Predicting Bullish Reversals with LightGBM: A Systematic Trading Strategy for Brazilian Equities

Wiliams Alves[1], Gustavo Callou[1]

[1]Universidade Federal Rural de Pernambuco (UFRPE)

{wiliams.alves,gustavo.callou}@ufrpe.br

Abstract. *Predicting stock market behavior is a complex challenge. This article presents an automated trading system based on machine learning, focused on identifying bullish reversals in the Brazilian financial market. The methodology combines technical and sectoral indicators, the LightGBM classifier, signal post-processing with non-maximum suppression, and backtesting with ATR-based risk management. Evaluated on assets representing different sectors of the economy (ABEV3, GGBR4, ITUB4, and VIVT3), the approach achieved PR-AUC values ranging from 0.955 to 0.979, win rates between 54.8% and 78.6%, and average returns per trade between 2.20% and 5.08%, consistently outperforming the buy-and-hold benchmark.*

1. Introduction

The application of reversal forecasting techniques in financial time series remains a recurring challenge in both academia and market practice. In environments characterized by high volatility, information asymmetry, and considerable price efficiency, identifying potential turning points following sharp downward movements can enable the development of investment strategies with a more favorable risk return profile.

Commonly, quantitative methods rely on technical indicators such as moving averages, the Relative Strength Index (RSI), and the Average True Range (ATR), which, when combined with statistical or econometric models, help identify patterns in price movements [Edwards et al. 2018]. Other methods employ deep learning models, such as Long Short-Term Memory (LSTM), for trend identification with the aid of technical analysis tools [Lopes de Amorim et al. 2024]. In this context, this study proposes a trading system that integrates bullish reversal detection with the Light Gradient Boosting Machine (LightGBM) algorithm, chosen for its computational efficiency, robustness to high-dimensional data, and ability to handle class imbalance. The process includes a supervised labeling stage for reversals, defined based on the occurrence of relevant drawdowns followed by price recovery within a predefined horizon, which extends the triple-barrier labeling approach [De Prado 2018].

The methodology employs time series cross validation, ensuring the preservation of the chronological order of data and preventing information leakage. Subsequently, a post-processing step is conducted to determine an optimal threshold based on the F1-score, along with the application of Non-Maximum Suppression (NMS), aiming to reduce redundant signals and improve the distribution of trade entries. The strategy is evaluated through a backtest that incorporates transaction costs, stop loss and take profit rules based on ATR, as well as a cooldown period. The experiments use data from 2010 to 2025 for four major assets traded on the Brazilian stock exchange: Ambev (ABEV3), Gerdau

(GGBR4), Itaú Bank (ITUB4), and Telefônica Brasil (VIVT3), covering the consumer staples, steel, financial, and communications sectors. The analysis includes both classification metrics and financial performance indicators, including a comparison with the Buy and Hold (B&H) strategy.

The main contributions of this work are concentrated in two central aspects: (i) the proposal of a modular and reproducible trading model, covering everything from feature engineering to the backtesting stage, incorporating contemporary machine learning and risk management techniques; and (ii) the demonstration that the use of post-processing with Non-Maximum Suppression, combined with F1-score-based thresholding, significantly improves the strategy's efficiency by reducing trade frequency and concentrating signals at more reliable moments.

The remainder of this paper is organized as follows. Section 2 discusses related work. Section 3 presents the necessary theoretical foundations. Section 4 describes the proposed methodology in detail. Section 5 presents and analyzes the obtained results. Finally, Section 6 provides conclusions and suggestions for future research.

2. Related Works

The literature includes numerous studies that explore the use of machine learning techniques for forecasting financial time series of assets traded on stock exchanges. For instance, [Zhao et al. 2023] introduces a cost-sensitive variant of LightGBM, which employs a customized loss function capable of dynamically adjusting the weight of each observation according to the complexity of its classification; the model is evaluated using data from the Shanghai, Hong Kong, and NASDAQ exchanges.

The application of the Directional Changes (DC) paradigm for predicting trend reversals in foreign exchange markets was investigated by [Adegboye and Kampouridis 2021]. In this approach, a classification model is used to determine whether a directional change will be followed by an overshoot, while a regression model is employed to estimate the duration of the subsequent movement. In turn, [Zhou et al. 2024] proposes hybrid models such as CNN-LightGBM-Fix and CNN-LightGBM-Wright, in which convolutional neural networks (CNNs) are used to extract local features, which are then combined with global patterns captured by LightGBM.

When comparing performance, [Guo et al. 2021] introduced the LSTM-LightGBM model for forecasting stock price fluctuations, showing that the hybrid approach outperforms standalone models in metrics such as RMSE and accuracy. Similarly, [Yang et al. 2021] demonstrated that combining XGBoost with LightGBM yields superior results compared to using each model independently. These findings highlight the effectiveness of LightGBM as a core component in financial forecasting experiments.

In a comparative study between LightGBM and ARIMA using technical indicators, [Yuvaraj et al. 2025] showed that LightGBM achieved coefficients of determination exceeding 99%, highlighting its ability to capture complex relationships in market data. In [Gupta and Kumar 2023], a framework is proposed that combines LightGBM with a Harris Hawk-based optimization algorithm to fine-tune hyperparameters and mitigate overfitting. The methodology also incorporates residual correction through a Markov chain, leading to improvements in error metrics. This emphasis on model generalization is closely related to the use of time-series cross-validation (TimeSeriesSplit) in the present study, ensuring that the model is evaluated across different temporal segments.

Overall, recent literature provides substantial evidence of LightGBM's potential in financial applications, particularly when integrated with error cost control mechanisms, feature fusion strategies, and hyperparameter optimization techniques. The present work advances this line of research by focusing on the specific phenomenon of bullish reversals following significant drawdowns, incorporating a comprehensive pipeline that spans from feature engineering with sector-based indicators to temporal validation and a realistic backtest framework using ATR, based stop-loss and take-profit rules, thereby addressing a gap not yet fully explored in an integrated manner in the literature.

3. Theoretical Framework

This section outlines the core concepts underlying the study: drawdown and reversal dynamics, the LightGBM algorithm, evaluation metrics for imbalanced classification, signal post-processing techniques, and backtesting procedures incorporating ATR-based risk management.

3.1. Drawdown and Uptrend Reversal

Drawdown quantifies the decline of an asset relative to its previous peak value, constituting a fundamental metric in the assessment of risk in investment strategies. Given a time series of prices P_t and a lookback window of W periods, the drawdown at time t is defined according to Equation 1.

$$DD_t = \frac{P_t}{\max_{k \in [t-W+1,\, t]} P_k} - 1 \tag{1}$$

This formulation yields non-positive values, such that $DD_t = 0$ when the current price equals the maximum within the considered window, and $DD_t < 0$ otherwise. A pronounced drawdown may signal that the asset has recently experienced a substantial loss, potentially indicating exhaustion of selling pressure and preceding a reversal, as discussed in the literature [Brandi 2021]. A bullish reversal, in turn, is defined based on Equation 2, representing the price recovery over a future horizon, measured by the maximum return achieved within the subsequent H trading sessions.

$$R_t^{\max} = \frac{\max_{j \in [t+1,\, t+H]} P_j}{P_t} - 1 \tag{2}$$

The occurrence of a reversal is characterized by the combination of a sufficiently negative drawdown and a sufficiently positive maximum future return. This definition, inspired by the triple-barrier method proposed by [De Prado 2018], enables supervised labeling of the data and formulates a binary classification problem aimed at predicting whether time t will be followed by a reversal.

3.2. Light Gradient Boosting Machine (LightGBM)

LightGBM [Ke et al. 2017] is a tree-based gradient boosting algorithm that leverages histogram-based techniques to reduce computational cost by discretizing continuous variables. It incorporates strategies such as Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB) to enhance efficiency and scalability in high-dimensional datasets. In financial applications, it stands out for its ability to handle imbalanced classes, its robustness against overfitting, and its use of binary cross-entropy as the loss function, leading to more accurate forecasts, especially when combined with time series cross-validation [Zhao et al. 2024].

3.3. Evaluation Metrics for Imbalanced Classes

In scenarios where positive-class events, such as reversals, are rare, accuracy may not adequately reflect model performance. In such cases, metrics derived from the confusion matrix, such as true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN), become more appropriate. Precision measures the proportion of correctly predicted positive instances, while recall evaluates the model's ability to identify actual events; the F1-score corresponds to the harmonic mean of these two metrics, balancing their trade-offs, as defined in Equation 3. Additionally, the precision–recall curve enables analysis of this trade-off across different decision thresholds, with its area (PR-AUC) being particularly suitable for imbalanced data [Davis and Goadrich 2006]. Finally, the ROC curve and its associated area (ROC-AUC) provide a more general assessment of the model's discriminative ability, regardless of the chosen threshold [Fawcett 2006].

$$Precision = \frac{TP}{TP + FP}, \quad Recall = \frac{TP}{TP + FN}, \quad \textit{F1-score} = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{3}$$

3.4. Signal Post-processing

After obtaining the predicted probabilities $\hat{y}_t \in [0, 1]$ for each time step, these values must be translated into actionable trading signals. This procedure involves two main steps: selecting an optimal threshold and applying non-maximum suppression.

First, regarding threshold selection, the cutoff value τ is determined by maximizing the F1-score on the training data. This strategy, commonly referred to as threshold tuning, is preferable to adopting a fixed value such as $\tau = 0.5$, as it accounts for class imbalance and the economic impact of classification errors [Provost and Fawcett 2001]. Recent studies further indicate that optimizing the threshold based on the F1-score is particularly effective when the goal is to balance precision and recall [Hancock et al. 2022].

Next, non-maximum suppression (NMS) is applied. In many cases, especially during periods of strong model confidence, multiple signals may be generated within a short time span. To prevent redundant consecutive entries, NMS is employed: within a window of L days, only the observation with the highest probability (provided it exceeds the threshold) is retained as the final signal. Originally developed in the field of computer vision [Neubeck and Van Gool 2006], this technique has proven effective in reducing the number of trades without compromising the detection of relevant events, and has more recently been adapted to financial time series applications [Adegboye and Kampouridis 2021].

3.5. Operations Simulation and Risk Management

The financial evaluation of the strategy is conducted through simulations (backtests) that replicate realistic market conditions. In this process, the test set is traversed chronologically, and for each valid signal, a long position is opened at the closing price of that day. To better reflect real, world conditions, a transaction cost proportional to the traded volume is considered, typically around 0.1% in the Brazilian market, applied to both entry and exit, as recommended in the literature [Bailey et al. 2014].

The definition of trade exits is based on stop loss and take profit levels, determined using the Average True Range (ATR) at the time of entry. Introduced by [Wilder 1978], ATR is widely used to measure volatility and define dynamic exit thresholds [Vezeris et al. 2018]. Its calculation follows Equation 4, which captures the true range of price variation.

$$ATR_t = \frac{1}{n}\sum_{i=0}^{n-1} TR_{t-i}, \quad \textit{where} \quad TR_t = \max(H_t - L_t, |H_t - C_{t-1}|, |L_t - C_{t-1}|) \tag{4}$$

Where ATR_t is the ATR at period t, n represents the number of periods used in the moving average of the True Range (TR), and TR_t corresponds to the TR at period t. The latter is calculated from the high price H_t, the low price L_t, and the previous period's closing price C_{t-1}.

The exit levels are then defined as multiples of the ATR at the entry point, as shown in Equation 5:

$$SL = P_{entry} - k_{SL} \cdot ATR_{entry}, \quad TP = P_{entry} + k_{TP} \cdot ATR_{entry} \tag{5}$$

Here, SL and TP denote the stop loss and take profit levels, respectively, P_{entry} is the trade entry price, k_{SL} and k_{TP} are the multipliers that determine the distance of these thresholds from the entry price, and ATR_{entry} indicates the ATR value calculated at the time of entry.

Using ATR multiples (commonly $k_{SL} = k_{TP} = 3$) makes the thresholds adaptive to recent volatility, preventing premature exits during turbulent periods and maintaining a consistent risk profile [Chan 2026]. Additionally, a cooldown period is enforced between consecutive trades to avoid excessive risk concentration and promote capital diversification.

At the end of the backtest, performance metrics are computed. The average return per trade is given by Equation 6:

$$\bar{r} = \frac{1}{N}\sum_{n=1}^{N} r_n \tag{6}$$

The win rate represents the proportion of trades with positive returns. The Sharpe ratio, defined in Equation 7, measures risk-adjusted return based on the mean and standard deviation of daily returns [Sharpe 1966]:

$$Sharpe = \frac{\bar{r}_{daily}}{\sigma_{daily}}\sqrt{252} \tag{7}$$

Finally, the maximum drawdown (MDD) quantifies the largest cumulative loss over time, as defined in Equation 8, where E_t represents the capital at time t:

$$MDD = \min_t \left(\frac{E_t - \max_{s \leq t} E_s}{\max_{s \leq t} E_s} \right) \tag{8}$$

For model validation, the chronological order of the data is preserved to prevent information leakage. The TimeSeriesSplit technique is used, where training data always precedes validation data. Recent studies highlight the importance of this approach in financial applications [Bergmeir et al. 2018]. In this work, 10 splits are performed, and features are normalized using a StandardScaler fitted only on the training data. This procedure provides a more realistic estimate of out-of-sample performance and improves robustness over time.

4. Methodology

The proposed methodology is structured into several stages, ranging from data collection and preprocessing to the training of the classification model and the simulation of the trading strategy, as illustrated in the methodology shown in Figure 1. In it, the process begins with data loading, followed by feature engineering and label definition. Next, classifier modeling and validation are performed, whose outputs undergo signal post-processing, where if there are valid signals, the backtest is carried out, otherwise, flow is interrupted and the process ends. The entire process was designed to ensure reproducibility and to prevent information leakage between the training and testing sets by strictly respecting the chronological order of events.

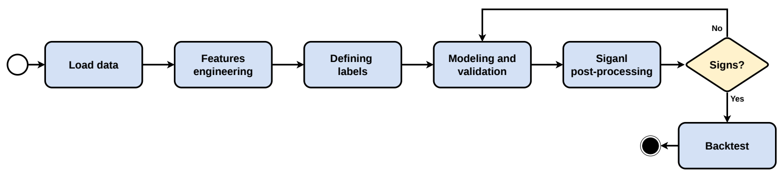

Figure 1. Methodology applied in the negotiation model.

4.1. Data and Feature Engineering

The data used in this research were obtained through the yfinance library, an API from Yahoo Finance, which enabled the download of daily time series for opening, high, low, closing prices, and trading volume of four highly liquid assets listed on the Ibovespa (IBOV): Ambev (ABEV3), Gerdau (GGBR4), Itaú Unibanco (ITUB4), and Telefônica Brasil (VIVT3). The selection of these assets considered both sectoral representativeness and the availability of long, reliable historical series. Additionally, the IBOV index itself was incorporated as a contextual variable, allowing the model to capture coordinated movements between each asset and the broader market.

In the feature engineering stage, the data were used to construct a set of predictive variables designed to capture asset-specific technical patterns as well as sectoral influences derived from the benchmark index. All variables were calculated exclusively using past information, employing rolling windows that do not incorporate future data, thereby avoiding bias or data leakage. Among the technical features are return rates, volatility, distances from moving averages, the Relative Strength Index (RSI), volume Z-score, and the Average True Range (ATR). To incorporate market context, additional sectoral features derived from the IBOV index were included: 1- and 5-day returns, capturing recent market direction; a 20-day rolling correlation between asset and index returns, indicating the degree of synchronization with overall market movements; and the ratio between the asset's trading volume and the moving average of the index volume, highlighting periods when investor interest in the asset stands out relative to the broader market.

4.2. Defining Labels

The definition of the target variable followed the approach presented in Section 3.1, which identifies upward reversal events based on the conjunction of a severe drawdown and a significant recovery. For each time step t, the drawdown was computed over a retrospective window of $W = 20$ days, as defined in Equation 1, and the maximum future return was calculated over a horizon of $H = 20$ days, as defined in Equation 2. The binary label y_t takes the value 1 when the drawdown is below $\theta_{dd} = -0.03$ and the maximum future return exceeds $\theta_{rev} = 0.03$; otherwise, $y_t = 0$.

It is important to note that the label construction uses past information for the drawdown calculation and future information exclusively for defining the target. The features used in the model are computed solely from data prior to time t, and observations for which a full future horizon is not available ($t > T - H$) are discarded. This procedure eliminates any risk of data leakage, ensuring that the model is trained and evaluated under conditions that replicate a real-time decision-making environment.

4.3. Modeling and Validation

With the data properly processed and labeled, a temporal split was performed into two subsets: 70% allocated for training and the remaining 30% for evaluation. This split preserves chronological order, ensuring that the model is evaluated on periods subsequent to those used for learning, a condition essential to simulate real-world performance.

To ensure model stability and generalization capability, time series cross-validation (TimeSeriesSplit) with 10 folds was applied. In each fold, the features were normalized using StandardScaler, fitted exclusively on the training set to avoid contamination of normalization statistics. The LightGBM classifier was used, whose fundamentals were detailed in Section 3.2. The cross-validation results provided a robust estimate of the expected performance on unseen data and guided potential adjustments prior to final training.

4.4. Signal Post-processing

At this stage, the objective is to transform the continuous probabilities generated by the model into effective trading signals, ensuring that only the most relevant events are selected. As discussed in Section 3.4, this process consists of two steps: determining the optimal probability threshold and applying non-maximum suppression.

The optimal threshold τ was defined as the one that maximizes the F1-score on the training data, where the true labels are known. After applying the threshold, non-maximum suppression (NMS) was performed using a temporal window of 20 days: within this window, only the point with the highest reversal probability was retained as the final signal, while the others were discarded. This approach prevents multiple redundant signals over short time intervals, concentrating entries at moments of highest model confidence.

4.5. Backtest

The trading strategy simulation (backtest) was conducted on the test set, strictly preserving chronological order and incorporating realistic trading elements, as established in Section 3.5. The process started with an initial capital of R$ 100,000.00, with each trade allocating the full available capital (full reinvestment). Transaction costs of 0.1% per trade (0.05% on entry and 0.05% on exit) were considered.

At each time step i, the presence of a signal was verified. When triggered, a buy operation was executed at the closing price, with transaction costs immediately deducted. Exit levels were defined based on the ATR at entry, using multipliers $k_{sl} = 3$ and $k_{tp} = 9$, as in Equation 5. Positions were closed when the stop loss, take profit, or the 20-day horizon was reached. After closing, capital was updated, transaction costs applied, and the equity curve recorded for performance evaluation. The backtest was performed separately for each asset, enabling comparative and consistency analysis.

5. Results and Discussion

This section presents the performance of the trading model proposed in the methodology, following the workflow illustrated in Figure 1. The results cover both the error metrics of the trained model for predicting upward reversals and the evaluation metrics obtained from the subsequent backtest.

5.1. Results of the Reversal Model Forecast

Table 1 confirms the predictive effectiveness of the proposed methodology. The PR-AUC values range from 0.955 to 0.979, while ROC-AUC values vary between 0.959 and 0.983 across the analyzed assets, indicating consistently strong discriminative capability. Precision, recall, and F1-score remain high across all cases, with values between 0.895 and 0.940. This demonstrates a favorable balance between correctly identified events and the control of false positives.

The optimal threshold τ varies significantly across assets, from 0.25 for ABEV3 to 0.91 for ITUB4. This variation reflects differences in the distribution of predicted probabilities and the model's calibration for each asset. Overall, these results highlight the robustness and consistency of the approach across different market sectors.

Table 1. Results of the evaluation metrics for the trained model.

Stock Symbols	PR-AUC	ROC-AUC	Threshold	Precision	Recall	F1-score
ABEV3	0.972	0.964	0.25	0.901	0.901	0.901
GGBR4	0.977	0.967	0.70	0.921	0.920	0.920
ITUB4	0.979	0.983	0.91	0.940	0.939	0.939
VIVT3	0.955	0.959	0.62	0.895	0.895	0.895

Beyond the quantitative metrics, the plots in Figure 2 illustrate the effect of non-maximum suppression (NMS) in removing redundant signals, retaining only those associated with the highest probability peaks. The combination of these metrics and the visual analysis reinforces the model's robustness and the reliability of the generated trading signals.

5.2. Backtest Results

The simulation results presented in Table 2 confirm the overall consistency of the strategy across different assets. GGBR4 achieved the highest average return (5.08%) over 33 trades, with a win rate of 69.69% and a Sharpe ratio of 0.5409. VIVT3 and ITUB4 also showed solid performance, with average returns of 4.05% and 3.70%, and win rates of 66.66% and 78.57%, respectively. Notably, ITUB4 stands out for its higher accuracy.

ABEV3 presented a more modest performance, with an average return of 2.20%, a Sharpe ratio of 0.4172, and a win rate of 54.83%. However, it exhibited the lowest maximum drawdown (-7.73%), indicating a more conservative risk profile. Overall, maximum drawdowns across assets ranged from -7.73% (ABEV3) to -14.11% (VIVT3), suggesting a controlled risk profile while maintaining consistent returns.

The analysis of the maximum drawdown plots in Figure 3 reveals important differences in the risk profiles of the assets over the evaluated period. ABEV3 exhibits the lowest maximum

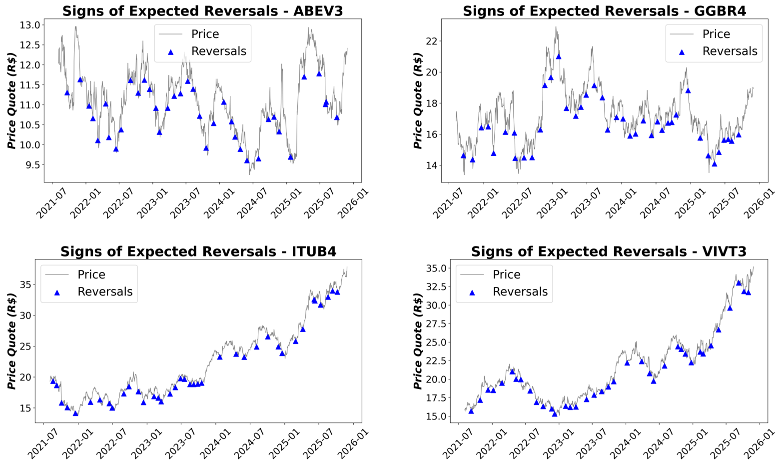

Figure 2. Performance charts for predicted signals using LightGBM, with the application of optimal threshold and NMS filters to the ABEV3, GGBR4, ITUB4, and VIVT3 stocks.

Table 2. Backtest metric results.

Stock Symbols	Trades	Average Return (%)	Sharpe Ratio	Win Rate (%)	Drawdown (%)
ABEV3	31	2.20	0.4172	54.83	-7.73
GGBR4	33	5.08	0.5409	69.69	-10.35
ITUB4	28	3.70	0.5766	78.57	-12.42
VIVT3	30	4.05	0.6401	66.66	-14.11

drawdown (-7.73%), indicating greater stability and lower exposure to severe losses, despite showing recurrent moderate declines. In contrast, VIVT3 presents the highest drawdown (-14.12%), mainly concentrated between 2022 and 2023, suggesting higher sensitivity to periods of market stress. GGBR4 (-10.35%) and ITUB4 (-12.43%) display intermediate behavior, with drawdowns more evenly distributed over time, including more pronounced events around 2022–2023. The synchronization of major drawdown periods across assets indicates the presence of systemic market shocks, while the variation in magnitude highlights different levels of resilience. This analysis is crucial for risk assessment and for interpreting the backtest results.

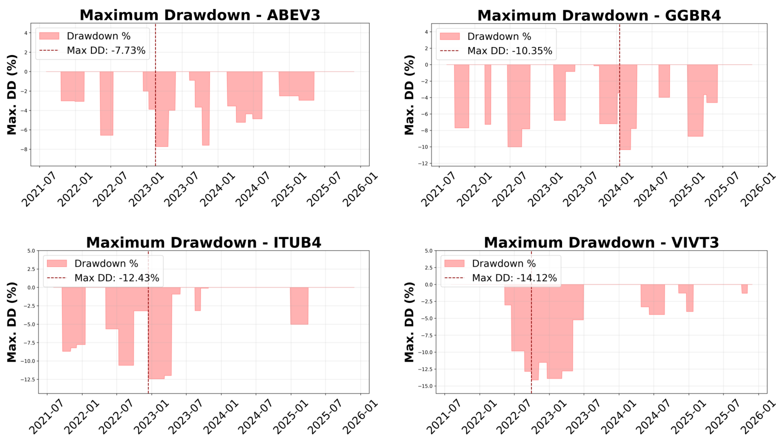

Figure 3. Drawdown charts for ABEV3, GGBR4, ITUB4, and VIVT3 stocks.

The superiority of the proposed strategy compared to a buy-and-hold (B&H) approach is evident in Table 3. The accumulated returns were significantly higher for all analyzed assets. GGBR4 stands out with a return of 349.74%, compared to only 12.40% under B&H. Similarly, ABEV3 achieved 88.40% versus 3.08%, while ITUB4 and VIVT3 reached 161.4% and 210.76%, outperforming their respective benchmarks of 95.7% and 121.38%.

Table 3. Comparison Results: Strategy vs. Buy and Hold.

Stock Symbols	Return Strategy (%)	Return B&H (%)
ABEV3	88.40	3.08
GGBR4	349.74	12.40
ITUB4	161.43	95.72
VIVT3	210.76	121.38

This superior performance is further supported by the visual analysis in Figure 4, which illustrates the higher accumulated capital achieved by the strategy compared to passive exposure. The equity curves in Figure 5 provide a consolidated view of these results, showing that the strategy consistently outperforms the B&H approach while maintaining controlled risk levels. Overall, these findings demonstrate the effectiveness of the proposed approach in capturing profitable market movements and consistently outperforming the traditional buy-and-hold strategy.

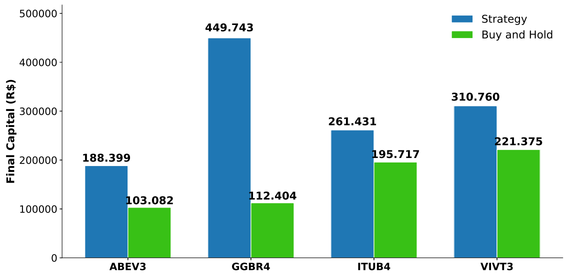

Figure 4. Comparison of final capital for the Strategy versus Buy and Hold (B&H) for the ABEV3, GGBR4, ITUB4, and VIVT3 assets.

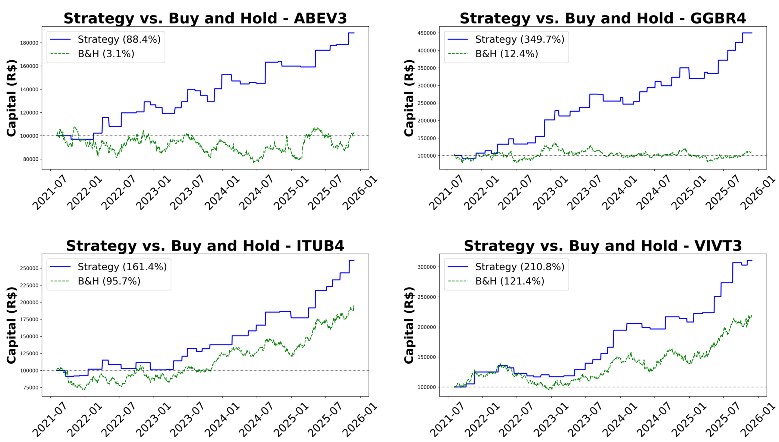

Figure 5. Equity curves comparing the Strategy versus Buy and Hold (B&H) for the ABEV3, GGBR4, ITUB4, and VIVT3 assets.

6. Conclusion

This work presented a methodology for detecting upward reversals based on LightGBM, feature engineering, pivot labeling, and post-processing with non-maximum suppression (NMS). The results demonstrated consistent and superior performance across all evaluated assets when compared to buy-and-hold (B&H), with classification metrics achieving PR-AUC and ROC-AUC above 0.95 and F1-score exceeding 0.89, indicating high discriminative capacity and low incidence of false positives. Temporal validation reinforced the model's robustness, confirming its stability and generalization capability, with results maintained at levels similar to those observed during training.

In the financial tests, the strategy consistently outperformed B&H across all analyzed assets. GGBR4 achieved the highest total return (349.74%), more than 28 times higher than B&H (12.40%), while ITUB4 obtained the highest win rate (78.57%). ABEV3 recorded a return of 88.40%, compared to 3.08% for the benchmark, and VIVT3 reached 210.76%, against 121.38% for B&H. Maximum drawdowns ranged from -7.73% to -14.11%, consistent with the strategy's risk profile, and Sharpe ratios, varying between 0.4172 and 0.6401, indicated attractive risk-adjusted returns.

Furthermore, the win rate between 54.8% and 78.6%, combined with an average return per trade ranging from 2.20% to 5.08%, reinforces the profit-generating potential of the methodology. Thus, the approach proves promising for systematic trading, opening possibilities for future work such as comparison with other quantitative strategies and other machine learning models, application of walk-forward validation with statistical significance analysis, and a more in depth overfitting analysis.

Acknowledgment

The authors would like to thank the National Council for Scientific and Technological Development (CNPq) - Brazil, Grant/Award Number 309338/2025-5, for the financial support for this research.

References

Adegboye, A. and Kampouridis, M. (2021). Machine learning classification and regression models for predicting directional changes trend reversal in fx markets. *Expert Systems with Applications*, 173:114645.

Bailey, D. H., Borwein, J. M., De Prado, M. L., and Zhu, Q. J. (2014). Pseudomathematics and financial charlatanism: The effects of backtest over fitting on out-of-sample performance. *Notices of the AMS*, 61(5):458–471.

Bergmeir, C., Hyndman, R. J., and Koo, B. (2018). A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis*, 120:70–83.

Brandi, V. R. (2021). Predictability of stock market indexes following large drawdowns and drawups. *Brazilian Review of Finance*, 19(1):1–23.

Chan, E. P. (2026). *Quantitative trading: how to build your own algorithmic trading business*. John Wiley & Sons.

Davis, J. and Goadrich, M. (2006). The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240.

De Prado, M. L. (2018). *Advances in financial machine learning*. John Wiley & Sons.

Edwards, R. D., Magee, J., and Bassetti, W. C. (2018). *Technical analysis of stock trends*. CRC press.

Fawcett, T. (2006). An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874.

Guo, Y., Li, Y., and Xu, Y. (2021). Study on the application of lstm-lightgbm model in stock rise and fall prediction. In *MATEC Web of Conferences*, volume 336, page 05011. EDP Sciences.

Gupta, V. and Kumar, E. (2023). H3o-lgbm: hybrid harris hawk optimization based light gradient boosting machine model for real-time trading. *Artificial Intelligence Review*, 56(8):8697–8720.

Hancock, J., Johnson, J. M., and Khoshgoftaar, T. M. (2022). A comparative approach to threshold optimization for classifying imbalanced data. In *2022 IEEE 8th International Conference on Collaboration and Internet Computing (CIC)*, pages 135–142. IEEE.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.

Lopes de Amorim, T., Callou, G., and Sena dos Santos, A. (2024). Compra e venda de ações negociadas na b3 utilizando lstm: análise de tendência e stop. *GESTÃO. Org: Revista Eletrônica de Gestão Organizacional*, 22(1).

Neubeck, A. and Van Gool, L. (2006). Efficient non-maximum suppression. In *18th international conference on pattern recognition (ICPR'06)*, volume 3, pages 850–855. Ieee.

Provost, F. and Fawcett, T. (2001). Robust classification for imprecise environments. *Machine learning*, 42(3):203–231.

Sharpe, W. F. (1966). Mutual fund performance. *The Journal of business*, 39(1):119–138.

Vezeris, D., Kyrgos, T., and Schinas, C. (2018). Take profit and stop loss trading strategies comparison in combination with an macd trading system. *Journal of risk and financial management*, 11(3):56.

Wilder, J. W. (1978). *New concepts in technical trading systems*. Greensboro, NC.

Yang, Y., Wu, Y., Wang, P., and Jiali, X. (2021). Stock price prediction based on xgboost and lightgbm. In *E3s web of conferences*, volume 275, page 01040. EDP Sciences.

Yuvaraj, S., Kumaran, J. C., Benschwartz, R., and Pradeepa, M. (2025). Technical indicator-driven stock prediction: A lightgbm and arima perspective. In *2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, pages 46–52. IEEE.

Zhao, X., Liu, Y., and Zhao, Q. (2023). Cost harmonization lightgbm-based stock market prediction. *Ieee Access*, 11:105009–105026.

Zhao, X., Liu, Y., and Zhao, Q. (2024). Improved lightgbm for extremely imbalanced data and application to credit card fraud detection. *IEEE Access*, 12:159316–159335.

Zhou, W., Jumahong, H., Cui, R., Wu, Y., Jing, C., and Lin, L. (2024). Predicting stock trends using web semantics and feature fusion. *International Journal on Semantic Web and Information Systems (IJSWIS)*, 20(1):1–25.