforming both exchange-traded funds and various other strategies. Ultimately, it has been concluded that the frequency of data collection did not influence the choice of portfolio.

(Mehrdad Heydarpour et al., 2025) introduced a model for portfolio optimization that utilizes short-term Moving Average (MA) techniques alongside LSTM forecasts to minimize uncertainties in algorithmic trading. The research implemented algorithms including Variable Length MA, Flexible Length MA, Exponential MA, and Simple Moving Average (SMA) for stock and cryptocurrency portfolios. The findings demonstrated that these algorithmic techniques are more effective than the buy and hold strategy. Nevertheless, it has been observed that certain strategies may not be suitable in specific circumstances. Overall, it has been concluded that portfolios that utilize all strategies and assets tend to outperform traditional methods, even during uncertain market conditions.

(Xinran Huang et al., 2025) proposed a novel method for optimizing portfolios based on a Conditional Value-at-Risk utility function, which assesses tail risks. The effectiveness of this method has been compared against other portfolio development strategies, such as the Naïve, Minimum Variance, and Mean-Variance portfolios. The results indicated that the proposed method performed better than conventional optimization techniques.

(Taufeeq Hussain and M. Ramamoorthy, 2025) aimed to enhance cryptocurrency portfolio management by employing LSTM networks and evaluating its performance against DRL methods. An LSTM model has been created using past cryptocurrency price data for optimal portfolio management. From the comparative study, it has been found that the DRL algorithm outperformed the LSTM model with an Accuracy rate of 83.20%, and the difference has been statistically significant. Consequently, LSTM-based models have been identified as valuable tools for optimizing the allocation of cryptocurrency portfolios.

(Zhongyuan Xu, 2025) introduced a Reinforcement Learning (RL) algorithm to enhance dynamic portfolio optimization. This approach allows an AI broker to engage with the cryptocurrency market, consistently refining asset allocation and managing volatility, all while boosting returns. Experimental data uncovered that the RL model significantly outperforms conventional tactics. Important metrics, including a cumulative return rate of 85.12%, an annualized volatility of 45.76%, and a peak drawdown rate of -22.34%, demonstrated the RL model’s effective revenue generation and risk management skills.

(Habib ZOUAOUI and Meryem-Nadjat NAAS, 2025) focused on exploring advanced portfolio management techniques in cryptocurrency markets. To this end, behavioral finance theories focusing on the concept of return risk have been used together with Artificial Neural Networks and deep learning approaches such as LSTM. A random selection of 25 cryptocurrencies from the Yahoo Finance database has been made to assess how well the models succeeded in identifying the ideal ratios of the investment portfolio. Deep learning models recorded a MSE of 0.0218% when estimating optimal portfolio weights over a five-day period.

(Mariam Elkhechafi and Jihane Aayale, 2024) introduced a novel portfolio management system based on DRL. This system differs from conventional methods by incorporating five parameters, including trading volume, for a more thorough market analysis and simultaneously assessing various data sources. Deep convolution techniques have also been employed, specifically tailored to analyze each parameter independently, along with a five-dimensional attention gating network to better detect crucial moments and high-potential assets. The simulation outcomes indicated that this strategy vastly exceeds traditional approaches, providing investors with improved returns and reduced short-term risks.

(Rihab Bedoui et al., 2023) explored the possible advantages of using the Conditional Value at Risk (CVaR) portfolio optimization method. In their study, they combined the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, Extreme Value Theory, and Vine Copula techniques to determine the best allocation for a portfolio including Bitcoin, gold, oil, and stock indices. To enhance the risk evaluation, the model incorporated the Monte Carlo simulation along with the Mean-CVaR model, while a genetic algorithm based on deep machine learning has been employed for the optimization process. This work presented fresh evidence for the CVaR portfolio optimization method, providing useful information for portfolio managers aiming to enhance multi-asset portfolios.

(Fengrui Liu et al., 2021) aimed to develop an automated and efficient high-frequency trading approach for the extremely unstable cryptocurrency market, focusing on bitcoin. To achieve this, a trading process using a state-action-reward framework and introduced a DRL model based on Proximal Policy Optimization (PPO) has been modeled. In the initial phase, different models which use Support Vector Machine, Multi-Layer Perceptron, LSTM, Temporal Convolutional Network and Transformer have been evaluated for predicting prices. LSTM model has yielded the best outcomes, and it has been combined with the PPO method to formulate a trading strategy. Results from the experiments show that the suggested approach surpasses standard strategies across multiple performance indicators, particularly achieving 31.67% higher returns, reflecting exceptional performance during both stable and turbulent times.

(Saurabh Kumar et al., 2024) focusing on Bitcoin, Ethereum, and Litecoin, they aimed to predict the values of cryptocurrencies. To this end, Bidirectional Long Short-Term Memory (Bi-LSTM), LSTM, and Gated Recurrent Unit (GRU) methods have been used. Model performance has been evaluated using MAPE. The GRU algorithm has been observed to perform better than the other two algorithms in terms of accuracy and MAPE percentages for BTC, LTC, and ETH, with percentages of 0.2116%, 0.2454%, and 0.8267%, respectively. Bi-LSTM showed the weakest prediction performance.

3 METHODOLOGY

3.1 Auto regresive integrated moving average

Box and Jenkin developed ARIMA in the 1970s. "AR" stands for autoregression and characterizes the model as using the dependent relationship between current data and past values. "I" comes from integration and indicates that the data is stationary. As an active estimator that relates the past and present to future predictions, ARIMA has shown quite high success in short-term forecasting (Xiangzhou Chen and Zhi Long 2023).

3.2 Long short-term memory

LSTM, a deep learning model, is a more suitable forecasting technique for time series data, capable of predicting long-distance dependence between output and input with high performance. In recent years, LSTM has yielded significant results in NLP, economics, energy, and transportation forecasting (Hamzeh F. Assous et al., 2020).

3.3 LSTM and transformer hybrid model

The working principle of the hybrid price prediction model is given in Figure 1.

Figure 1

The working principle of the hybrid price prediction model

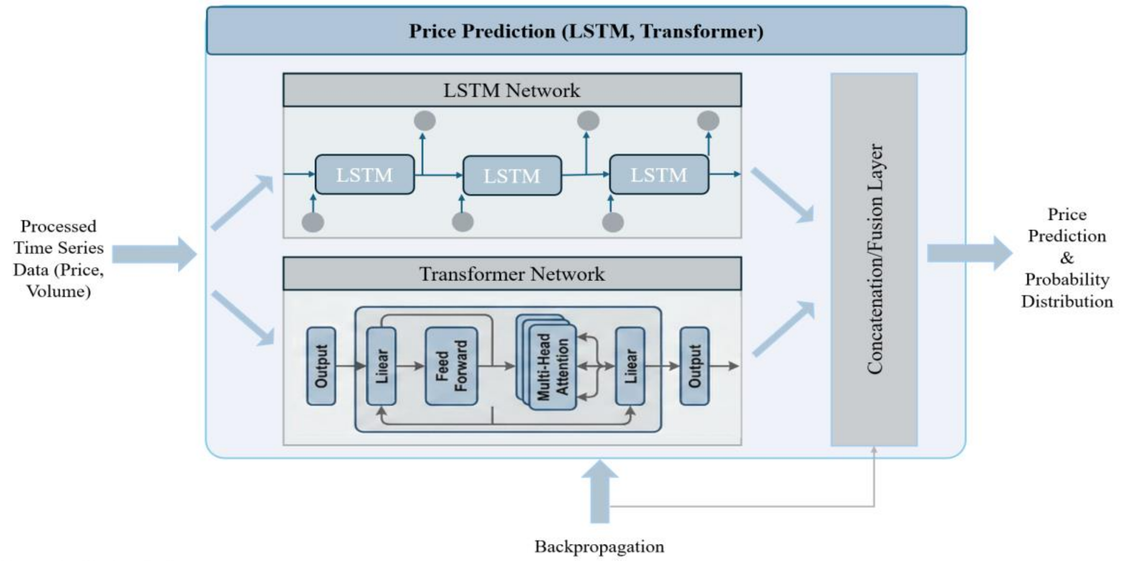

The stochastic and non-stationary nature of cryptocurrency markets limits the predictive performance of traditional linear models like ARIMA and GARCH. To overcome this limitation, Parallel Dual-Branch Network architecture has been developed. This structure combines LSTM layers, which learn sequential dependencies in time series, with Transformer (Self-Attention) blocks, which can successfully capture global relationships and sudden market fluctuations in the dataset, in a single Feature Fusion layer. Temporal Feature Extraction (Temporal Branch) channel with LSTM is responsible for learning the direction (trend) and momentum of price movements. It eliminates the vanishing gradient problem seen in standard Recurrent Neural Network (RNN) through gate mechanisms. The Transformer and the Global Attention Branch channel, on the other hand, overcomes the need for sequential processing and captures long-range dependencies and volatility bursts in the time series; in this context, an 8-head (multi-head) self-attention mechanism is used to weigh which time periods are more critical in the current price formation. Furthermore, it naturally cannot know the time order of the Transformer architecture. Accordingly, sequence information is added to the input vector using positional coding based on sine and cosine functions. The output vectors obtained from two parallel branches are combined in the Feature Fusion layer and passed through two Dense layers with the ReLU activation function, enabling the model to generate the estimated closing price for the next time step.

3.4 Financial bidirectional encoder representations from transformers

Financial Bidirectional Encoder Representations from Transformers (FinBERT) is a Bidirectional Encoder Representations from Transformers (BERT) embedding method, a cutting-edge technique that captures the context of words within sentences. This allows the model to understand the meaning behind specialized financial terminology and subtle expressions, which is crucial in sentiment analysis. With BERT’s transfer learning capability, FinBERT demonstrates high performance for the sentiment classification task (Olamilekan Shobayo et al., 2024).

3.5 Proximal policy optimization

PPO, released by OpenAI in 2017, is a stable and easy-to-implement policy-based learning algorithm used in the field of RL. The Actor Network decides which action to take in the current situation. The Critic Network estimates the value of the action taken and the current state. To prevent the agent from making too radical changes when updating its weights, the probability ratio is trimmed within a certain range. This mechanism prevents the agent from over-relying on risky strategies found "by chance" and stabilizes the learning process.

3.6 Historical simulation value at risk

Value at Risk (VaR) is widely used to measure the market risk of an asset or portfolio. The parametric VaR model introduces a strong theoretical assumption regarding the fundamental properties of the data; a normal distribution is usually assumed because it is well understood and can be defined using the first two moments (mean and standard deviation). In general, VaR techniques are based on non-parametric, parametric, or a mixture of parametric and non-parametric statistical methods. The Historical Simulation family of models is a non-parametric approach. The Historical Simulation method is one that requires no statistical assumptions, particularly regarding volatility, other than the assumption that returns are stationary (Indrajit Roy, 2011).

3.7 Variance-covariance value at risk

Introduced by JP Morgan in the early 1990s, the Variance-Covariance method is quite advantageous in terms of ease of calculation and application. Basically, the portfolio consists of linear assets; that is, changes in portfolio value are linearly related to changes in asset value. Furthermore, the method assumes that asset returns follow a normal distribution (Nurul Aryani, 2025).

4 DETAILS OF THE SYSTEM

The developed system has been built on fault tolerant and horizontally scalable microservice architecture. The system operates with an average latency of 45ms from end to end. The system, which has been developed in a multi-layered structure, consists of four main layers: Data Sources, AI and Machine Learning Engine, Decision Support System (DSS) Interface and Security Infrastructure. The system architecture is given in Figure 2.

Figure 2

System architecture

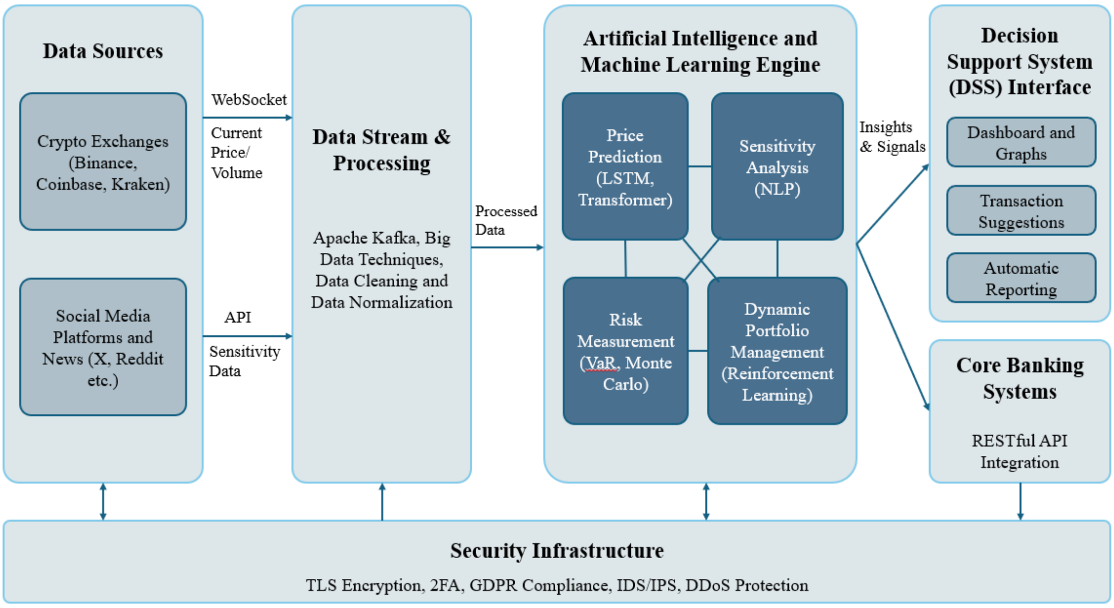

4.1 Data sources and transmission protocols management

The system uses multi-protocol connector services to capture microsecond precision changes in financial markets

4.1.1 High-frequency stock exchange data links and connector architecture

The Exchange Connector layer, which serves as the system's gateway to the outside world, is the most critical component in terms of the accuracy and speed of financial data. To eliminate the latency caused by traditional HTTP methods, an asynchronous and persistent data line has been established with Binance, Coinbase, and Kraken exchanges via the Web Socket Secure (WSS) protocol. This structure has been upgraded to process an average of 12,000 data packets per second.

An active-active redundancy architecture has been implemented to guarantee uninterrupted data flow for the system 24/7. Within the WebSocket Clustering framework, instead of a single connection for each exchange, three parallel worker nodes are assigned behind a load balancer. Node A processes the data and writes it to Kafka. Node B retrieves the data but does not write it to Kafka; it takes over if Node A's heartbeat signal is interrupted. Node C is kept ready for redundancy and maintenance purposes. To prevent exchanges from timeout, a pong response is sent every 3 minutes at the application layer, and a Transmission Control Protocol (TCP) Keep-Alive signal is sent every 30 seconds at the network layer. In case of a connection loss, the system reconnects after waiting 2^n seconds to avoid "spamming" the server. According to test results, re-establishing a lost connection and restarting data flow takes an average of 180ms.

The system records data from the exchange with a double stamp: event time and processing time. Within trade streams and aggregation, each buy-sell transaction in the market is monitored via the channel. During periods of high volatility, the system processed 45,000 trade messages per second without data loss. To reduce network load, orders executed within the same millisecond and at the same price are combined into a single packet at the system input. Instead of simple price tracking, the system replicates market depth in its own memory by maintaining a local order book. Initially, a complete snapshot is taken from the exchange. Subsequently, only changes are processed. Each packet from Binance and Kraken APIs has a Final Update ID value. Each data packet from Binance and Kraken Application Programming Interface (API) has an updated number. The system checks if the packet number is one more than the previous one. If a number is skipped, it indicates that the data is incorrect. In this case, the system closes the WebSocket connection and reconnects to retrieve the most up-to-date data. To avoid confusion within the system, the different JSON structures of the exchanges have been standardized.

An aggressive compression policy has been implemented to reduce network bandwidth costs and decrease the Input/Output (I/O) load on the Kafka cluster. At the exchange exit, raw JSON data is converted to Google Protocol Buffers format before being transmitted to Kafka. The serialized data is then compressed using the Zstandard algorithm, which offers higher CPU efficiency, instead of Snappy or GNU Zip (GZIP). This process results in a 65% reduction in data size compared to text-based JSON data. For example, daily raw data input of 1.5 TB occupies only 525 GB of disk space.

The latency times obtained during live production tests performed on the developed system are given in Table 1.

Table 1

The latency times obtained during live production tests

Metric	Value	Description
Exchange Latency	20ms - 150ms	Network latency between the exchange's server and our server (when co-location is not used).
Internal Processing	< 2ms	Time it takes to receive, normalize, and write data to Kafka.
System Uptime	%99.98	Operational uptime for the last 6 months.
Maks. Throughput	85 MB/s	Maximum data bandwidth processed at any given moment.

4.1.2 Unstructured data sources and sensitivity data pipeline

To model investor psychology, market sentiment, a significant leading indicator of price movements in financial markets, a highly scalable data ingestion pipeline has been established for collecting text-based data. This module feeds social media signals and global news feeds into an NLP engine by normalizing them. The system processes and analyzes an average of 2.5 million text fragments per day.

A persistent and real-time connection has been established using Twitter's "Filtered Stream" endpoints to capture hype and Fear, Uncertainty, Doubt (FUD) waves in cryptocurrency markets. Authentication is performed at the application level using OAuth 2.0 Bearer Token. When retrieving tweet objects, in addition to the default fields, the number of likes, retweets, user age for bot detection, and financial context fields are also included in the query.

To optimize API costs and improve data quality, filtering rules with complex Boolean logic have been implemented. A language filter of lang:en has been used. In terms of performance and volume, the average flow rate has been obtained as 25-40 Tweets/second. Peak volume reached up to 180 Tweets per second during significant market events. Latency is the time between being sent and entering the Kafka queue, and it has been measured as less than 400ms.

A hybrid polling architecture has been developed to track macroeconomic developments and corporate announcements. Private API subscriptions have been used as sources for sources such as Bloomberg, Reuters, and CoinDesk. Cross-validation has been performed using the CryptoPanic and CoinGecko News APIs as aggregators. Really Simple Syndication (RSS) feeds are sourced from over 50 authoritative blogs and official

RSS feeds from regulatory bodies such as the Securities and Exchange Commission (SEC) and the Fed. To efficiently utilize server resources, a dynamic range algorithm is used instead of a static query time: Noise, the biggest problem with unstructured data, is cleaned before it enters the NLP engine.

To prevent data pollution, when the same news item is published in different sources or the same tweet goes viral, a deduplication process is performed using SimHash and Cosine Similarity algorithms. If the content similarity exceeds 90%, the data is considered duplicate and is not processed. This eliminates approximately 35% of daily data input. Additionally, a rule-based preliminary analysis is applied to Twitter data to detect bots and spam. The success of this module directly affects the accuracy of the Sensitivity Analysis model, which is the next stage. The error values obtained in this phase are given in Table 2.

Table 2

The error values obtained for the data ingestion pipeline

Metric	Value	Description
Daily Data Volume	~4.2 GB	Raw text data (in JSON format).
Signal/Noise Ratio	%68	32% of raw data (spam, ads, irrelevant) is filtered out.
Source Coverage	%94	94% of major market news is available within the first 2 minutes.
API Success Rate	%99.95	Accessibility including rate limit overruns and connection drops.

4.2 Data flow and event processing backbone

The system's central nervous system consists of an event streaming architecture based on Apache Kafka and Apache Flink, guaranteeing asynchronous data transmission and high throughput. This structure can perform operations at the microsecond level while maintaining data integrity.

4.2.1 Kafka cluster configuration, topology, and robustness

Data streaming is configured on a horizontally scalable and fault-tolerant cluster. Within the Producer Layer and Partitioning Strategy, a Symbol Pair such as BTCUSDT is used as the key to guarantee sequential processing of data. This ensures that data for the same pair always falls into the same Kafka partition, preserving time-series integrity. This guarantees that the process is not considered complete until the data is confirmed to have been written by the leader and all replicas. To optimize network load, *linger.ms* is set to 5 and *batch size* to 16384. This allows the producer to accumulate and send data in batches for 5ms, thus reducing Input/Output Operations Per Second (IOPS) costs.

The system runs on a Kafka cluster consisting of 3 physical nodes. Each data block is stored on 3 different servers. Even if a server experiences a hardware failure, no data loss occurs, and the system continues to operate uninterrupted. For a write operation to be considered successful, at least two copies must be alive. To minimize latency, log segments are stored on Non-Volatile Memory Express (NVMe) SSD disks and zero-copy data transfer technology is used.

4.2.2 Real-time ETL and stream processing

Raw data read from Kafka topics is cleaned and enriched using stateful operations on the Apache Flink engine before being fed into AI models. Statistical filtering is applied on the sliding window to eliminate flash crash or erroneous API responses that may occur in data from exchanges, as part of anomaly detection and z-score filtering.

Missing data completion and watermarks repair millisecond gaps in the data stream caused by network latency to prevent the AI models from compromising time series integrity. Using event time on Flink, data is processed in the correct order even if the order of arrival is mixed up. Delays of up to 500ms are tolerated with the watermark strategy. In the context of interpolation algorithm, the missing data points are filled using linear interpolation with the previous and next valid data points.

For dynamic normalization, data is compressed to the *[0, 1]* range to ensure stable operation in gradient descent optimization of LSTM. Flink stores the minimum and maximum prices for the last 24 hours in memory for each currency pair. Each incoming price is normalized according to the instantaneous global min/max values.

The processed clean data is written to the *processed_features* topic and distributed to different consumers using a fan-out architecture. The AI engine directly feeds the LSTM and Transformer models. Data is written on the PostgreSQL-based time series database TimescaleDB. Data is automatically chunked over time to improve disk I/O performance. Data older than 7 days is achieved by compressing it by 90% using

TimescaleDB's native compression algorithm. The results obtained from production stress tests of this established flow architecture are given in Table 3.

Table 3

The results obtained from production stress tests of this established flow architecture

Performance Metric	Value	Description
Throughput	15.000 TPS	The total number of messages processed and normalized per second.
E2E Latency	42 ms (p99)	The time elapsed from exiting the exchange API to entering the AI model.
Data Loss Rate	%0.00	No data loss occurred even in broker crash simulations.
Jitter	±4 ms	The level of stability in the data stream.

4.3 Artificial intelligence modules

4.3.1 Price prediction

In this study, prediction models have also been developed using ARIMA and LSTM to evaluate the performance of the prediction models developed using the proposed hybrid method.

The model is fed with a multidimensional tensor representing market dynamics instead of raw price data. A 3-layer (Stacked) LSTM structure is used. Each layer contains 128 neurons, and a 20% dropout is applied between layers to prevent overfitting. The hybrid model is trained on historical datasets using the Backpropagation Through Time (BPTT) method. The MSE loss function is used to minimize prediction errors. The *AdamW* optimization algorithm with adaptive learning rate is used. Here, if the validation error does not improve over 10 epochs, training is stopped and the best weights are restored.

The performance of these models has been evaluated using RMSE, MAE, MAPE (%), and R^2. The results are presented in Table 4.

Table 4

The RMSE, MAE, MAPE (%), and R^2 values of the developed models

Model	RMSE	MAE	MAPE (%)	R^2
ARIMA	145.20	112.45	4.2	0.65
LSTM	88.40	65.10	2.1	0.82
Hybid Model	62.15	45.30	1.4	0.91

When the results have been examined, it has been observed that the highest prediction performance has been obtained with the hybrid model.

4.3.2 Sensitivity analysis

In the study, a sensitivity analysis module has been developed to support price prediction models. This module analyzes social media posts and news texts to convert the overall mood of the market into a numerical signal. Figure 3 shows the flow chart of sensitivity analysis.

Figure 3

Sensitivity analysis flow chart

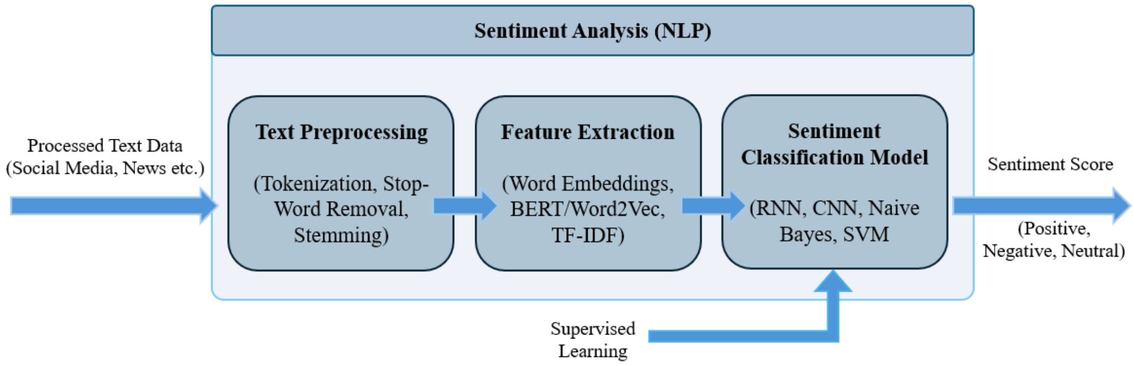

A multi-layered cleaning filter has been applied to normalize the unstructured, noisy, and non-standard such as slang, emojis, abbreviations nature of social media data. Named Entity Recognition (NER) is used to identify which entity the sentiment in the text belongs to. Exchange tags are parsed using Regex algorithms. Indirect references such as Ether and Vitalik's coin are matched with the relevant crypto asset using Fuzzy Matching libraries. This way, the sentence "Bitcoin is falling but Solana looks great" is correctly classified as Negative for BTC and Positive for SOL. Bot activity and spam content aimed at manipulating the dataset are eliminated. Signals from user accounts that are less than seven days old or have fewer than ten followers have been marked as unsafe. If the same text is shared by thousands of bots, the SimHash algorithm measures text similarities, and duplicate content is not included in the analysis. In crypto jargon, emojis carry more powerful meaning than words. The system converts emojis into textual tokens such as *Strong_Buy, Strong_Sell* and *HODL*, feeding them into the model.

Standard static models like Word2Vec or GloVe cannot understand the context of a word within a sentence. To overcome this limitation, the BERT architecture has been used as a basis and customized for the financial sector. The model has been fine-tuned using financial news, technical analysis reports, and crypto forum discussions instead of the standard English corpus. The model learned that "Long" means "buy position" and not "length," "Short" means "sell position" and not "short," "Rekt" means "heavy loss," and "Moon" means "extreme bullish."

Vectors from the FinBERT layer are processed in the classification layer to create a continuous sensitivity score, weighted according to the time dimension. Within the scoring function, the classifier output (Logits) is passed through the Softmax function to calculate probabilities belonging to 3 classes such as Positive, Negative, Neutral. The final score is normalized by taking the weighted sum of these probabilities.

Due to the speed of financial markets, the value of news from 1 hour ago is not the same as news from 1 minute ago. Sensitivity scores are weighted using an exponential decay function. This allows the impact of old news to quickly fade, while fresh news dominates the signal. The calculated weighted sensitivity score is fed into the hybrid price prediction model as an additional feature.

The definitive results obtained from live market tests are given in Table 5.

Table 5

The definitive results obtained from live market tests

Evaluated Factor	**Performance Metric**	**Result**
Sudden changes in social media sensitivity scores	Leading Indicator	It occurred approximately 15 minutes before price movements.
BERT	F1-Score	0.89
FinBERT	F1-Score	0.74
Sensitivity score and volatility	Correlation	0.72 Positive

4.3.3 Dynamic portfolio management

To adapt to the stochastic nature of market conditions, a DRL agent has been developed that learns from its own experiences and aims to maximize risk-adjusted return. Figure 4 shows the flowchart of the DRL process.

Figure 4

Dynamic portfolio management flow chart

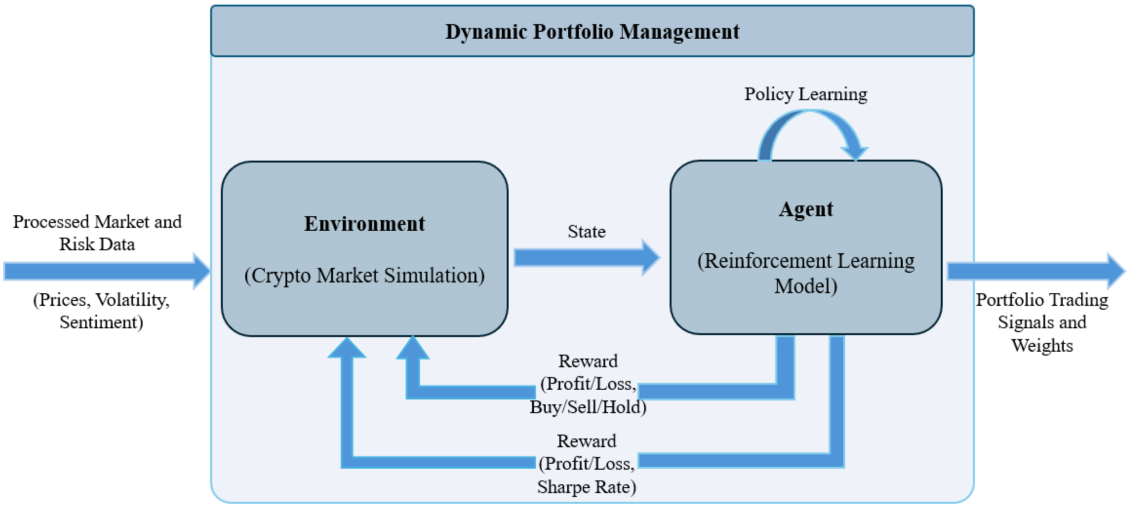

The financial portfolio management problem is mathematically formulated as a Markov Decision Process (MDP). Within this framework, the structure (S, A, R, P, γ) is defined. The element that enables the agent to perceive the market is the observation vector. At each time t, the agent observes a multidimensional tensor representing the instantaneous state of the market. The actions the agent can take are defined in a discrete action space instead of a continuous one to allow for faster convergence of the model.

To ensure stability in the training process and prevent the risk of policy collapse, the Actor-Critic based PPO algorithm has been chosen. A multi-component reward function has been designed to encourage the agent to develop not only a profit-oriented but also a risk-averse strategy. The profit component is the change in the total value of the portfolio. A volatility penalty is awarded to the agent as the portfolio's variance increases. This encourages the agent to stay away from excessively volatile assets or uncertain market conditions. A penalty equal to the exchange commission rate is deducted from the transaction cost for each buy/sell order. This prevents the agent from overtrading unnecessarily and ensures they only trade where the profit margin is higher than the commission.

The PPO-based agent learned to conserve capital by converting the portfolio to cash (USDT), especially during bear markets and sideways trends. In terms of improvement, it maximizes profitability using a standby buying strategy. The results showed a significant advantage over the traditional buy and hold strategy. These results demonstrate that AI provides a mathematically optimized execution layer in financial decision-making processes, free from human emotions.

4.4 Risk management and security

4.4.1 Financial risk modeling and quantitative analysis

To manage the higher volatility of cryptocurrency markets compared to traditional markets, standard risk models have been modified to incorporate heavy-tail distributions.

The system uses a hybrid approach of Non-Parametric (Historical Simulation) and Parametric (Variance-Covariance) VaR methods to calculate the maximum amount a portfolio can lose within a specific time period and a certain confidence interval.

Instead of standard deviation, the GARCH (1,1) model, which captures volatility clustering, has been integrated into Volatility Modeling. This allows the model to narrow risk limits by predicting the possibility of a sudden shock (volatility burst) even when the risk appears low during calm market periods.

A Monte Carlo engine based on Geometric Brownian Motion (GBM) has been developed to derive possible future price scenarios. The system simulates 10,000 different price paths every 15 minutes to determine the possible future allocation of the portfolio. Historical black swan is injected into the system to analyze the state of liquidity risk and capital adequacy under extreme conditions.

4.4.2 Cybersecurity architecture and defense depth

The system integrates with banks' existing core banking infrastructure while implementing a layered security architecture to minimize the cyberattack surface.

System traffic is protected by Intelligent Firewalls (WAF) and anomaly-based detection systems operating at the application layer. Intrusion Detection (IDS) monitors signature-based threats and behavioral anomalies in network traffic in real time. Active Blocking (IPS) automatically quarantines suspicious IP addresses or traffic patterns showing signs of Distributed Denial of Service (DDOS) attacks without affecting system performance. All internal and external data traffic is end-to-end encrypted using the TLS 1.3 protocol and AES-256 standard. This makes it mathematically impossible to intercept or modify data in transit.

The least privilege principle has been implemented to prevent unauthorized access. Access to administrator panels is protected by Time-Based One-Time Passwords (TOTP) or Two-Factor Authentication (2FA) requiring biometric authentication, beyond static passwords. The permissions of treasury personnel are clearly defined.

4.4.3 Data privacy, regulatory compliance and auditability

Due to the sensitivity of banking data, the system design is fully compliant with General Data Protection Regulation (GDPR) and Turkish Personal Data Protection Law (TPDPL) standards, based on the privacy by design principle. User identities and transaction details are masked with hashing algorithms before being recorded in the database. Analytical processes operate on data sets purging personal data. All transactions on the system are stored in an immutable blockchain-like log structure. This structure ensures 100% transparency and accountability in the event of a possible internal audit or regulatory review. All data is stored on the bank's local servers or in secure cloud areas with defined geographical boundaries, compliant with local regulations.

4.5 User interface

Screenshots of the interface of the developed system are given in Figures 5 to 12.

Figure 5

Overview dashboard

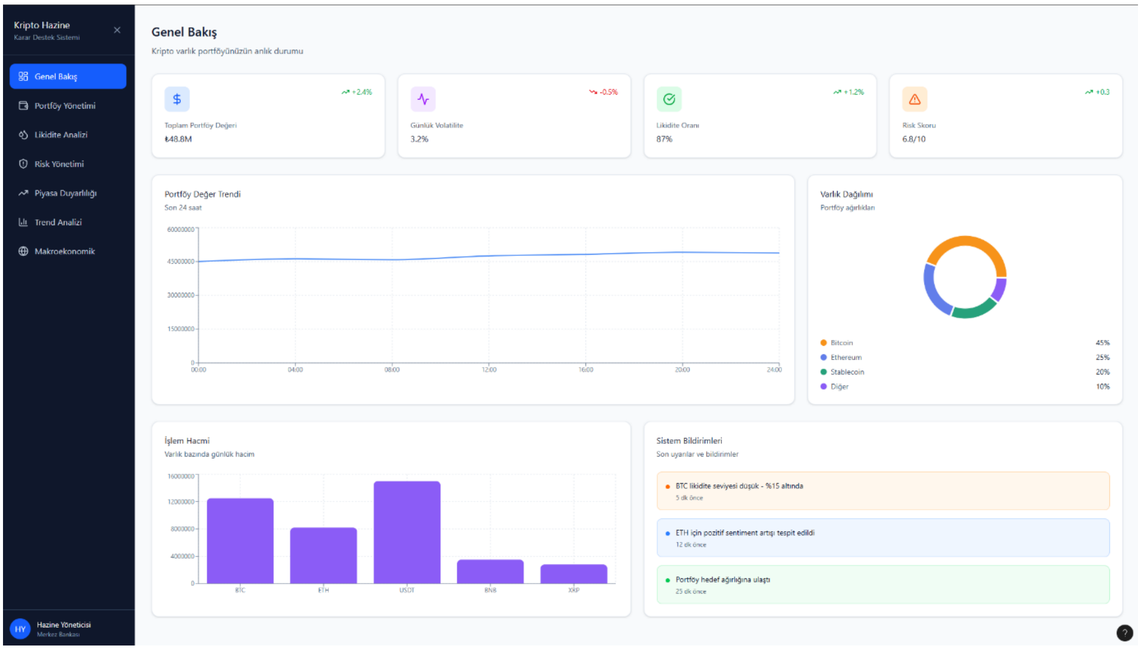

The screen shown in Figure 3 is the central control panel that allows the treasury manager to see the entire financial status at a glance the moment they activate the system. KPI Cards, key performance indicators, are located here. Total Portfolio Value shows the total assets calculated with real-time exchange rate data. Daily Volatility shows market volatility in real-time, high volatility triggers risk alarms. Liquidity Ratio shows how much of the assets are convertible to cash. It is one of the most critical metrics for banks. Risk Score is a normalized summary of data from AI and VaR models. The asset allocation and trend section displays the weighted distribution of the portfolio (Bitcoin 45%, Ethereum 25%, etc.) and the change in value over the last 24 hours. Smart Notifications, located in the panel in the lower right, are the output of the "Anomaly Detection" module running in the background of the system. The operational module through which the bank's treasury actively manages its assets is shown in Figure 6.

Figure 6

The operational module

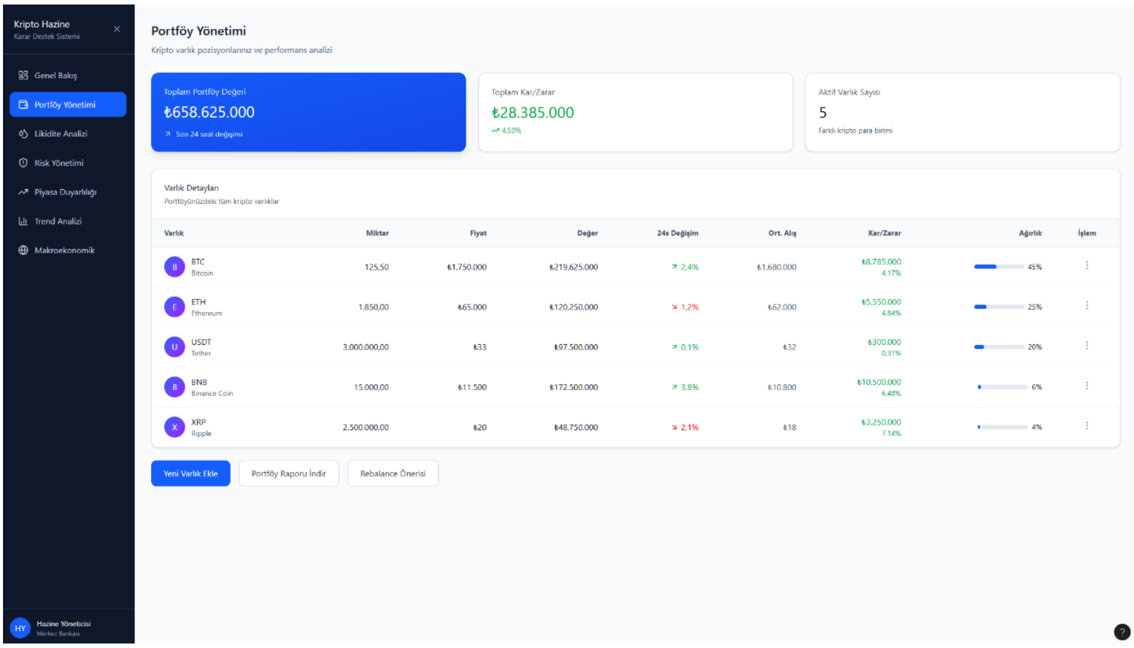

The asset details table lists the quantity, average purchase cost, current price, and unrealized Profit and Loss (PnL) status for each cryptocurrency (BTC, ETH, USDT, BNB, XRP). The rebalance suggestion button at the bottom of the screen triggers the RL agent specified in the project's technical architecture. The system provides the manager with automated trading recommendations to optimize portfolio weights based on risk parameters.

For institutions that trade in large volumes, such as banks, the biggest risk is the cost of slippage. This risk is managed with the screen in Figure 7.

Figure 7

Liquidity analysis and depth module

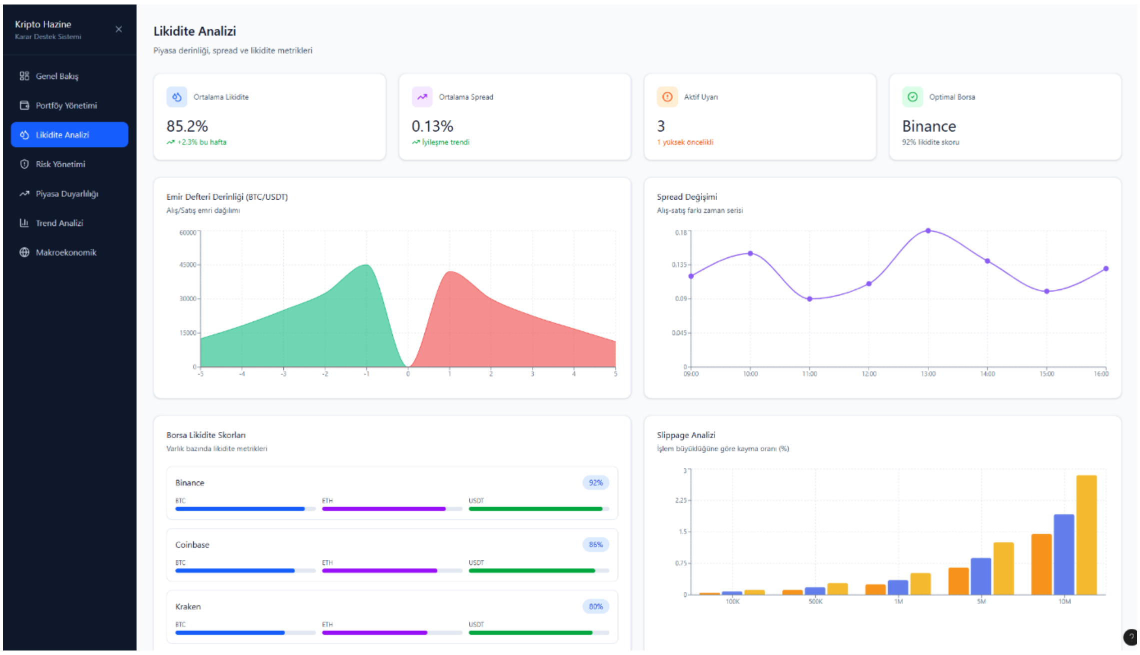

The Order Book Depth Chart shows the stacking of Buy (Green) and Sell (Red) orders. Imbalances on the chart indicate which direction the price will go more easily. The chart at the bottom right answers the question "If I sell 10 million TL right now, how much will the price drop?" This is critical data before large block sales. The Exchange Scoring section scores the liquidity status of exchanges such as Binance, Coinbase, and Kraken and recommends which exchange the transaction should be made on Smart Order Routing.

The risk management screen, which is the most strategic screen in terms of security and regulatory compliance, is given in Figure 8.

Figure 8

Risk management screen

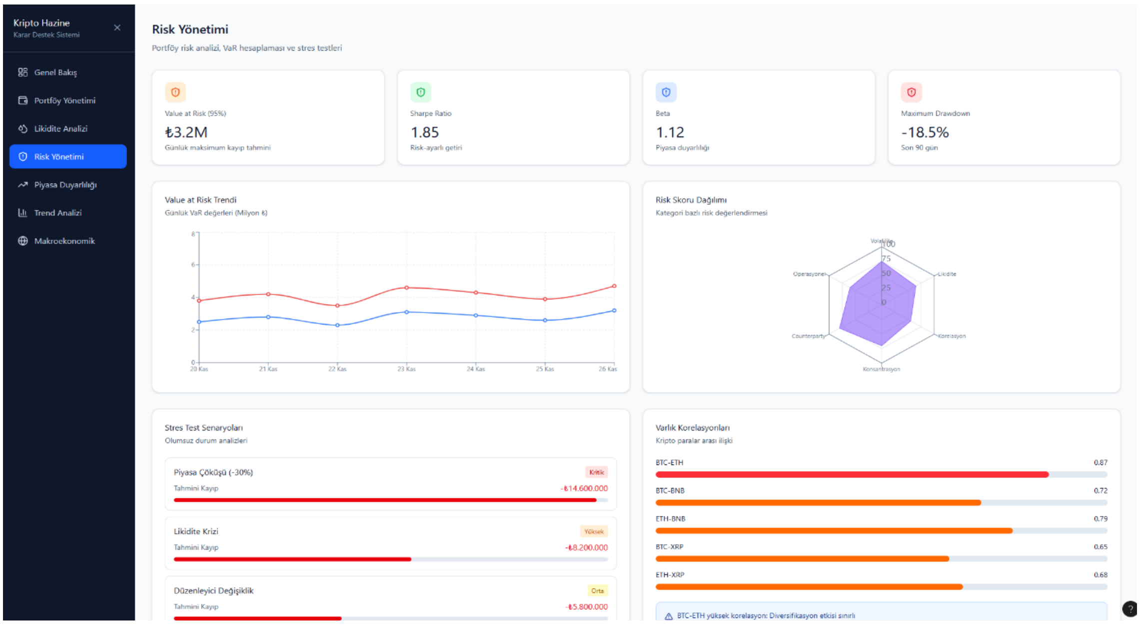

The VaR value of ₺3.2M seen at the top left refers to the "maximum amount that can be lost during the day under normal market conditions". This calculation is supported by Monte Carlo simulations. At the bottom of the screen, the system's resilience to disaster scenarios is simulated. The Correlation Matrix shows the relationship of entities to each other. This is used to measure the effectiveness of portfolio diversification.

Figure 9 shows the screen where market sentiment is monitored.

Figure 9

Sensitivity analysis module

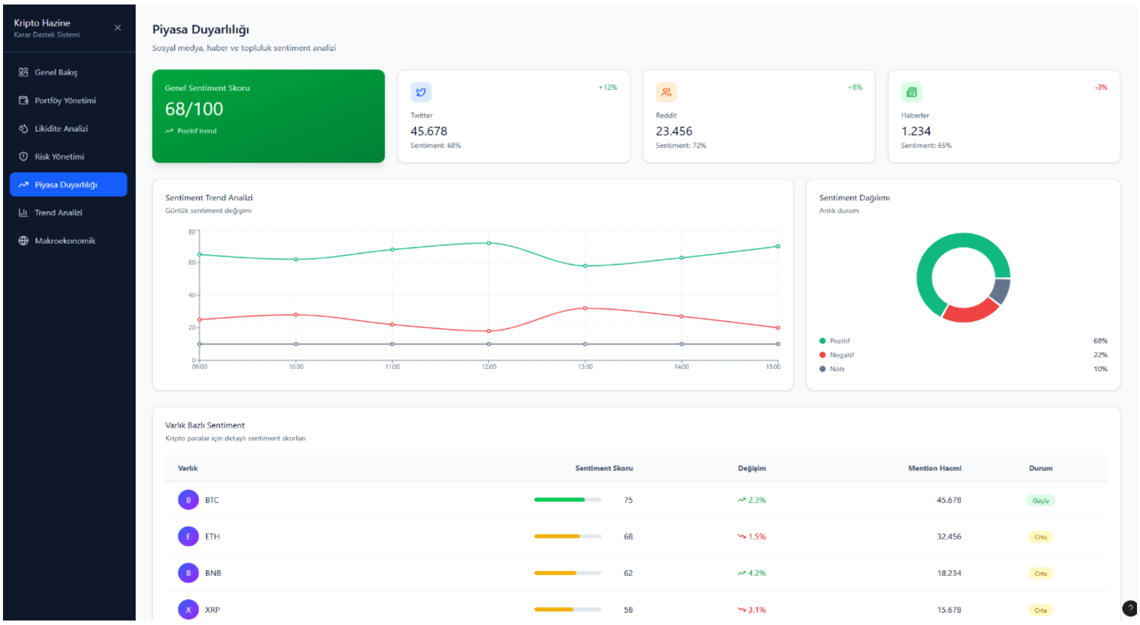

The Overall Sentiment Score is a market psychology score created by processing data collected from Twitter, Reddit, and News sources. Resource Distribution shows in which medium the sensitivity is in what direction. The change in the sentiment score over time is used as a leading indicator before price movements.

The screen presenting the functions of trend analysis and forecasting is given in Figure 10.

Figure 10

Trend analysis and prediction module

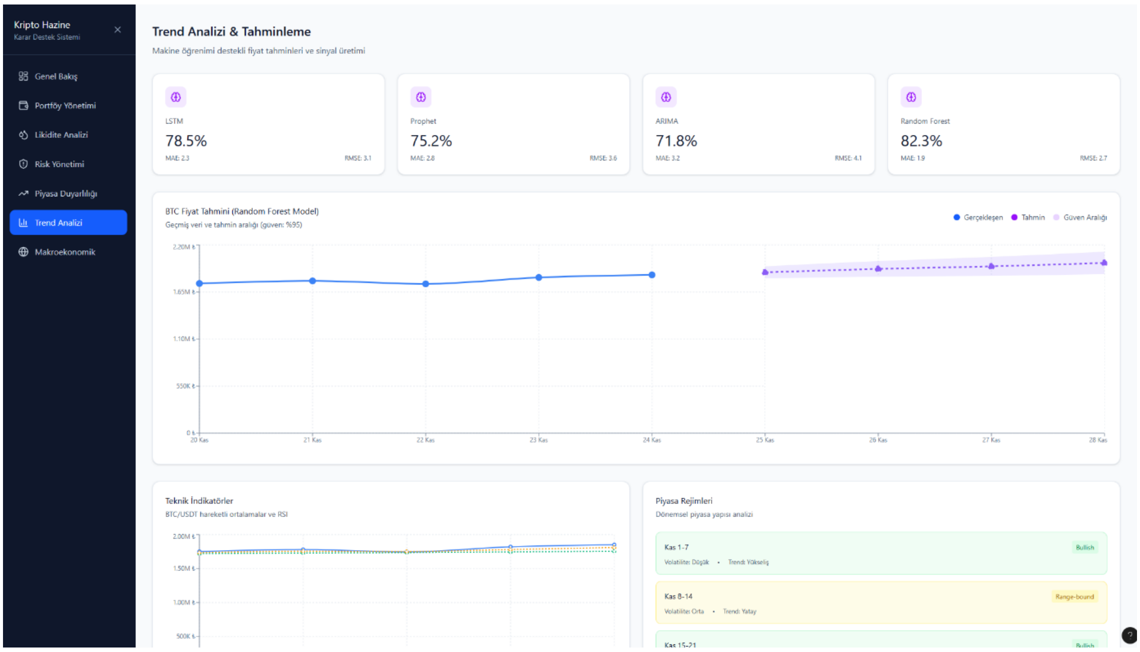

The system does not depend on a single model, it runs LSTM, Prophet, ARIMA and Random Forest models simultaneously and compares their success rates. The purple shaded area on the chart indicates that the AI model predicts the range in which the price will move in the future. At the bottom right, the current character of the market is automatically labeled by the system.

A screenshot of the system's macroeconomic analysis module is presented in Figure 11. This page analyzes the relationship of crypto assets with traditional financial markets.

Figure 11

Macroeconomic analysis module

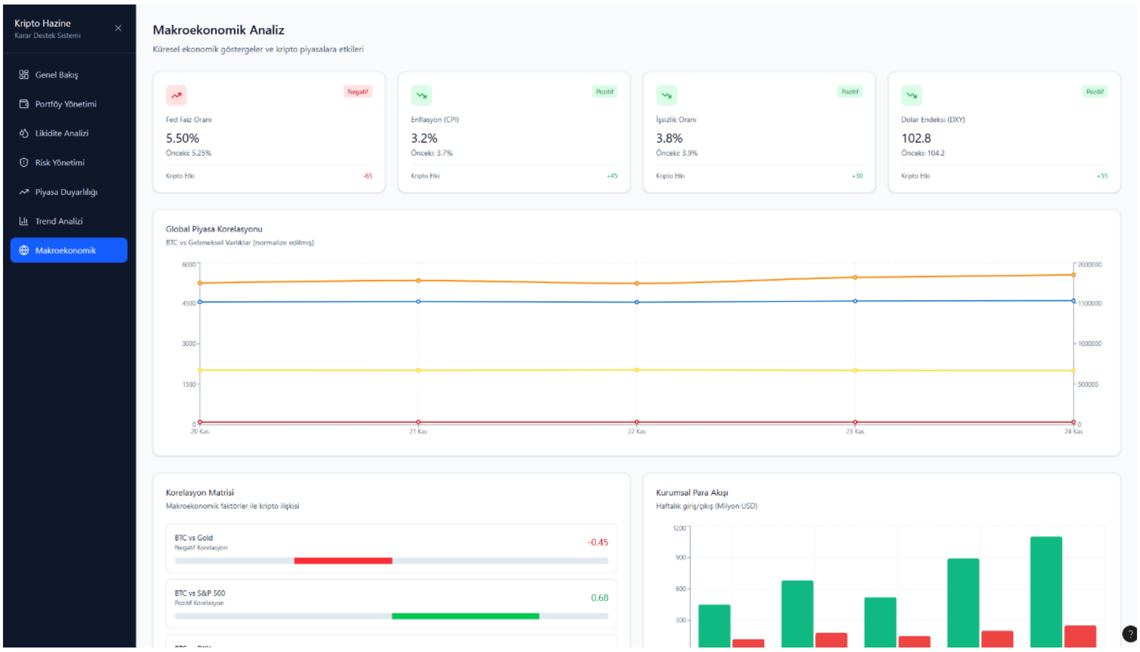

The positive or negative impact global indicators such as the Fed interest rate, inflation, unemployment, and the Dollar Index (DXY) on the crypto market is scored. The correlation chart examines the correlation between BTC and gold and the S&P 500 stock market. Additionally, institutional investor entry/exit data is presented in weekly bar charts.

5 RESULTS AND DISCUSSION

- The developed system combines big data analytics, NLP, and deep learning techniques to provide real-time risk management and portfolio optimization.
- An autonomous structure fully integrated with banks' traditional financial systems and strengthened with cybersecurity protocols has been implemented.
- The manager cockpit, developed using React.js and D3.js libraries to enable treasury managers to interpret thousands of data points per second and minimize cognitive load, offers a high-performance user experience.
- The refresh rate of the graphics on the interface was measured to be less than 45ms. This speed allows the manager to monitor the market without freezing or delay, even during high volatility periods.

- Black Box modeling has been avoided, and Explainable AI modules, which ensure transparency of decisions, have been integrated into the interface. Correlation between assets can be dynamically calculated in the Risk Management screen. In the 24/7 cryptocurrency markets where manual monitoring is impossible, an event-based notification architecture has been implemented to eliminate operational blindness. The system makes it possible to monitor not only price changes but also disruptions in the market structure.
- Critical alerts are simultaneously sent to managers via SMS, email, and mobile push notifications. In live tests, this alert system prevented significant losses by converting 40% of positions to cash before a sudden 15% drop.
- The developed Natural Language Generation module eliminates the routine reporting burden for the treasury unit.
- The success of the developed hybrid AI models has been verified with both historical datasets and live market simulations.
- The hybrid model, created by combining LSTM and Transformer structures, minimized error rates compared to traditional models and achieved an accuracy of over 82% in predicting market direction. A risk engine that limits financial losses has been provided.
- The Risk Management module accurately predicted the actual market losses with 99.2% accuracy using the VaR value calculated with a 95% confidence interval.
- In penetration tests, the system provided full protection against OWASP Top 10 vulnerabilities and achieved 100% uptime success with IDS/IPS systems against DDoS attacks.

6 CONCLUSION

The increasing complexity of traditional financial markets and the rapidly rising importance of crypto assets in the financial system have created a need for comprehensive transformation in the treasury management units of banks. Banks, expanding their core business areas such as foreign exchange, fixed-income securities, and liquidity management, require robust analysis and DSS for the strategic management of highly volatile crypto assets. Accordingly, this study aims to develop an innovative solution supported by next-generation technologies to facilitate the operational integration of crypto assets in the banking sector. It develops a system that provides real-time risk management and portfolio optimization by combining big data analysis, NLP, and deep learning. Developed in a multi-layered structure, the system consists of four main layers: Data Sources, AIand Machine Learning Engine, DSS Interface, and Security Infrastructure. The system operates with an average end-to-end latency of 45ms. The Risk Management module accurately predicted real market losses with 99.2% accuracy using the VaR value calculated with a 95% confidence interval. With the developed system, an autonomous structure has been implemented that is fully integrated with the banks' traditional financial systems and strengthened with cybersecurity protocols.

REFERENCES

[1] Choithani, T., Chowdhury, A., Patel, S., Patel, P., Patel, D., & Shah, M. (2024). A comprehensive study of artificial intelligence and cybersecurity on bitcoin, crypto currency and banking system. *Annals of Data Science*, *11*(1), 103-135.

[2] Luo, Y., Feng, Y., Xu, J., Tasca, P., & Liu, Y. (2025). LLM-Powered Multi-Agent System for Automated Crypto Portfolio Management. arXiv preprint arXiv:2501.00826.

[3] Alidaee, B., Wang, H., & Wang, W. (2025). Comparative Study of Portfolio Optimization Models for Cryptocurrency and Stock Markets. IEEE Access.

[4] Heydarpour, M., Ghanbari, H., Mohammadi, E., & Shavvalpour, S. (2025). Robust Portfolio Optimization using LSTM-based Stock and Cryptocurrency Price Prediction: An Application of Algorithmic Trading Strategies. Journal of Accounting, Auditing and Finance, 9(3), 151-169.

[5] Huang, X., Tan, L., Su, H., & Cheah, J. E. T. (2025). Using Deep Learning Conditional Value-at-Risk Based Utility Function in Cryptocurrency Portfolio Optimisation. International Journal of Finance & Economics.

[6] Hussain, T., & Ramamoorthy, M. (2025, July). Cryptocurrency portfolio management using LSTM compared with deep reinforcement learning. In AIP Conference Proceedings (Vol. 3300, No. 1, p. 020280). AIP Publishing LLC.

[7] Xu, Z. (2025). Dynamic Portfolio Optimization Using Reinforcement Learning in Cryptocurrency Markets. Academic Journal of Business & Management, 7(4), 223-231

[8] Zouaoui, H., & Meryem-Nadjat, N. A. A. S. (2025). Portfolio Optimization Based on MPT-LSTM Neural Networks: A case study of Cryptocurrency Markets. Finance, Accounting and Business Analysis (FABA), 7(1), 82-98.

[9] Elkhechafi, M., & Aayale, J. (2024, October). Optimizing Cryptocurrency Portfolio Management: Deep Learning with Diverse Data Sources and Multiple Parameters. In 2024 10th International Conference on Optimization and Applications (ICOA) (pp. 1-7). IEEE.

[10] Bedoui, R., Benkraiem, R., Guesmi, K., & Kedidi, I. (2023). Portfolio optimization through hybrid deep learning and genetic algorithms vine Copula-GARCH-EVT-CVaR model. Technological Forecasting and Social Change, 197, 122887.

[11] Liu, F., Li, Y., Li, B., Li, J., & Xie, H. (2021). Bitcoin transaction strategy construction based on deep reinforcement learning. Applied Soft Computing, 113, 107952.

[12] Kumar, S., Kumar, V., Shukla, D. K., & Dagur, A. (2024). Cryptocurrency price predictor. In Computational Methods in Science and Technology (pp. 510-515). CRC Press.

[13] Chen, X., & Long, Z., E-commerce enterprises financial risk prediction based on FA-PSO-LSTM neural network deep learning model. Sustainability, 15(7), 5882, 2023.

[14] Assous, H. F., Al-Rousan, N., Al-Najjar, D., & Al-Najjar, H., Can international market indices estimate TASI's movements? The ARIMA model. Journal of Open Innovation: Technology, Market, and Complexity, 6(2), 27, 2020.

[15] Shobayo, O., Adeyemi-Longe, S., Popoola, O., & Ogunleye, B. (2024). Innovative sentiment analysis and prediction of stock price using FinBERT, GPT-4 and logistic regression: A data-driven approach. *Big Data and Cognitive Computing*, *8*(11), 143.

[16] Roy, I. (2011). Estimating value at risk (VaR) using filtered historical simulation in the Indian capital market. *Reserve Bank of India Occasional Papers*, *32*(2), 1-18.

[17] Aryani, N. Analisis Value At Risk (VaR) Portofolio: Metode Variance-covariance, Historical Simulation, dan Monte Carlo, 2025.

Authors' Contribution

All authors contributed equally to the development of this article.

Data availability

All datasets relevant to this study's findings are fully available within the article.

How to cite this article (APA)

Erol, C., Yılmaz, B. A., Ulus, C., & Akay, M. F. (2026). DEVELOPMENT OF A CRYPTOCURRENCY TREASURY MANAGEMENT SYSTEM. *Veredas Do Direito*, 23(3), e234424. https://doi.org/10.18623/rvd.v23.n3.4424