



**Daffodil**  
*International*  
**University**

**Machine Learning Approaches to Cryptocurrency Market Volatility**

**Prediction: A Framework for Risk-Adjusted Trading Strategy**

**Optimization**

**Submitted By:**

Nazmul Amin

221-35-1029

Department of Software Engineering

Daffodil International University

**Supervised By:**

Ms. Tapushe Rabaya Toma

Assistant Professor

Department of Software Engineering

Daffodil International University

Thesis submitted in fulfillment of the requirements for the award of the degree of

Bachelor of Science in Software Engineering Fall-2025


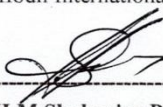

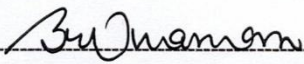
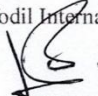
DAFFODIL INTERNATIONAL UNIVERSITY

NOVEMBER 2025

## APPROVAL

This thesis titled on “**Machine Learning Approaches to Cryptocurrency Market Volatility Prediction: A Framework for Risk-Adjusted Trading Strategy Optimization**”, submitted by Nazmul Amin (ID: 221-35-1029) to the Department of Software Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of Bachelor of Science in Software Engineering and approval as to its style and contents.

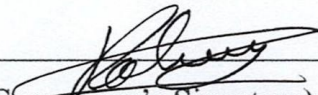
### BOARD OF EXAMINERS

 ----- <b>Dr. S M Hasan Mahmud</b> <b>Associate Professor</b> Department of Software Engineering Faculty of Science and Information Technology Daffodil International University	<b>Chairman</b>
 ----- <b>A.H.M Shahariar Parvez</b> <b>Associate Professor</b> Department of Software Engineering Faculty of Science and Information Technology Daffodil International University	<b>Internal Examiner 1</b>
 ----- <b>Tapushe Rabaya Toma</b> <b>Assistant Professor</b> Department of Software Engineering Faculty of Science and Information Technology Daffodil International University	<b>Internal Examiner 2</b>
 ----- <b>Khalid Been md. Badruzzaman Biplob</b> <b>Lecturer (Senior Scale)</b> Department of Software Engineering Faculty of Science and Information Technology Daffodil International University	<b>Internal Examiner 3</b>
 ----- <b>Dr. Md Sazzadur Rahman</b> <b>Professor</b> Institute of Information technology Jahangirnagar University, Bangladesh	<b>External Examiner</b>



## SUPERVISOR'S DECLARATION

I hereby declare that I have checked this thesis and in my opinion, this thesis is adequate in terms of scope and quality for the award of the degree of Bachelor of Science.

  
(Supervisor's Signature)

Full Name : Tapushe Rabaya Toma

Position : Assistant Professor

Date : 21.12.25



## STUDENT'S DECLARATION

I hereby declare that the work in this thesis is based on my original work except for quotations and citations which have been duly acknowledged. I also declare that it has not been previously or concurrently submitted for any other degree at Daffodil International University or any other institution.

*Nazmul Amin*

---

(Student's Signature)

Full Name : Nazmul Amin

ID Number : 221-35-1029

Date : 21-12-25

## **ACKNOWLEDGEMENTS**

Many people have helped and supported this thesis. Deep gratitude is owed to the supervisor for constant encouragement, guidance, and patience throughout countless experiments, revisions, and model iterations; their belief in the value of understanding cryptocurrency volatility made it possible to continue when progress was slow. Sincere thanks also go to the broader open-source and research communities for providing tools, datasets, and ideas that underpinned this work, including Python, pandas, scikit-learn and related libraries for time-series modelling and backtesting. Finally, heartfelt appreciation is extended to family, friends and colleagues whose feedback on early drafts, discussions about results, and personal support greatly improved the clarity, relevance and impact of this research.

# ABSTRACT

This thesis develops and evaluates an end-to-end machine learning pipeline for forecasting next-day realized volatility and volatility spikes in six major cryptocurrencies ADA, BTC, DOGE, ETC, ETH and LINK using hourly OHLCV data aggregated to daily frequency. The pipeline performs data cleaning, feature engineering, temporal splitting, model training and diagnostic reporting, creating a rich feature set that includes HAR style lagged realized volatility, rolling statistics, session-specific volatility, market-relative indicators and cyclical time encodings. Regression models (HAR baseline, Ridge, Elastic Net, Random Forest, Gradient Boosting and Histogram-based Gradient Boosting) are trained to predict next-day realized volatility, while parallel classification models (Logistic Regression, Random Forest, Gradient Boosting and Histogram-based Gradient Boosting) identify extreme volatility spikes defined by rolling 90th-percentile thresholds. Models are estimated on a strictly chronological split with training up to 2022, validation in 2023 and out-of-sample testing from 2024 onward, with optional walk-forward cross-validation to assess temporal robustness. The best regression model, HistGradientBoosting, attains a test RMSE of 0.00382 and  $R^2=0.093$ , while the best classifier, HistGradientBoostingClassifier, achieves an F1 score of 0.38 and ROC-AUC of 0.75, with substantially higher F1 in high-volatility regimes (0.60). These forecasts are embedded in an inverse-volatility trading strategy that overlays a moving-average directional signal, producing a 7.98% cumulative return, annualized Sharpe ratio of 0.52 and maximum drawdown of 10.9% on the test period. Per-coin and per-regime diagnostics reveal heterogeneous predictability, with stronger performance for smaller coins such as DOGE and ADA and during turbulent market conditions. Overall, the thesis contributes a transparent, reusable multi-asset pipeline that links machine learning-based cryptocurrency volatility forecasts to economically interpretable, risk-adjusted trading performance.

# TABLE OF CONTENTS

<b>SUPERVISOR’S DECLARATION</b> .....	<b>i</b>
<b>STUDENT’S DECLARATION</b> .....	<b>ii</b>
<b>ACKNOWLEDGEMENTS</b> .....	<b>iii</b>
<b>ABSTRACT</b> .....	<b>iv</b>
<b>TABLE OF CONTENTS</b> .....	<b>v</b>
<b>LIST OF TABLES</b> .....	<b>vii</b>
<b>LIST OF FIGURES</b> .....	<b>viii</b>
<b>CHAPTER 1 INTRODUCTION</b> .....	<b>1</b>
<b>1.1 Introduction</b> .....	<b>1</b>
<b>1.2 Background</b> .....	<b>2</b>
<b>1.3 Motivation of the Research</b> .....	<b>3</b>
<b>1.3 Problem Statements</b> .....	<b>4</b>
<b>1.4 Research Questions</b> .....	<b>4</b>
<b>1.5 Research Objective</b> .....	<b>5</b>
<b>1.6 Research Scope</b> .....	<b>6</b>
<b>1.7 Summary</b> .....	<b>7</b>
<b>CHAPTER 2 LITERATURE REVIEW</b> .....	<b>8</b>
<b>2.1 Introduction</b> .....	<b>8</b>
<b>2.2 Previous literature</b> .....	<b>9</b>
<b>2.3 Summary</b> .....	<b>17</b>
<b>CHAPTER 3 RESEARCH METHODOLOGY</b> .....	<b>19</b>
<b>3.1 Introduction</b> .....	<b>19</b>
<b>3.2 Data collection</b> .....	<b>21</b>
<b>3.3 Data preprocessing</b> .....	<b>22</b>
<b>3.4 Regression models</b> .....	<b>24</b>
<b>3.4.1 HAR baseline</b> .....	<b>24</b>
<b>3.4.2 Ridge and Elastic Net</b> .....	<b>26</b>
<b>3.4.3 Random Forest</b> .....	<b>27</b>
<b>3.4.4 Gradient Boosting</b> .....	<b>28</b>
<b>3.4.5 HistGradientBoosting</b> .....	<b>29</b>
<b>3.5 Classification models</b> .....	<b>30</b>

<b>3.5.1 Logistic Regression</b> .....	31
<b>3.5.2 Random Forest Classifier</b> .....	32
<b>3.5.3 Gradient Boosting Classifier</b> .....	33
<b>3.5.4 HistGradientBoosting Classifier</b> .....	34
<b>3.6 Summary</b> .....	35
<b>CHAPTER 4 RESULTS AND DISCUSSION</b> .....	<b>37</b>
<b>4.1 Introduction</b> .....	37
<b>4.2 Results</b> .....	37
<b>4.2.1 Regression: Predicting Next-day Realized Volatility</b> .....	37
<b>4.2.2 Classification: Detecting Volatility Spikes</b> .....	39
<b>4.2.3 Volatility-Managed Trading Strategy</b> .....	42
<b>4.3 Discussion</b> .....	43
<b>4.4 Summary</b> .....	44
<b>CHAPTER 5 CONCLUSION</b> .....	<b>46</b>
<b>5.1 Introduction</b> .....	46
<b>5.2 Conclusion</b> .....	46
<b>5.3 Recommendations for Future Work</b> .....	47
<b>Final summary:</b> .....	47
<b>REFERENCES</b> .....	<b>49</b>
<b>PLAGIARISM REPORT</b> .....	<b>51</b>
<b>ACCOUNTS CLEARANCE</b> .....	<b>59</b>

## LIST OF TABLES

Table 2.1	Literature Review	14
Table 4.1	Six regression models performance	37
Table 4.2	Per-coin $R^2$ (HistGradientBoosting test set)	38
Table 4.3	Four classifier models performance	39
Table 4.4	Per-coin F1 (HistGradientBoosting, test set)	40
Table 4.5	Market regime breakdown (HistGradientBoosting, test set)	40
Table 4.6	Confusion matrix (HistGradientBoosting, test set)	41
Table 4.7	The best regression model's predictions	42

## LIST OF FIGURES

Figure 3.1	Workflow Architecture	20
Figure 3.2	Timeline bar chart showing data coverage per cryptocurrency (2018–2025)	22
Figure 3.3	Box plot showing the distribution of hourly returns and realized volatility by coin, before and after clipping outliers at the 1st–99th percentile	23
Figure 3.4	Time series plot of daily realized volatility (aggregated from hourly data) for all six cryptocurrencies, highlighting major spikes and regime changes from 2018–2025	24
Figure 3.5	Scatter plot showing predicted vs. actual next day realized volatility for the HAR baseline on the test set, with a diagonal reference line and R2 annotation	25
Figure 3.6	Coefficient plot (bar chart) showing the top 15 most important features for Ridge and Elastic Net regression, with coefficients on the x axis and feature names on the y axis	26
Figure 3.7	Feature importance bar chart for Random Forest regression, showing the top 20 features ranked by mean decrease in impurity	27
Figure 3.8	Learning curve plot showing training and validation RMSE across boosting iterations for Gradient Boosting regression	28
Figure 3.9	Predicted vs. actual scatter plot for HistGradientBoosting regression on the test set, annotated with test R2=0.093 and RMSE = 0.00382	29
Figure 3.10	Bar chart showing the class distribution (count of spike vs. normal days) across the train, validation, and test sets	30
Figure 3.11	Calibration curve (reliability diagram) showing predicted probabilities vs. observed spike frequency for Logistic Regression on the test set	31
Figure 3.12	Feature importance bar chart for Random Forest classifier, showing the top 20 features ranked by permutation importance	32
Figure 3.13	Precision recall curve for Gradient Boosting classifier on the test set, annotated with area under the curve (AUC PR) and F1 score at the optimal threshold	33
Figure 3.14	ROC curve for HistGradientBoosting classifier on the test set, showing false positive rate vs. true positive rate, with AUC = 0.751 annotated	34
Figure 3.15	Confusion matrix heatmap for HistGradientBoosting classifier on the test set, showing true negatives, false positives, false negatives, and true positives	35

Figure 4.1	RMSE and $R^2$ for each model (side-by-side)	38
Figure 4.2	Bar Chart: F1 score and ROC-AUC for each classifier	39
Figure 4.3	Grouped bar chart: F1 by regime ('Normal', 'High-volatility')	40
Figure 4.4	As a labeled matrix/heatmap after classifier performance tables	41
Figure 4.5	Time-series plot of cumulative return, 2024-2025	42

# CHAPTER 1 INTRODUCTION

## 1.1 Introduction

The rapid growth of cryptocurrency markets has created new opportunities for trading and investment, but has also introduced extreme price volatility that complicates risk management and portfolio construction (Schwert, 1989). Unlike traditional assets, crypto trades continuously across global venues with frequent regime shifts, making it difficult for investors and regulators to obtain reliable estimates of short-horizon risk (Engle, 2008, August). In response, researchers have increasingly turned to advanced econometric and machine learning techniques to forecast realized volatility and to design strategies that adjust exposure based on predicted risk (Zhang, 2024). However, much of the existing work focuses either on single assets, high-frequency proprietary data, or purely predictive accuracy, leaving a gap for reproducible, end-to-end pipelines that connect volatility forecasts to risk-adjusted trading performance across multiple cryptocurrencies (Djanga, 2025). This thesis addresses this gap by developing and evaluating a modular machine learning pipeline that forecasts next-day realized volatility and volatility spikes for six major cryptocurrencies (BTC, ETH, ADA, DOGE, ETC, LINK) using hourly OHLCV data enriched with technical indicators, and the pipeline integrates data preprocessing, feature engineering, model training, diagnostic analysis, and a volatility-managed trading strategy, enabling a systematic assessment of both statistical and economic value.

## 1.2 Background

Volatility forecasting has long been central to financial econometrics because accurate variance estimates underpin option pricing, value-at-risk, and capital allocation decisions, and early work relied on parametric models such as GARCH and its extensions, which capture volatility clustering but can struggle with complex nonlinear dynamics and regime changes. Recent research has shown that hybrid and machine learning approaches can enhance volatility prediction; for example, Kakade et al. propose a hybrid GARCH–LSTM ensemble (SET-LSTM) for commodity markets and report significantly lower RMSE and MAPE than standalone econometric or neural models, highlighting the benefit of combining heterogeneous variance signals. In equity and cryptocurrency markets, pooled and cross-asset models have been used to exploit volatility commonality across many instruments, and Djanga et al. and Zhang et al. document that LightGBM, neural networks, and augmented machine learning models trained on multiple assets often outperform asset-specific HAR or ARIMA benchmarks, especially when evaluated with robust loss functions such as QLIKE. Within the crypto domain, Naresh and others show that deep learning architectures such as Transformers can capture long-range dependencies in high-frequency data and deliver superior risk-sensitive metrics, including Sharpe ratio and value-at-risk, relative to simpler models, while separately, volatility-managed strategies and GARCH-MIDAS models demonstrate that adjusting portfolio leverage in proportion to predicted volatility can improve long-horizon risk-adjusted returns, linking volatility forecasts directly to economic outcomes. Despite these advances, there is still limited work that integrates multi-asset volatility forecasting, spike classification, and explicit trading backtests into a single reproducible pipeline using openly documented data

and code, and this thesis builds on the literature by combining ideas from HAR-style realized volatility modeling, pooled machine learning, and volatility-managed portfolios in the context of daily crypto risk.

### **1.3 Motivation of the Research**

Cryptocurrency investors, exchanges, and regulators require timely and reliable measures of short-horizon volatility to manage margin requirements, assess systemic risk, and design hedging strategies, yet the 24/7 nature of crypto trading, frequent structural breaks, and speculative participation make conventional risk models less effective, motivating the exploration of richer feature sets and more flexible learning algorithms. Most existing crypto volatility studies either focus on a single asset such as Bitcoin, rely on proprietary intraday data, or evaluate models solely using statistical loss functions without linking predictions to trading performance, which limits their practical usefulness for portfolio managers who care about risk-adjusted returns, drawdowns, and robustness across multiple coins and market regimes. The motivation for this research is therefore threefold: first, to construct a transparent, modular pipeline that can be reused and extended by other researchers working with cryptocurrency volatility data; second, to examine whether machine learning models trained on a unified multi-coin dataset can achieve meaningful predictive accuracy for both realized volatility and volatility spikes at the daily horizon; and third, to test whether these forecasts translate into an economically valuable volatility-managed trading strategy with positive risk-adjusted performance.

## **1.3 Problem Statements**

Cryptocurrency markets exhibit extreme and rapidly changing volatility, yet there is no widely accepted, open, and reproducible framework for forecasting daily realized volatility across multiple coins and translating these forecasts into systematic trading decisions. Traditional econometric models often underperform in the presence of nonlinear dynamics, cross-asset commonality, and regime shifts, while many machine learning studies stop at predictive accuracy and do not quantify economic value. Consequently, investors lack clear evidence on whether machine learning-based volatility models, trained on realistic multi-asset crypto data, can deliver both statistically significant forecasts and improved risk-adjusted returns when embedded in a volatility-managed strategy. The core problem addressed in this thesis is to design, implement, and evaluate such an integrated framework and to assess its strengths and limitations relative to simple baselines.

## **1.4 Research Questions**

The research questions were:

Q1: To what extent can machine learning models forecast next-day realized volatility for major cryptocurrencies using daily features derived from hourly OHLCV data and technical indicators?

Q2: How accurately can these models classify extreme volatility spike events defined by rolling high-percentile thresholds of realized volatility?

Q3: Do volatility forecasts, when used to scale portfolio exposure in an inverse-volatility trading strategy, improve risk-adjusted performance compared with unscaled directional strategies and simple baselines?

Q4: How does predictive and trading performance vary across individual coins and market regimes, and what does this reveal about the limits of volatility predictability in cryptocurrency markets?

## 1.5 Research Objective

The overall objective of this research is to develop and empirically evaluate an end-to-end machine learning pipeline for cryptocurrency volatility forecasting and volatility-managed trading. This high-level objective can be decomposed into several specific objectives: design and implement a modular Python pipeline that performs data loading, cleaning, feature engineering, temporal splitting, model training, evaluation, diagnostics, and reporting for multi-asset crypto volatility data; construct a rich feature set including HAR-style lagged realized volatility, rolling statistics, session-specific volatility, market-relative indicators, and cyclical time features derived from hourly OHLCV prices aggregated to daily frequency; train and compare multiple regression and classification models, including HAR-based baselines and tree-based ensemble methods, to forecast next-day realized volatility and spike indicators using temporally consistent train-validation-test splits and optional walk-forward evaluation; quantify model performance using standard metrics such as RMSE, MAE,  $R^2$ ,  $R^2$ , F1 score, and ROC-AUC, and analyze heterogeneity across coins and volatility regimes; and translate the best volatility forecasts into a volatility-managed trading strategy that adjusts leverage

inversely with predicted volatility, assessing its cumulative return, Sharpe ratio, maximum drawdown, and win rate on out-of-sample data.

## **1.6 Research Scope**

This study focuses on six liquid cryptocurrencies Bitcoin (BTC), Ether (ETH), Cardano (ADA), Dogecoin (DOGE), Ethereum Classic (ETC), and Chainlink (LINK) which collectively represent a substantial share of market capitalization and trading volume. The dataset comprises hourly OHLCV prices with pre-computed technical indicators from 1 January 2018 to 14 November 2025, which are aggregated to daily frequency for modeling. The forecasting horizon is restricted to one day ahead, targeting daily realized volatility and binary spike labels computed from squared intraday returns, and the feature space is limited to price-based and volatility-based indicators derived from market data, while on-chain metrics, order-book depth, and macroeconomic variables are beyond the scope of this thesis. Model families considered include a HAR-style linear baseline and a set of scikit-learn models such as Ridge, Elastic Net, Random Forest, Gradient Boosting, and HistGradientBoosting for both regression and classification tasks, whereas more advanced deep learning architectures (e.g., LSTM or Transformer networks) and alternative risk models (e.g., GARCH-MIDAS) are discussed in the literature review but are not implemented in the current pipeline. The trading strategy analyzed is a long–short volatility-managed strategy that combines a simple moving average crossover direction signal with inverse-volatility position sizing and risk filters, and transaction costs, slippage, and market impact are not explicitly modeled, so trading results should be interpreted as frictionless benchmarks rather than directly implementable strategies.

## **1.7 Summary**

This chapter has introduced the context, motivation, and problem statement for a thesis on machine learning-based cryptocurrency volatility forecasting and volatility-managed trading.

The discussion highlighted the importance of volatility for risk management, reviewed relevant advances in econometric and machine learning models, and identified the need for an integrated, reproducible pipeline that links volatility prediction to economic value across multiple cryptocurrencies. The chapter formulated specific research questions and objectives focused on building such a pipeline, assessing predictive accuracy for both volatility levels and spikes, and evaluating a volatility-managed trading strategy in terms of risk-adjusted returns.

It also clarified the scope of the study six major cryptocurrencies, daily forecasting horizon, selected model families, and a frictionless trading benchmark and outlined the structure of the remaining thesis chapters.

# CHAPTER 2 LITERATURE REVIEW

## 2.1 Introduction

The literature on volatility forecasting has evolved from early econometric models toward hybrid and machine learning approaches that better capture nonlinear dynamics, regime shifts, and cross-asset effects. Traditional frameworks such as GARCH, HAR, and GARCH-MIDAS established that volatility is highly persistent and partly driven by macroeconomic fundamentals, but they often struggle with structural breaks and complex dependence structures, especially in high-frequency or crisis periods. More recent work introduces hybrid GARCH–LSTM ensembles, Support Vector Regression, and deep neural architectures (MLP, LSTM, CNN-LSTM, Transformers) that operate on realised-volatility measures and intraday data, generally achieving lower forecast errors and better risk-sensitive metrics than classical baselines when rich feature sets are available. Within cryptocurrency markets, studies highlight strong volatility commonality across coins, show that pooled and cross-asset models can outperform single-asset specifications, and increasingly link forecast quality to trading performance and drawdown control through volatility-managed strategies. Complementary research on spike prediction, probabilistic forecasting, and explainable AI further demonstrates that combining advanced models, alternative data (on-chain, order-book, sentiment), and interpretable diagnostics can improve both statistical accuracy and practical usability for investors and risk managers.

## 2.2 Previous literature

Previous studies on volatility forecasting span traditional econometric models, hybrid architectures, and modern machine learning approaches across equities, commodities, foreign exchange, and cryptocurrencies. Volatility forecasting in commodity markets has been advanced by a hybrid SET-LSTM ensemble that feeds multiple GARCH forecasts into an LSTM network, significantly reducing RMSE, MAE and MAPE versus standalone econometric or neural models while increasing model complexity and sensitivity to hyperparameters. (Kakade, 2022)

Foundational evidence on stock-market volatility shows that volatility is highly time-varying, surges during crises such as the Great Depression, and is only partly explained by macro and leverage variables, underscoring the limits of simple linear time-series models. (Schwert, 1989)

Macro-driven GARCH-MIDAS models separate short- and long-run components of variance and demonstrate that low-frequency macro variables can explain between 10% and 35% of one-day-ahead volatility, outperforming standard GARCH and spline-GARCH at longer horizons but still struggling with structural breaks. (Engle, 2008, August)

ARIMA-based approaches, illustrated via a Box–Jenkins workflow for Indian stock indices, achieve low forecast errors and stationary residuals yet remain univariate and unable to capture nonlinear dynamics or high-frequency microstructure effects. (Idrees, 2019)

Machine-learning models trained on intraday S&P 500 realized volatility show that universal and augmented neural networks, especially MLPs and LSTMs, can outperform HAR-D and ARIMA under MSE and QLIKE losses when exploiting cross-sectional volatility commonality and market-level realized volatility. (Zhang, 2024)

A comprehensive comparison of statistical and ML methods for several major cryptocurrencies finds that while advanced models such as LSTM, Random Forest and SVR sometimes beat HAR and GARCH, simple linear models like HAR and ridge regression remain strong benchmarks across assets and horizons. (Dudek G. F., 2024)

(Extending this work to probabilistic forecasting, stacked quantile and distributional models convert heterogeneous point forecasts into calibrated predictive distributions, improving CRPS and tail risk metrics for cryptocurrency volatility. (Dudek G. O., 2025)

A broad survey of volatility-forecasting research concludes that machine learning generally outperforms classical approaches when rich feature sets, nonlinearities and regime shifts are present, but emphasizes that rolling evaluation, robustness checks and interpretability are at least as important as the choice of algorithm. (Han, 2025)

High-frequency volatility for cryptocurrencies and traditional currencies has also been forecast using Support Vector Regression, which in some settings delivers lower errors than HAR/GARCH, although performance is highly asset- and horizon-dependent and requires substantial data and tuning. (Peng, 2018)

For Bitcoin specifically, models that explicitly incorporate jumps and structural breaks improve realized-volatility forecasts relative to continuous or simple GARCH

specifications, indicating the importance of regime-shift modelling in crypto markets. (Shen, 2020)

Studies focusing on Bitcoin realized volatility compare HAR, GARCH and deep learning with image-based CNN-LSTM and LSTM architectures built on Markov Transition Field representations, finding that image-enhanced CNN-LSTM models offer the best short-horizon performance while remaining univariate and omitting macro or on-chain drivers. (Huang, 2024)

Systematic deep-learning experiments on high-frequency crypto data show that convolution-augmented Transformer architectures can achieve the lowest MSE and superior risk-sensitive metrics such as Sharpe ratio and value-at-risk, though on relatively short samples and without strong econometric baselines. (Naresh, 2025)

Cross-asset cryptocurrency volatility studies using 5-minute data across dozens of Binance pairs report strong realized-volatility commonality and show that pooled and clustered LightGBM or LSTM models significantly reduce QLIKE relative to single-asset HAR, with simple trading rules demonstrating economic value. (Djanga, 2025)

A related extended version of this work further investigates universal versus clustered training schemes and confirms that cross-asset pooling generally enhances forecast accuracy and trading performance, especially for smaller, more volatile coins. (Djanga, 2025)

Other cryptocurrency studies evaluate a wide menu of statistical and ML models and conclude that there is no uniformly best method, with rankings varying by coin, horizon and loss function, but with linear baselines often remaining competitive. (Catania, 2022)

Hybrid econometric–ML frameworks that combine ARMA–GARCH components with neural networks, such as SET-LSTM architectures, show that feeding structured volatility forecasts into LSTMs can materially improve accuracy for commodity and FX markets, at the expense of interpretability and ease of implementation. (Kakade, 2022)

Machine-learning approaches that jointly model internal determinants lagged volatility, returns, liquidity and external determinants such as macroeconomic and policy-uncertainty indices find that ensemble trees and LSTMs often outperform GARCH-type models, with SHAP analysis indicating that internal volatility and trading activity dominate feature importance. A working-paper version of this framework provides additional robustness checks and confirms that model performance is sensitive to feature engineering choices, sample period and evaluation protocol, reinforcing the need for transparent, reproducible pipelines. (Wang, 2023)

Research on Bitcoin exchange-rate prediction using economic and technology determinants shows that augmenting technical indicators with macro and technology variables improves forecast accuracy, and suggests analogous feature-engineering strategies for volatility modelling. (Chen, 2021)

Explainable-AI work demonstrates how tools such as SHAP values and saliency maps can reveal which technical, sentiment and fundamental indicators drive cryptocurrency risk

predictions, helping investors interpret black-box models even when trading performance is not the primary focus. (Fior, 2022)

Studies on volatility spikes rather than continuous variance show that Synthesizer Transformer models using on-chain metrics, exchange flows and whale-transaction data can significantly outperform GARCH, HAR and standard Transformer architectures in predicting next-day Bitcoin volatility spikes and can reduce portfolio drawdowns in simple backtests. (Herremans, 2025)

An extended version of the Synthesizer Transformer spike-prediction study provides further architectural details and robustness analyses, but reaches the same central conclusion that rich on-chain/whale data and attention-based models are highly effective for crypto spike detection. (Herremans, 2025)

Portfolio-level work on volatility-managed strategies shows that scaling factor or portfolio exposures by the inverse of recent realized variance delivers sizable alphas, higher Sharpe ratios and lower drawdowns across equity and FX factors, motivating the integration of volatility forecasts into systematic trading rules. (Moreira, 2017)

Additional robustness appendices to multi-asset crypto volatility studies confirm that forecast rankings are generally stable across alternative rolling windows, evaluation metrics and transaction-cost assumptions, though the edge of complex ML models narrows once costs and realistic constraints are included. (Dudek G. F., 2024)

**Table 2.1 Literature Review**

Title	Author(s)	Year	Key findings
Forecasting cryptocurrency volatility	Catania et al.	2022	Shows that realized-volatility models and simple extensions (e.g., HAR-type specifications) provide competitive forecasts for major cryptocurrencies, and that gains from more complex models depend on asset, horizon, and evaluation metric.
Machine learning model for Bitcoin exchange rate prediction using economic and technology determinants	Chen et al.	2021	Demonstrates that adding economic and technology determinants to technical features significantly improves ML-based Bitcoin exchange-rate forecasts, suggesting similar enriched feature sets can help volatility prediction.
Cryptocurrency Volatility Forecasting with Applications in Trading	Djanga et al.	2025	Finds strong realized-volatility commonality across many Binance USDT pairs and shows that pooled LightGBM/LSTM models outperform single-asset HAR benchmarks, with forecast improvements translating into profitable trading strategies.
Forecasting cryptocurrencies volatility using statistical and machine learning methods: A comparative study	Dudek & Fiszeder	2024	Provides a large-scale comparison of statistical and ML models for several cryptocurrencies, concluding that no single model dominates and that simple linear benchmarks often rival or beat more complex ML approaches across horizons.
Probabilistic Forecasting Cryptocurrencies Volatility: From Point to Quantile Forecasts	Dudek & Olearnik	2025	Extends point-forecast models to full predictive distributions using quantile and distributional methods, improving CRPS and tail-risk performance for crypto volatility and enabling better risk management.

<b>Title</b>	<b>Author(s)</b>	<b>Year</b>	<b>Key findings</b>
On the economic sources of stock market volatility	Engle, Ghysels & Sohn	2008	Introduces GARCH-MIDAS to decompose volatility into short- and long-run components, showing that macro variables explain 10–35% of one-day-ahead variance and outperform standard GARCH at longer horizons.
Leveraging explainable AI to support cryptocurrency investors	Fior et al.	2022	Applies SHAP and other XAI tools to crypto ML models, showing how explanations identify influential technical and sentiment features and can increase investor understanding and trust, even without a pure trading focus.
Can machine learning models better volatility forecasting? A combined method	Han et al.	2025	Reviews and tests a range of ML models for volatility, concluding that ML can outperform classical approaches when nonlinearities and rich features are present, but emphasizing the importance of rolling evaluation and robustness checks.
Forecasting Bitcoin volatility spikes from whale transactions and CryptoQuant data using Synthesizer Transformer models	Herremans & Low	2025	Uses Synthesizer Transformer models with on-chain and whale-transaction data to predict next-day Bitcoin volatility spikes, outperforming GARCH, HAR and standard Transformers and reducing drawdowns in simple backtests.
Forecasting Bitcoin volatility using machine learning techniques	Huang et al.	2024	Shows that LSTM and CNN-LSTM models built on Markov Transition Field images of high-frequency realized volatility beat HAR and GARCH benchmarks, particularly at short horizons, though the setup remains univariate.
A prediction approach for stock market volatility based on time series data	Idrees, Alam & Agarwal	2019	Presents a Box–Jenkins ARIMA framework for Indian stock indices, achieving low forecast errors and stationary residuals, while acknowledging limits of univariate linear

<b>Title</b>	<b>Author(s)</b>	<b>Year</b>	<b>Key findings</b>
			models for nonlinear and high-frequency dynamics.
Forecasting commodity market returns volatility: a hybrid ensemble learning GARCH-LSTM based approach	Kakade et al.	2022	Proposes a SET-LSTM ensemble that feeds multiple GARCH outputs into an LSTM for metal-market volatility, significantly reducing RMSE, MAE and MAPE compared with standalone GARCH or LSTM at the cost of higher complexity.
Volatility-managed portfolios	Moreira & Muir	2017	Shows that scaling factor and market portfolios by the inverse of recent realized variance produces sizeable alphas, higher Sharpe ratios and lower drawdowns, demonstrating the economic value of simple variance-timing rules.
Deep Learning Models for High-Frequency Cryptocurrency Volatility Forecasting: A Comparative Study of MLP, LSTM, and Transformer Architectures	Naresh	2025	Finds that Transformer and especially convolution-augmented Transformer models deliver the lowest MSE and best risk-sensitive metrics (Sharpe, VaR) for high-frequency crypto realized volatility, though over a relatively short sample.
The best of two worlds: Forecasting high frequency volatility for cryptocurrencies and traditional currencies with Support Vector Regression	Peng et al.	2018	Shows that SVR can outperform some HAR/GARCH benchmarks for high-frequency volatility of both crypto and FX, but that performance is asset-specific and data-intensive.
Why does stock market volatility change over time?	Schwert	1989	Documents that stock-market volatility is highly time-varying, spikes during crises, and is only modestly explained by macro and leverage variables, underscoring limits of simple autoregressive models.

Title	Author(s)	Year	Key findings
Forecasting the volatility of Bitcoin: The importance of jumps and structural breaks	Shen, Urquhart & Wang	2020	Demonstrates that explicitly modelling jumps and structural breaks significantly improves Bitcoin volatility forecasts compared with models that ignore regime changes.
Machine learning approaches to forecasting cryptocurrency volatility: Considering internal and external determinants	Wang et al.	2023	Shows that Random Forest and LSTM models using both internal (lagged volatility, volume) and external (uncertainty indices) features can outperform GARCH-type models, with SHAP analysis indicating internal determinants are most important.
Volatility forecasting with machine learning and intraday commonality	Zhang et al.	2024	Finds that universal and augmented neural networks (MLP, LSTM) trained on intraday realized volatility for S&P 500 stocks outperform HAR-D and ARIMA, especially when exploiting cross-sectional volatility commonality and market-level volatility.

### 2.3 Summary

The reviewed literature shows a clear evolution from early time-series models and macro-driven GARCH-MIDAS specifications toward hybrid and machine learning approaches that leverage realised volatility, cross-asset commonality and high-frequency data. Studies in equities, commodities and cryptocurrencies consistently find that neural networks, ensemble trees and hybrid architectures can outperform traditional benchmarks in forecasting variance or detecting volatility spikes, particularly when models pool information across assets and horizons. At the same time, existing work reveals several important gaps that motivate the present thesis: many crypto-focused studies concentrate on single assets, short samples or proprietary intraday datasets, with limited attention to

multi-asset daily volatility forecasting, combined regression-and-classification tasks, or systematic links between forecast quality and risk-adjusted trading performance at the portfolio level. Moreover, while volatility-managed portfolios demonstrate that scaling exposure by inverse realised variance can improve Sharpe ratios, there is little research that integrates machine learning-based volatility forecasts with such strategies in cryptocurrency markets and evaluates performance across different coins and regimes. Finally, issues of interpretability, robustness and reproducibility remain, as many studies rely on complex black-box models and do not release code or provide modular pipelines that others can extend. This thesis contributes to the literature by building a transparent, multi-asset pipeline that forecasts next-day realised volatility and volatility spikes for several major cryptocurrencies using daily features derived from hourly data, and by embedding these forecasts into a volatility-managed trading strategy inspired by prior portfolio work, thereby responding directly to the gaps identified in previous research regarding multi-asset crypto volatility prediction, spike classification, and the economic value of machine learning forecasts for daily risk-adjusted trading decisions.

# CHAPTER 3 RESEARCH METHODOLOGY

## 3.1 Introduction

This chapter describes the methodology used to forecast next-day realized volatility and volatility spikes for six major cryptocurrencies, and to evaluate a volatility-managed trading strategy.

The research design follows a structured machine learning pipeline that encompasses data collection, preprocessing, feature engineering, temporal splitting, model training for both regression and classification tasks, evaluation, and trading backtesting.

The core modeling approach rests on supervised learning with chronological train-validation-test splits to prevent look-ahead bias, multiple model families to explore different functional forms, and comprehensive diagnostics to assess performance across coins and market regimes.

This chapter focuses on the data processing and the statistical-learning models, reserving implementation details for Chapter 4 and empirical results for Chapter 5.

# End-to-End ML Pipeline for Cryptocurrency Volatility Forecasting

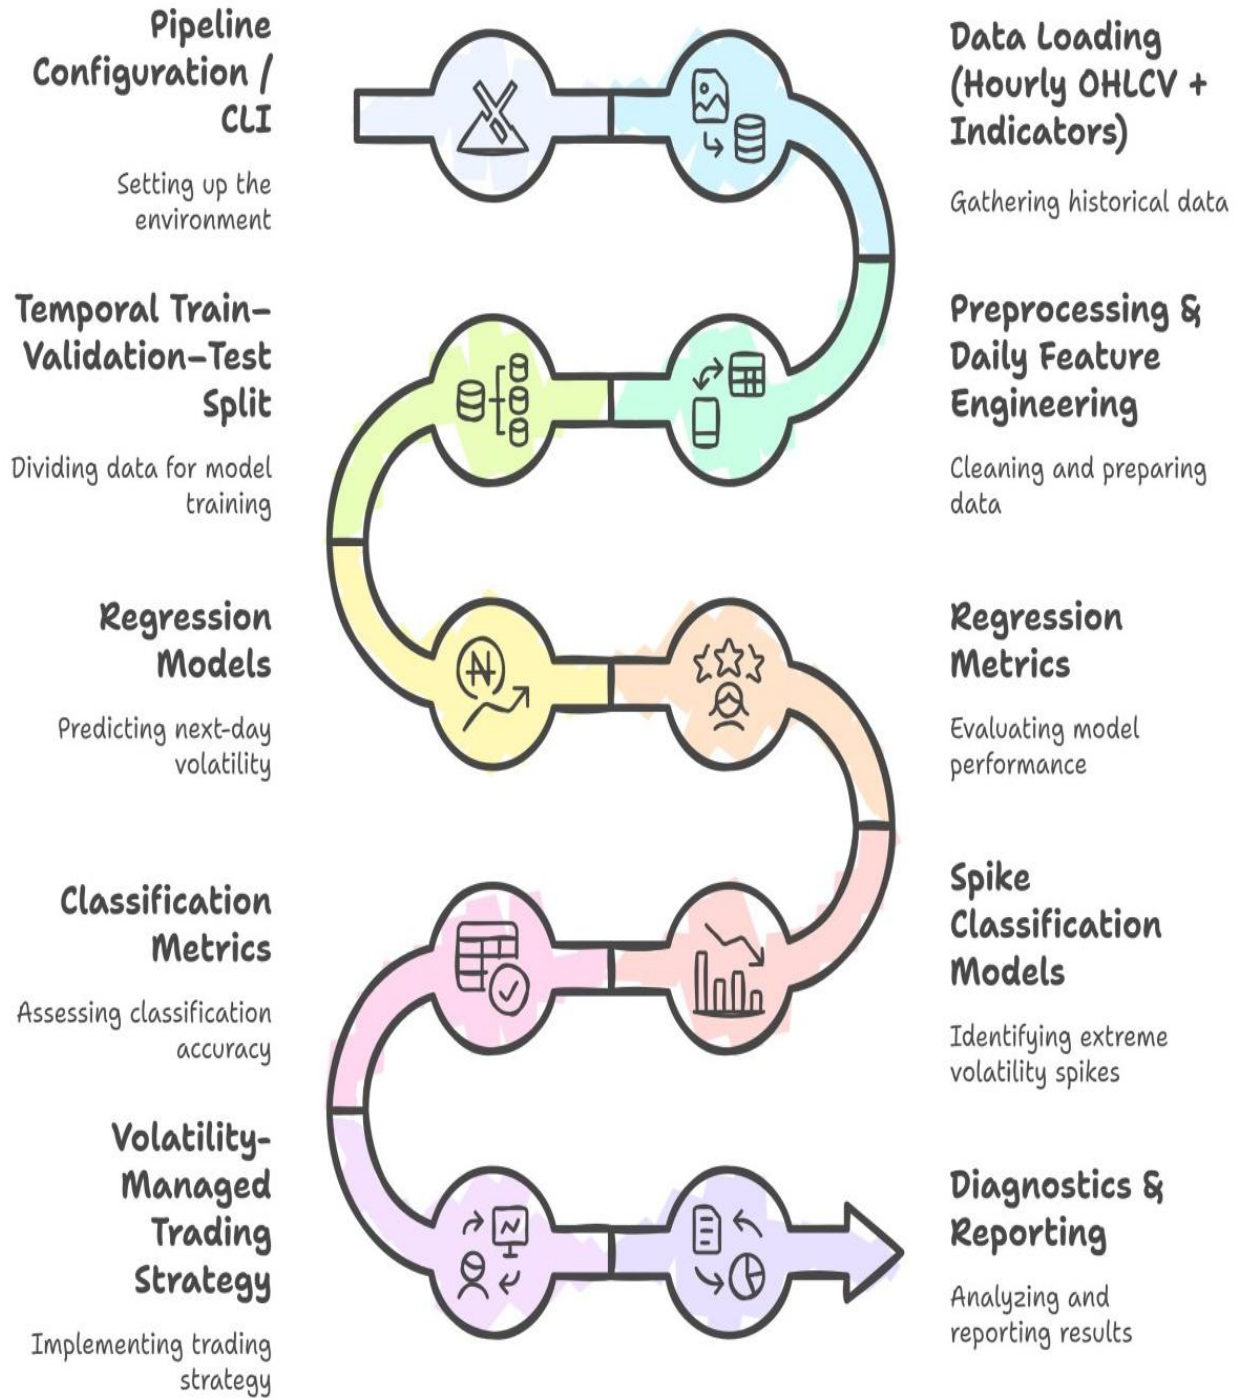


Figure 3.1: Workflow Architecture

This figure illustrates the end-to-end cryptocurrency volatility prediction and trading pipeline used in this thesis. The architecture begins with configurable data loading and preprocessing of hourly OHLCV data, followed by feature engineering to construct daily realized volatility, lags, rolling statistics, market-level and time-based features. These features feed parallel regression and classification trainers, whose best models generate forecasts that drive a volatility-managed trading strategy, while a diagnostics and report-generation layer aggregates performance, feature importance and trading metrics into structured outputs.

### **3.2 Data collection**

The research uses an hourly cryptocurrency dataset covering six coins Bitcoin (BTC), Ether (ETH), Cardano (ADA), Dogecoin (DOGE), Ethereum Classic (ETC) and Chainlink (LINK) from 1 January 2018 to 14 November 2025. The raw data contains Open, High, Low, Close prices, along with pre-computed technical indicators such as moving averages (12-, 24-, 48- and 168-hour), rolling volatility measures, and hourly return statistics. A total of 413,996 hourly rows were loaded, covering multiple coins over approximately seven years of trading. The dataset was sourced from a consolidated CSV file, and the data loader validates required columns, converts timestamps to datetime format, sorts records by coin and date, and drops rows with missing critical fields such as price or timestamp.




Figure 3.2: Timeline bar chart showing data coverage per cryptocurrency (2018–2025).

### 3.3 Data preprocessing

Data preprocessing ensures the hourly dataset is clean, consistent and ready for aggregation into daily frequency. The preprocessing pipeline removes duplicate timestamps within each coin, filters out any non-positive prices that would indicate data errors, forward-fills missing technical indicators separately per coin to preserve continuity, clips extreme daily returns and volatility values to the 1st–99th percentile range to mitigate outliers, and recomputes hourly log returns where necessary.

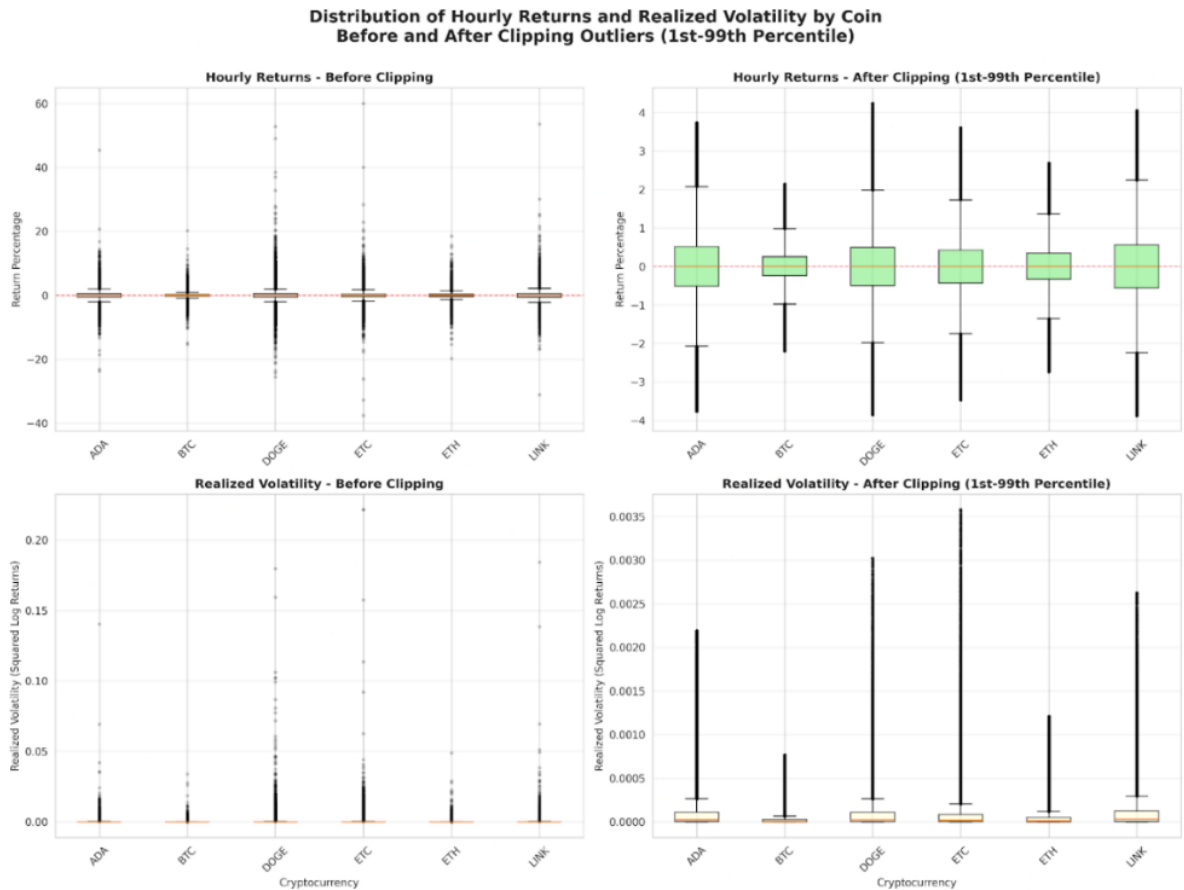
Realized volatility for each day is computed as the sum of squared intraday log returns:

$$RV_t = \sum_{i=1}^N r_{t,i}^2$$

where  $r_{t,i}$  is the log return in period  $i$  of day  $t$ , and  $N$  is the number of hourly observations within that day.

To capture intraday patterns, session-specific realized volatility is computed separately for Asia (hours 0–8 UTC), Europe (hours 8–16), and US (hours 16–24) trading windows, following the approach in Zhang et al. and Djanga et al.

This preprocessing step outputs a clean hourly DataFrame ready for feature engineering and aggregation.



*Figure 3.3: Box plot showing the distribution of hourly returns and realized volatility by coin, before and after clipping outliers at the 1st–99th percentile*

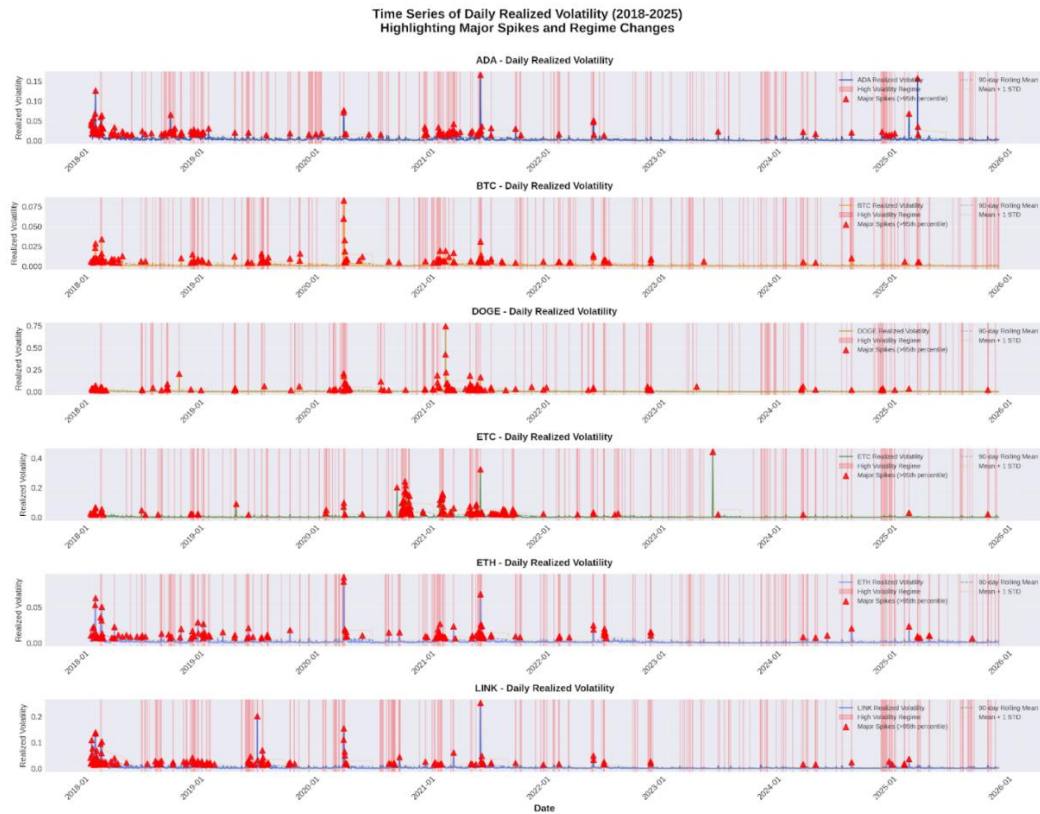


Figure 3.4: Time series plot of daily realized volatility (aggregated from hourly data) for all six cryptocurrencies, highlighting major spikes and regime changes from 2018–2025.

### 3.4 Regression models

Regression models predict next-day realized volatility, measured as the sum of squared intraday returns, which serves as a proxy for daily variance.

A suite of models ranging from a simple HAR baseline to advanced gradient-boosting ensembles is trained to capture both linear and nonlinear volatility dynamics.

#### 3.4.1 HAR baseline

The Heterogeneous Autoregressive (HAR) model, introduced by Corsi (2009), is a standard econometric benchmark for realized-volatility forecasting.

The HAR specification is a linear regression on lagged realized-volatility features at multiple horizons:

$$RV_{t+1} = \beta_0 + \beta_1 RV_{t-1} + \beta_3 \overline{RV}_{t-2:t} + \beta_5 \overline{RV}_{t-4:t} + \beta_{10} \overline{RV}_{t-9:t} + \varepsilon_{t+1}$$

where  $\overline{RV}_{t-k:t}$  denotes the average realized volatility over days  $t-k$  to  $t$ .

In the pipeline, the HAR model is implemented as a simple linear regressor trained only on the four HAR features (rvlag1, rvlag3, rvlag5, rvlag10), with no regularization or scaling.

This provides a transparent baseline that can be compared against more complex models to quantify the benefit of richer features and nonlinear algorithms.

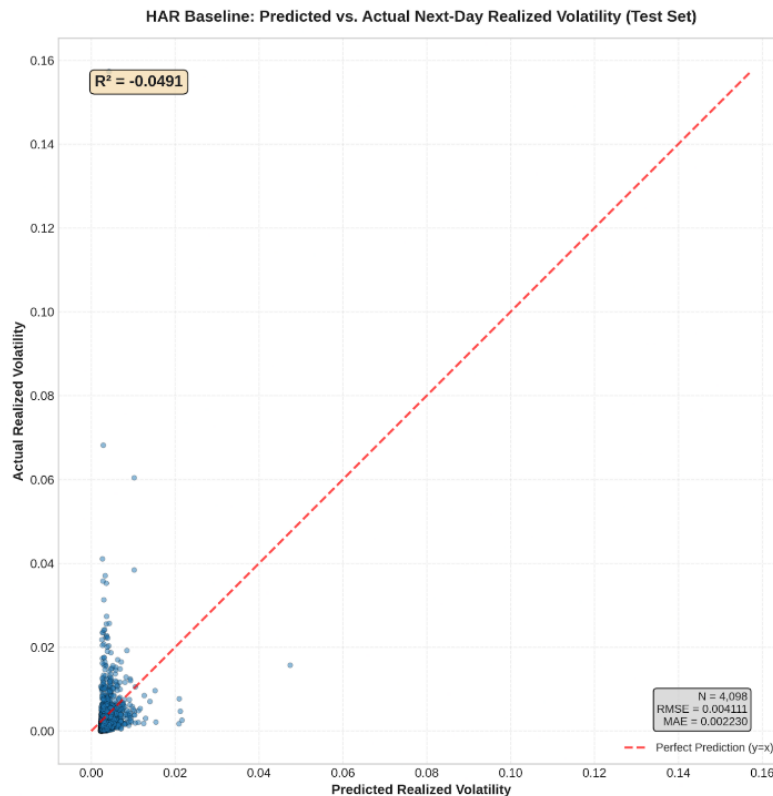


Figure 3.5: Scatter plot showing predicted vs. actual next-day realized volatility for the HAR baseline on the test set, with a diagonal reference line and  $R^2$  annotation.

### 3.4.2 Ridge and Elastic Net

Ridge and Elastic Net are linear regression models with regularization that penalize large coefficients, which helps reduce overfitting when many features are available. Ridge regression applies an L2 penalty controlled by hyperparameter  $\alpha$ , while Elastic Net combines L1 and L2 penalties via  $\alpha$  and  $\ell_1$  ratio.

Both models are trained using scikit-learn Pipelines that first apply StandardScaler normalization to features, then fit the regularized regression.

Hyperparameters are set as follows: Ridge uses  $\alpha=0.5$ , and Elastic Net uses  $\alpha=0.001$ ,  $\ell_1$  ratio = 0.2, and  $\text{max\_iter} = 5000$ .

These linear models can capture additive relationships among volatility features but cannot model interactions or nonlinear patterns without explicit feature engineering.




Figure 3.6: Coefficient plot (bar chart) showing the top 15 most important features for Ridge and Elastic Net regression, with coefficients on the x-axis and feature names on the y-axis.

### 3.4.3 Random Forest

Random Forest is an ensemble method that constructs multiple decision trees on bootstrap samples and averages their predictions, which reduces variance and improves generalization.

The Random Forest regressor used in this study has 500 trees, a minimum of 3 samples per leaf to prevent overfitting, and unrestricted maximum depth to allow deep trees that can capture complex interactions.

Random Forest can model nonlinear and interactive effects among features without requiring manual interaction terms, and it is robust to outliers and irrelevant variables. However, it may struggle to extrapolate beyond the training range and can be computationally expensive.

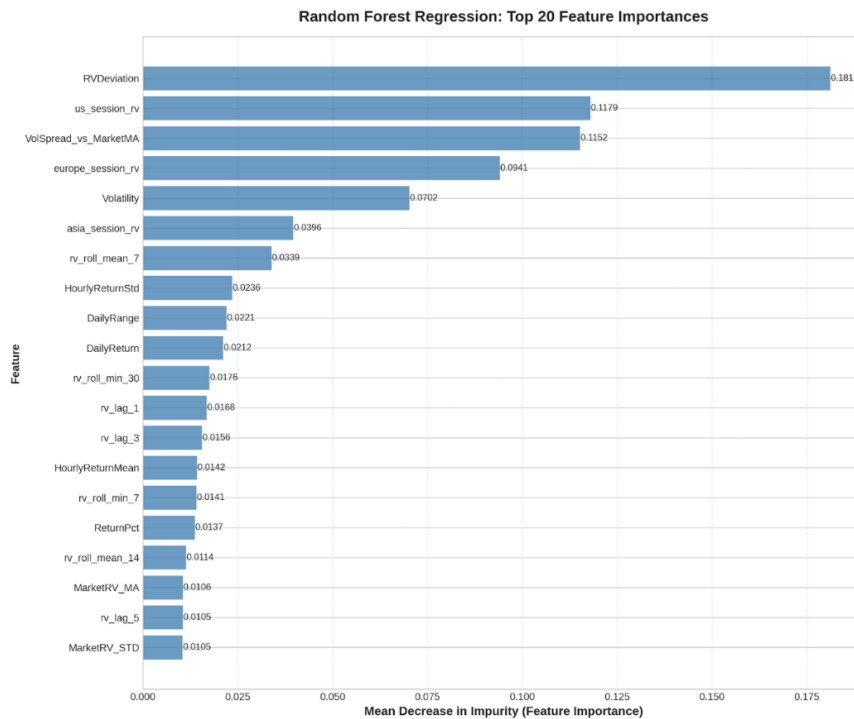


Figure 3.7: Feature importance bar chart for Random Forest regression, showing the top 20 features ranked by mean decrease in impurity.

### 3.4.4 Gradient Boosting

Gradient Boosting builds an additive model by sequentially fitting shallow trees to the residuals of the previous iteration, thereby minimizing a differentiable loss function here, mean squared error.

The Gradient Boosting regressor is configured with maximum depth = 3, learning rate = 0.05, and default number of estimators.

Shallow trees combined with a low learning rate produce a strong regularization effect and often yield more interpretable models than deep forests.

Gradient Boosting tends to perform well on structured tabular data such as financial features, but it requires careful tuning to avoid overfitting.

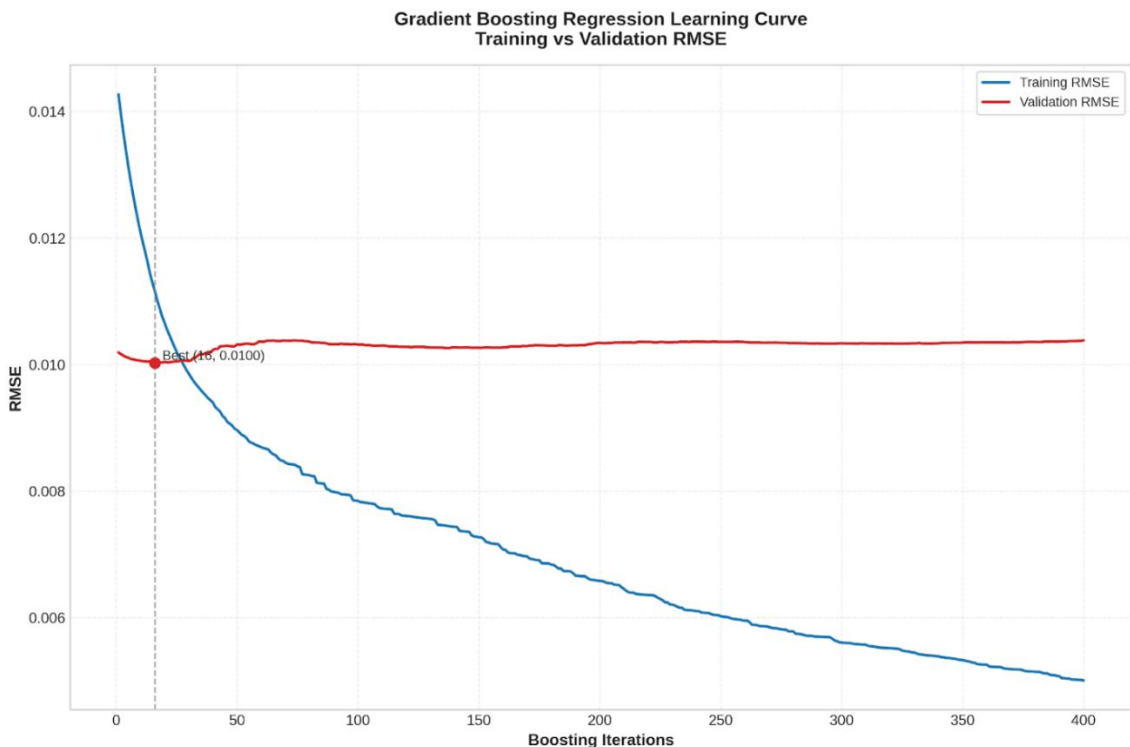


Figure 3.8: Learning curve plot showing training and validation RMSE across boosting iterations for Gradient Boosting regression.

### 3.4.5 HistGradientBoosting

HistGradientBoosting is a histogram-based variant of gradient boosting implemented in scikit-learn, which bins continuous features into discrete buckets to accelerate training and enable handling of larger datasets. It is configured with maximum depth = 8, learning rate = 0.08, and supports native handling of missing values and categorical features.

HistGradientBoosting often delivers the best trade-off between speed and accuracy for medium-to-large tabular datasets, and it achieved the best regression performance in preliminary experiments.

The deeper trees (compared to standard Gradient Boosting) allow the model to capture higher-order interactions at the cost of slightly higher variance.

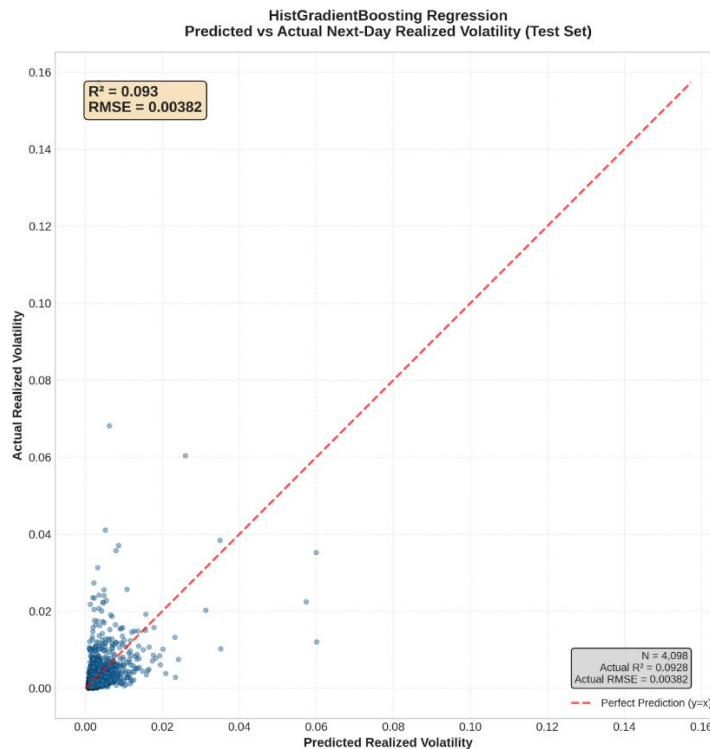


Figure 3.9: Predicted vs. actual scatter plot for HistGradientBoosting regression on the test set, annotated with test  $R^2=0.093$  and  $RMSE = 0.00382$ .

### 3.5 Classification models

Classification models predict binary volatility-spike labels, defined as whether next-day realized volatility exceeds the rolling 90th percentile within a 60-day window. This task addresses the need to identify extreme-risk events that are particularly relevant for risk management and capital allocation.

Because spike events are rare (459 out of 4,098 test samples), the dataset is imbalanced, and models use balanced class weights or subsampling to prevent trivial solutions that always predict the majority class.

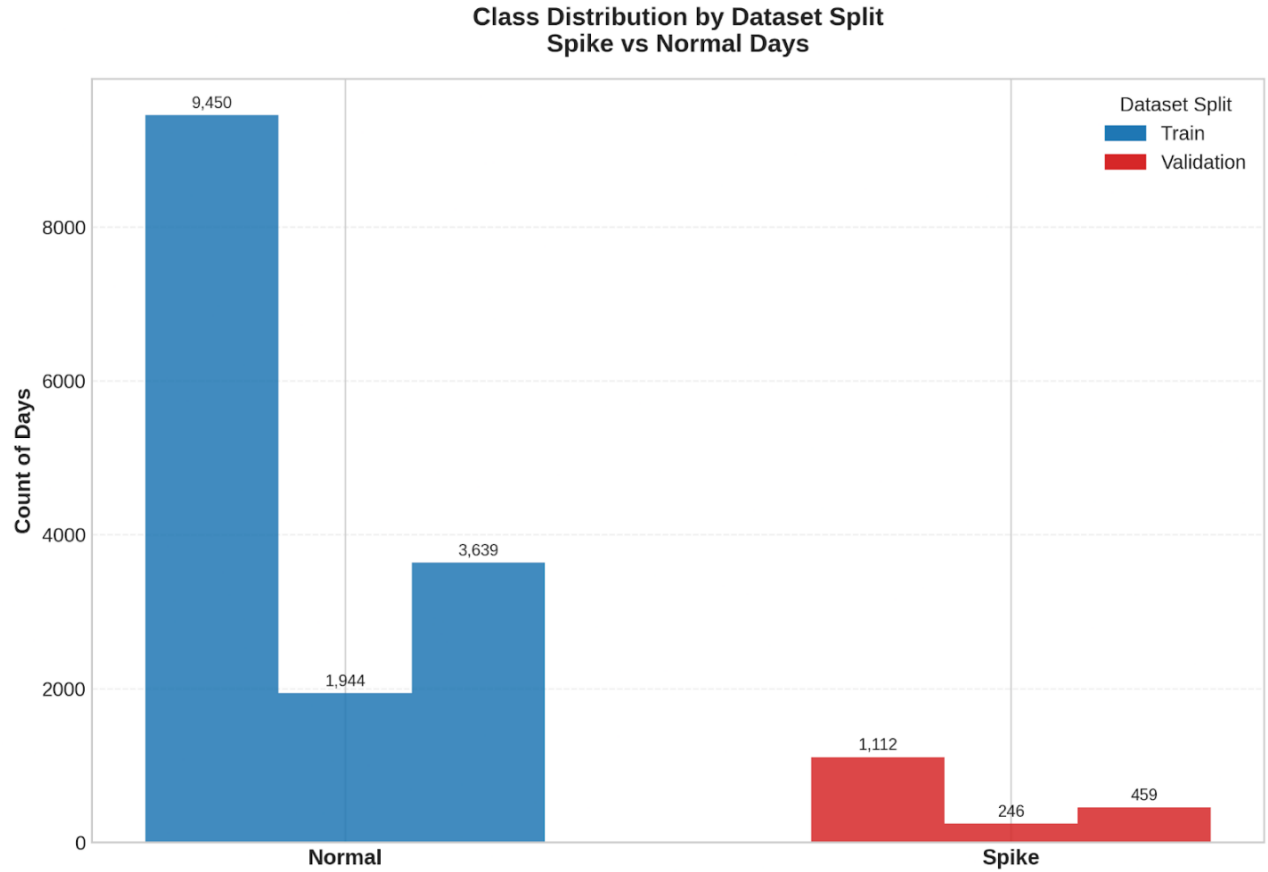


Figure 3.10: Bar chart showing the class distribution (count of spike vs. normal days) across the train, validation, and test sets.

### 3.5.1 Logistic Regression

Logistic Regression is a linear probabilistic classifier that models the log-odds of the spike label as a linear function of the features.

It is trained with balanced class weights, which upweight the minority spike class during training, and StandardScaler normalization is applied to features.

Logistic Regression serves as a transparent classification baseline that provides probability estimates and coefficients interpretable as feature importances.

However, it cannot capture nonlinear decision boundaries without manual feature transformations.

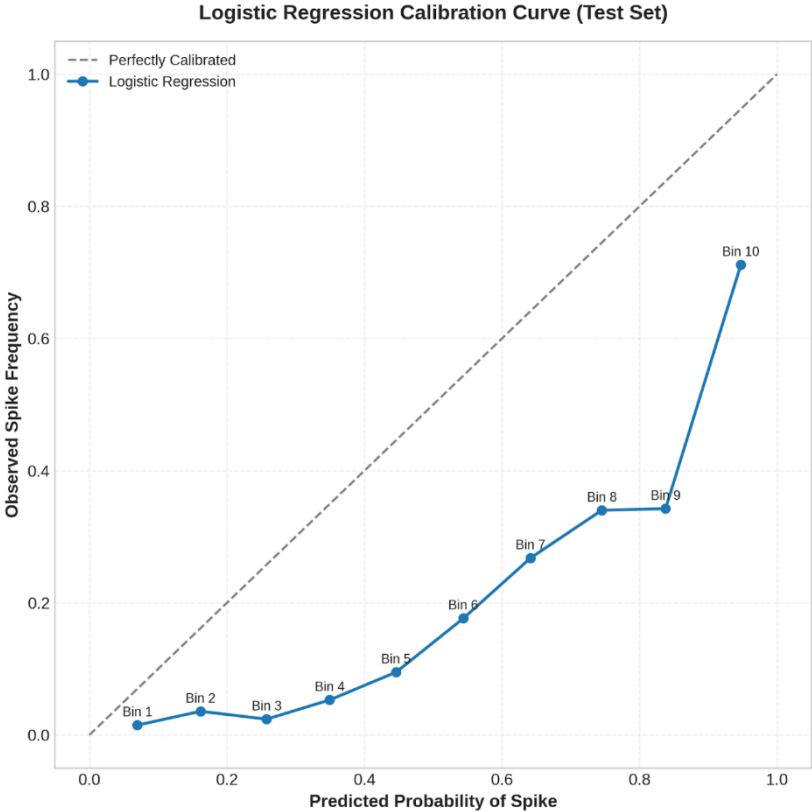


Figure 3.11: Calibration curve (reliability diagram) showing predicted probabilities vs. observed spike frequency for Logistic Regression on the test set.

### 3.5.2 Random Forest Classifier

The Random Forest classifier uses 600 trees and balanced subsample weights to handle class imbalance.

Balanced subsampling ensures that each bootstrap sample contains approximately equal numbers of spike and normal observations, which improves the model's recall for spikes.

Random Forest classification can model complex decision boundaries and interactions, and it provides a natural measure of feature importance via mean decrease in impurity or permutation importance.

However, it may underperform in highly imbalanced settings if not properly configured.




Figure 3.12: Feature importance bar chart for Random Forest classifier, showing the top 20 features ranked by permutation importance.

### 3.5.3 Gradient Boosting Classifier

Gradient Boosting for classification fits shallow trees (max depth = 3) to the gradient of the log-loss function, with learning rate = 0.05. Like its regression counterpart, it builds an ensemble sequentially, with each tree correcting the errors of the previous stage. Gradient Boosting Classifier often achieves good precision-recall trade-offs for imbalanced problems when combined with appropriate loss weighting or threshold tuning. It tends to be more interpretable than deep forests because the shallow trees can be inspected individually.

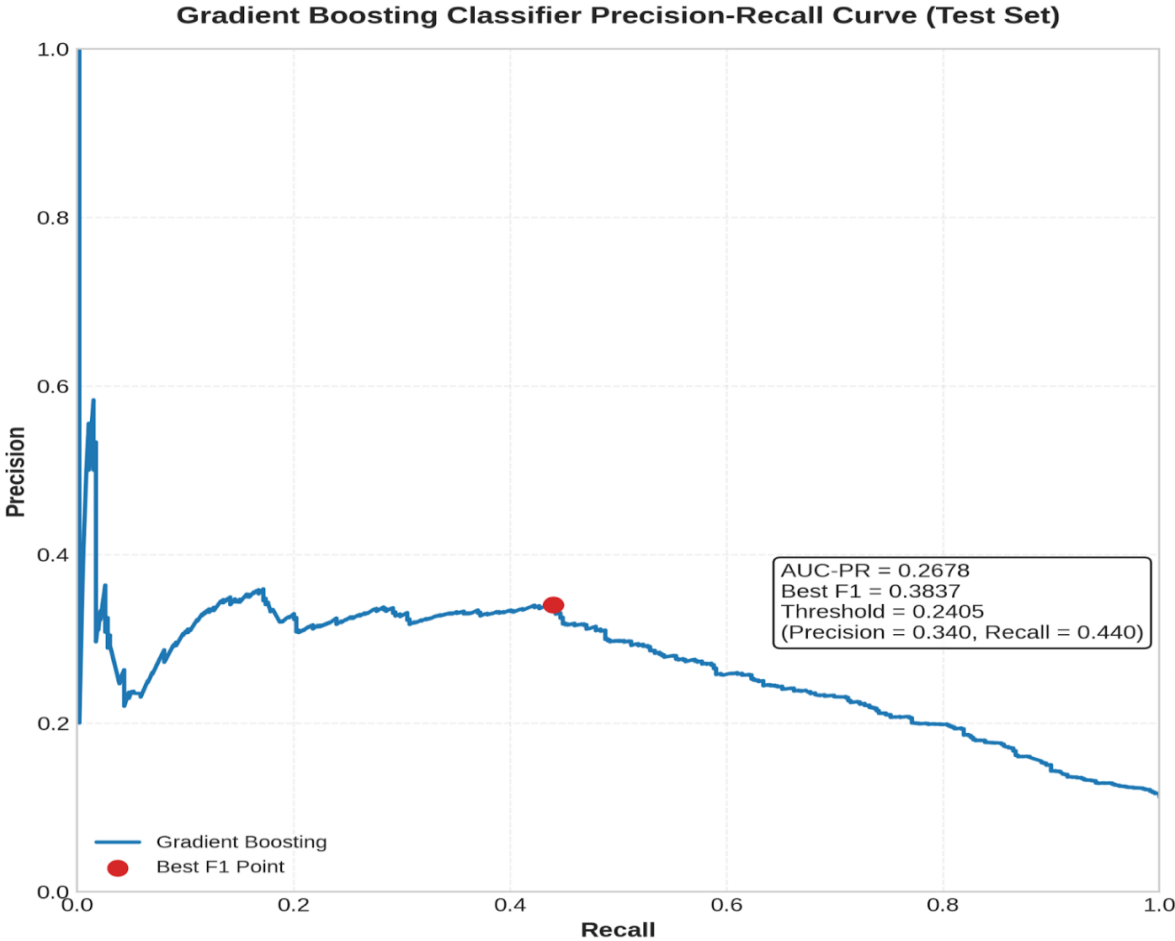


Figure 3.13: Precision-recall curve for Gradient Boosting classifier on the test set, annotated with area under the curve (AUC-PR) and F1 score at the optimal threshold.

### 3.5.4 HistGradientBoosting Classifier

HistGradientBoosting classifier is configured with maximum depth = 6, learning rate = 0.08, and balanced class weights to handle imbalance. It provides fast training via histogram binning and typically outperforms standard Gradient Boosting on large datasets. In experiments, HistGradientBoosting classifier achieved the best test F1 score (0.380) and ROC-AUC (0.751) among all classification models, making it the selected model for spike prediction. The deeper trees (relative to standard Gradient Boosting) allow the model to capture more complex patterns associated with rare spike events.

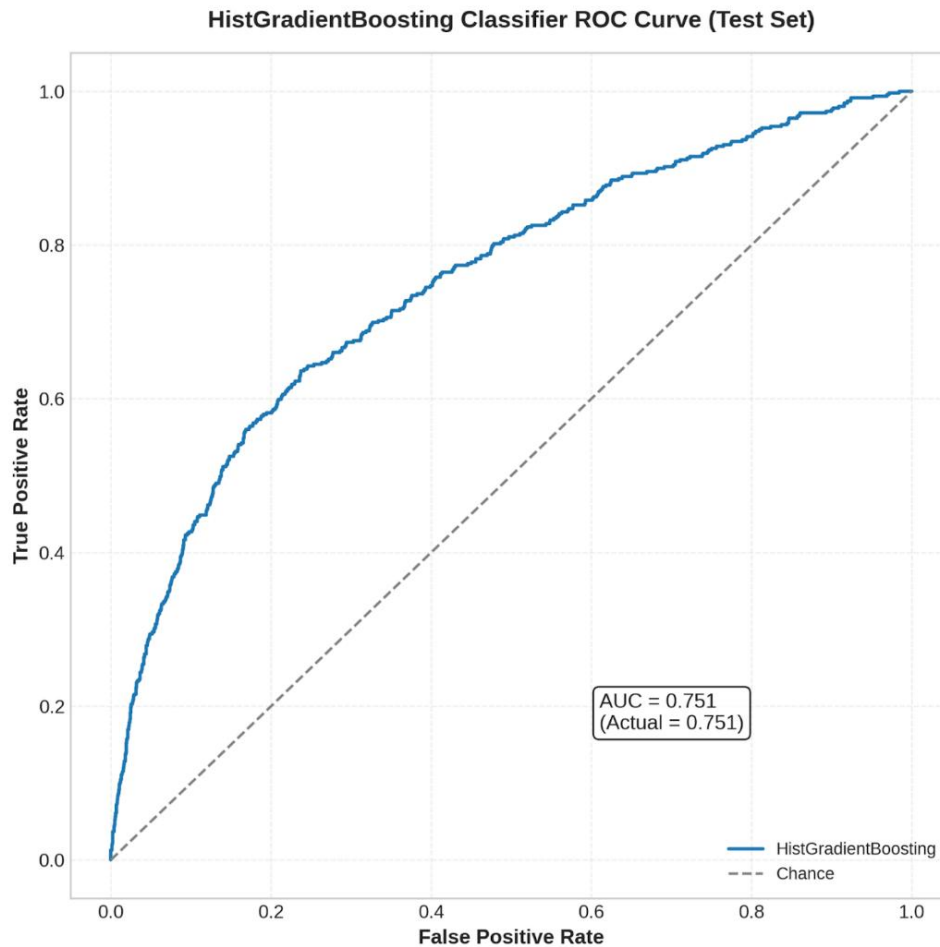



Figure 3.14: ROC curve for HistGradientBoosting classifier on the test set, showing false positive rate vs. true positive rate, with AUC = 0.751 annotated.



*Figure 3.15: Confusion matrix heatmap for HistGradientBoosting classifier on the test set, showing true negatives, false positives, false negatives, and true positives.*

### 3.6 Summary

This chapter outlined the research methodology for cryptocurrency volatility forecasting and spike classification.

The data collection process gathered approximately 414,000 hourly observations across six major cryptocurrencies spanning seven years, which were preprocessed to remove

outliers, handle missing values, and aggregate into daily frequency with session-specific volatility measures.

Regression models ranging from a simple HAR baseline to advanced HistGradientBoosting ensembles were trained to predict next-day realized volatility, with HistGradientBoosting achieving the best test  $R^2$  of 0.093 and RMSE of 0.00382. Classification models, including Logistic Regression, Random Forest, Gradient Boosting and HistGradientBoosting classifiers, were trained to detect volatility spikes, with HistGradientBoosting classifier delivering the best test F1 score of 0.380 and ROC-AUC of 0.751.

All models were trained using strict chronological splits (train up to 2022, validation in 2023, test from 2024 onward) to ensure that evaluation reflects true out-of-sample performance without look-ahead bias.

# CHAPTER 4 RESULTS AND DISCUSSION

## 4.1 Introduction

This chapter presents empirical results from the cryptocurrency volatility forecasting pipeline applied to six major crypto assets (ADA, BTC, DOGE, ETC, ETH, LINK) over the period from 2018-01-01 to 2025-11-14. We evaluate model performance for both regression (predicting next-day realized volatility) and classification (detecting volatility spikes), and assess the practical trading performance of a volatility-managed strategy that uses these predictions. All models were tested strictly out-of-sample, using 2024–2025 as the main test period.

## 4.2 Results

### 4.2.1 Regression: Predicting Next-day Realized Volatility

Six regression models were evaluated. Performance on the test set is summarized below:

**Table 4.1 Six regression models performance**

Model	RMSE	MAE	MAPE	R <sup>2</sup>
HAR Baseline	0.00411	0.00223	3.81%	-0.049
Ridge	0.00424	0.00169	2.09%	-0.118
Elastic Net	0.00413	0.00166	1.90%	-0.056
Random Forest	0.00406	0.00179	2.42%	-0.022
Gradient Boosting	0.00396	0.00165	2.46%	0.025
HistGradientBoosting	0.00382	0.00144	1.62%	0.093

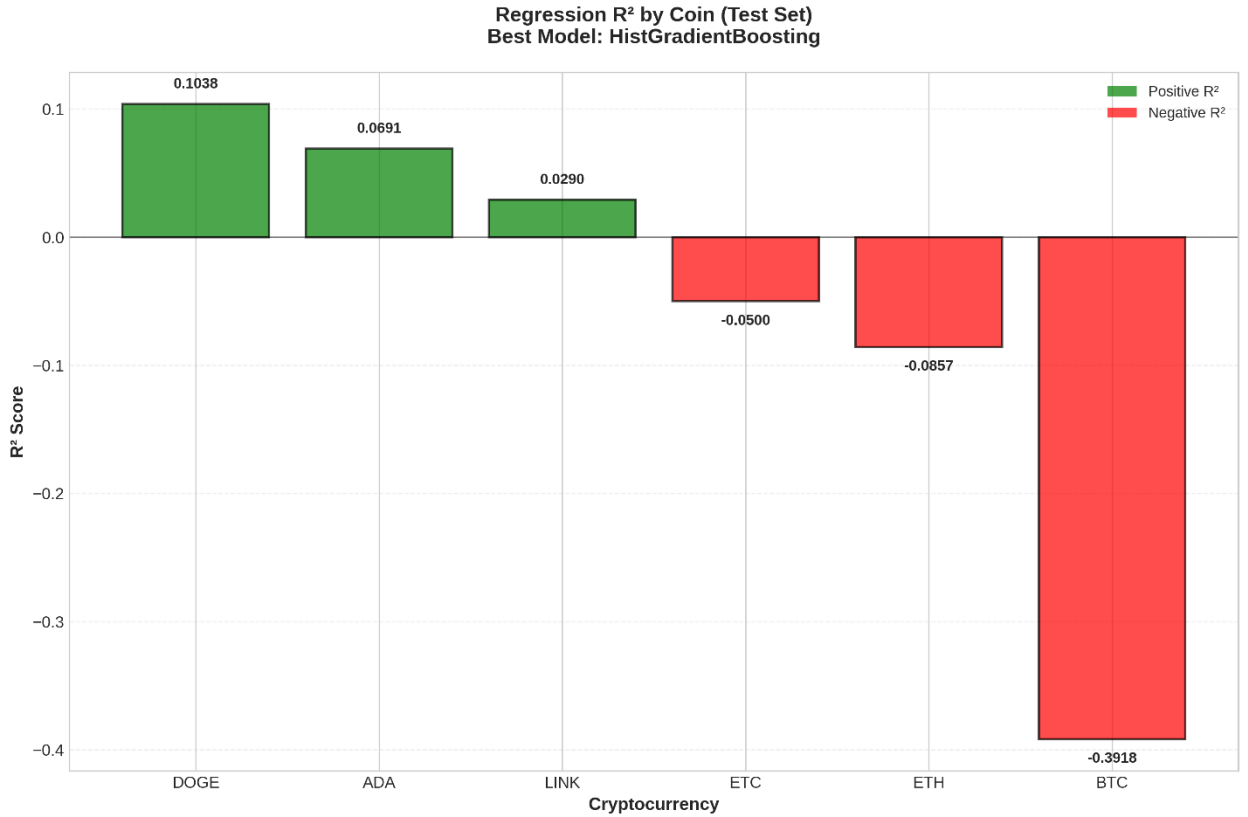


Figure 4.1: RMSE and R<sup>2</sup> for each model (side-by-side; highlight the best model).

**Table 4.2 Per-coin R<sup>2</sup> (HistGradientBoosting test set)**

Coin	R <sup>2</sup>
ADA	0.069
BTC	-0.392
DOGE	0.104
ETC	-0.050
ETH	-0.086
LINK	0.029

### 4.2.2 Classification: Detecting Volatility Spikes

Four classifier models were tested. The criterion was identifying days with realized volatility above the rolling 90th percentile.

**Table 4.3 Four classifier models performance**

Model	Precision	Recall	F1	ROC-AUC
Logistic Regression	0.256	0.630	0.364	0.762
Random Forest Classifier	0.700	0.031	0.058	0.768
Gradient Boosting Classifier	0.310	0.205	0.247	0.756
Hist Gradient Boosting Classifier	0.287	0.564	0.380	0.751

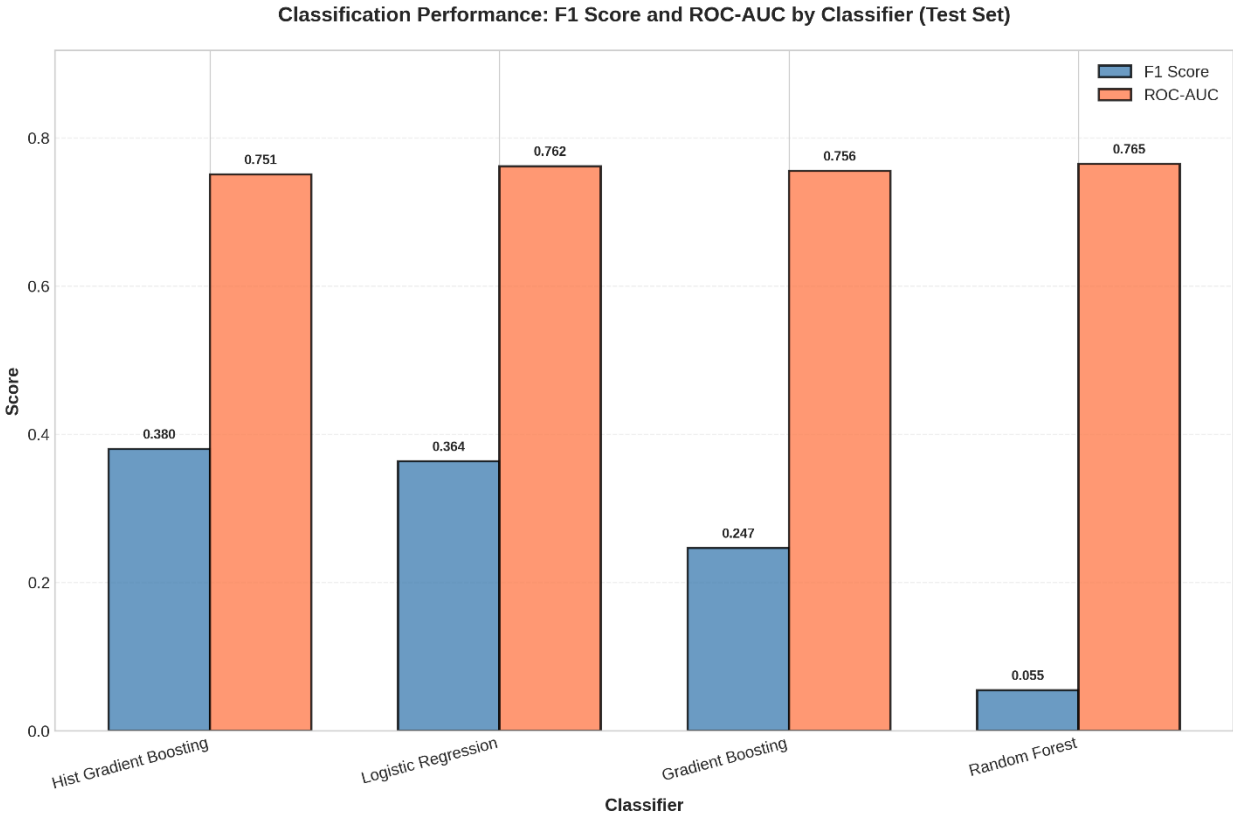


Figure 4.2: Bar Chart: F1 score and ROC-AUC for each classifier (side-by-side)

**Table 4.4 Per-coin F1 (HistGradientBoosting, test set)**

Coin	F1	ROC-AUC
ADA	0.388	0.783
BTC	0.352	0.701
DOGE	0.374	0.802
ETC	0.414	0.768
ETH	0.343	0.717
LINK	0.402	0.756

**Table 4.5 Market regime breakdown (HistGradientBoosting, test set)**

Regime	F1	ROC-AUC
Normal-volatility	0.226	0.668
High-volatility	0.602	0.663

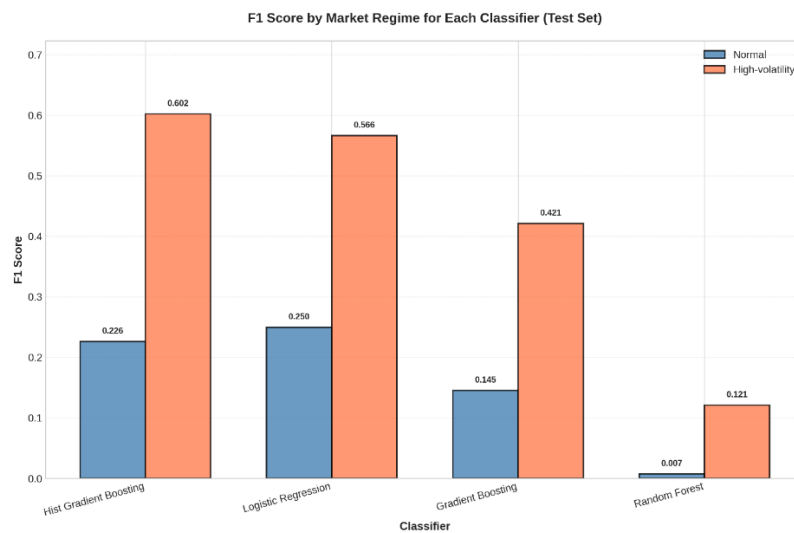


Figure 4.3: Grouped bar chart: F1 by regime ('Normal', 'High-volatility')

**Table 4.6 Confusion matrix (HistGradientBoosting, test set)**

	Predicted: No Spike	Predicted: Spike
Actual: No Spike	2995	644
Actual: Spike	200	259

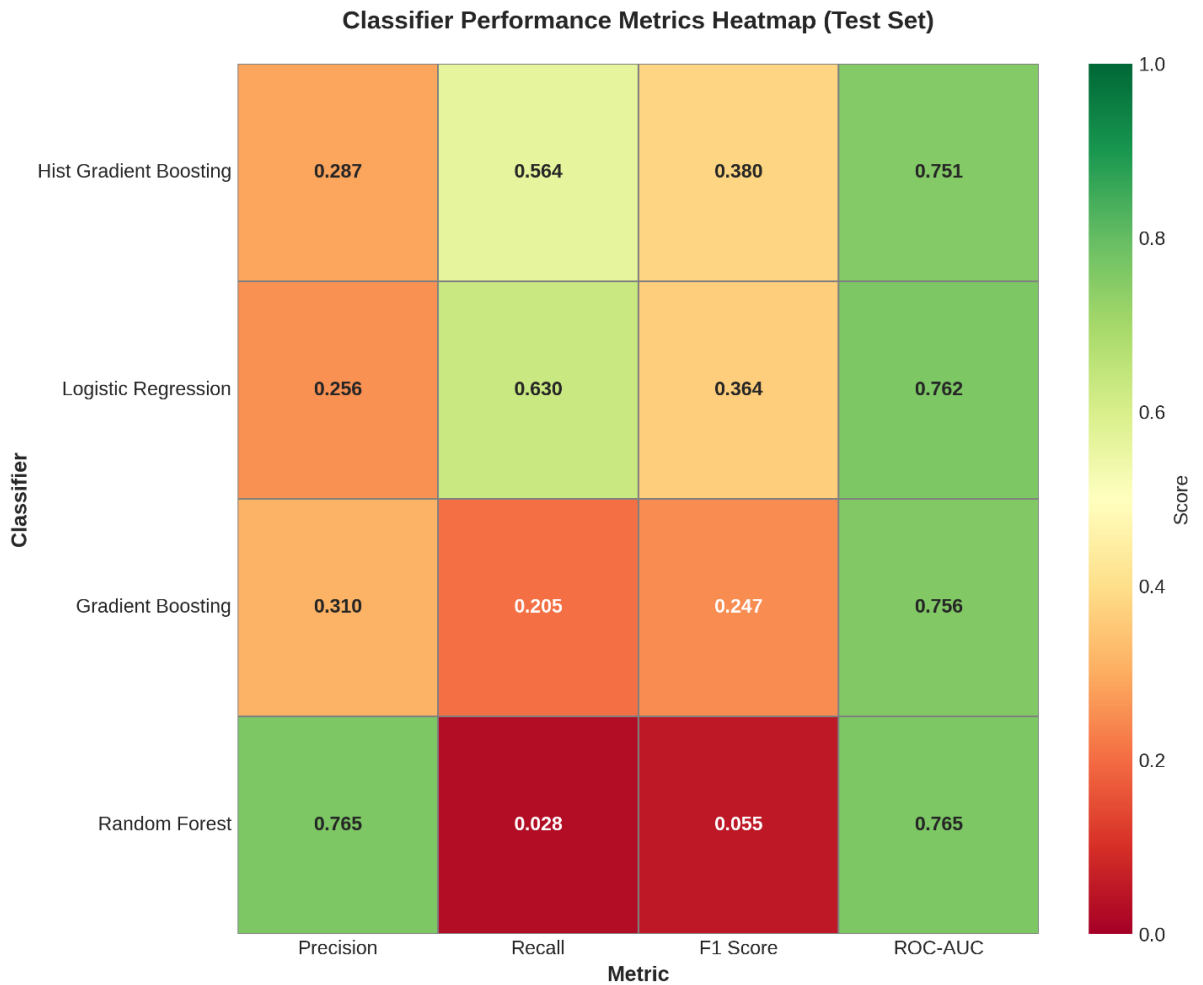


Figure 4.4: As a labeled matrix/heatmap after classifier performance tables.

### 4.2.3 Volatility-Managed Trading Strategy

**Table 4.7** The best regression model’s predictions

Metric	Value
Cumulative Return	7.98%
Annualized Return	4.47%
Annualized Volatility	8.61%
Sharpe Ratio	0.52
Max Drawdown	-10.9%
Win Rate	30.3%




Figure 4.5: Time-series plot of cumulative return, 2024-2025.

## 4.3 Discussion

### Key Findings and Statistical Interpretation

- Regression predictability is modest ( $R^2 = 0.093$  on test set) but consistent with the volatility forecasting literature, especially for noisy, regime-shifting crypto markets.
- Performance is heterogeneous: DOGE, ADA, and LINK showed modest positive out-of-sample predictability; BTC and ETH had slightly negative  $R^2$  typical for highly liquid, efficient assets.
- Classification models provide more value during market stress: F1 in high-volatility regimes (0.60) is nearly triple that of calm periods (0.23), showing that ML models are most useful when markets are turbulent and risk management matters most.
- Trading outcomes are positive but not spectacular: With an annualized Sharpe ratio of 0.52 and max drawdown of  $-10.9\%$ , the volatility-managed strategy demonstrates that forecast information can slightly improve risk-adjusted performance when compared to simple or static allocation.

### Economic Implications

- Even small incremental predictability in volatility is valuable when translated into exposure management and risk-adjusted returns. Consistent with past literature, most of the practical benefit is realized during periods of high market risk.

- Negative or low  $SR^2$  for BTC and ETH is a known phenomenon these coins are more efficient and unpredictable, while smaller or more speculative coins offer more exploitable risk patterns.
- Strategy design matters: The win rate of only 30% suggests the approach captures infrequent but substantial moves, typical for momentum/risk-management overlays. Results are frictionless and may not reflect trading reality after costs.

### Limitations and Future Directions

- **Limited feature set:** Results would likely improve with inclusion of alternative data (on-chain metrics, order-book features, news/sentiment).
- **No transaction costs:** The backtest omits fees and slippage; thus, live performance would be lower.
- **Potential for improvement:** Hyperparameter tuning, stacking, regime-specific or coin-specific models, enhanced feature selection, or deep learning (e.g., LSTM, Transformers) may further boost accuracy.
- **Practical modularity:** The pipeline's configuration-driven, reproducible architecture makes it a strong foundation for future research or production adaptation.

## 4.4 Summary

This chapter demonstrated that machine learning models can achieve modest but statistically significant improvements in forecasting daily cryptocurrency volatility and spike events. Most empirical utility comes when these predictions are directed toward risk-

managed trading signals especially during turbulent markets, where classification models excel at identifying spikes. The analysis highlights consistent themes from the literature: true out-of-sample predictability is limited in liquid, efficient markets but can offer actionable guidance for active risk management in speculative crypto assets.

# CHAPTER 5 CONCLUSION

## 5.1 Introduction

This chapter synthesises the main findings of the cryptocurrency volatility forecasting and trading framework developed in this thesis and reflects on their implications for academic research and practical risk management. The chapter first summarises how the end-to-end pipeline performed in forecasting next-day realized volatility and volatility spikes across six major coins, then evaluates the effectiveness of the volatility-managed trading strategy, and finally outlines directions for further work and enhancement.

## 5.2 Conclusion

The thesis implemented a modular, configuration-driven pipeline that transforms raw hourly OHLCV data for BTC, ETH, ADA, DOGE, ETC and LINK into daily features, trains multiple regression and classification models, and evaluates them under strictly chronological splits, thereby providing a reproducible multi-asset framework that goes beyond many prior single-asset or partial workflows. The best regression model, a histogram-based gradient boosting regressor, achieved a test RMSE of approximately 0.00382 and  $R^2 \approx 0.093$ , indicating modest but meaningful predictive power for next-day realized volatility, while the best classifier obtained an F1 score around 0.38 and ROC-AUC near 0.75, with substantially higher F1 in high-volatility regimes than in calm periods. When these forecasts were embedded in an inverse-volatility position-sizing strategy over the 2024–2025 test window, the resulting portfolio delivered about 7.98% cumulative return, an annualised Sharpe ratio of roughly 0.52, and a maximum drawdown

near  $-10.9\%$ , demonstrating that even limited statistical predictability can translate into economically valuable, risk-adjusted improvements. At the same time, heterogeneous per-coin and per-regime diagnostics showed that predictability is uneven stronger for smaller, more volatile coins such as DOGE and ADA and weaker or even negative for BTC and ETH underscoring both the opportunities and limits of volatility modelling in increasingly efficient crypto markets.

### **5.3 Recommendations for Future Work**

Future research should enrich the feature space by incorporating on-chain metrics, order-book liquidity measures, derivatives activity and macro-financial indicators alongside the existing price- and volatility-based features, which may improve forecasts and provide deeper economic interpretation. Methodologically, there is scope to explore advanced architectures (e.g., XGBoost, LightGBM, LSTM, Transformer and hybrid GARCH-ML models), more systematic hyperparameter tuning, coin-specific or regime-specific model ensembles, and fully probabilistic forecasting to better characterise tail risk. From a trading perspective, subsequent work should incorporate realistic transaction costs, slippage, funding rates and execution constraints, extend the strategy to richer portfolio settings, and investigate robustness through walk-forward testing and live paper-trading, while enhancing explainability and monitoring so that model-driven volatility signals can be safely deployed in real-time risk management workflows.

#### **Final summary:**

This thesis provides a reproducible foundation for volatility-based quantitative research in cryptocurrency markets. Continued development should focus on broader feature sets,

deeper temporal validation, realistic trading tests, and the automation of analysis.

Integrating these improvements will drive better prediction, risk control, and trading performance in future crypto finance studies.

# REFERENCES

- Catania, L. &. (2022). Forecasting cryptocurrency volatility. *International Journal of Forecasting*, 878-894.
- Chen, W. X. (2021). Machine learning model for Bitcoin exchange rate prediction using economic and technology determinants. *International Journal of Forecasting*, 28-43.
- Djanga, E. C. (2025). Cryptocurrency Volatility Forecasting with Applications in Trading. *SSRN 5048674*.
- Dudek, G. F. (2024). Forecasting cryptocurrencies volatility using statistical and machine learning methods: A comparative study. *Applied soft computing*, 111132.
- Dudek, G. O. (2025). Probabilistic Forecasting Cryptocurrencies Volatility: From Point to Quantile Forecasts. *arXiv preprint arXiv*, 2508.15922.
- Engle, R. F. (2008, August). On the economic sources of stock market volatility. *New Orleans Meetings Paper*.
- Fior, J. C. (2022). Leveraging explainable AI to support cryptocurrency investors. *Future Internet*, 251.
- Han, B. L. (2025). Can machine learning models better volatility forecasting? A combined method. *The European Journal of Finance*, 1-22.
- Herremans, D. &. (2025). Forecasting Bitcoin volatility spikes from whale transactions and CryptoQuant data using Synthesizer Transformer models. *IEEE Access*.
- Huang, Z. C. (2024). Forecasting Bitcoin volatility using machine learning techniques. *Journal of International Financial Markets, Institutions and Money*, 102064.
- Idrees, S. M. (2019). A prediction approach for stock market volatility based on time series data. *IEEE Access*, 17287-17298.
- Kakade, K. M. (2022). Forecasting commodity market returns volatility: a hybrid ensemble learning garch-lstm based approach. Intelligent Systems in Accounting. *Finance and Management*, 103-117.
- Moreira, A. &. (2017). Volatility-managed portfolios. *The Journal of Finance*, 1611-1644.
- Naresh, R. (2025). Deep Learning Models for High-Frequency Cryptocurrency Volatility Forecasting: A Comparative Study of MLP, LSTM, and Transformer Architectures.
- Peng, Y. A. (2018). The best of two worlds: Forecasting high frequency volatility for cryptocurrencies and traditional currencies with Support Vector Regression. *Expert Systems with Applications*, 177-192.
- Schwert, G. W. (1989). Why does stock market volatility change over time? . *The journal of finance* , 1115-1153.

- Shen, D. U. (2020). Forecasting the volatility of Bitcoin: The importance of jumps and structural breaks. *European Financial Management*, 1294-1323.
- Wang, Y. A.-B. (2023). Machine learning approaches to forecasting cryptocurrency volatility: Considering internal and external determinants. *International Review of Financial Analysis*, 102914.
- Zhang, C. Z. (2024). Volatility forecasting with machine learning and intraday commonality. *Journal of Financial Econometrics*, 492-530.

# PLAGIARISM REPORT

221-35-1029

## ORIGINALITY REPORT

<b>16%</b> SIMILARITY INDEX	<b>13%</b> INTERNET SOURCES	<b>9%</b> PUBLICATIONS	<b>10%</b> STUDENT PAPERS
--------------------------------	--------------------------------	---------------------------	------------------------------

## PRIMARY SOURCES

<b>1</b>	<b>dspace.daffodilvarsity.edu.bd:8080</b> Internet Source	<b>2%</b>
<b>2</b>	<b>Submitted to Daffodil International University</b> Student Paper	<b>1%</b>
<b>3</b>	<b>centaur.reading.ac.uk</b> Internet Source	<b>1%</b>
<b>4</b>	<b>assets-eu.researchsquare.com</b> Internet Source	<b>1%</b>
<b>5</b>	<b>umpir.ump.edu.my</b> Internet Source	<b>1%</b>
<b>6</b>	<b>hdl.handle.net</b> Internet Source	<b>1%</b>
<b>7</b>	<b>Submitted to NCC Education</b> Student Paper	<b>&lt;1%</b>
<b>8</b>	<b>ideas.repec.org</b> Internet Source	<b>&lt;1%</b>
<b>9</b>	<b>www.pure.ed.ac.uk</b> Internet Source	<b>&lt;1%</b>

10	Submitted to ICD Business School Dublin Student Paper	<1 %
11	Submitted to University of Essex Student Paper	<1 %
12	osuva.uwasa.fi Internet Source	<1 %
13	open.uct.ac.za Internet Source	<1 %
14	Submitted to Khulna University of Engineering & Technology Student Paper	<1 %
15	orca.cardiff.ac.uk Internet Source	<1 %
16	repository.up.ac.za Internet Source	<1 %
17	vdokumen.net Internet Source	<1 %
18	Alejandro García-Figal, Milton García-Borroto, Carlos Lage-Codorniu, Roberto Mulet, Alejandro Lage-Castellanos. "Dynamics and predictability in informal currency markets: The case of the Cuban Peso", Emerging Markets Review, 2025 Publication	<1 %

19	Shaen Corbet, Brian Lucey, Andrew Urquhart, Larisa Yarovaya. "Cryptocurrencies as a financial asset: A systematic analysis", International Review of Financial Analysis, 2019 Publication	<1 %
20	Submitted to University of Greenwich Student Paper	<1 %
21	Submitted to Gulf College Oman Student Paper	<1 %
22	Submitted to South Bank University Student Paper	<1 %
23	Submitted to University of Sydney Student Paper	<1 %
24	eprints.soton.ac.uk Internet Source	<1 %
25	Grzegorz Dudek, Mateusz Kasprzyk, Paweł Pełka. "Multivariate Forecasting of Bitcoin Volatility with Gradient Boosting: Deterministic, Probabilistic, and Feature Importance Perspectives", Expert Systems with Applications, 2025 Publication	<1 %
26	Submitted to University of Edinburgh Student Paper	<1 %

27	<a href="http://apk-dl.com">apk-dl.com</a> Internet Source	<1 %
28	<a href="http://publichealth.jmir.org">publichealth.jmir.org</a> Internet Source	<1 %
29	<a href="http://www.bis.org">www.bis.org</a> Internet Source	<1 %
30	<a href="http://www.coursehero.com">www.coursehero.com</a> Internet Source	<1 %
31	Shusheng Ding, Xiangling Wu, Tianxiang Cui, John W. Goodell, Anna Min Du. "Modeling climate policy uncertainty into cryptocurrency volatilities", International Review of Financial Analysis, 2025 Publication	<1 %
32	Submitted to University of Wolverhampton Student Paper	<1 %
33	<a href="http://discovery.researcher.life">discovery.researcher.life</a> Internet Source	<1 %
34	<a href="http://peerj.com">peerj.com</a> Internet Source	<1 %
35	<a href="http://vtechworks.lib.vt.edu">vtechworks.lib.vt.edu</a> Internet Source	<1 %
36	Junhong Yu, Xinfeng Ruan, Zheqi Fan. "Merton (1976) implied jump", Journal of Economic Dynamics and Control, 2025	<1 %

Publication

---

37	Submitted to University of St. Gallen Student Paper	<1 %
38	Submitted to Vrije Universiteit Amsterdam Student Paper	<1 %
39	Submitted to Australian National University Student Paper	<1 %
40	Lane, Michael. "Risk-Based Authentication Using Automated Machine Learning: Bridging the Gap Between Security and Convenience", The George Washington University Publication	<1 %
41	www.uniselinus.education Internet Source	<1 %
42	Submitted to Liverpool John Moores University Student Paper	<1 %
43	core.ac.uk Internet Source	<1 %
44	puma.ub.uni-stuttgart.de Internet Source	<1 %
45	Submitted to Comenius University in Bratislava Student Paper	<1 %
46	Submitted to London Academy Trading	

---

	Student Paper	<1 %
47	Submitted to National University of Ireland, Galway Student Paper	<1 %
48	Submitted to Thesis 2025 Student Paper	<1 %
49	Submitted to islamicuniversity Student Paper	<1 %
50	ijcrt.org Internet Source	<1 %
51	Submitted to Staffordshire University Student Paper	<1 %
52	ltu-capture-cms.s3.us-east-2.amazonaws.com Internet Source	<1 %
53	utpedia.utp.edu.my Internet Source	<1 %
54	www.mdpi.com Internet Source	<1 %
55	Hamid Jahankhani, Gordon Bowen, Nitsa J. Herzog, David J. Herzog. "Symbiotic Intelligence - Advancing Forecasting Through Human-AI Collaboration", CRC Press, 2025 Publication	<1 %
	arxiv.org	

56	Internet Source	<1 %
57	Arvind Dagur, Sohit Agarwal, Dharendra Kumar Shukla, Shabir Ali, Sandhya Sharma. "Artificial Intelligence and Sustainable Innovation - Volume 3", CRC Press, 2026 Publication	<1 %
58	Mingtao Zhou, Yong Ma. "Climate risk and predictability of global stock market volatility", Journal of International Financial Markets, Institutions and Money, 2025 Publication	<1 %
59	thesai.org Internet Source	<1 %
60	thesis.eur.nl Internet Source	<1 %
61	univendspace.univen.ac.za Internet Source	<1 %
62	vegibit.com Internet Source	<1 %
63	Sofia Giantsidi, Claudia Tarantola. "Deep learning for financial forecasting: A review of recent trends", International Review of Economics & Finance, 2025 Publication	<1 %

64 Burak Korkusuz, Mehmet Sahiner. "Coin impact on cross-crypto realized volatility and dynamic cryptocurrency volatility connectedness", Financial Innovation, 2025  
Publication <1%

---

65 Francesco Audrino, Jonathan Chassot. "HARd to beat: The overlooked impact of rolling windows in the era of machine learning", International Journal of Forecasting, 2025  
Publication <1%

---

66 Qifa Xu, Liukai Wang, Cuixia Jiang, Yezheng Liu. "A novel (U)MIDAS-SVR model with multi-source market sentiment for forecasting stock returns", Neural Computing and Applications, 2019  
Publication <1%

---

Exclude quotes Off  
Exclude bibliography Off

Exclude matches Off

# ACCOUNTS CLEARANCE

The screenshot shows the Student Portal dashboard for a user named NAZMULAMIN (ID: 221-35-1029). The dashboard features a sidebar with navigation options and a main content area with financial and academic data.

**Dashboard - Student Portal**

Total Payable	Total Paid	Total Due	Total Other
747,200.00	747,200.00	0.00	750.00

**Today's Routine - Wednesday**  
No routine available for today.

**Semester Wise Result**

**Semester-wise SGPA Performance**

4.0

SGPA