

Article

Temporal Fusion Transformer-Based Trading Strategy for Multi-Crypto Assets Using On-Chain and Technical Indicators

Ming Che Lee 

Applied Artificial Intelligence Department, Ming Chuan University, Taoyuan 320, Taiwan;
leemc@mail.mcu.edu.tw

Abstract: Cryptocurrency markets are characterized by high volatility, nonlinear dependencies, and limited transparency, making short-term forecasting particularly challenging for both researchers and practitioners. To address these complexities, this study introduces a Temporal Fusion Transformer (TFT)-based forecasting framework that integrates on-chain and technical indicators to improve predictive performance and inform tactical trading decisions. By combining multi-source features—such as Spent Output Profit Ratio (SOPR), Total Value Locked (TVL), active addresses (AA), exchange net flow (ENF), Realized Cap HODL Waves, and the Crypto Fear and Greed Index—with classical signals like Relative Strength Index (RSI) and moving average convergence divergence (MACD), the model captures behavioral patterns, investor sentiment, and price dynamics in a unified structure. Five major cryptocurrencies—BTC, ETH, USDT, XRP, and BNB—serve as the empirical basis for evaluation. The proposed TFT model is benchmarked against LSTM, GRU, SVR, and XGBoost using standard regression metrics to assess forecasting accuracy. Beyond prediction, a signal-based trading strategy is developed by translating model outputs into daily buy, hold, or sell signals, with performance assessed through a comprehensive set of financial metrics. The results suggest that integrating attention-based deep learning with domain-informed indicators provides an effective and interpretable approach for multi-asset cryptocurrency forecasting and real-time portfolio strategy optimization.

Keywords: Temporal Fusion Transformer; multi-crypto-asset trading; on-chain indicators; technical analysis



Academic Editor: Vladimír Bureš

Received: 13 May 2025

Revised: 7 June 2025

Accepted: 13 June 2025

Published: 16 June 2025

Citation: Lee, M.C. Temporal Fusion Transformer-Based Trading Strategy for Multi-Crypto Assets Using On-Chain and Technical Indicators. *Systems* **2025**, *13*, 474. <https://doi.org/10.3390/systems13060474>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The cryptocurrency market has continued to evolve rapidly over the past decade, transitioning from a niche technological experiment to a mainstream component of the global financial landscape. As of early 2025, the total market capitalization of cryptocurrencies has stabilized above USD 2.5 trillion, with Bitcoin (BTC), Ethereum (ETH), and a growing number of decentralized finance (DeFi) and Layer-2 assets sustaining dominant positions [1]. This continued relevance is driven by the maturation of blockchain infrastructure, rising institutional participation, and the diversification potential that digital assets contribute to modern investment portfolios [2]. As crypto adoption deepens, there is increasing academic and industry interest in developing predictive tools that can effectively navigate such high-risk, high-reward environments.

Despite the enticing prospects, the cryptocurrency market is notorious for its extreme volatility and unpredictability. Factors such as a 24/7 trading cycle, regulatory developments, technological innovations, and the pervasive influence of social media contribute to abrupt and significant price fluctuations [3]. This inherent instability poses considerable challenges for traders and investors striving to devise effective strategies that can navigate

the tumultuous landscape of digital assets. Traditionally, market participants have relied on technical analysis (TA) to inform their trading decisions. TA involves the study of past market data, primarily price and volume, to forecast future price movements [4]. Commonly employed indicators include moving averages, RSI, and MACD [5]. While these tools have demonstrated utility in various financial markets, their efficacy in the cryptocurrency domain is often limited due to the market's unique characteristics and the exclusion of fundamental factors inherent to blockchain-based assets [6]. This gap highlights the need for more advanced forecasting frameworks that can incorporate both technical signals and blockchain-native features to improve predictive accuracy in dynamic conditions.

In response to the limitations of technical analysis, on-chain analysis has emerged as a pivotal approach in understanding cryptocurrency markets. On-chain metrics are derived directly from blockchain data and offer insights into network activity, investor behavior, and overall market sentiment [7]. Notable on-chain indicators include the Spent Output Profit Ratio (SOPR), which assesses the profitability of spent outputs and provides a gauge of market sentiment; Total Value Locked (TVL), reflecting the aggregate value of assets deposited in decentralized finance (DeFi) protocols and serving as a barometer for ecosystem growth; and the Crypto Fear and Greed Index, which amalgamates various data sources to quantify market sentiment on a scale from extreme fear to extreme greed [8–10]. Integrating these on-chain metrics with traditional technical indicators presents a more holistic view of the market, yet such integration remains underexplored in existing algorithmic trading strategies [11].

The rise in deep learning has further transformed financial forecasting, allowing models to capture complex nonlinearities, long-range dependencies, and temporal dynamics within financial data. Recurrent neural networks (RNNs) and long short-term memory (LSTM) models have been widely used for time-series prediction tasks across domains, including finance. However, they often face limitations in capturing long-term dependencies and lack scalability when dealing with heterogeneous inputs from multiple data sources [12]. The emergence of transformer architecture, initially introduced for natural language processing, addressed many of these challenges. Self-attention mechanisms allow for flexible modeling of relationships between time steps and variables without recurrence, providing a scalable and interpretable solution for time-series modeling [13]. Among transformer-based models, the Temporal Fusion Transformer (TFT) stands out for its effectiveness in multi-horizon forecasting and its ability to integrate static covariates, known and unknown time-varying features, and temporal attention mechanisms. Originally introduced by Lim et al. (2021) [14], TFT has been successfully applied in domains such as energy demand prediction, retail sales forecasting, and economic indicator tracking. It offers the dual advantages of strong predictive performance and model interpretability through attention layers and variable selection gates [15–17]. In the financial domain, recent studies have applied TFT to market prediction tasks, showing that the integration of economic indicators and sentiment analysis can enhance forecasting accuracy compared to traditional time-series models [18–20].

While the TFT model has been increasingly adopted in traditional financial markets, its application in the cryptocurrency sector remains relatively nascent. Recent studies have explored the use of TFT and other neural models for forecasting the price of digital assets [21]. For example, Ref. [22] proposed a windowed ensemble framework, WETT, to improve financial time series forecasting using temporal transformers. Their method combines multiple models, TFT and SeTT, trained on sliding windows and integrates them via averaging and meta-learning. Tested on 20 DJIA stocks across different market regimes, the ensemble significantly outperformed standalone TFT models, demonstrating that model diversity and sliding-window training enhance forecasting under financial volatility. Prior studies [23,24] have demonstrated the efficacy of deep learning in capturing volatile price

trends in financial and commodity markets, but few have incorporated crypto-native signals, such as on-chain data, into these architectures.

Moreover, most existing studies focus narrowly on single-asset price prediction, ignoring the interdependence among cryptocurrencies. However, recent research shows that digital assets are becoming increasingly interlinked, with correlations between BTC, ETH, and DeFi tokens affected by liquidity shifts, macroeconomic announcements, and user behavior across blockchains [25]. Effective modeling of these dynamics is critical for managing multi-asset portfolios and capturing spillover effects that influence returns and risk exposure. Ref. [26] proposed a multi-asset deep reinforcement learning framework for crypto portfolio management, but it lacked interpretability and did not incorporate blockchain-derived data.

At the same time, model explainability has become essential in high-stakes financial applications. Users and regulators require transparency in algorithmic decision-making. The attention mechanisms and gating structures within TFT allow analysts to discern which features, such as TVL surges or SOPR shifts, are influencing specific predictions, enabling better understanding and trust in model outputs. This level of explainability sets TFT apart from “black box” models like vanilla LSTMs and convolutional architectures, making it more suitable for real-world deployment in trading systems. Recognizing these gaps, this study proposes a comprehensive framework for developing a multi-cryptocurrency trading strategy using TFT, integrated with both technical and on-chain indicators. The goal is not merely to predict future prices but to translate rich time-series signals into actionable trading decisions. This study constructs a dataset incorporating multiple cryptocurrencies—such as BTC, ETH, BNB, and selected DeFi tokens—across several indicators, including RSI, MACD, SOPR, TVL, and the Crypto Fear and Greed Index. These features are used to train a TFT model that captures temporal patterns, asset interactions, and indicator relevance through interpretable attention scores.

This research contributes to the field in several ways. First, it bridges the gap between deep learning forecasting and trading execution by proposing a model-to-signal pipeline grounded in both theory and practice. Second, it demonstrates that combining technical and on-chain metrics improves signal quality and model robustness across volatile market conditions. We hypothesize that a TFT model, when trained on a multi-source feature set composed of technical indicators, on-chain behavioral data, and sentiment measures, can outperform traditional and deep learning baselines in both predictive accuracy and trading performance. Third, it extends the literature on multi-asset trading in cryptocurrency markets, a domain still dominated by single-token studies. Lastly, it validates the proposed strategy through rigorous backtesting, using financial performance metrics such as cumulative returns, Sharpe ratio, and maximum drawdown, thereby offering a practical and scalable solution for automated crypto trading systems.

The remainder of this paper is organized as follows. Section 2 presents a review of related work, including the literature on deep learning models for financial forecasting, the use of technical and on-chain indicators in crypto trading, and prior applications of the TFT. Section 3 outlines the research methodology, including data preprocessing, feature engineering, model training, and trading signal design. Section 4 discusses the experimental setup and empirical results, while Section 5 concludes with insights, limitations, and future directions for research.

2. Literature Review

Numerous studies have explored the use of deep learning models, particularly recurrent neural networks such as LSTM and GRU, for cryptocurrency price forecasting. These models are known for capturing complex temporal dependencies and nonlinear market behaviors. Jay et al. (2020) [27] introduced stochastic variants of MLP and LSTM models by injecting pertur-

bations into the activation layers to model financial uncertainty. Their approach demonstrated that learnable stochasticity parameters can outperform deterministic models across various cryptocurrencies. Expanding on deep learning's capabilities, Li et al. (2022) [28] proposed a hybrid VMD-LMH-BiLSTM model that combines variational mode decomposition and BiLSTM for short-term Bitcoin prediction. The model exhibited high directional accuracy and trading profitability, outperforming ARIMA, SVR, and traditional BiLSTM baselines.

Subsequent research further benchmarked LSTM, GRU, and transformer models in comparative frameworks. While LSTM and GRU remain prominent, attention-based architecture such as the TFT have also been introduced for their interpretability and feature selection capabilities. In a comprehensive benchmark, Murray et al. (2023) [29] included TFT alongside recurrent and convolutional models. Although theoretically superior in structure, TFT underperformed LSTM due to limited covariate usage and dataset constraints, suggesting that its potential is conditional on input richness and architectural tuning. Syed et al. (2024) [30] extended this direction by integrating bias correction mechanisms into LSTM, GRU, and BiLSTM structures. The results show that GRU-BC performs best for BTC and USDT, while BiLSTM-BC excels in XRP and BNB forecasting. The study demonstrates that bias correction significantly improves accuracy, with GRU-BC achieving MAE as low as 0.0006 and MAPE under 3% on real market data, especially when incorporating macroeconomic variables and precious metals as exogenous signals. Zhao et al. (2024) [31] developed a hybrid framework integrating LSTM time series prediction and deep reinforcement learning (DRL) for cryptocurrency market trend forecasting and risk management. The experiments utilize data from BTC, ETH, DOGE, and ADA covering the period from 2015 to 2021. The LSTM model captures long-term temporal dependencies, while DRL optimizes investment decisions using the forecasted trends. Similarly, Alnami et al. (2025) [32] employed a hybrid approach integrating ensemble learning with deep neural networks and Z-score anomaly detection, achieving excellent R^2 and early-warning capabilities across multiple cryptocurrencies. The experimental results show that Random Forest consistently performs best, while deep learning shows potential under nonlinear regimes. The study demonstrates the value of combining ensemble learning with statistical anomaly detection in volatile crypto markets.

In parallel, sentiment and emotion-based modeling have emerged as an important complement to price and technical data. Parekh et al. (2022) [33] developed DL-GuesS, which integrates LSTM and GRU models with sentiment signals from tweets related to multiple cryptocurrencies. The model is trained using normalized time-series data and structured tweet sentiment inputs. Experiments on Dash and Bitcoin Cash show that DL-GuesS consistently outperforms baseline models in MSE, MAE, and MAPE under 1-, 3-, and 7-day windows, confirming the effectiveness of cross-asset dependency modeling and sentiment fusion. Their multi-input design showed that price co-movement and social sentiment are both useful for improving prediction accuracy. Further extending sentiment analysis, Feizian and Amiri (2023) [34] incorporated author-level influence, showing that follower-weighted sentiment scores improved forecast accuracy when used with LSTM and ARIMA models. Shahiki Tash et al. (2024) [35] employed SenticNet to extract fine-grained emotions (e.g., joy, fear) and demonstrated strong price-emotion correlations, particularly for Fantom platform. Building on aspect-level sentiment understanding, Jahanbin and Chahooki (2025) [36] proposed a BiGRU-based attention model trained on tweets from crypto influencers. Their hybrid system, enhanced by tweet-level importance and transfer learning, achieved high F1-scores and trend prediction accuracy across eight major cryptocurrencies.

Traditional machine learning methods and technical indicators also continue to provide useful baselines and components for hybrid strategies. Sebastião and Godinho (2021) [37] used linear regression, ordinary SVM, and RF models trained on technical

and blockchain indicators under bearish regimes. Their ensemble strategies-maintained profitability and high Sharpe ratios even in adverse markets. Lin et al. (2023) [38] combined six classical indicators—such as KD, RSI, and OBV—with XGBoost, LSTM, and SVM models, showing that countertrend operations were more effective than momentum-based ones. Naganjaneyulu et al. (2023) [39] constructed a multi-indicator framework using EMA, PSAR, BB, and RSI, revealing that hybrid rule-based strategies can significantly outperform single-indicator systems in cumulative returns.

Several review papers have systematically organized these findings. Otabek and Choi (2024) [40] categorized existing forecasting models into statistical, ML, sentiment-based, and hybrid groups. They emphasized that accurate forecasting remains central to algorithmic trading success. Pečiulis et al. (2024) [41] conducted a bibliometric study on 168 high-quality papers, identifying LSTM, SVM, and GARCH as the dominant models, and pointing to emerging methods like TVM-aGAS and sentiment-enhanced volatility modeling. This review fills a gap in structured bibliometric synthesis and provides a detailed map of research evolution and emerging trends in cryptocurrency forecasting. Nguyen and Chan (2024) [42] expanded this view by mapping 622 studies and highlighting a trend toward automation, explainability, and multi-modal input integration. It highlights the use of machine learning, neural networks, and hybrid models in decision-making, along with popular inputs such as technical indicators, social sentiment, macroeconomic data, and engineered features. Notably, LSTM, GRU, SVM, and XGBoost emerged as the dominant models.

Complementing the above, deep reinforcement learning (DRL) and attention-based architecture have gained traction for optimizing trading policies. Jing and Kang (2024) [43] introduced a candlestick image-based DRL system using ResNet18 and PPO agents. Experiments on BTC/USDT from Binance show that the PPO model with candlestick images outperformed raw-data DRL models and baseline heuristics in both bullish and bearish markets. The study highlights the interpretability advantage of image-based trading systems and demonstrates superior cumulative returns (up to +26.13%) over buy-and-hold strategies. Peng et al. (2024) [44] designed a CNN-LSTM model with dual attention layers and a triple-class trend label, reducing trade frequency by 90% while improving excess return. Kochliaridis et al. (2024) [45] proposed UNSURE, a hybrid system incorporating TCN-based price bound forecasting, unsupervised clustering, and PPO/IMPALA agents. Using Binance hourly data (2017–2022), the system achieves high Sharpe/Sortino ratios with low drawdowns. UNSURE outperforms PPO, technical indicator baselines, and prior DRL models in both profit and stability. The framework proves effective in high-noise environments, offering robustness, explainability, and superior risk-adjusted returns across nine cryptocurrencies. Huang and Su (2024) [46] trained a Deep Q-Learning (DQL) trading strategy for multiple cryptocurrencies (BTC, ETH, VET, ADA, TRX, XRP) across diverse market trends (uptrend, horizontal, downtrend). The method showed that specialized go-long or go-short agents could yield significantly higher returns than buy-and-hold strategies.

In addition to modeling techniques, numerous studies have proposed hybrid or anomaly aware forecasting frameworks. García-Medina and Aguayo-Moreno (2024) [47] proposed hybrid LSTM-GARCH models for high-frequency volatility forecasting in cryptocurrency portfolios. The study compares GARCH family models, vanilla MLP, LSTM, and LSTM-GARCH hybrids using 5 min price and volume data for 10 cryptocurrencies during COVID-19 turbulence. The study concludes that simple models may offer better computational efficiency and comparable accuracy in volatile crypto environments. Feng et al. (2024) [48] proposed a daily dynamic tuning strategy for regularized regressors using 24 macroeconomic (e.g., EFFR, CPI), blockchain, and technical indicators. Their analysis highlighted hash rate and EFFR as key predictors, and their model outperformed HAR-RV across metrics and economic value simulations. Pellicani et al. (2025) [49] developed CAR-

ROT, a multi-target LSTM anomaly detection system that clusters assets using DTW and labels anomalies through local extrema detection. The model labels anomalies via curve-shifting and integrates 61 features, including 48 technical indicators and the Crypto Fear and Greed Index. Evaluated via 18-fold time series cross-validation (2020–2021), CARROT outperformed single-target LSTM, CNN, and MLP models across all settings, achieving up to 20% improvement in macro F1-score and greater robustness under volatility.

Several works also focus on integrating blockchain-native indicators into forecasting and trading pipelines. King et al. (2024) [50] constructed “blockchain ribbons” by applying smoothing and statistical adjustments to 21 raw blockchain metrics, such as user count, transaction cost, and hash rate, these indicators extend the Hash Ribbon concept using moving averages and statistical refinements. These refined indicators significantly outperformed the classic Hash Ribbon in both direct price prediction and strategy profitability and demonstrated the statistical advantage of incorporating blockchain metrics into trading systems under the Adaptive Market Hypothesis. Beyond signals and features, some researchers have examined uncertainty as a direct forecasting input. Ah Mand (2025) [51] proposed the Cryptocurrency Uncertainty Index (UCRY) and employed wavelet coherence analysis to evaluate its lead-lag relationship with crypto returns from 2013 to 2023. Using Wavelet Coherence (WTC), Cross-Wavelet Transform (XWT), and Partial Wavelet Coherence (PWC), the study confirms that UCRY consistently leads cryptocurrency returns across all frequencies (short-, medium-, and long-term). UCRY was found to consistently lead price changes across multiple timeframes and asset types, outperforming traditional indices such as EPU and VIX. These findings support UCRY’s use in market timing and portfolio diversification.

Table 1 provides a comparative summary of selected studies in cryptocurrency forecasting, focusing on their model types, input feature sets, and target assets. The table is curated to illustrate key differences and similarities across methodologies, highlighting the role of technical indicators, sentiment signals, and on-chain data. This literature synthesis supports the motivation for integrating diverse features within the proposed TFT-based framework.

Conceptually, compared with traditional recurrent models like LSTM and GRU, which are effective in capturing temporal dependencies and nonlinear dynamics, transformer-based models such as the TFT offer improved interpretability through attention mechanisms and enable the incorporation of known future inputs—making them particularly suitable for multi-source financial time series forecasting. However, these models often require extensive hyperparameter tuning and are sensitive to feature quality. In contrast, classical machine learning models, such as Support Vector Regression (SVR) and XGBoost, are efficient to train and perform well on small, structured datasets, but their lack of sequential modeling capabilities and temporal awareness limits their effectiveness in dynamic financial environments. These architectural trade-offs form the basis for benchmarking TFT against both deep learning and traditional baselines in this study.

Furthermore, from a design perspective, although prior studies have explored a variety of modeling techniques and data sources, many remain limited in scope, focusing on single-asset prediction or overlooking the synergistic integration of technical and on-chain indicators. Additionally, the application of attention-based models like TFT is still relatively underexplored in multi-asset cryptocurrency trading systems. To address these research gaps, this study proposes a TFT-based trading framework that integrates domain-specific technical signals (e.g., RSI, MACD) with blockchain-derived metrics (e.g., SOPR, TVL, AA, ENF, HODL Waves, and the Fear and Greed Index). This integrative approach aims to improve short-term forecasting accuracy while enhancing the robustness and interpretability of multi-asset trading strategies.

Table 1. Comparative summary of selected studies on cryptocurrency forecasting models, features, and applications.

Study	Model	Features	Assets	Notes
Jay et al. (2020) [27]	Stochastic LSTM	RSI, MACD	BTC, ETH	Models uncertainty with stochastic units
Li et al. (2022) [28]	VMD-BiLSTM	Daily trading data	BTC	Decomposition + BiLSTM boosts short term accuracy
Murray et al. (2023) [29]	TFT vs. RNN/CNN	Daily trading data	BTC	TFT underperforms with limited covariates
Syed et al. (2024) [30]	GRU-BC, BiLSTM-BC	Technical + Macro	BTC, USDT, XRP, BNB	Bias correction improves MAPE and MAE
Zhao et al. (2024) [31]	LSTM + DRL	Daily trading data	BTC, ETH, DOGE	Combines LSTM with RL for trading actions
Parekh et al. (2022) [33]	DL-GuesS	Sentiment	Dash, BCH	Sentiment fusion boosts prediction accuracy
Jahanbin and Chahooki (2025) [36]	BiGRU + Attention	Sentiment	BTC, ETH, LTC	High F1 and trend accuracy across platforms
Kochliaridis et al. (2024) [45]	UNSURE	Technical + Market States	9 Tokens	TCN + DRL, strong risk-adjusted returns
This Study	TFT	Technical + On-chain + Sentiment	BTC, ETH, USDT, XRP, BNB	First to unify multi-source features with TFT for trading decisions under multi-asset training

3. Research Method

This section details the design of the forecasting and trading framework, including data sources, feature structure, model architecture, and signal transformation. The entire pipeline is structured to align with the architectural strengths of the TFT, which serves as the core forecasting model. Unlike traditional recurrent models, TFT is designed to process both observed and known time-varying inputs while capturing short- and long-term dependencies through gating and attention mechanisms.

The architecture of the full system is illustrated in Figure 1. These input features are organized into observed inputs, known future covariates, and static covariates, and then passed to the TFT model. The TFT architecture includes a variable selection network for dynamic feature weighting, gated residual networks (GRN) to capture nonlinear relationships, and a temporal attention layer to focus on key time points during sequence modeling. To fully leverage these capabilities, this study constructs a multi-asset time series dataset wherein historical OHLCV (open, high, low, close, and volume) data is treated as observed inputs. A set of precomputable features are incorporated as known covariates, including technical indicators such as RSI and MACD, as well as a comprehensive set of blockchain-based and sentiment-aware indicators: Spent Output Profit Ratio (SOPR), Total Value Locked (TVL), active addresses (AA), exchange net flow (ENF), Realized Cap HODL Waves, and the Crypto Fear and Greed Index. All features are processed through Z-score, logarithmic, or logit normalization as appropriate. The model outputs the next-day predicted closing prices for each asset. These predictions are passed to a signal-based trading strategy module that calculates standardized returns and applies a z-score threshold to generate buy, hold, or sell signals. A daily backtesting engine simulates trading execution and portfolio rebalancing under a fixed transaction cost. Forecasting accuracy is evaluated using RMSE, MAE, MAPE, and R^2 , while trading performance is assessed through cumulative return, Sharpe ratio, maximum drawdown, and hit ratio.

3.1. Technical and On-Chain Indicators

To enable robust and interpretable forecasting across multiple cryptocurrency assets using TFT, this study constructs a structured dataset composed of market-derived technical

indicators, blockchain-level behavioral metrics, and high-level sentiment signals. The selected assets include the top five cryptocurrencies by market capitalization—Bitcoin (BTC), Ethereum (ETH), Tether (USDT), Ripple (XRP), and Binance Coin (BNB)—representing a diverse spectrum of coin types and market functions.




Figure 1. Workflow of the TFT-Based Forecasting and Trading Framework.

In alignment with the architecture of TFT, the input features in this study are systematically categorized into three distinct types. First, observed inputs encompass time-dependent variables whose values are only available up to the forecast creation point, such as closing prices, trading volumes, and price volatility measures. Second, known future inputs refer to features that are fully known in advance of the prediction horizon, including technical indicators and on-chain behavioral metrics. Third, static covariates represent asset-specific characteristics that remain constant over time, including token classification, market capitalization rank, and historical volatility percentiles. The subsequent subsections provide a comprehensive overview of the construction, and normalization of each category of input variables, ensuring their compatibility with the TFT model’s temporal encoding structure.

3.1.1. Technical Indicators

Technical indicators are calculated from OHLCV data to capture price momentum, volatility, and trend signals. For each asset, the Relative Strength Index (RSI) is computed as

$$RSI_t = 100 - \left(\frac{100}{1 + RS_t} \right), \quad RS_t = \frac{U_t}{D_t}, \quad (1)$$

where t denotes the time index in the daily time series, corresponding to each trading day in the forecasting dataset, U_t and D_t are the average upward and downward price changes over a 14-day window. The moving average convergence divergence (MACD) is expressed as the difference between two exponential moving averages:

$$\begin{aligned} MACD_t &= EMA_t^{(12)} - EMA_t^{(26)}, \\ EMA_t^{(k)} &= \alpha_k x_t + (1 - \alpha_k) EMA_{t-1}^{(k)}, \quad \alpha_k = \frac{2}{k+1}, \end{aligned} \quad (2)$$

where α_k is the smoothing factor for the EMA with period k . In addition to RSI and MACD, the system computes short- and medium-term EMAs (e.g., 7-day and 21-day), Bollinger Bands, and Average True Range (ATR) for volatility modeling. All technical indicators are treated as known time-varying covariates, as they can be computed prior to the forecast period. To ensure model stability, each indicator is normalized via the Z-score transformation:

$$z_t = \frac{x_t - \mu}{\sigma}, \quad (3)$$

where x_t is the raw indicator value at time t , and μ and σ are its mean and standard deviation over the training window.

The selected technical indicators—RSI and MACD—are particularly well-suited to short-term forecasting tasks. RSI captures recent momentum by quantifying overbought or oversold market conditions, while MACD reveals trend reversals through moving average crossovers. Both are responsive to price volatility over short horizons and are

widely adopted in cryptocurrency trading. Given that this study focuses on next-day price prediction and signal-based strategy execution, such indicators offer timely and interpretable cues aligned with the forecast horizon.

3.1.2. On-Chain Behavioral Indicators

To enhance the model's understanding of market behavior beyond price-based technical indicators, this study incorporates five carefully selected on-chain indicators. These metrics capture aspects of investor psychology, capital flow, and network-level participation. Each indicator is treated as a known time-varying covariate in the model and normalized before input. The following subsections detail the definition, intuition, and normalization of each feature.

Spent Output Profit Ratio (SOPR)

The Spent Output Profit Ratio (SOPR) measures the average profit or loss realized by investors when spending their coins. It is defined as the ratio between the value at which outputs are spent and the price at which those coins were initially acquired:

$$\text{SOPR}_t = \frac{\sum_{i=1}^{N_t} P_{i,t}^{\text{out}}}{\sum_{i=1}^{N_t} P_{i,t}^{\text{in}}}. \quad (4)$$

Here, N_t is the number of outputs spent at time t , $P_{i,t}^{\text{out}}$ is the output value when spent, and $P_{i,t}^{\text{in}}$ is the coin's original cost basis. A SOPR value greater than 1 suggests that most coins are being sold at a profit, typically reflecting bullish sentiment. Values below 1 imply loss realization, often associated with capitulation. Given that SOPR is strictly positive and exhibits non-normal distribution with volatility spikes, this study applies a logarithmic transformation followed by Z-score normalization:

$$x_t^{\log} = \log(\text{SOPR}_t + \epsilon), \quad z_t = \frac{x_t^{\log} - \mu_{\text{SOPR}}}{\sigma_{\text{SOPR}}}, \quad (5)$$

where $\epsilon = 10^{-6}$, μ_{SOPR} and σ_{SOPR} are computed from the training data. The constant ϵ is added to avoid undefined values when SOPR_t approaches zero, ensuring numerical stability in the logarithmic transformation. The choice of ϵ is standard in time series preprocessing and has minimal impact on the scale or behavior of the transformed feature. The standardized feature z_t captures directional shifts in aggregate investor profitability.

Total Value Locked (TVL)

Total Value Locked (TVL) represents the aggregate amount of capital deployed in DeFi protocols across a given network. It reflects the level of user engagement, trust in decentralized systems, and risk-on positioning. Higher TVL generally indicates bullish market sentiment and increased protocol usage. Let TVL_t denote the total USD value locked in DeFi contracts at time t , calculated as

$$\text{TVL}_t = \sum_{j=1}^{M_t} V_{j,t}, \quad (6)$$

where M_t is the number of DeFi contracts and $V_{j,t}$ is the locked value in contract j . To correct for scale disparities and skewed distribution, a log transformation is applied:

$$x_t^{\log} = \log(\text{TVL}_t + \epsilon), \quad z_t = \frac{x_t^{\log} - \mu_{\text{TVL}}}{\sigma_{\text{TVL}}}. \quad (7)$$

The transformed and standardized series allows the model to detect shifts in macro-level capital inflows and outflows.

Active Addresses (AA)

The number of Active Addresses (AA) quantifies user engagement by counting the unique wallet addresses involved in transactions during a given day. Unlike price or

volume, AA captures organic blockchain activity, serving as a proxy for network usage and sentiment among retail participants. Let AA_t represent the count of active addresses at time t . Because this variable can spike dramatically and has a highly skewed distribution, this study normalizes it using

$$x_t^{\log} = \log(AA_t + \epsilon), z_t = \frac{x_t^{\log} - \mu_{AA}}{\sigma_{AA}}. \quad (8)$$

This processed indicator reflects shifts in retail involvement and transaction intensity across time.

Exchange Net Flow (ENF)

Exchange Net Flow (ENF) captures the net volume of coins transferred to or from centralized exchanges (CEX) in each day. It is defined as

$$ENF_t = \text{Inflow}_t - \text{Outflow}_t, \quad (9)$$

where Inflow_t and Outflow_t denote total token volumes sent into and out of exchange wallets, respectively. Positive ENF indicates potential selling pressure, while negative ENF suggests accumulation. Because ENF is a signed variable with wide dispersion, this study applies a signed log transformation:

$$x_t^{\log} = \text{sign}(ENF_t) \cdot \log(|ENF_t| + \epsilon), z_t = \frac{x_t^{\log} - \mu_{ENF}}{\sigma_{ENF}}. \quad (10)$$

This transformation preserves directionality while reducing variance, allowing the model to interpret sudden liquidity movements more effectively.

Realized Cap HODL Waves

Realized Cap HODL Waves represent the distribution of held coins by age, expressed as a fraction of the total realized market capitalization. They are used to infer holding behavior, risk tolerance, and long-term conviction. A growing share of older HODL bands indicates strong hands and lower liquidity, while the dominance of younger coins may signal speculative activity. Let $R_t^{(a)}$ be the value held by coins aged within band a , and RC_t be the total realized cap:

$$\text{Wave}_t^{(a)} = \frac{R_t^{(a)}}{RC_t}. \quad (11)$$

This study focuses on coins held for 1 to 12 months, which often signal mid-term investor behavior. Since wave values are bounded in $(0,1)$, this study applies the logit transformation:

$$x_t^{\text{logit}} = \log\left(\frac{\text{Wave}_t^{(a)} + \epsilon}{1 - \text{Wave}_t^{(a)} + \epsilon}\right), z_t = \frac{x_t^{\text{logit}} - \mu_{\text{HODL}}}{\sigma_{\text{HODL}}}. \quad (12)$$

The final standardized variable is sensitive to structural changes in market holding patterns and serves as a long-term sentiment indicator.

The on-chain indicators included in this study are chosen for their ability to reflect behavioral, liquidity, and network-level dynamics in the crypto ecosystem. SOPR quantifies investor profitability and sentiment; TVL indicates capital inflow into DeFi protocols; AA captures real-time user engagement; ENF measures fund movement to and from exchanges, often signaling accumulation or selling pressure; and Realized Cap HODL Waves provide insights into coin holding patterns and long-term conviction. These indicators complement market-derived signals by embedding behavioral context into the model, making them especially valuable for enhancing short-term predictions under high volatility.

3.1.3. Sentiment Indicator

While on-chain indicators reflect structural behavior and wallet-level dynamics, they do not fully capture the real-time psychological state of the market. To bridge this gap, this study incorporates a macro sentiment metric designed to quantify fear, greed, and investor mood in a unified and interpretable manner. This metric complements the behavioral signals discussed in Section 3.1.2 by explicitly modeling the emotional dimension of market activity.

The Crypto Fear and Greed Index is a daily sentiment score ranging from 0 (extreme fear) to 100 (extreme greed). It is published by Alternative.me and incorporates multiple components, including price momentum, volatility, social media activity, Google search volume, and market dominance. As a composite index, it is designed to represent short-term crowd sentiment, making it especially useful for capturing rapid shifts in market mood that may not be reflected in transactional or price data alone. Let $FG_t \in [0, 100]$ denote the raw index value at time t , to prepare it for input into the TFT model, this study first normalizes it to a $[0, 1]$ scale:

$$x_t = \frac{FG_t}{100}. \quad (13)$$

Given that this variable is bounded and may exhibit saturation effects near 0 or 1, this study applies a logit transformation to map it to the real number line:

$$x_t^{\text{logit}} = \log\left(\frac{x_t + \epsilon}{1 - x_t + \epsilon}\right). \quad (14)$$

This study then standardizes the transformed values using Z-score normalization:

$$z_t = \frac{x_t^{\text{logit}} - \mu_{FG}}{\sigma_{FG}}, \quad (15)$$

where μ_{FG} and σ_{FG} are computed over the training window. The resulting standardized signal z_t is treated as a known time-varying covariate and supplied to the model alongside technical and on-chain features. By including the Crypto Fear and Greed Index, the model is equipped to learn nonlinear interactions between crowd psychology and market behavior, particularly during times of extreme volatility, irrational exuberance, or panic selling. This emotional context strengthens the model's ability to generalize across diverse regimes and complements the rational structure of price and blockchain-derived indicators.

3.2. Algorithm Design

This section presents the algorithmic framework that underpins the forecasting and trading components of the proposed system. The architecture is composed of two integrated modules: a deep learning model designed to perform multi-asset price prediction, and a rule-based algorithm that translates model outputs into executable trading signals.

The core forecasting engine is based on the Temporal Fusion Transformer (TFT), a sequence-to-sequence model optimized for multi-horizon prediction with heterogeneous time-series inputs. In this study, the TFT model is configured to perform a univariate regression task, predicting the future closing price of each asset on a daily basis. This design allows the model to capture both short-term and long-term dependencies using gated residual connections and interpretable attention mechanisms.

To convert model forecasts into actionable investment decisions, a trading strategy module is implemented. This module interprets predicted price changes and generates trading signals under a rule-based framework. The strategy logic supports multi-day holding, dynamic entry and exit, and optional constraints on position size and turnover. The two modules are integrated sequentially, with model inference serving as the basis for signal generation and strategy evaluation.

The following subsections detail the architecture, training procedure, and output structure of the TFT model, followed by the signal-based trading algorithm used to backtest and evaluate the practical performance of the forecasting system.

3.2.1. TFT-Based Price Forecasting Model

To generate accurate short-term price forecasts for multiple cryptocurrencies, this study adopts the TFT as our primary regression model. TFT is designed for multivariate time-series forecasting with both observed and known future inputs. In our setup, historical OHLCV data is treated as observed variables, while technical and on-chain indicators are

incorporated as known covariates. The model predicts the next-day closing price for each asset, leveraging attention mechanisms and gated residual networks to capture both short- and long-term temporal dependencies. The training samples are constructed using a sliding window approach with a fixed encoder length and single-step decoder target. The detailed forecasting procedure is outlined in Algorithm 1 below.

Algorithm 1: TFT-Based Multi-Cryptocurrency Price

Input:

X_t : Historical price features [$Open_t, High_t, Low_t, Close_t, Volume_t$]

I_t : Technical indicators and on-chain indicators [RSI, MACD, SOPR, TVL, AA, ENF, HODL, Fear and Greed]

T : Encoder window length

$\tau = 1$: Forecast horizon (next-day prediction)

η : Learning rate

ϵ : Smoothing constant

$\mathcal{A}$: Asset set $\{a_1, a_2, \dots, a_N\}$

Output:

$\hat{y}_{t+1}^{(n)} \in \mathbb{R}$: Predicted next-day closing price for asset a_n

1. Data Processing

1.1 Normalize X_t and I_t using Z-score and appropriate transformations

1.2 Construct joint feature set:

$$z_t = [X_t || I_t] \in \mathbb{R}^d$$

1.3 Apply sliding window to create sequential training samples:

$$\{z_{t-T+1:t}^{(n)}\}, \forall a_n \in \mathcal{A}$$

2. Model Construction

2.1 Encoder Layer

Process observed and known inputs using Variable Selection Network:

$$\tilde{Z}_t = \text{GRN}(z_t, c_t) \odot \text{Softmax}(W_v z_t)$$

2.2 Temporal Encoding

Apply gated residual network (GRN) with LayerNorm:

$$E_t = \text{GRN}(\tilde{Z}_t)$$

2.3 Attention-Based Decoder

Compute multi-head attention over past encoder outputs:

$$\alpha_t = \text{Softmax}(QK^T / \sqrt{d_k})$$

$$\text{AttentionOutput}_t = \alpha_t V$$

2.4 Prediction Layer

Feed attention output into final GRN and linear projection:

$$\hat{y}_{t+1}^{(n)} = W_o \cdot \text{GRN}(\text{AttentionOutput}_t) + b_o$$

3. Loss and Optimization

3.1 Compute mean squared error over all assets:

$$\mathcal{L}_{\text{total}} = \frac{1}{N} \sum_{n=1}^N \left(y_{t+1}^{(n)} - \hat{y}_{t+1}^{(n)} \right)^2$$

3.2 Update parameters:

$$\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}_{\text{total}}$$

The proposed model predicts the next-day closing price $\hat{y}_{t+1}^{(n)}$ for each asset a_n , by processing multivariate sequential inputs constructed from concatenated historical price features X_t and technical/on-chain indicators I_t . These combined feature vectors $z_t \in \mathbb{R}^d$ are normalized and organized into fixed-length sequences of length T using a sliding window approach. Within the TFT architecture, variable selection networks dynamically learn feature relevance, while gated residual networks (GRNs) capture nonlinear temporal dynamics. Multi-head attention mechanisms enable the model to selectively focus on

informative time steps during decoding, improving short-term prediction accuracy. The model is trained to minimize the Mean Squared Error (MSE) between predicted and actual closing prices across all assets, and parameters are optimized via gradient descent using the Adam optimizer. This design allows TFT to effectively model heterogeneous input sources and temporal dependencies in cryptocurrency price forecasting.

3.2.2. Signal-Based Trading Strategy Algorithm

To operate the forecasts generated by TFT, this study designs a statistically grounded trading strategy that translates model outputs into actionable decisions. The strategy is built upon standardized return signals, which reflect not only the directional magnitude of expected price changes but also their statistical relevance within the context of each asset's recent volatility. Given the one-day-ahead price forecast $\hat{y}_{t+1}^{(n)}$ for asset a_n , and the current closing price $y_t^{(n)}$, this study defines the predicted return as

$$r_t^{(n)} = \frac{\hat{y}_{t+1}^{(n)} - y_t^{(n)}}{y_t^{(n)}}. \quad (16)$$

To normalize this return and account for asset-specific behavior, this study computes its z-score based on a rolling historical distribution:

$$z_t^{(n)} = \frac{r_t^{(n)} - \mu_r^{(n)}}{\sigma_r^{(n)}}, \quad (17)$$

where $\mu_r^{(n)}$ and $\sigma_r^{(n)}$ represent the mean and standard deviation of predicted returns over a fixed-length lookback window. Trading actions are taken only when the standardized return deviates meaningfully from the mean, ensuring that signals are both significant and volatility adjusted. The full decision procedure is summarized in Algorithm 2, which outlines the transformation from model output to trading signal.

Algorithm 2: Signal-Based Trading Strategy

Input:

- $\hat{y}_{t+1}^{(n)}$: Forecasted closing price for asset a_n
- $y_t^{(n)}$: Observed closing price at time t
- $\mu_r^{(n)}, \sigma_r^{(n)}$: Rolling mean and standard deviation of forecasted returns
- θ : Statistical threshold
- $\mathcal{P}_{t-1}^{(n)}$: Previous position (Buy = +1, Hold = 0, Sell = -1)

Output:

- $s_t^{(n)} \in \{-1, 0, +1\}$: Current trading signal (Sell, Hold, Buy)

1. Compute forecasted return:

$$r_t^{(n)} \leftarrow (\hat{y}_{t+1}^{(n)} - y_t^{(n)}) / y_t^{(n)}$$

2. Standardize return using z-score:

$$z_t^{(n)} \leftarrow (r_t^{(n)} - \mu_r^{(n)}) / \sigma_r^{(n)}$$

3. Generate trading signal:

- If $z_t^{(n)} > \theta$, then $s_t^{(n)} \leftarrow +1$
- Else if $z_t^{(n)} < -\theta$, then $s_t^{(n)} \leftarrow -1$
- Else $s_t^{(n)} \leftarrow 0$

4. Signal stability check:

- If $s_t^{(n)} = \mathcal{P}_{t-1}^{(n)}$, maintain current position
- Else, update position to new signal $s_t^{(n)}$

5. Execute trade at market close of day t

- Open or adjust position according to $s_t^{(n)}$
-

This rule-based strategy is designed to be interpretable and adaptable to varying market conditions. By standardizing predicted returns using z-scores, the algorithm filters out insignificant signals and focuses on high-confidence directional movements. The threshold θ governs sensitivity and can be tuned via backtesting. The signal $s_t^{(n)}$ maps directly to trade actions +1 for initiating or holding a long position, -1 for a short position, and 0 for remaining in cash or holding the current state. The inclusion of a signal stability check (Step 4) helps reduce transaction costs and spurious reversals by avoiding frequent position flipping in response to marginal signal shifts. This creates a more realistic trading simulation when applied in backtesting.

3.3. Backtesting and Evaluation Metrics

To assess the practical viability of the TFT-based prediction and trading system, this study conducts historical backtesting using a rolling evaluation framework. In this setup, the model is trained over a fixed historical window and evaluated on the immediately following day. The process repeats with the window sliding forward chronologically. For each trading day t , the predicted closing price $\hat{y}_{t+1}^{(n)}$ for asset a_n is compared with the actual price $y_t^{(n)}$ to compute a return-based trading signal, which is then executed at the close of day t . All positions are closed the next day to match the single-step prediction horizon. To ensure fidelity to realistic trading environments, this study simulates execution with a fixed round-trip transaction cost of 0.2%. Trades are executed without leverage and with equal-weighted capital distribution across selected assets. The portfolio is rebalanced daily, and all positions are reset at the end of each trading day.

3.3.1. Cumulative Return

The cumulative return CR quantifies the total capital growth over the backtest period. It is calculated by compounding daily portfolio returns across the evaluation horizon. Let $r_t^{(p)}$ denote the portfolio return at time t . The cumulative return is defined as

$$CR = \left(\prod_{t \in T_{\text{test}}} (1 + r_t^{(p)}) \right) - 1. \quad (18)$$

This metric measures the overall profitability of the strategy over the simulation window T_{test} .

3.3.2. Sharpe Ratio

To evaluate risk-adjusted performance, this study computes the annualized Sharpe ratio SR , which measures the average return per unit of volatility. Let $\mu_{r^{(p)}}$ be the mean, and $\sigma_{r^{(p)}}$ be the standard deviation of daily portfolio returns. The Sharpe ratio is expressed as

$$SR = \frac{\mu_{r^{(p)}}}{\sigma_{r^{(p)}}} \cdot \sqrt{252}. \quad (19)$$

The factor $\sqrt{252}$ annualizes the metric under the assumption of 252 trading days per year.

3.3.3. Maximum Drawdown

Maximum drawdown MDD measures the largest loss experienced from a peak to a subsequent trough in portfolio value. Let P_t be the portfolio value at time t . The drawdown is computed by tracking the historical maximum and measuring the decline from the peak. Formally, it is defined as

$$MDD = \max_{t \in T} \left(\frac{\max_{s \leq t} P_s - P_t}{\max_{s \leq t} P_s} \right). \quad (20)$$

This risk measure captures the worst-case drop in capital, which is critical for evaluating downside exposure.

3.3.4. Hit Ratio

The hit ratio evaluates the directional accuracy of the trading strategy by measuring how often the model correctly predicts the direction of the market. For each time step t , this study compares the sign of the portfolio return $r_t^{(p)}$ with that of the actual market return $r_t^{(b)}$:

$$HR = \frac{1}{|T_{\text{test}}|} \sum_{t \in T_{\text{test}}} \left[\text{sign}\left(r_t^{(p)}\right) = \text{sign}\left(r_t^{(b)}\right) \right]. \quad (21)$$

In this expression, the bracketed term returns to a value of 1 if the condition inside is true (i.e., the signs of the strategy and market returns match), and 0 otherwise. This metric reflects the frequency with which the strategy makes correct directional predictions, regardless of the magnitude of returns.

3.3.5. Strategy Stability

In addition to return and risk-based metrics, this study analyzes the turnover and volatility clustering of the strategy. High turnover may indicate excessive sensitivity to marginal signals, while low turnover may reflect excessive conservatism. Volatility clustering is evaluated by examining sequences of large positive or negative returns, providing insight into the strategy's robustness under regime shifts. Together, these evaluation metrics offer a comprehensive assessment of both the predictive quality of the TFT model and the reliability of the trading system in operational scenarios.

4. Experimental Results and Discussion

This chapter presents the empirical evaluation of the proposed forecasting and trading framework. The experimental design aims to assess both the predictive capability of the TFT and its effectiveness when integrated into a signal-based trading strategy. The results are reported in two stages. First, this study evaluates the regression accuracy of the TFT model in comparison with several widely used machine learning baselines, including LSTM, GRU, XGBoost, Support Vector Regression (SVR), and Linear Regression. These comparisons are conducted across five major cryptocurrencies using standard performance metrics such as RMSE, MAE, MAPE, and R^2 . The evaluation emphasizes out-of-sample prediction quality based on rolling test windows.

Second, this study analyzes the practical trading performance of the proposed strategy when it operates on TFT-generated signals. Backtesting simulations are used to quantify real-world performance across return, risk, and consistency dimensions. This two-phase evaluation provides a comprehensive understanding of the model's forecasting ability as well as its viability in live-market conditions.

4.1. Experimental Setup

To evaluate the predictive performance of the proposed TFT framework, this study conducts experiments across five widely traded cryptocurrencies: Bitcoin (BTC), Ethereum (ETH), Tether (USDT), Ripple (XRP), and Binance Coin (BNB). The dataset includes daily OHLCV data and a range of computed indicators, covering both technical and on-chain dimensions. These indicators include RSI, MACD, SOPR, TVL, AA, ENF, Realized Cap, HODL Waves, as well as sentiment proxies such as the Crypto Fear and Greed Index. The raw data spans from 1 January 2022 to 31 December 2024, yielding a total of 1095 daily observations per asset. This time range was deliberately chosen to capture a period of heightened volatility and regime shifts in the cryptocurrency market. From 2022 to 2024, the market experienced a full correction, a prolonged bearish trend, and the early phase

of a new bull cycle. These dynamics provide a suitable environment for testing model responsiveness to rapid market changes. Even for stablecoins like USDT, meaningful variations in on-chain activity and net flows during this period offer valuable signals for evaluation. All continuous features are normalized using Z-score standardization, and skewed variables are preprocessed using log or logit transformations where applicable. Features are aligned temporally, and missing values are forward filled. Each model is trained to perform one-step-ahead regression, predicting the next-day closing price for each asset.

This study benchmarks the TFT model against five baseline models commonly used in time series regression: LSTM, GRU, SVR, XGBoost Regressor, and Ordinary Least Squares Linear Regression. These models are selected to represent a broad methodological spectrum encompassing both traditional machine learning and deep learning architectures. Specifically, LSTM and GRU are recurrent neural networks widely adopted for financial forecasting due to their ability to capture nonlinear dependencies and temporal dynamics. SVR and XGBoost are chosen for their proven efficiency in small to medium-sized datasets and their robustness under structured input conditions. Including these baselines allows us to compare TFT's performance against both memory-based and tree-based learners, offering a well-rounded assessment of its forecasting capabilities. All deep learning models (TFT, LSTM, GRU) are trained using a sliding-window setup with an encoder window size of 30 days. Training is performed for 200 epochs with early stopping and a batch size of 64. The Adam optimizer is used with an initial learning rate of 0.005, and a cosine annealing learning rate scheduler is applied for stability. Model evaluation follows a walk-forward testing approach. The dataset is split chronologically with 80% used for training and 20% for out-of-sample testing. To ensure consistency, all models are trained and evaluated on the same feature space and prediction task. No future information is leaked into training windows.

All experiments are executed on a workstation running Windows 11 Pro, equipped with an Intel® Core™ i9-12900K CPU, 32 GB of RAM, and an NVIDIA GeForce RTX 4080 GPU. The hardware configuration ensures efficient handling of sequence modeling, backpropagation, and time-windowed forecasting tasks. This experimental setup is designed to fairly assess the regression capabilities of each model across multiple assets under the same temporal and informational conditions, providing a consistent basis for cross-model comparison in Section 4.2.

4.2. Experimental Results

This section presents the core experimental findings of the proposed framework. The evaluation is conducted in two stages: first, by assessing the predictive accuracy of the TFT and several baseline models in a one-step-ahead regression task across multiple cryptocurrencies; and second, by analyzing the performance of a signal-based trading strategy driven by the predicted outputs. The results provide insights into both model-level forecasting precision and the practical utility of integrating such models into trading operations.

4.2.1. Regression-Based Forecasting Analysis

This section presents the regression-based evaluation of the proposed TFT model, benchmarked against five representative forecasting approaches: LSTM, GRU, XGBoost, SVR, and Linear Regression (LR). Each model is trained to predict the next-day closing price of the target cryptocurrencies using the same input features and data splits as described in Section 4.1.

To quantify forecasting accuracy, this study adopts four standard regression metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage

Error (MAPE), and the Coefficient of Determination (R^2) across five major cryptocurrencies: BTC, ETH, USDT, XRP, and BNB. These metrics are defined as follows.

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T (\hat{y}_t - y_t)^2}, \quad (22)$$

$$\text{MAE} = \frac{1}{T} \sum_{t=1}^T |\hat{y}_t - y_t|, \quad (23)$$

$$\text{MAPE} = \frac{100\%}{T} \sum_{t=1}^T \left| \frac{\hat{y}_t - y_t}{y_t} \right|, \quad (24)$$

$$R^2 = 1 - \frac{\sum_{t=1}^T (\hat{y}_t - y_t)^2}{\sum_{t=1}^T (y_t - \bar{y}_t)^2}. \quad (25)$$

In these equations, $\hat{y}_t$ is the predicted closing price at time t , y_t is the true closing price, $\bar{y}_t$ is the mean of the true values over the test period, and T is the total number of predictions, representing the number of out-of-sample time steps over which the model's regression performance is evaluated. RMSE penalizes large errors more heavily due to squaring, while MAE reflects the average absolute deviation. MAPE provides a scale-invariant error percentage, particularly useful when assets differ in price range. The R^2 score quantifies the proportion of variance explained by the model. The full results are summarized in Table 2, while Table 3 provides average performance scores across all assets. Overall, TFT consistently outperforms the other models across all metrics on average, with particularly notable advantages in forecasting stability and relative accuracy.

Table 2. Regression results of different models across five major cryptocurrencies.

Asset	Metric	TFT	LSTM	GRU	XGBoost	SVR	LR
BNB-USD	RMSE	356.3	353.3	387.6	421.4	516.4	538.7
	MAE	238.9	251.5	268.3	260.7	267.3	307.8
	MAPE (%)	3.670	4.670	4.730	4.990	5.950	5.50
	R^2	0.948	0.899	0.902	0.928	0.841	0.881
BTC-USD	RMSE	390.4	428.1	432.6	458.9	506.1	517.8
	MAE	256.1	278.3	284.9	298.6	323.9	335.2
	MAPE (%)	3.410	3.950	3.960	4.240	4.710	4.880
	R^2	0.937	0.905	0.899	0.876	0.846	0.834
ETH-USD	RMSE	371.2	406.2	419.7	438.2	477.3	499.6
	MAE	242.8	255.4	258.6	272.1	296.8	305.9
	MAPE (%)	3.630	4.010	4.060	4.330	4.820	5.060
	R^2	0.944	0.922	0.916	0.892	0.869	0.862
USDT-USD	RMSE	112.8	140.5	142.7	159.4	172.4	180.2
	MAE	80.30	98.20	99.10	110.3	118.7	123.7
	MAPE (%)	1.180	1.430	1.490	1.610	1.880	1.990
	R^2	0.961	0.941	0.936	0.918	0.891	0.879
XRP-USD	RMSE	405.7	448.9	462.1	472.7	518.7	538.3
	MAE	271.2	295.7	302.1	318.8	345.2	352.6
	MAPE (%)	4.010	4.620	4.870	5.210	5.780	5.930
	R^2	0.926	0.882	0.871	0.852	0.821	0.812

Table 3. Average regression metrics across all assets.

Model	RMSE	MAE	MAPE (%)	R ²
TFT	327.28	217.86	3.180	0.9432
LSTM	355.40	235.82	3.736	0.9098
GRU	368.94	242.60	3.822	0.9048
XGBoost	390.12	252.10	4.076	0.8932
SVR	438.18	270.38	4.628	0.8536

Table 2 presents the regression performance of all models—TFT, LSTM, GRU, XGBoost, SVR, and Linear Regression—evaluated across five major cryptocurrencies: BNB, BTC, ETH, USDT, and XRP. Each asset-specific block reports four metrics: RMSE, MAE, MAPE (%), and R². These rows indicate the forecasting accuracy of each model on an individual asset. Table 3 shows the average forecasting performance of each model across all assets. Each column corresponds to a regression metric (RMSE, MAE, MAPE, R²), allowing for direct comparison of overall model accuracy.

In terms of RMSE, which penalizes large errors more heavily, TFT achieves the lowest average (327.28), showing a 25.3% improvement over SVR, which performed the worst (438.18). Across individual assets, TFT yields the lowest RMSE in ETH and USDT, and also ranks among the best for BTC and BNB. XRP remains the most difficult asset to predict, with all models displaying significantly higher RMSE values. Interestingly, while Linear Regression performed worst overall, LSTM achieves a slightly better result than TFT for BNB (353.3 vs. 356.3), reflecting the competitive nature of temporal neural models under certain conditions.

When comparing MAE, which provides a more interpretable sense of average error magnitude, TFT again ranks first with an average of 217.86. This reflects a 19.4% reduction compared to SVR, the weakest model. The model performs particularly well on stable assets like USDT, with an MAE of only 80.3, and maintains low error margins for BTC, ETH, and BNB. XRP once again exhibits the largest residual errors, where LSTM slightly edges out TFT. These results suggest that TFT’s architecture is especially effective in high-signal, low-noise settings, while other models may be more robust in chaotic regimes.

MAPE, which adjusts for relative price scale, offers another perspective on predictive consistency. TFT achieves the best average score of 3.18%, significantly outperforming SVR (4.628%) by 31.3%. This suggests TFT maintains a high percentage-level accuracy across assets of varying price levels. USDT stands out with a remarkably low MAPE of 1.18%, further confirming the model’s ability to handle low-volatility series. XRP and BNB, conversely, show higher MAPE values across all models, indicating persistent forecasting difficulty in high-volatility assets.

Finally, the R² metric reveals how much variance in the actual price can be explained by model predictions. TFT records an average R² of 0.9432, indicating high explanatory power. In comparison, SVR and Linear Regression trail significantly with values of 0.8536 and 0.854, respectively. TFT produces particularly strong R² scores for USDT (0.961), ETH (0.944), and BNB (0.948), while XRP again lags behind. This consistent pattern across metrics confirms that TFT generalizes well across different asset types, particularly when asset behaviors are distinguishable and feature-rich.

These regression results demonstrate that TFT, implemented as a unified multi-asset model, is able to outperform both deep learning and traditional baselines not only in aggregated metrics, but also in maintaining strong and stable performance across diverse cryptocurrency assets. This performance advantage provides a solid foundation for further assessing the model’s practical applicability in real-world trading scenarios, which will be analyzed in the next section.

4.2.2. Trading Strategy Performance Analysis

To evaluate the practical utility of the regression models in real-world financial decision-making, this study conducts a backtesting experiment using a signal-based trading strategy as described in Section 3.2.2. The strategy converts model-predicted price changes into actionable positions based on the z-score of the predicted return, which is derived by comparing the forecasted next-day closing price with the current price and normalizing it over a rolling historical distribution. For this experiment, this study applies a 1-day holding policy, whereby the trading signal is generated at the close of day t , and the corresponding position is opened immediately and closed at the end of day $t+1$. This short-term structure avoids position compounding and aligns directly with the model's one-step-ahead forecast horizon. The position sizing is fixed at one unit per signal, and a round-trip transaction fee of 0.2% is included to simulate realistic market friction. This study compares the trading performance of the TFT-based strategy with the same strategy applied to three benchmark models: LSTM, GRU, and SVR. All models generate price forecasts using the same input features and data splits as in Section 4.1. To ensure fairness, the z-score threshold, signal generation logic, and backtesting pipeline are kept consistent across models.

The trading performance is evaluated based on the methodology established in Section 3.3, utilizing several standard financial metrics to comprehensively assess strategy effectiveness. Specifically, the cumulative return (CR) captures the total return accumulated over the full backtest period, providing a measure of overall profitability. The Sharpe ratio is computed as the average daily return divided by the standard deviation of daily returns, and annualized over 252 trading days to facilitate comparison with conventional investment performance standards. Maximum drawdown (MDD) quantifies the greatest peak-to-trough decline experienced during the backtesting period, highlighting the magnitude of potential downside risk. In addition to risk-adjusted and absolute returns, this study evaluates directional accuracy through the hit ratio, which measures the proportion of trades where the predicted price movement correctly matches the actual market outcome. The average daily return provides insight into the consistency and profitability of signals on a per-trade basis, while volatility, measured as the standard deviation of daily returns, indicates the level of return variability and risk exposure associated with each trading strategy.

Using the full dataset from 1 January 2022 to 31 December 2024, this study simulates each strategy across all five target cryptocurrencies: BTC, ETH, USDT, XRP, and BNB. Positions are opened and closed independently per asset, allowing for disaggregation as well as portfolio-level analysis. Performance metrics are calculated both individually and as portfolio averages.

Table 4 summarizes the trading performance of the TFT-based strategy in comparison with several benchmarks, including LSTM, GRU, SVR, Buy-and-Hold, and Naive Momentum strategies. All metrics are computed based on multi-asset average results across BTC, ETH, USDT, XRP, and BNB over the 2022–2024 period.

Table 4. Performance comparison of signal-based trading strategies across models.

Model	Cumulative Return (%)	Sharpe Ratio	Max Drawdown (%)	Hit Ratio	Average Daily Return (%)	Volatility (%)
TFT	38.6	1.06	−22.4	0.56	0.13	1.21
LSTM	34.2	0.98	−19.8	0.58	0.12	1.18
GRU	31.5	0.91	−20.9	0.57	0.11	1.20
SVR	18.7	0.61	−27.2	0.53	0.07	1.30
Buy-and-Hold	28.1	0.79	−25.4	0.51	0.09	1.27
Naive Momentum	26.4	0.75	−23.8	0.52	0.08	1.25

In terms of cumulative return, the TFT strategy achieved the highest return at 38.6%, outperforming LSTM (34.2%) and GRU (31.5%) by approximately 12.9% and 22.5%, respectively. Compared to the traditional SVR model, which accumulated only 18.7%, the TFT strategy generated 106.4% higher returns, demonstrating its superior ability to capture profitable price movements across different assets. The Buy-and-Hold strategy resulted in a cumulative return of 28.1%, meaning that TFT still outperformed passive investing by 37.3%. Meanwhile, the Naive Momentum strategy lagged behind at 26.4%.

When examining the Sharpe Ratio, which reflects the risk-adjusted return, the TFT strategy again leads with a ratio of 1.06. This represents a 9.2% improvement over LSTM (0.98) and a 16.5% improvement over GRU (0.91). Compared to SVR (0.61), TFT's Sharpe Ratio is 73.8% higher, highlighting a significantly better reward-to-risk trade-off. However, it is important to note that although TFT has the highest Sharpe Ratio, the absolute magnitude remains moderate, reflecting the challenging volatility environment characteristic of cryptocurrency markets.

Interestingly, in terms of maximum drawdown, TFT shows a deeper risk exposure. TFT's maximum drawdown reached -22.4% , which is larger than that of LSTM (-19.8%) and GRU (-20.9%). This indicates that while TFT captures more upside potential, it also endures relatively larger temporary losses during market downturns. Compared to SVR's maximum drawdown of -27.2% , however, TFT still exhibits better downside protection. This suggests that TFT, while aggressive in capturing trends, remains relatively contained in extreme market conditions.

Hit ratio analysis further reveals the nuances of the models' behavior. TFT achieved a hit ratio of 0.56, which is slightly lower than LSTM (0.58) and GRU (0.57). This means that while TFT generates higher overall returns, its directional accuracy is not the highest among the compared models. In practical terms, TFT compensates for a slightly lower hit ratio with larger gains on correct trades, demonstrating a "higher reward per correct signal" behavior. LSTM, although achieving a slightly higher directional accuracy, sacrifices some profitability in exchange for increased signal consistency.

Regarding average daily return and volatility, the TFT strategy maintained an average daily return of 0.13% with a volatility of 1.21%. This volatility level is slightly higher than LSTM (1.18%) but lower than SVR (1.30%) and Buy-and-Hold (1.27%). These results suggest that while TFT operates with a moderate level of risk, it does so efficiently relative to the returns achieved. In contrast, SVR's high volatility coupled with lower returns underscores its instability in short-term forecasting for cryptocurrency assets.

The analysis demonstrates that the TFT-based trading strategy effectively balances return generation with risk management, despite minor trade-offs in hit ratio and drawdown severity. It outperforms both traditional machine learning models and naive benchmarks across most key metrics, validating its applicability in dynamic and volatile multi-crypto-asset trading environments. The trading performance results clearly demonstrate the superiority of the TFT-based strategy in cumulative return and risk-adjusted profitability across multiple cryptocurrency assets. Although the model accepts slightly greater drawdowns and a marginally lower hit ratio compared to LSTM and GRU, it consistently delivers better overall reward-to-risk trade-offs. These outcomes reflect inherent characteristics of deep attention-based forecasting models. TFT's architecture, which emphasizes dynamic variable selection and sequence-to-sequence learning, is inherently more sensitive to short-term market changes. This heightened signal sensitivity enables the model to more aggressively capture emerging trends, which contributes to higher cumulative returns. However, it also makes the strategy more vulnerable to short-term market noise, explaining the relatively larger maximum drawdowns observed during volatile periods. Moreover, the hit ratio analysis illustrates that higher directional accuracy does not necessarily equate

to superior trading performance. The TFT strategy compensates for its slightly lower hit rate by ensuring that correct signals yield substantially larger gains than the losses incurred from incorrect signals. This behavior aligns with modern trading theory, where profitability is often more dependent on the magnitude of gains versus losses rather than sheer prediction frequency.

Overall, the results validate the practical effectiveness of deploying multi-asset deep learning models such as TFT in volatile trading environments. Nevertheless, the observed trade-offs between sensitivity and robustness suggest that future refinements could focus on adaptive signal thresholding, dynamic holding periods, and portfolio optimization strategies aimed at enhancing stability while maintaining profitability.

5. Conclusions and Future Research

This study demonstrates the potential of integrating deep learning with financial domain knowledge to enhance cryptocurrency forecasting and trading. By leveraging a TFT with technical indicators and on-chain metrics (SOPR, TVL, AA, etc.), this study constructs a multi-asset framework capable of modeling short-term price movements with high accuracy.

The TFT model achieves the strongest predictive performance among all baselines, with the lowest average RMSE (327.28), MAE (217.86), and MAPE (3.18%), and the highest R^2 score (0.9432). Its robustness across diverse crypto assets indicates strong generalization, though forecasting remains more challenging for highly volatile tokens like XRP. In addition, a signal-based trading strategy derived from TFT predictions achieves a cumulative return of 38.6% and a Sharpe ratio of 1.06 over a three-year backtest, surpassing both traditional models and passive benchmarks. While the strategy experiences slightly higher drawdowns, it compensates with greater returns per successful signal, reflecting a favorable reward-to-risk profile. These findings underscore the viability of combining interpretable deep learning architectures with rich, behavior-driven financial inputs to support data-driven decision-making in volatile and high-noise environments such as crypto-asset markets.

However, several practical limitations remain. The current framework assumes one-day holding periods, which may not fully capture market complexity. In particular, real-world trading constraints such as liquidity limitations, slippage, and latency in trade execution are not explicitly modeled, potentially affecting the feasibility of deploying the strategy at scale. Future research could extend this framework in several specific directions. First, integrating macro-financial indicators such as interest rate shifts, inflation expectations, or geopolitical risk indices could help improve model robustness under different market regimes. Second, exploring adaptive signal thresholding and variable holding durations may better align trading actions with changing volatility and investor sentiment. Third, incorporating order book dynamics and execution latency models would make the system more microstructure-aware, enabling practical deployment in real-world trading environments. Lastly, future work can experiment with advanced transformer variants such as Autoformer or FEDformer, especially for longer-term horizon forecasting, to complement the short-term focus of this study.

Funding: This work was supported by the National Science Council for providing the research grant NSTC 113-2221-E-130-006.

Data Availability Statement: The data presented in this study are available in Yahoo Finance at <https://reurl.cc/gek5Gp>; <https://reurl.cc/aeMKnY>; <https://pse.is/7hz5ay>; <https://pse.is/7hz5cs>; <https://pse.is/7hz5eb>; and <https://www.coinglass.com/>, access date: 1 May 2025.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. CoinMarketCap. Global Cryptocurrency Market Capitalization—Q1 2025. Available online: <https://coinmarketcap.com/charts> (accessed on 10 April 2025).
2. Chainalysis. *The 2025 Geography of Cryptocurrency Report*; Chainalysis: New York, NY, USA, 2025. Available online: <https://www.chainalysis.com/reports> (accessed on 10 April 2025).
3. Fang, F.; Ventre, C.; Basios, M.; Kanthan, L.; Martinez-Rego, D.; Wu, F.; Li, L. Cryptocurrency trading: A comprehensive survey. *Financ. Innov.* **2022**, *8*, 13. [\[CrossRef\]](#)
4. Nti, I.K.; Adekoya, A.F.; Weyori, B.A. A systematic review of fundamental and technical analysis of stock market predictions. *Artif. Intell. Rev.* **2020**, *53*, 3007–3057. [\[CrossRef\]](#)
5. Pramudya, R.; Ichsani, S. Efficiency of technical analysis for the stock trading. *Int. J. Financ. Bank. Stud.* **2020**, *9*, 58–67.
6. Zatwarnicki, M.; Zatwarnicki, K.; Stolarski, P. Effectiveness of the relative strength index signals in timing the cryptocurrency market. *Sensors* **2023**, *23*, 1664. [\[CrossRef\]](#)
7. Demosthenous, G.; Georgiou, C.; Polydorou, E. From On-chain to Macro: Assessing the Importance of Data Source Diversity in Cryptocurrency Market Forecasting. Proceedings of the VLDB Endowment. ISSN 2150-8097. Available online: <https://vldb.org/workshops/2024/proceedings/FAB/FAB-6.pdf> (accessed on 12 June 2025).
8. Casella, B.; Paletto, L. Predicting cryptocurrencies market phases through on-chain data long-term forecasting. In Proceedings of the 2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), Dubai, United Arab Emirates, 1–5 May 2023; pp. 1–4.
9. Metelski, D.; Sobieraj, J. Decentralized finance (DeFi) projects: A study of key performance indicators in terms of DeFi protocols' valuations. *Int. J. Financ. Stud.* **2022**, *10*, 108. [\[CrossRef\]](#)
10. Grande, M.; Borondo, J. Trust as a driver in the DeFi market: Leveraging TVL/MCAP bands as confidence indicators to anticipate price movements. *Financ. Res. Lett.* **2025**, 106705. [\[CrossRef\]](#)
11. Cong, L.W.; Prasad, E.S.; Rabetti, D. Financial and informational integration through oracle networks. *Natl. Bur. Econ. Res.* **2025**, 33639. [\[CrossRef\]](#)
12. Al-Selwi, S.M.; Hassan, M.F.; Abdulkadir, S.J.; Muneer, A.; Sumiea, E.H.; Alqushaibi, A.; Ragab, M.G. RNN-LSTM: From applications to modeling techniques and beyond—Systematic review. *J. King Saud Univ. Comput. Inf. Sci.* **2024**, *36*, 102068. [\[CrossRef\]](#)
13. Sajun, A.R.; Zualkernan, I.; Sankalpa, D. A historical survey of advances in transformer architectures. *Appl. Sci.* **2024**, *14*, 4316. [\[CrossRef\]](#)
14. Lim, B.; Arik, S.Ö.; Loeff, N.; Pfister, T. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *Int. J. Forecast.* **2021**, *37*, 1748–1764. [\[CrossRef\]](#)
15. Eşki, D.; Kaya, T. Retail Demand Forecasting Using Temporal Fusion Transformer. In *International Conference on Intelligent and Fuzzy Systems*; Springer Nature: Cham, Switzerland, 2024; pp. 165–170.
16. Zheng, P.; Zhou, H.; Liu, J.; Nakanishi, Y. Interpretable building energy consumption forecasting using spectral clustering algorithm and temporal fusion transformers architecture. *Appl. Energy* **2023**, *349*, 121607. [\[CrossRef\]](#)
17. Han, Y.; Tian, Y.; Yu, L.; Gao, Y. Economic system forecasting based on temporal fusion transformers: Multi-dimensional evaluation and cross-model comparative analysis. *Neurocomputing* **2023**, *552*, 126500. [\[CrossRef\]](#)
18. Balara, V.; Mach, M.; Machova, K. The Impact of Sentiment in S&P 500 volatility prediction with the use of Deep Learning. In Proceedings of the 2023 21st International Conference on Emerging eLearning Technologies and Applications (ICETA), Stary Smokovec, Slovakia, 26–27 October 2023; pp. 25–30.
19. Hajek, P.; Novotny, J. Beyond Sentiment in Stock Price Prediction: Integrating News Sentiment and Investor Attention with Temporal Fusion Transformer. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*; Springer Nature: Cham, Switzerland; pp. 30–43.
20. Kumar, D.; Kumar, P.; Anandhi, S. Augmented Gold Price Forecasting via Image Processing and Machine Learning Fusion. In Proceedings of the 2024 Second International Conference on Advances in Information Technology (ICAIT), Chikkamagaluru, Karnataka, India, 24–27 July 2024; Volume 1, pp. 1–6.
21. Farooq, A.; Uddin, M.I.; Adnan, M.; Alarood, A.A.; Alsolami, E.; Habibullah, S. Interpretable multi-horizon time series forecasting of cryptocurrencies by leverage temporal fusion transformer. *Heliyon* **2024**, *10*, e40142. [\[CrossRef\]](#)
22. Olorunnimbe, K.; Viktor, H. Ensemble of temporal Transformers for financial time series. *J. Intell. Inf. Syst.* **2024**, *62*, 1087–1111. [\[CrossRef\]](#)
23. Mintarya, L.N.; Halim, J.N.; Angie, C.; Achmad, S.; Kurniawan, A. Machine learning approaches in stock market prediction: A systematic literature review. *Procedia Comput. Sci.* **2023**, *216*, 96–102. [\[CrossRef\]](#)

24. Oyewole, A.T.; Adeoye, O.B.; Addy, W.A.; Okoye, C.C.; Ofodile, O.C.; Ugochukwu, C.E. Predicting stock market movements using neural networks: A review and application study. *Comput. Sci. IT Res. J.* **2024**, *5*, 651–670. [\[CrossRef\]](#)
25. Josué, T. Assessing Cryptomarket Risks: Macroeconomic Forces, Market Shocks and Behavioural Dynamics. *Mark. Shock. Behav. Dyn.* **2025**.
26. Cheng, L.C.; Sun, J.S. Multiagent-based deep reinforcement learning framework for multi-asset adaptive trading and portfolio management. *Neurocomputing* **2024**, *594*, 127800. [\[CrossRef\]](#)
27. Jay, P.; Kalariya, V.; Parmar, P.; Tanwar, S.; Kumar, N.; Alazab, M. Stochastic neural networks for cryptocurrency price prediction. *IEEE Access* **2020**, *8*, 82804–82818. [\[CrossRef\]](#)
28. Li, Y.; Jiang, S.; Li, X.; Wang, S. Hybrid data decomposition-based deep learning for bitcoin prediction and algorithm trading. *Financ. Innov.* **2022**, *8*, 31. [\[CrossRef\]](#)
29. Murray, K.; Rossi, A.; Carraro, D.; Visentin, A. On forecasting cryptocurrency prices: A comparison of machine learning, deep learning, and ensembles. *Forecasting* **2023**, *5*, 196–209. [\[CrossRef\]](#)
30. Syed, S.; Talha, S.M.; Iqbal, A.; Ahmad, N.; Alshara, M.A. Seeing Beyond Noise: Improving Cryptocurrency Forecasting with Linear Bias Correction. *AI* **2024**, *5*, 2829–2851. [\[CrossRef\]](#)
31. Zhao, F.; Zhang, M.; Zhou, S.; Lou, Q. Application of Deep Reinforcement Learning for Cryptocurrency Market Trend Forecasting and Risk Management. *J. Ind. Eng. Appl. Sci.* **2024**, *2*, 48–55.
32. Alnami, H.; Mohzary, M.; Assiri, B.; Zangoti, H. An Integrated Framework for Cryptocurrency Price Forecasting and Anomaly Detection Using Machine Learning. *Appl. Sci.* **2025**, *15*, 1864. [\[CrossRef\]](#)
33. Parekh, R.; Patel, N.P.; Thakkar, N.; Gupta, R.; Tanwar, S.; Sharma, G.; Davidson, I.E.; Sharma, R. DL-GuesS: Deep learning and sentiment analysis-based cryptocurrency price prediction. *IEEE Access* **2022**, *10*, 35398–35409. [\[CrossRef\]](#)
34. Feizian, F.; Amiri, B. Cryptocurrency Price Prediction Model Based on Sentiment Analysis and Social Influence. *IEEE Access* **2023**, *11*, 142177–142195. [\[CrossRef\]](#)
35. Shahiki Tash, M.; Ahani, Z.; Tash, M.; Kolesnikova, O.; Sidorov, G. Analyzing Emotional Trends from X Platform Using SenticNet: A Comparative Analysis with Cryptocurrency Price. *arXiv* **2024**, arXiv:2405.03084. [\[CrossRef\]](#)
36. Jahanbin, K.; Chahooki, M.A.Z. Cryptocurrency Trend Prediction Through Hybrid Deep Transfer Learning. *Int. J. Intell. Syst.* **2025**, *2025*, 4211799. [\[CrossRef\]](#)
37. Sebastião, H.; Godinho, P. Forecasting and trading cryptocurrencies with machine learning under changing market conditions. *Financ. Innov.* **2021**, *7*, 1–30. [\[CrossRef\]](#)
38. Lin, T.M.; Yu, J.L.; Chen, J.W.; Huang, C.S. Application of machine learning with news sentiment in stock trading strategies. *Int. J. Financ. Res.* **2023**, *14*. [\[CrossRef\]](#)
39. Prashanth, G.; VSSKRNaganjaneyulu, G.; Revanth, M.; Narasimhadhan, A.V. Multi indicator based hierarchical strategies for technical analysis of crypto market paradigm. *Int. J. Electr. Comput. Eng. Syst.* **2023**, *14*, 765–780.
40. Otabek, S.; Choi, J. From prediction to profit: A comprehensive review of cryptocurrency trading strategies and price forecasting techniques. *IEEE Access* **2024**. [\[CrossRef\]](#)
41. Pečiulis, T.; Ahmad, N.; Menegaki, A.N.; Bibi, A. Forecasting of cryptocurrencies: Mapping trends, influential sources, and research themes. *J. Forecast.* **2024**, *43*, 1880–1901. [\[CrossRef\]](#)
42. Nguyen, D.T.A.; Chan, K.C. Cryptocurrency trading: A systematic mapping study. *Int. J. Inf. Manag. Data Insights* **2024**, *4*, 100240. [\[CrossRef\]](#)
43. Jing, L.; Kang, Y. Automated cryptocurrency trading approach using ensemble deep reinforcement learning: Learn to understand candlesticks. *Expert Syst. Appl.* **2024**, *237*, 121373. [\[CrossRef\]](#)
44. Peng, P.; Chen, Y.; Lin, W.; Wang, J.Z. Attention-based CNN–LSTM for high-frequency multiple cryptocurrency trend prediction. *Expert Syst. Appl.* **2024**, *237*, 121520. [\[CrossRef\]](#)
45. Kochliaridis, V.; Papadopoulou, A.; Vlahavas, I. UNSURE-A machine learning approach to cryptocurrency trading. *Appl. Intell.* **2024**, *54*, 5688–5710. [\[CrossRef\]](#)
46. Huang, C.S.; Su, Y.S. Trading Strategy of the Cryptocurrency Market Based on Deep Q-Learning Agents. *Appl. Artif. Intell.* **2024**, *38*, 2381165. [\[CrossRef\]](#)
47. García-Medina, A.; Aguayo-Moreno, E. LSTM–GARCH hybrid model for the prediction of volatility in cryptocurrency portfolios. *Comput. Econ.* **2024**, *63*, 1511–1542. [\[CrossRef\]](#)
48. Feng, L.; Qi, J.; Lucey, B. Enhancing cryptocurrency market volatility forecasting with daily dynamic tuning strategy. *Int. Rev. Financ. Anal.* **2024**, *94*, 103239. [\[CrossRef\]](#)
49. Pellicani, A.; Pio, G.; Ceci, M. CARROT: Simultaneous prediction of anomalies from groups of correlated cryptocurrency trends. *Expert Syst. Appl.* **2025**, *260*, 125457. [\[CrossRef\]](#)

50. King, J.C.; Dale, R.; Amigó, J.M. Blockchain metrics and indicators in cryptocurrency trading. *Chaos Solitons Fractals* **2024**, *178*, 114305. [\[CrossRef\]](#)
51. Ah Mand, A. Cryptocurrency returns and cryptocurrency uncertainty: A time–frequency analysis. *Financ. Innov.* **2025**, *11*, 52. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.