Stock Market Forecasting Using Machine Learning Models Through Volatility-Driven Trading Strategies

Ivan Letteri[a]
Department of Life, Health and Environmental Sciences, University of L'Aquila, Italy

Keywords: Machine Learning, Statistical Analysis, Algorithmic Trading, K-Means++, Granger Causality Test.

Abstract: The purpose of our research was to explore volatility-based trading strategies in financial markets to leverage market dynamics for capital gain. We sought to introduce a strategy that integrated statistical analysis with machine learning to predict stock market trends. Our method involved using the k-means++ clustering algorithm to examine the mean volatility of the nine largest stocks in both the NYSE and Nasdaq markets. The clusters formed the basis for understanding relationships among stocks based on their volatility patterns. We further subjected the mid-volatility clustered dataset to the Granger Causality Test, which helped identify stocks with strong predictive connections. These stocks were crucial in formulating our trading strategy, serving as trend indicators for decisions on target stock trades. Our empirical approach included thorough backtesting and performance analysis. Our findings demonstrated the effectiveness of our method in exploiting profitable trading opportunities. This was achieved through predictive insights derived from volatility clusters and Granger causality relationships among stocks. In conclusion, our research contributed to the field of volatility-based trading strategies by offering a methodology that combined a statistical approach with machine learning. This enhanced the predictability of stock market trends.

1 INTRODUCTION

The field of finance is witnessing a growing interest in volatility-based trading, which capitalises on market dynamics. Artificial Intelligence (AI) plays a crucial role in this, providing robust tools for analysing and leveraging market volatility. Specifically, AI's ability to estimate mean volatility offers valuable insights into the uncertainty and risk associated with specific securities or the overall market (Letteri et al., 2022).

In our work, the key research questions include examining the effectiveness of k-means++ clustering in analyzing the mean volatility of major stocks, understanding relationships among stocks based on distinctive volatility patterns, and utilizing the Granger Causality Test to assess predictive influences between stocks. The study aims to formulate trading strategies based on identified predictive connections, leveraging influential stocks as trend indicators. Rigorous backtesting and performance analysis validate the reliability of the proposed volatility-driven trading strategy.

To answer the aforementioned research questions, we created an AI trading strategy using k-means++ clustering of average volatility data (Arthur and Vassilvitskii, 2007) from nine major stock markets. Initially, we aim to identify distinct volatility patterns in the market and group assets accordingly. We then utilize the Granger Causality Test (GCT) (Kirchgässner and Wolters, 2007) to pinpoint stocks that significantly predict others in our analysis, establishing buy, sell, or hold trading decisions.

In this study, we used the AITA framework (Letteri, 2023a) to rigorously analyse the historical performance of the proposed strategy, employing multiple performance metrics to evaluate its profitability, effectiveness, and resilience.

Previously, our focus on technical trading strategies emphasised technical indicators (Letteri et al., 2022),(Letteri et al., 2023), particularly for investment timing. We now explore Historical Volatility estimators as a dataset for identifying medium volatility and selecting stocks for the Granger Causality Test asset cointegration approach (Engle and Granger, 1987).

This paper is organised as follows: Section 2 introduces foundational concepts within our AITA framework, highlighting the Volatility Trading System (VolTS)(Letteri, 2023b) module and Aita Back-

[a] https://orcid.org/0000-0002-3843-386X

Letteri, I.
Stock Market Forecasting Using Machine Learning Models Through Volatility-Driven Trading Strategies.
DOI: 10.5220/0012607200003717

In *Proceedings of the 6th International Conference on Finance, Economics, Management and IT Business (FEMIB 2024)*, pages 96-103
ISBN: 978-989-758-695-8; ISSN: 2184-5891

Testing (AitaBT). Section 3 outlines the methodology within the VolTS module, which analyses securities' volatility averages and establishes predictive relationships. It then delves into the implementation of the trading strategy and includes a thorough empirical analysis of its performance and robustness. Section 4 presents practical findings achieved through backtesting with the AitaBT module, followed by a discussion. Finally, Section 5 concludes the study by summarising the effectiveness and applicability of the proposed method.

2 BACKGROUND

2.1 Price Action

The price action (PA) influences Historical Volatility (HV), and in turn, HV can provide insights into future PA. When the PA exhibits strong price movements, such as wide trading ranges, breakouts, or rapid directional changes, it tends to increase.

VolTS, a module within the AITA framework, adheres to these principles. Low HV signifies a period of consolidation or low price volatility, indicating a potential upcoming spike in volatility or a shift in the PA. On the other side, high HV suggests a higher probability of sharp market movements or trend changes.

Within VolTS, the PA is encoded as OHLC, i.e., the open, high, low, and close prices of the assets, as represented in the candlesticks charts (see figure 1). For each timeframe t, the OHLC of an asset is represented as a 4-dimensional vector $X_t = (x_t^{(o)}, x_t^{(h)}, x_t^{(l)}, x_t^{(c)})^T$, where $x_t^{(l)} > 0$, $x_t^{(l)} < x_t^{(h)}$ and $x_t^{(o)}, x_t^{(c)} \in [x_t^{(l)}, x_t^{(h)}]$.

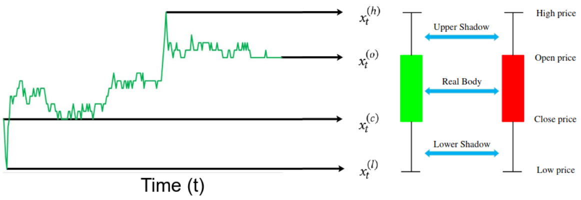

Figure 1: Example of candlestick chart.

2.2 Historical Volatility Module

The construction of the dataset is designed to use the following HV estimators:

- The *Parkinson* (PK) estimator incorporates the stock's daily *high* and *low* prices as follow:

$$PK = \sqrt{\frac{1}{4Nln(2)}\sum_{i=1}^{N}\left(ln\frac{x_t^{(h)}}{x_t^{(l)}}\right)^2}.$$

It is derived from the assumption that the true volatility of the asset is proportional to the logarithm (ln) of the ratio of the high $x_t^{(h)}$ and low $x_t^{(l)}$ prices of N observations.

- The *Garman-Klass* (GK) estimator assumes that price movements are log-normally distributed calculated as follows: $\sqrt{\frac{1}{N}\sum_{i=1}^{N}\frac{1}{2}\left(ln\frac{x_t^{(h)}}{x_t^{(l)}}\right)^2 - \sum_{i=1}^{N}(2ln(2)-1)\left(ln\frac{x_t^{(c)}}{x_t^{(o)}}\right)^2}$

- The *Rogers-Satchell* (RS) estimator uses the range of prices within a given time interval as a proxy for the volatility of the asset as follows: $RS = \sqrt{\frac{1}{N}\sum_{t=1}^{N} ln\left(\frac{x_t^{(h)}}{x_t^{(c)}}\right) ln\left(\frac{x_t^{(h)}}{x_t^{(o)}}\right) + ln\left(\frac{x_t^{(l)}}{x_t^{(c)}}\right) ln\left(\frac{x_t^{(l)}}{x_t^{(o)}}\right)}$
RS assumes that the range of prices within the interval is a good proxy for the volatility of the asset, additionally, the estimator may be sensitive to outliers and extreme price movements.

- The *Yang-Zhang* (YZ) estimator (Yang and Zhang, 2000) incorporates OHLC prices as follows: $YZ = \sqrt{\sigma^2_{OvernightVol} + k\sigma^2_{OpenToCloseVol} + (1-k)\sigma^2_{RS}}$, where $k = 0.34/1.34 + \frac{N+1}{N-1}$, $\sigma^2_{OpenToCloseVol} = \frac{1}{N-1}\sum_{i=1}^{N}\left(ln\frac{x_t^{(c)}}{x_t^{(o)}} - ln\frac{x_t^{(c)}}{x_t^{(o)}}\right)^2$, and $\sigma^2_{OvernightVol} = \frac{1}{N-1}\sum_{i=1}^{N}\left(ln\frac{x_t^{(o)}}{x_{t-1}^{(c)}} - ln\frac{x_t^{(o)}}{x_{t-1}^{(c)}}\right)^2$.
Empirical studies have demonstrated that the YZ estimator exhibits notable performance across a broad spectrum of scenarios, including those characterised by jumps and non-normality in the data. However, this estimator is not without its limitations, and its effectiveness may be constrained in certain contexts.

In this research, our attention is centred on mid-volatility. This focus allows us to either close open positions or refrain from entering a position when the anticipated volatility coefficient is high, thereby mitigating the risk of losses. On the other hand, if the expected volatility is too low, it does not offer any potential for gains.

2.3 Trading Strategies

Three distinct trading strategy classes are implemented in AITA framework:

- *Buy and Hold* (B&H) strategy is used as a benchmark to compare the performance of the two strategies below. It involves buying one single share on the first date of the period studied on the market close and selling the share at the market close on the last date as follows: $V_t = Q \cdot P_t$, where V_t is the value of the investment at time t. Q is the quantity of the asset purchased at time $t = 0$, and P_t is the price of the asset at time t with P_0 the initial price.
- *Trend Following* (TF) strategy is one way to engage in trend trading, where a trader initiates an order in the direction of the breakout after the price surpasses the resistance line as follows: let P_t the price at time t, and let MA denote the Moving Average of the asset price over a certain period. If $P_t \geq MA_t$ indicates an upward trend to take a long position otherwise it is a downward trend to take a short position.
- *Mean Reversion* (MR) strategy suggests that a security's maximum and minimum prices are temporary, and the security will eventually move towards its mean as follows: let P_t the price of the asset at time t, and let μ and σ represent the mean and standard deviation of the asset price, respectively. The entry/exit conditions for a long/short position are given by: $P_t < \mu - k \cdot \sigma$ and $P_t > \mu - k \cdot \sigma$, respectively where k is a constant representing the number of standard deviations from the mean at which the entry condition is triggered.

For the sake of brevity, in this study, the experiment is focused on the trend-follow strategy and we compare it with the B&H considered as a benchmark. It is important to note that both trend-following and mean reversion strategies, which are theoretically opposing concepts, can be applied to the same stock without conflicting with each other. Nonetheless, we find it beneficial to apply the mean reversion strategy when dealing with mid-volatility assets.

2.4 Backtesting Module

AitaBT module considers both profit and risk metrics as crucial factors in trading, in order to evaluate the potential profitability of investments and manage risk exposure.

- (i) *Drawdown (DD)* is a measure of the peak-to-trough decline in the value of a trading account before a new peak is attained. DD is defined as follows: $DD = \frac{P-T}{T}$, where P is the highest value or peak of the portfolio. T is the lowest value or trough after the peak. Maximum Drawdown (MDD) is the most significant loss from peak to trough during a specific period calculated as follows: $MDD = max_i\left(\frac{P_i - min_{j:i\leq j\leq N} T_j}{P_i}\right)$, where P_i is the highest value or peak of the portfolio time i. T_j is the lowest value or trough after the peak up to time j. N is the total number of data points.
- (ii) The *Sortino ratio (SoR)* is a risk-adjusted profit measure, which refers to the return per unit of deviation as follows: $SoR = \frac{R_p - R_f}{\sigma_d}$, where R_p is the expected portfolio return, R_f the risk-free rate of return, and σ_d denotes the downside deviation of the portfolio returns.
- (iii) The *Sharpe ratio (SR)* is a variant of the risk-adjusted profit measure, which applies σ_p as a risk measure: $SR = \frac{R_p - R_f}{\sigma_p}$ where σ_p is the standard deviation of the portfolio return.
- (iv) The *Calmar ratio (CR)* is another variant of the risk-adjusted profit measure, which applies MDD as risk measure: $CR = \frac{R_p - R_f}{MDD}$.

To check the goodness of trades, we mainly focused on the *Total Returns* $TR_k(t)$ for each stock $(k = 1, \ldots, p)$ in the time interval $(t = 1, \ldots, n)$ with the price P, defined as follows:

$$TR_k(t) = \frac{P_k(t + \Delta t) - P_k(t)}{P_k(t)}.$$

Furthermore, we analyzed the standardized returns $r_k = (TR_k - \mu_k)/\sigma_k$, with $(k = 1, \ldots, p)$, where σ_k is the standard deviation of TR_k, e μ_k denote the average overtime for the studied period.

3 METHOD

3.1 Asset Collections

AITA automatically downloads the OHLC prices via an internal Python library connected to an API, using the MetaTrader5 (MT5)[1] directly associated with the broker TickMill[2]. The data collected for this study includes the OHLC prices of the stocks listed in Table 1.

3.2 Anomalies Filtering

AITA framework starts to examine the price time series of the assets to determine the time window without considerable anomalies. The criterion implemented is based on the anomaly score calculated

[1] https://www.metatrader5.com/
[2] https://tickmill.eu

Table 1: List of the main 9 stocks selected for the experimentation.

Ticker	Company	Market
MSFT	Microsoft Corporation	Nasdaq
GOOGL	Alphabet Inc.	Nasdaq
MU	Micron Technology, Inc.	Nasdaq
NVDA	NVIDIA Corporation	NYSE
AMZN	Amazon.com, Inc.	NYSE
META	Meta Platforms, Inc.	NYSE
QCOM	QUALCOMM Incorporated	Nasdaq
IBM	Int. Business Machines Corp.	NYSE
INTC	Intel Corporation	NYSE

by a K-Nearest Neighbors (KNN) model (Wahid and Chandra Sekhara Rao, 2020). One of the key advantages of KNN is its ability to handle non-linear and complex relationships between data points (Letteri et al., 2021a)(Letteri et al., 2020a). The KNN model is fit to the time series data and the anomaly score is calculated based on the distance between the points and their k nearest neighbours.

The threshold (th) for detecting anomalies is then determined based on the mean (μ) and standard deviation (σ) of the anomaly scores. The criterion can be expressed as follows: let x_t be the value of the time series at time t, and k be the number of Nearest Neighbours to use in the KNN model with the Euclidean distance between x_t and x_i, where x_i is the i^{th} nearest neighbour (NN) of x_t. The anomaly score (as_t) for x_t is defined as follow:

$$as_{t,i} = \frac{1}{k} \sum \sqrt{(x_t - x_i)^2}, \forall i \in NN(x_t, k).$$

The threshold th for detecting anomalies as follows: $th = \mu + 3 \cdot \sigma$. Data points with anomaly scores greater than the threshold are considered to be anomalies.

Figure 2 shows only one critical anomaly during March 2020 (the global pandemic), so we decided to use only the time window in the period after instead of simply removing it, starting from 1st May 2020 to 1st May 2023.

3.3 Historical Volatility Dataset

The History Volatility Clustering process of our approach determines the stocks with intermediate volatility. First calculate the average of historical volatility time series among the aforementioned estimators (see sect. 2.2). Next, the resulting volatility series are clusterized using the KMeans++ algorithm with the Dynamic Time Warping (DTW) metric (Niennattrakul and Ratanamahatana, 2007). DTW is used to compare couples of time series that may have different lengths and speeds of variation, which makes it well-suited for this type of clustering. In particular, we split into three clusters ($K = 3$) *high*, *middle*, and *low* volatility. The centroids are selected using the maximum DTW distance with respect to the previous centroid.

Figure 3 shows the results displayed through a plot of the time series belonging to the middle cluster where we are focused on our strategy. It is worth noting that, the main region is in the time window from 1st November 2022 to 1st May 2023. So, we use this interval as the dataset, and then from the intermediate cluster, the candidate assets selected are TSLA with the highest, AMZN and META in the middle, with QCOM and IBM with the lowest values, respectively.

3.4 Regression Analysis

AITA performs regression analysis to determine whether one time series can predict another. Initially, it uses linear regression to model how one variable (independent variable) explains or predicts changes in another variable (dependent variable) considering F-statistic and Durbin-Watson statistics.

- *F-statistic (F-stat)* is used by VolTS to evaluate the overall adequacy of the model by comparing the full model with a null model (without any independent variables) by determining whether at least one of the independent variables contributes significantly to explaining the variations in the dependent variable.
- *Durbin-Watson statistic* calculated by VolTS detects autocorrelation in the model residuals because it can influence the interpretation of the results. A value close to 2 indicates no autocorrelation, while values significantly different from 2 suggest the presence of autocorrelation.

3.5 Cointegration and Causality

Cointegration refers to the long-term equilibrium relationship between two or more time series. If two time series are cointegrated, it means there exists a stable linear combination between them, even if the individual series may be non-stationary.

In the context of volatility-based trading, the VolTS module performs the GCT to examine the relationship between the lagged volatility of one asset and the future volatility of another asset by applying the following steps:

- **Step 1.** *Significant Granger causality:* Let X and Y be the pair stocks time series volatility to check, where X represents the potential causal

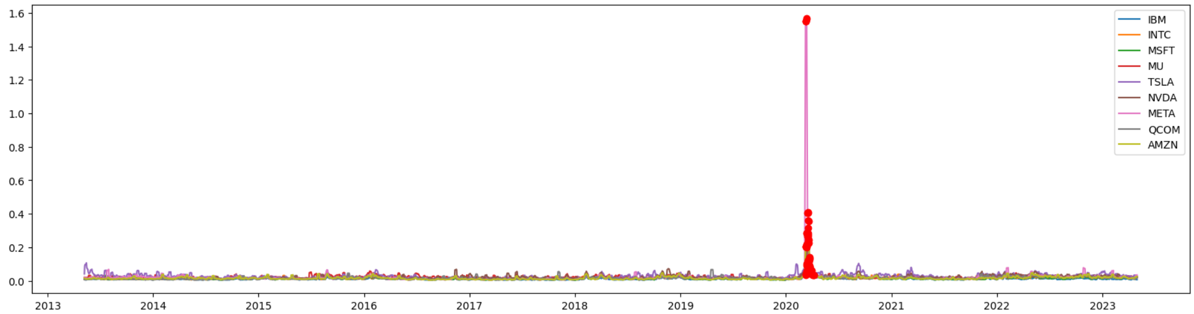

Figure 2: Red dots highlight the anomalies detected in the interval analyzed from 2020/05/01 to 2023/05/01.

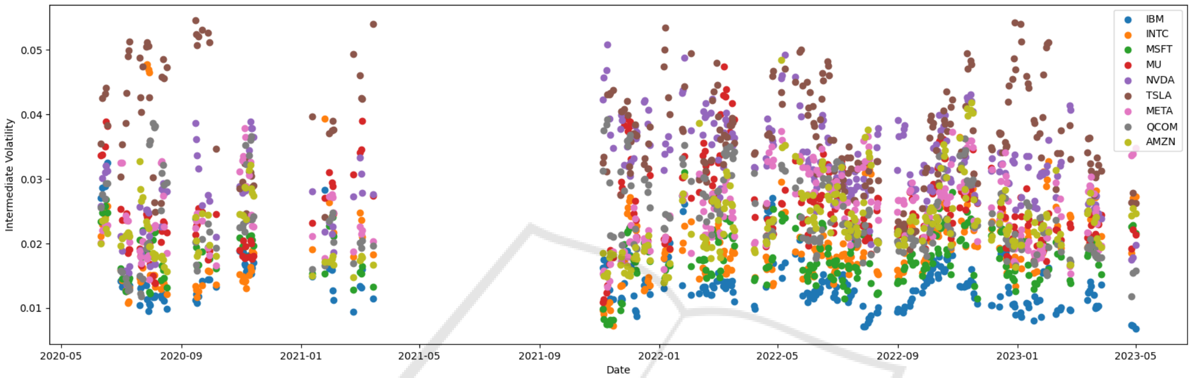

Figure 3: Kmeans++ clusters with $k = 3$ of the Historical Volatility estimators dataset, from 1st May 2020 to 1st May 2023.

variable and Y represents the potential effect variable. The null hypothesis (H0) states that X does not Granger cause Y, while the alternative hypothesis (H1) states that X does Granger cause Y. The F-test is defined as follows:

$$F-test = \frac{[(RSS_{Y(t)} - RSS_{YX_{(t)}})/p]}{[(RSS_{YX_{(t)}})/(n-p-k)]},$$

where RSS is the *Residual Sum of Squares* for the two AutoRegressive models: $Y(t) = c_Y + \beta_{Y_1} * Y(t-1) + \beta_{Y_2} * Y(t-2) + \cdots + \beta_{Y_p} * Y(t-p) + \varepsilon_{Y(t)}$, and $X : Y(t) = c_{YX} + \beta_{YX_1} * X(t-1) + \beta_{YX_2} * X(t-2) + \cdots + \beta_{YX_p} * X(t-p) + \varepsilon_{YX(t)}$, with p the lag order, n the number of observations, and k the number of parameters in the models.

- **Step 2.** *F-statistic* comparison with the critical value from the F-distribution where the significance level has $\alpha = 0.05$. If the F-statistic is greater than the critical value, reject the null hypothesis (H0) and conclude that X Granger causes Y with statistical significance. If the F-statistic is not greater than the critical value, fail to reject the null hypothesis (H0) and conclude that there is no significant Granger causality between X and Y.

- **Step 3.** *Direction of causality:* If the volatility of Stock X Granger causes the volatility of Stock Y, it suggests that changes in Stock X's volatility can be used to predict changes in Stock Y's volatility.

3.6 The Algorithm

- **Regression Step:** For each pair of time series (X_i, Y_j), where $i \neq j$, we construct a linear regression model:

$$X_i = \beta_{0,ij} + \beta_{1,ij} Y_j + \varepsilon_{ij}$$

where $\beta_{0,ij}$ is the intercept, $\beta_{1,ij}$ is the regression coefficient, and ε_{ij} is the error term. We calculate the F-statistic and the Durbin-Watson statistic to evaluate the overall adequacy of the model and detect autocorrelation in the residuals, respectively.

- **GCT Step:** For each pair of time series (X_i, Y_j), we perform the Granger causality test. The model for the Granger test can be expressed as $X_i(t) = \alpha_{ij} + \sum_{k=1}^{n} \beta_{k,ij} X_i(t-k) + \sum_{k=1}^{n} \gamma_{k,ij} Y_j(t-k) + \varepsilon_{ij}(t)$, where $X_i(t)$ is the current value of X_i, $X_i(t-k)$ and $Y_j(t-k)$ are the lagged values of X_i and Y_j, respectively, and $\varepsilon_{ij}(t)$ is the error term. If the coefficients $\gamma_{k,ij}$ are statistically different from zero, we reject the null hypothesis and conclude that Y_j Granger causes X_i.

4 RESULTS AND DISCUSSIONS

4.1 The Experiment

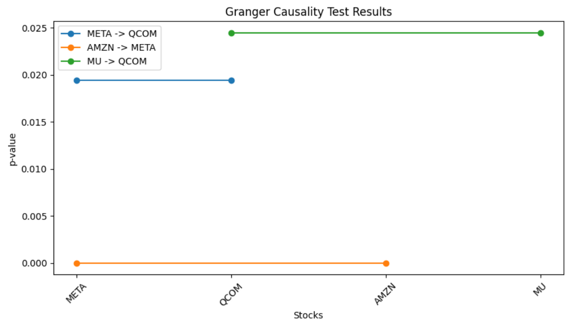

Figure 4: Co-integration via GCT.

The VolTS algorithm iterates the daily lags in a range from 2 to 30 days to determine the best result. In this experiment, the best result is achieved with lags=5, where '*best*' is considered when there is direction coherency among the stocks with the maximum cardinality of the set of stocks. In other words, the GCT direction does not generate the acyclic graph in the connection among the highest number of nodes, as shown in figure 5.

In figure 7, we can see how the GCT suggests buying QCOM when META has a positive trend and vice versa, the same thing with MU. Furthermore, when AMZN price increases, it is time to buy META and so on.

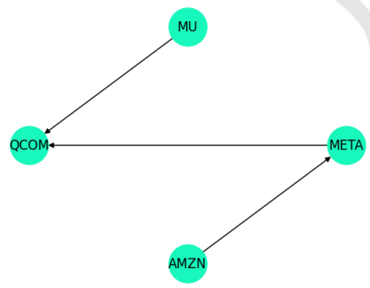

Figure 5: The best Acyclic Graph of the co-integration.

Figure 6 shows the scatter plot to visualise whether there are patterns in the residuals that suggest autocorrelation and to assess the overall adequacy of a regression model. We can see how the stocks META, QCOM, and AMZN confirmed their autocorrelation.

The experiment results indicate that the volatility-based trading strategy has performed well during the tested period from 8th April 2023 to 1st June 2023. The strategy resulted in a total gain of 231.77$ in 40 days of market opening, starting with an initial budget of 1000$ per stock. The exposure time of the positions being open was quite high at 88.89% for all the stocks, indicating active trading and frequent changes in the portfolio.

Tab. 2 contains further details about the performance metrics of the strategy and shows how the total amount in the portfolio is increased to 3231.77$ (7.725%), which is a positive sign of profitable trading, also considering the fixed commission of 9$ per trade. Notice that, the managing of the budget is set in compounded mode, so the full amount is reused for each trade.

4.2 The Backtesting

The analysis of individual stocks' performance is presented in figure 8 about META co-integration. The trades of META bought following the AMZN trend resulted in a Profit and Loss (PnL) of 1.281%, with a return of 9.721%. This return outperforms the B&H strategy, which would have yielded a return of 6.684%.

Figure 9 shows the trades of QCOM bought following the META trend showed a PnL of 2.774%, with a return of 12.866% compared to the B&H return of 9.235%. Furthermore, the trades of QCOM bought following the MU trend resulted in a PnL of 1.562%, with a return of 6.302% as opposed to the B&H return of 3.969%.

We compare, our backtesting trades to the optimal portfolio derived against 10000 possible portfolios constructed, in the same testing period, using the Markowitz Efficient Frontier (MEF), with the same 3000$ of budget and the constraint of 1000$ invested in META and $2 \times$ 1000$ in QCOM. MEF identifies the best portfolio with the highest TR of 3164.94$ when the volatility, measured with the standard deviation (72.41), is in the average. This confirms our idea to exploit the mid-volatility and highlights that our trading approach wins with a TR of 3231.77$, so 2.23% more than the optimal portfolio.

5 CONCLUSION

In this work, we propose an effective method to handle volatility in trading strategy and combine causality by the Historical Volatility Granger Causality Test implemented in the AITA framework with the module VolTS. The innovation of our system lies in selecting moderately volatile assets using Historical Volatility Estimators on market data and determining the most profitable stock pairings using K-means++ combined with a statistical method to choose the predictive property of our approach.

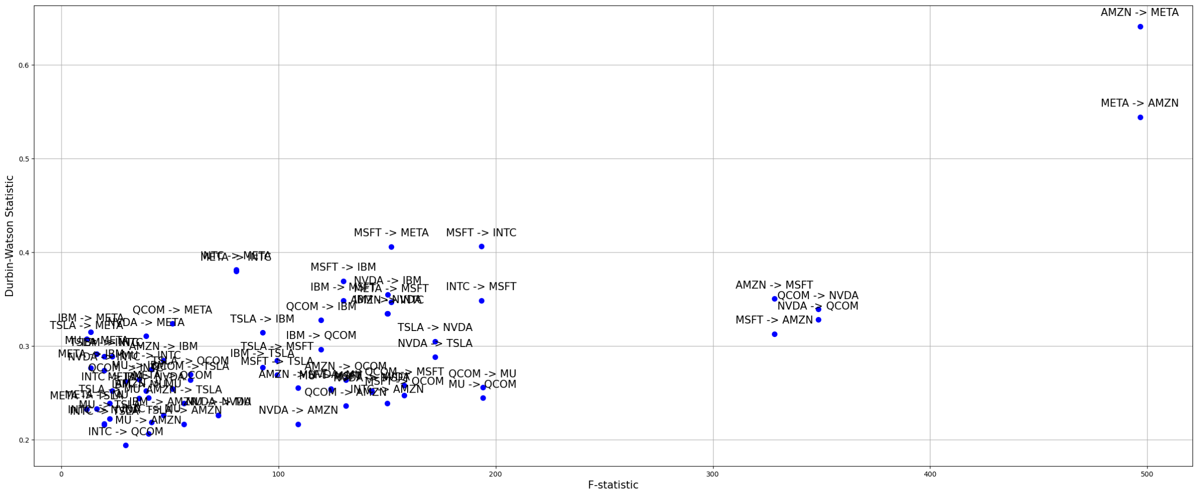

Figure 6: F-Statistic for regression model evaluation and Durbin-Watson statistics for autocorrelation in residuals.

Table 2: Results of the backtesting in the experiment.

Stock	Trades	Win rate (%)	TR ($)	SR	SoR	CR	MDD (%)
AMZN ->META	16	37.5	1045.01	1.1784	4.6421	14.264	1.77
META ->QCOM	16	43.75	1110.11	3.8511	44.4613	248.342	-1.33
MU ->QCOM	16	56.25	1076.65	1.2130	6.3624	23.687	-7.65

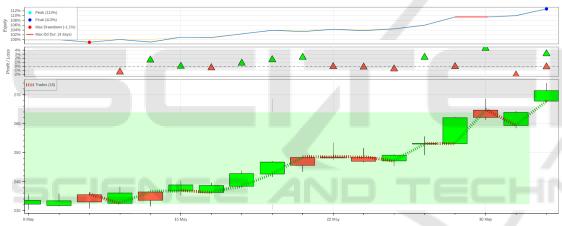

Figure 7: Co-integration AMZN to META without spurious correlation.

Figure 8: Co-integration META to QCOM without spurious correlation.

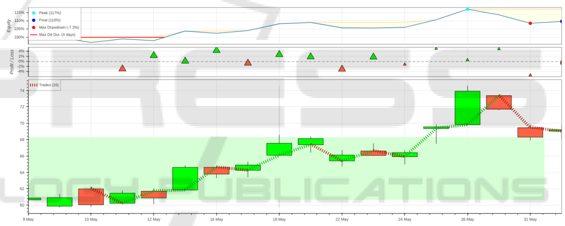

Figure 9: Co-integration MU to QCOM without spurious correlation.

Our future research areas include improving text data handling techniques through dataset optimization approaches (Letteri et al., 2020b),(Letteri et al., 2021b) and incorporating domain expert knowledge to enhance the model's understanding of price and volume data. Furthermore, we will expose the AITA framework's API as a secure service to thwart botnet attacks using Deep Learning models (Letteri et al., 2019b)(Letteri et al., 2019a). To enhance resilience, we plan to create a Multi-agent System which features transparent Ethical Agents for customer service (Dyoub et al., 2020) or combines logic constraint and DRL (Gasperis et al., 2023). We will evaluate dialogues (Dyoub et al., 2021) with guidance from an ethical teacher (Dyoub et al., 2022), also in other contexts like technology-enhanced learning (Angelone et al., 2023).

REFERENCES

Angelone, A., Letteri, I., and Vittorini, P. (2023). First evaluation of an adaptive tool supporting formative assessment in data science courses. In *Methodologies and Intelligent Systems for Technology Enhanced Learning, 13th International Conference, MIS4TEL 2023, Guimaraes, Portugal, 12-14 July 2023*, volume 764 of *Lecture Notes in Networks and Systems*, pages 144–151. Springer.

Arthur, D. and Vassilvitskii, S. (2007). K-means++: The advantages of careful seeding. In *Proceedings of the*

Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '07, page 1027–1035, USA. Society for Industrial and Applied Mathematics.

Dyoub, A., Costantini, S., and Letteri, I. (2022). Care robots learning rules of ethical behavior under the supervision of an ethical teacher (short paper). In *Joint Proceedings of the 1st International Workshop on HYbrid Models for Coupling Deductive and Inductive ReAsoning (HYDRA 2022) and the 29th RCRA Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion (RCRA 2022) co-located with the 16th International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR 2022), Genova Nervi, Italy, September 5, 2022*, volume 3281 of *CEUR Workshop Proceedings*, pages 1–8. CEUR-WS.org.

Dyoub, A., Costantini, S., Letteri, I., and Lisi, F. A. (2021). A logic-based multi-agent system for ethical monitoring and evaluation of dialogues. In *Proceedings 37th International Conference on Logic Programming (Technical Communications), ICLP Technical Communications 2021, Porto (virtual event), 20-27th September 2021*, volume 345 of *EPTCS*, pages 182–188.

Dyoub, A., Costantini, S., Lisi, F. A., and Letteri, I. (2020). Logic-based machine learning for transparent ethical agents. In *Proceedings of the 35th Italian Conference on Computational Logic - CILC 2020, Rende, Italy, October 13-15, 2020*, volume 2710 of *CEUR Workshop Proceedings*, pages 169–183. CEUR-WS.org.

Engle, R. F. and Granger, C. W. J. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, 55(2):251–276.

Gasperis, G. D., Costantini, S., Rafanelli, A., Migliarini, P., Letteri, I., and Dyoub, A. (2023). Extension of constraint-procedural logic-generated environments for deep q-learning agent training and benchmarking. *J. Log. Comput.*, 33(8):1712–1733.

Kirchgässner, G. and Wolters, J. (2007). *Granger Causality*, pages 93–123. Springer Berlin Heidelberg, Berlin, Heidelberg.

Letteri, I. (2023a). AITA: A new framework for trading forward testing with an artificial intelligence engine. In *Proceedings of the Italia Intelligenza Artificiale - Thematic Workshops co-located with the 3rd CINI National Lab AIIS Conference on Artificial Intelligence (Ital IA 2023), Pisa, Italy, May 29-30, 2023*, volume 3486 of *CEUR Workshop Proceedings*, pages 506–511. CEUR-WS.org.

Letteri, I. (2023b). Volts: A volatility-based trading system to forecast stock markets trend using statistics and machine learning.

Letteri, I., Cecco, A. D., Dyoub, A., and Penna, G. D. (2020a). A novel resampling technique for imbalanced dataset optimization. *CoRR*, abs/2012.15231.

Letteri, I., Cecco, A. D., Dyoub, A., and Penna, G. D. (2021a). Imbalanced dataset optimization with new resampling techniques. In Arai, K., editor, *Intelligent Systems and Applications - Proceedings of the 2021 Intelligent Systems Conference, IntelliSys 2021, Amsterdam, The Netherlands, 2-3 September, 2021, Volume 2*, volume 295 of *Lecture Notes in Networks and Systems*, pages 199–215. Springer.

Letteri, I., Cecco, A. D., and Penna, G. D. (2020b). Dataset optimization strategies for malwaretraffic detection. *CoRR*, abs/2009.11347.

Letteri, I., Cecco, A. D., and Penna, G. D. (2021b). New optimization approaches in malware traffic analysis. In *Machine Learning, Optimization, and Data Science - 7th International Conference, LOD 2021, Grasmere, UK, October 4-8, 2021, Revised Selected Papers, Part I*, volume 13163 of *Lecture Notes in Computer Science*, pages 57–68. Springer.

Letteri, I., Penna, G. D., and Caianiello, P. (2019a). Feature selection strategies for HTTP botnet traffic detection. In *2019 IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2019, Stockholm, Sweden, June 17-19, 2019*, pages 202–210. IEEE.

Letteri, I., Penna, G. D., and Gasperis, G. D. (2019b). Security in the internet of things: botnet detection in software-defined networks by deep learning techniques. *Int. J. High Perform. Comput. Netw.*, 15(3/4):170–182.

Letteri, I., Penna, G. D., Gasperis, G. D., and Dyoub, A. (2022). Dnn-forwardtesting: A new trading strategy validation using statistical timeseries analysis and deep neural networks.

Letteri, I., Penna, G. D., Gasperis, G. D., and Dyoub, A. (2023). Trading strategy validation using forwardtesting with deep neural networks. In Arami, M., Baudier, P., and Chang, V., editors, *Proceedings of the 5th International Conference on Finance, Economics, Management and IT Business, FEMIB 2023, Prague, Czech Republic, April 23-24, 2023*, pages 15–25. SCITEPRESS.

Niennattrakul, V. and Ratanamahatana, C. A. (2007). On clustering multimedia time series data using k-means and dynamic time warping. In *2007 International Conference on Multimedia and Ubiquitous Engineering (MUE'07)*, pages 733–738.

Wahid, A. and Chandra Sekhara Rao, A. (2020). An outlier detection algorithm based on knn-kernel density estimation. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.

Yang, D. and Zhang, Q. (2000). Drift-independent volatility estimation based on high, low, open, and close prices. *The Journal of Business*, 73(3):477–492.