ANNALS OF DATA ANALYSIS

Intelligent Buy and Hold: Buy Signaling for Perennial Stocks with Random Forest

Gabriel Canuto de Alencar [1] * - https://orcid.org/0009-0008-5372-9264
João Gabriel Gama Vila Nova [2] * - https://orcid.org/0000-0002-6920-2650
Marcus Cezar Moreira Ferraz [3] * - https://orcid.org/0009-0003-2909-3140

[1] Department of Computer Science, Maceió University Center, Maceió, Brazil
Email: gabrielcanuto2011@hotmail.com
* corresponding author

[2] Department of Computer Science, Maceió University Center, Maceió, Brazil
Email: joao.nova@unima.edu.br

[3] Department of Computer Science, Maceió University Center, Maceió, Brazil
Email: marcusferraz98@gmail.com

RESEARCH ARTICLE

Annals of Data Analysis
https://periodicos.ufpe.br/revistas/ADA/
ISSN Online: xxxx-xxxx

Submitted on: 12.12.2025.
Accepted on: 01.08.2026.
Published on: 01.08.2026.

Abstract

The optimization of entry points is a critical factor in enhancing long-term returns within the Buy and Hold strategy. This project addresses this challenge by developing and evaluating a framework that utilizes Random Forest models to predict opportune buying moments in a portfolio of blue-chip stocks. The effectiveness of the signals generated by the models is validated through a backtesting system, which compares the performance of dynamic allocation strategies against passive benchmarks. The results indicate that Machine Learning-guided timing can offer significant advantages, and the project also contributes its methodological structure for testing quantitative strategies.

Keywords

Machine Learning, Portfolio Optimization, Random Forest, Buy and Hold, Backtesting

1. Introduction

Building an efficient investment portfolio is directly related to the careful selection of assets and the adoption of strategies that maximize risk-adjusted returns. In Brazil, despite the growing interest in the financial market, the number of individual investors is still relatively low, representing only 11% of the population (B3, 2024). This scenario is influenced by a lack of financial knowledge and the country's economic uncertainty, factors that make it difficult to choose solid assets and define long-term investment strategies (Pixley, 2004). The volatility and instability intrinsic to financial markets represent a central concern, especially for strategies focused on the short term. Conversely, literature often points to long-term horizon investing as an approach to mitigate the risks associated with these daily fluctuations. In this sense, long-term strategies, such as Buy and Hold, are a way for investors to benefit from structural economic growth, minimizing the impact of market noise (da Silva SANTOS et al., 2021).

Given this scenario of uncertainty, one of the most consolidated strategies for the individual investor is Buy and Hold, which stands out as a way to benefit from the yields, results, and appreciation that companies may present in the future, as pointed out by Damodaran (1997). This investment philosophy consists of acquiring high-quality and resilient assets with the intention of keeping them in the portfolio for long periods, ignoring short-term movements, whose predictability is often noted as low compared to long-term strategies (Haryanto, 2024). Its fundamental premise is that the value of solid companies tends to grow consistently over time, and that trying to predict shorter-term movements is often an inefficient and costly task for the non-professional investor.

Financial market forecasting is widely recognized as a challenging problem, given the nature of its time series, which are dynamic, non-linear, and chaotic (Deboeck, 1994), often influenced by macroeconomic factors and investor behavior. Faced with this complexity, a growing body of research has been dedicated to applying Machine Learning (ML) models, a subfield of Artificial Intelligence (AI) focused on developing models capable of learning patterns from data (El Naqa and Murphy, 2015). This approach has proven effective in asset analysis and forecasting, seeking to capture patterns that traditional methods might not identify (Ma et al., 2021). Several architectures have shown promising results in this field. Ensemble models such as Random Forest (RF) (Ballings et al., 2015), for example, are employed to forecast stock price trends. Other models involving deep learning techniques have been successfully applied in portfolio optimization based on return prediction, as demonstrated by Ma et al. (2021).

In the vast universe of the stock market, a fundamental challenge for the investor is the allocation of capital among a large number of available assets. A common approach to managing this complexity is asset pre-selection, a process in which the investment universe is deliberately restricted to a set of companies that meet established criteria (Silva et al., 2024).

That said, the present study proposes the development of a systematic investment framework that unites the Buy and Hold philosophy with the intelligence of ML techniques to optimize allocation in perennial assets. The methodology starts from a careful pre-selection of companies that simultaneously belong to the theoretical portfolios of the BOVESPA (IBOV) and Dividend Index (IDIV) indices, restricting the analysis universe to companies with high liquidity, stability, and a history of dividends. On this asset base, instead of a purely passive approach, an RF model is used to qualify the accumulation strategy, an application supported by works that demonstrate the model's potential to optimize market entry points (Zheng et al., 2024). The model is trained to identify potentially more favorable moments for market entry, optimizing capital allocation by capturing non-linear patterns that traditional methods might ignore. In this way, the contribution of this work is to create an optimized Buy and Hold strategy, which seeks to boost long-term returns through more strategic entries into the market.

2. Related Works

The application of ML in finance has proven to be a powerful tool for capturing complex patterns in large volumes of data, surpassing traditional statistical methods in some cases. Studies such as that of Aziz et al. (2022), for example, map the vast range of ML applications in the financial sector, which includes everything from portfolio optimization to sentiment analysis in news to predict market trends.

Within the vast field of ML techniques, RF is frequently employed due to its robustness against overfitting and its ability to handle high-dimensional data (Shaik and Srinivasan, 2018). This effectiveness in complex scenarios is confirmed in financial applications. In a comparative study, Ma et al. (2021) demonstrated the superiority of RF in predicting asset returns, where the algorithm not only outperformed deep learning models, achieving the lowest mean absolute error (MAE of 1.85×10-2) and mean squared error (MSE of 8.08×10-4), but the portfolio strategy based on its predictions (RF+MVF) reached a total return of 1780.34% in the Chinese market between 2012 and 2015 (disregarding transaction fees).

In addition to predicting returns, RF also proves effective in trend classification tasks. Zheng et al. (2024), for example, corroborate the potential of RF by using it to classify the trend of major American stocks (upward, sideways, or downward), achieving an accuracy of up to 98% in long-term predictions and surpassing other models such as SVM (Support Vector Machine). While these works validate the high predictive power of RF for predicting returns or general trends, our study differs by applying this technique to optimize entry points within a specific Buy and Hold strategy for a portfolio of perennial stocks.

3. Methodology

The methodology adopted in this project followed a structured approach in four main stages: data collection, data analysis, model training, and Trade System architecture. Below, in Figure 1, is the diagram of the processes carried out in the research methodology phase.

Figure 1: Methodology Process

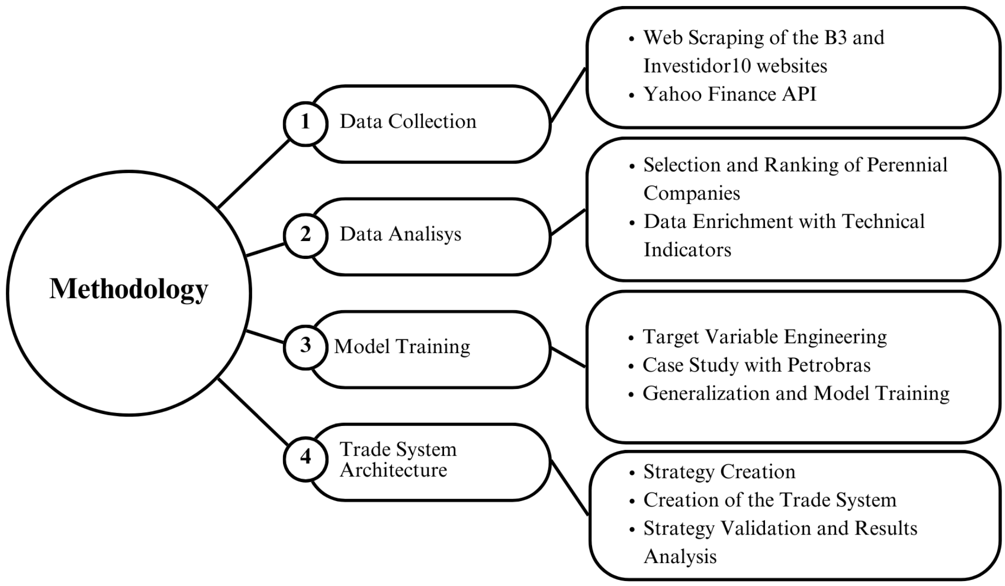

3.1. Data Collection

Initially, the Yahoo Finance Application Programming Interface (API) was used—a resource that allows automated communication between systems, to obtain information such as opening and closing prices, highs and lows, as well as trading volume. However, due to limitations of this API in providing extensive historical data, complementary data obtained via Web Scraping was incorporated, an automated data extraction technique from the web (Navlani et al., 2021), from platforms such as B3, Status Invest, and Investidor10 (B3, 2025a; Status Invest, 2025; Investidor10, 2025).

3.2. Data Analysis

The second stage consisted of data processing and preparation, using the Pandas library, which is widely used for data cleaning and formatting (Navlani et al., 2021). From this, the methodology followed these phases:

3.2.1. Selection and Ranking of Perennial Companies

This stage consisted of defining and ranking a universe of "perennial stocks." To this end, only companies whose shares simultaneously composed the theoretical portfolios of the IBOV and the IDIV were selected. This double selection criterion created an initial filter that prioritizes assets with high liquidity and market relevance and, at the same time, with a consistent history of dividend distribution.

To quantify the relevance of each asset, a customized score was created. Initially, the assets were ranked separately according to their percentage share in the IBOV and the IDIV. The arithmetic mean of these two rankings generated an average rank, and the inverse of this value was used to create a raw score, where a higher number indicates more relevance. Finally, the score was normalized to become a relative percentage weight within the group of selected assets.

The complete result of this ranking process, which formed the basis for the selection and weighting of assets for subsequent strategies, is consolidated in Table 1, presented below. It details the individual rankings, the raw score, and the final relative weight for each of the selected perennial assets.

3.2.2. Data Enrichment with Technical Indicators

With the universe of perennial stocks defined by the ranking stage, the next step consisted of enriching the data for each selected company. Technical indicators from multiple categories were calculated, such as trend and volatility (SMA, EMA, BB, ATR, MACD), momentum oscillators (RSI), and past performance (return). The selection of these indicators followed consolidated practices in financial technical analysis (Achelis and Achelis, 2001; Oguz et al., 2019), forming the set of input variables that fed the predictive modeling phase.

3.3. Model Training

3.3.1. Target Variable Engineering

The target variable was designed to directly reflect the defined intelligent Buy and Hold strategy, seeking to optimize monthly contributions. First, for each historical data point, the percentage return over a future time horizon of 30 days was calculated. The choice of a 30-day horizon is directly aligned with the frequency of contributions sought in the strategies. Thus, for the model to focus only on statistically relevant appreciation opportunities, a performance threshold was defined, corresponding to the upper quartile (Q3) of the previously calculated return distribution, representing the top 25% of the asset's historical results. Thus, with the threshold defined, the target variable was discretized into two classes: Class 1 (Buy) and Class 0 (Hold). A consequence of this formulation is the creation of an unbalanced dataset.

Table 1: List of selected perennial assets, ordered by perennial share.

Asset	Share IBOV (%)	Share IDIV (%)	Position IBOV	Position IDIV	*Score*	Perennial Share (%)
PETR4	8,31	5,35	1	2	0,6667	20,84
PETR3	4,95	4,74	2	5	0,2857	8,93
BBAS3	3,43	4,00	3	6	0,2222	6,95
GGBR4	1,22	6,55	8	1	0,2222	6,95
VALE3	11,74	5,06	7	4	0,1818	5,68
CMIG4	1,05	5,08	10	3	0,1538	4,81
JBSS3	2,16	3,14	6	10	0,1250	3,91
BBSE3	1,15	3,94	9	7	0,1250	3,91
ITSA4	2,61	2,65	5	15	0,1000	3,13
KLBN11	0,87	3,47	12	9	0,0952	2,98
VIVT3	0,99	2,80	11	12	0,0870	2,72
STBP3	0,55	3,52	15	8	0,0870	2,72
BBDC4	3,06	2,27	4	21	0,0800	2,50
ISAE4	0,47	2,80	16	11	0,0741	2,32
TIMS3	0,61	2,63	14	16	0,0667	2,08
CXSE3	0,40	2,74	19	13	0,0625	1,95
BBDC3	0,81	2,36	13	20	0,0606	1,89
CSNA3	0,39	2,66	20	14	0,0588	1,84
EGIE3	0,47	2,48	17	19	0,0556	1,74
GOAU4	0,37	2,52	21	17	0,0526	1,65
SANB11	0,43	2,08	18	22	0,0500	1,56
TAEE11	0,36	2,49	22	18	0,0500	1,56
CMIN3	0,30	2,07	23	23	0,0435	1,36
FLRY3	0,30	2,03	24	24	0,0417	1,30
CPFE3	0,30	2,03	25	25	0,0400	1,25
BRAP4	0,22	1,48	26	26	0,0385	1,20
USIM5	0,15	1,02	27	27	0,0370	1,16
AURE3	0,13	0,90	28	28	0,0357	1,12

3.3.2. Case Study with Petrobras

A case study phase was defined to optimize the modeling approach, with the PETR4 asset as the representative due to its prominent position in the previously selected set of perennial assets, as well as its resilience and cyclical characteristics. The main focus of this stage was to refine a predictive model based on the RF algorithm, seeking the best combination of input variables and hyperparameters.

As a first step, and because it is a time series problem, the dataset was divided chronologically in a proportion of 80% for training and 20% for testing. This is a mandatory approach to prevent lookahead bias, in which the model would be trained with future information to predict the past (Botache et al., 2023). Additionally, to guarantee reproducibility and ensure that there is no data leakage between the training and testing process, the entire modeling flow was encapsulated in a Pipeline object from the Scikit-Learn library (Scikit-learn developers, 2024). The use of the Pipeline is a practice that ensures that pre-processing steps are adjusted only to the training dataset and subsequently to the testing and validation sets.

The Pipeline developed for this process is composed of two sequential stages, as illustrated in Figure 2.

Figure 2: Pipeline Diagram

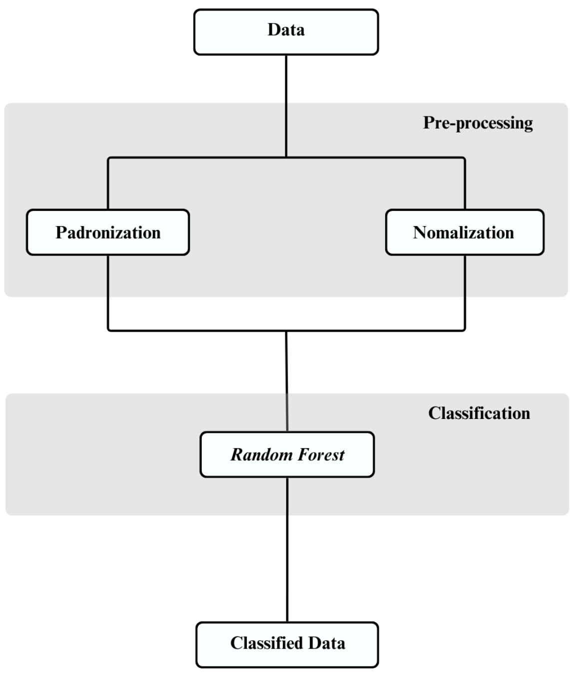

After defining the Pipeline, an iterative process was carried out to select input variables and adjust hyperparameters, using cross-validation for time series to ensure robustness. The optimization was configured to maximize the recall of class 1 (buy), a strategic choice for the Buy and Hold approach, where the cost of a false negative (missing a purchase) is greater than that of a false positive (buying mistakenly). To address the strong class imbalance, three approaches were compared: no balancing, resampling with SMOTE (Synthetic Minority Oversampling Technique), and weight adjustment. According to Table 2, the treatment strategy was selected based on which one best maximized the recall of class 1, aligning with the primary objective of the model.

Table 2: Results of the tests of different balancing techniques

Treatment Strategy	**Overall Precision (%)**	***Recall* (Class 1) (%)**
Scenario 3: Weight Adjustment	79,50	60,00
Scenario 2: Resampling (SMOTE)	72,38	21,00
Scenario 1: No Balancing	68,93	5,00

3.3.3. Generalization and Model Training

The model configuration optimized in the previous phase with PETR4 was used as a base model. This architecture and hyperparameters were then applied in a standardized manner to individually train predictive models for all other previously determined assets.

3.4. Trade System Architecture

3.4.1. Strategy Creation

An experimental design was implemented to conduct a robust comparative analysis between different investment strategies. These strategies were defined by varying along three main axes.

3.4.1.1. Timing Strategies (Buy Trigger)

Defines the logic that triggers a buy order. Both rules based on model usage and benchmark rules were implemented:

A. **Model-based strategies:** Model-based strategies were designed to test two distinct logics for acting on signals. The first, first_buy_signal, adopts an agile approach, executing the buy order for an asset as soon as the first positive signal is generated in the month, limited to one monthly contribution per asset. The second, last_day_top3_abs, uses an evidence accumulation logic, waiting until the last business day of the month to allocate capital to the three assets that received the highest count of buy signals throughout the period.
B. **Benchmark Strategies:** To establish a baseline for comparison, passive strategies that ignore model signals were implemented. They execute purchases based on simple temporal rules: on the first business day of the month (first_day), on the last business day (last_day), or on a randomly selected day (random_day). A crucial mechanism applied to all of them is that, each day, the evaluation sequence of assets for purchase is randomly reordered, mitigating order bias that could systematically favor certain assets in capital allocation.

3.4.1.2. Capital Allocation Methods

Determines how available capital is distributed among assets selected for purchase.

A. **Equal division (equal):** Divides capital equally among candidate assets.
B. **Signal-based division (signal-weighted):** This method was designed to distribute capital non-equally, prioritizing assets with greater signal "strength." The weighting logic, however, adapts to the buy trigger used: for the last_day_top3_abs strategy, strength is measured by the count of buy signals accumulated in the month, with capital being distributed proportionally to this count. For the first_buy_signal strategy, as the purchase is immediate, strength is determined by the asset's raw score, allocating more capital to companies with a higher ranking among the evaluated perennial companies.

3.4.1.3. Operational Constraint Parameters

A. **Minimum Lot (min_shares):** Imposes a minimum lot of shares to avoid operations of negligible value. Minimum lot values were determined empirically for the strategies.
B. **Maximum Allocation (max_allocation):** Limits capital concentration in a single asset to a defined percentage of the portfolio, acting as a risk management tool. The allocation percentage values were defined empirically.

This multifaceted experimental design allows for isolating the impact of each strategy component and assessing, in a quantitative way, whether the intelligence generated by ML models offers a real advantage over passive or simpler investment approaches. The detailed configurations of each tested strategy are summarized in Table 3, below.

Table 3: List of strategies defined for the Backtesting process

Strategy	**Type**	**Buy Trigger**	**Capital Allocation**	**Min. Lot**	**Max. Alloc. (%)**
IA-FS-R1	RF-based	*first_buy_signal*	*signal_weighted*	1	100
IA-FS-R2	RF-based	*first_buy_signal*	*signal_weighted*	10	100
IA-FS-R3	RF-based	*first_buy_signal*	*signal_weighted*	1	25
IA-FS-R4	RF-based	*first_buy_signal*	*signal_weighted*	10	25
IA-L3-R1	RF-based	*last_day_top3_abs*	*signal_weighted*	1	100
IA-L3-R2	RF-based	*last_day_top3_abs*	*signal_weighted*	10	100
IA-L3-R3	RF-based	*last_day_top3_abs*	*signal_weighted*	1	25
IA-L3-R4	RF-based	*last_day_top3_abs*	*signal_weighted*	10	25
BM-FD-R1	*Benchmark*	*first_day*	*equal*	1	100
BM-FD-R2	*Benchmark*	*first_day*	*equal*	10	100
BM-FD-R3	*Benchmark*	*first_day*	*equal*	1	25
BM-FD-R4	*Benchmark*	*first_day*	*equal*	10	25
BM-LD-R1	*Benchmark*	*last_day*	*equal*	1	100
BM-LD-R2	*Benchmark*	*last_day*	*equal*	10	100
BM-LD-R3	*Benchmark*	*last_day*	*equal*	1	25
BM-LD-R4	*Benchmark*	*last_day*	*equal*	10	25
BM-RD-R1	*Benchmark*	*random_day*	*equal*	1	100

BM-RD-R2	*Benchmark*	*random_day*	*equal*	10	100
BM-RD-R3	*Benchmark*	*random_day*	*equal*	1	25
BM-RD-R4	*Benchmark*	*random_day*	*equal*	10	25

3.4.2. Creation of the Trade System

The Trade System architecture is event-oriented, processing market data sequentially. To ensure a robust evaluation, the models were trained and tested with data from January 1, 2010, to December 31, 2019. The simulation, in turn, operated on a distinct and subsequent validation dataset, covering the period from January 1, 2020, to December 31, 2024, ensuring a purely out-of-sample evaluation and assessing the system's true generalization capacity. At the start of the simulation, the RF models, previously trained for each asset, are loaded. The backtesting engine then executes a temporal loop that iterates day by day over a unified date index, ensuring that all decisions are based only on information available up to the moment of the transaction and thus eliminating lookahead bias.

For portfolio state management, the system dynamically maintains two central structures: one that stores the state of each position (quantity, average price, etc.), and a variable that represents available liquidity, adjusted by periodic contributions. A crucial point for the impartiality of the simulation is the randomization of the asset evaluation order daily, mitigating the bias that could systematically favor certain assets in capital allocation. At the end of each trading session, the system performs portfolio reconciliation, calculating its total market value to record the equity evolution.

3.4.3. Strategy Validation and Results Analysis

After defining the strategies and creating the Trade System, the final step of the methodology consisted of empirical validation through backtesting, a standard approach for evaluating the effectiveness of strategies based on historical data (Campbell, 2005). This phase is crucial to ensure the robustness of the results, mitigate risks of overfitting, and, fundamentally, evaluate the practical value of the signals generated by the asset models in an environment that simulates real market conditions.

4. Results

The training of the models was conducted using a historical dataset of the companies mentioned above covering the period from January 1, 2010, to December 31, 2019. For the case study phase with the Petrobras asset (PETR4), this set was subjected to a chronological split of 80% for training and 20% for testing, in order to preserve the temporal integrity of the data. Validation followed an iterative and multiphase process, which started from a baseline model to subsequently conduct a comparative analysis of balancing techniques; hyperparameter optimization focused on the recall of class 1 (Buy) and a selection of input variables guided by the importance of each was also performed. The final configuration was then assessed through cross-validation to ensure a reliable performance estimate. This validation and optimization process culminated in the selection of a final and optimized set of predictive variables.

This optimized configuration, derived from the case study with PETR4, was then replicated for the training of models for all other assets. The set of variables was defined as: maximum price

(HIGH), minimum price (LOW), closing price (CLOSE), Exponential Moving Averages (EMA) of 20 and 50 days, and three lagged return variables (1, 3, and 5 days). In this analysis, "Overall Precision" measures the model's percentage of total correct hits, while "Recall of Class 1" evaluates its ability to identify the buy opportunities that actually existed in the testing period. The consolidated performance metrics of the models are presented in Table 4.

Table 4: Metrics of the trained model for each asset.

Asset	Overall Precision (%)	*Recall* (Class 1) (%)
ITSA4	30,75	100,00
GGBR4	26,99	100,00
KLBN11	24,46	100,00
SANB11	29,29	100,00
GOAU4	29,29	100,00
FLRY3	36,19	100,00
BRAP4	29,08	100,00
USIM5	20,92	100,00
CMIG4	71,13	94,00
PETR3	75,52	75,00
PETR4	75,31	72,00
CSNA3	48,74	71,00
BBAS3	80,75	48,00
JBSS3	38,49	48,00
TIMS3	74,06	45,00
CPFE3	68,83	42,00
BBDC4	76,36	38,00
BBDC3	70,29	37,00
STBP3	45,61	30,00
BBSE3	70,06	28,00
VIVT3	77,62	28,00
VALE3	67,47	26,00
TAEE11	54,60	15,00
ISAE4	69,60	0,00
EGIE3	64,44	0,00
CXSE3	-	-
CMIN3	-	-
AURE3	-	-

Once the training phase was completed, the performance of each strategy was empirically validated through backtesting. The simulation was conducted in a customized trading system, covering the period from 01/01/2020 to 12/31/2024. Each simulated strategy started with a capital of R$ 1,000.00 and was fed with monthly contributions of R$ 300.00. The comparative results, which demonstrate the equity evolution of each approach, are displayed in Table 5.

Table 5: Strategy results ordered by return.

Strategy	Total Return (%)	Annual Return (%)	Final Cost (R$)	Final Balance (R$)	Final Cash (R$)
IA-FS-R2	91,11	18,46	18.685,56	35.709,80	325,76
IA-FS-R4	85,42	17,30	18.675,21	34.627,77	325,02
BM-RD-R2	82,79	16,77	16.834,99	30.773,31	2.173,81
BM-LD-R2	81,21	16,45	16.416,94	29.749,61	2.591,46
BM-FD-R2	80,36	16,28	16.802,87	30.305,84	2.204,43
BM-LD-R4	78,83	15,97	16.415,44	29.356,31	2.553,78
BM-RD-R4	77,91	15,78	16.778,41	29.850,48	2.207,05
BM-FD-R4	76,43	15,48	16.833,14	29.699,26	2.159,36
BM-LD-R3	56,09	11,36	18.249,60	28.466,02	758,37
BM-LD-R1	56,02	11,35	18.241,22	28.459,81	759,27
BM-FD-R1	55,91	11,33	18.559,68	28.935,72	451,25
BM-FD-R3	55,82	11,31	18.558,25	28.918,20	456,32
BM-RD-R3	55,70	11,28	18.551,35	28.885,14	452,27
BM-RD-R1	53,90	10,92	18.544,41	28.539,35	453,77
IA-FS-R1	46,83	9,49	18.993,22	27.887,07	4,69
IA-FS-R3	46,18	9,36	19.005,42	27.785,97	2,50
IA-L3-R2	33,27	6,74	18.460,76	24.602,61	533,32
IA-L3-R4	29,10	5,89	18.385,02	23.734,82	608,30
IA-L3-R3	25,64	5,19	18.558,44	23.316,24	443,66
IA-L3-R1	19,56	3,96	18.576,90	22.211,47	435,51

Table 5 summarizes the performance of the 20 tested strategies, highlighting IA-FS-R2 as the one with the highest total return (91.11%). To dissect the factors that led to this superior performance, the following analysis focuses on a direct comparison between this strategy and the one with the best result from the benchmark group, BM-RD-R2 (82.79%). The investigation will delve into how each approach constructed its portfolio. The following tables detail the final distribution of assets in the portfolio for each of these two strategies.

Table 6: Distribution and performance of the IA-FS-R2 strategy portfolio.

Asset	Quantity	Average Price (R$)	Final Price (R$)	Final Return (%)	Portfolio Allocation (%)
PETR3	199	7,15	38,32	435,94	21,35
PETR4	186	7,28	35,13	382,55	18,30
BBAS3	277	12,72	23,57	85,30	18,28
ITSA4	733	7,68	8,18	6,51	16,79
GGBR4	142	12,60	17,89	41,98	7,11

GOAU4	134	8,90	10,10	13,48	3,79
KLBN11	52	15,69	22,88	45,83	3,33
CMIG4	118	2,90	10,06	246,90	3,32
USIM5	129	6,65	5,32	-20,00	1,92
BBDC3	68	10,79	9,89	-8,34	1,88
BBDC4	54	11,35	10,76	-5,20	1,63
BBSE3	10	15,96	36,18	126,69	1,01
BRAP4	15	8,40	15,71	87,02	0,66
STBP3	10	3,09	12,86	316,18	0,36
CSNA3	10	9,49	8,86	-6,64	0,25

Table 7: Distribution and performance of the BM-RD-R2 strategy portfolio.

Asset	Quantity	Average Price (R$)	Final Price (R$)	Final Return (%)	Portfolio Allocation (%)
CMIG4	950	5,20	10,06	93,46	31,06
STBP3	606	5,23	12,86	145,89	25,32
ITSA4	599	7,34	8,18	11,44	15,92
PETR3	72	5,90	38,32	549,49	8,97
GOAU4	228	5,54	10,10	82,31	7,48
USIM5	345	6,37	5,32	-16,48	5,96
PETR4	26	4,90	35,13	616,94	2,97
GGBR4	23	5,12	17,89	249,41	1,34
CSNA3	34	5,83	8,86	51,97	0,98

The comparison between the final portfolios of the top-performing strategies in each category, IA-FS-R2 and BM-RD-R2, reveals a fundamental difference in their structure. The first and most notable is the level of diversification: the RF-guided strategy built a broader portfolio with 15 assets, while the benchmark concentrated its capital in only 9 assets. Theoretically, a more diversified portfolio tends to be more resilient during periods of high volatility, yet it is subject to having its gains diluted, preventing the exceptional performance of a single asset from dominating the total result.

5. Discussion

The analysis of the final model metrics allows for the extraction of important conclusions regarding the optimization process of an ML model. The main one is that the strategy of generalizing the characteristics and hyperparameters from the PETR4 case study to the other assets proved ineffective. This highlights that, even within a portfolio of perennial assets, each company possesses distinct price dynamics and behaviors, meaning that a configuration optimized for one asset is not transferable to others, especially across different sectors.

Two contrasting failure patterns are observed. On one hand, models such as those for ITSA4 and GGBR4 showed a recall of 100% for the buy class, combined with extremely low overall accuracy (between 20% and 30%). Such behavior indicates that the model developed an extreme prediction bias, classifying almost all instances as "Buy" and generating a massive amount of false positives. In contrast, for ISAE4 and EGIE3, the model became excessively conservative, predicting only the "Hold" class and resulting in a recall of 0%. Additionally, assets with a shorter data history or issues returning data via Yahoo Finance (CXSE3, CMIN3, and AURE3) did not generate valid models.

The models for PETR4 and PETR3, which served as the basis for the study, presented a balanced performance between recall and overall accuracy, validating the success of the focused optimization process. Other assets with reasonable results, such as CMIG4 and BBAS3, share characteristics with Petrobras, such as being large and consolidated companies with state participation, which may justify superior performance compared to the others.

The differences in individual model performance, discussed previously, directly impact the performance of the investment strategies. The analysis of the backtesting results allows for the quantification of how model behaviors translated into financial return. Thus, it is observed that the IA-FS-R2 strategy emerged as the top performer, with a total return of 91.11% (18.46% annual). Its configuration combines buy execution at the first signal from the model, with no allocation limit per asset and a minimum lot of 10 shares at the time of purchase. This result suggests that agility in capital allocation is a critical factor and appears to have been fundamental in capturing the beginning of upward movements.

To contextualize the magnitude of this result, it is essential to compare it with the main index of the Brazilian stock market, the IBOVESPA. In the same simulation period (from January 2020 to December 2024), the IBOVESPA had a total accumulated return of approximately 7.6% (B3, 2025b), which represents an annualized return of about 1.52%. This disparity shows that the IA-FS-R2 strategy generated a return superior to the market by more than 16 percentage points per year. This represents a significant opportunity cost—that is, the potential return lost by keeping capital in one option over another—for the investor who chose to follow the index, validating that the proposed active strategy succeeded in adding more value than the purely passive approach.

It is notable that the benchmark strategies, which do not make use of the RF models, presented competitive performance. Strategies such as BM-RD-R2 (82.79%), BM-LDR2 (81.21%), and BM-FD-R2 (80.36%) outperformed most of the variants using the models. However, the high final cash balance in these strategies, above R$ 2,000.00, which represents between 12% and 16% of the invested value depending on the strategy, indicates that their configurations frequently failed to invest all available capital. This unallocated capital failed to benefit from asset appreciation, which also represents a significant opportunity cost, indicating that these strategies were inefficient compared to the strategies that used the models in this regard.

In contrast to the success of the other strategies, the IA-L3 group of strategies presented notably inferior performance, with total returns between 19% and 33%. Although they used the signals generated by the models, the rule of waiting until the end of the month to allocate capital proved counterproductive. This weak performance reinforces a central hypothesis already suggested by the success of the best strategy: latency in execution is detrimental. The wait caused the system to miss the ideal timing of upward movements, which apparently occur well before the close of the monthly period.

Although the total return results are revealing, it is essential to dissect how each strategy built its portfolio to understand the origin of its performance. The analysis of the portfolio composition for the best AI strategy (IA-FS-R2) and the best benchmark (BM-RD-R2) exposes fundamental differences. The first of these is diversification: while the AI strategy built a portfolio with 15 assets, the benchmark limited itself to only 9.

This difference deepens in the concentration analysis. The BM-RD-R2 strategy proved extremely concentrated, with only two assets (CMIG4 and STBP3) accounting for more than 56% of the total portfolio. This means that its excellent result was largely dependent on the performance of a very restricted number of companies. In contrast, the IA-FS-R2 strategy, while also having significant concentration at the top (57.9% in the top three assets), distributed the remaining capital much more evenly among 12 other assets.

However, the decisive factor for the superiority of the model-based strategy was not diversification, but rather the quality of selection. IA-FS-R2 achieved the best return because it was able to identify and allocate capital to assets with the most explosive growth in the period. The return of PETR3 (435.94%) and PETR4 (382.55%) was of a magnitude much higher than that of the benchmark's best assets (STBP3 with 145.89%). This indicates the model's ability, when well-optimized for a specific asset, to effectively signal buy moments preceding exponential upward movements, validating the potential of the "Intelligent Buy and Hold" approach.

6. Conclusion

This work proposed and evaluated a framework to improve the Buy and Hold strategy through an algorithmic trading system, which uses RF as a predictive model to optimize buy moments in a portfolio of perennial assets in the Brazilian market. The studies conducted through the backtesting process generated fundamental conclusions about predictive modeling and the implementation of the proposed investment strategies.

Although the scope of this work did not include a direct comparative study with other architectures, such as deep learning models, the choice of RF was deliberately based on evidence from the literature. Works such as that of Ma et al. (2021) have already demonstrated its superiority in return forecasting tasks compared to more complex models, validating it as a robust and effective starting point for the optimization of financial strategies.

The results demonstrated, first, that generalizing input variables and hyperparameters from a model optimized for one asset (PETR4) is an ineffective approach for a diversified portfolio. The unique dynamics of each asset require individualized optimization to create reliable predictive models. Second, the analysis of backtesting strategies revealed that timing in the execution of contributions is a critical factor for higher returns.

It is fundamental, however, to recognize the limitations of the present study. Although it served to validate the main thesis, the generalization of model characteristics is a methodological simplification. Furthermore, the simulation did not consider some elements of the real market, such as transaction costs, bid-ask spreads, dividends, and bonuses.

These limitations, however, open clear paths for the evolution and follow-up of this work. Perspectives for continuity include the individualized optimization of the predictive model for each asset and the exploration of more complex model architectures, among others.

In summary, the modular and scalable framework developed serves as a basis for future quantitative research, allowing for the testing and implementation of new ideas in the dynamic landscape of the Brazilian financial market.

Acknowledgements

No funding sources to disclose.

Conflict of Interest Declaration

The authors have no conflicts of interest to declare.

References

1. Achelis, S. B. & Achelis, S. (2001). Technical Analysis from A to Z. McGraw Hill New York.
2. Aziz, S., Dowling, M., Ghandour, A., & Trigonometricio, P. (2022). Machine learning in finance: A topic modeling approach. European Financial Management, 28(3), 744–770.
3. B3. (2024). Pessoas Físicas na B3. Acessado em 13 de agosto de 2024, de https://www.b3.com.br/pt_br/noticias/pessoas-fisicas-na-b3.htm
4. B3. (2025a). B3 - Brasil, Bolsa, Balcão. Disponível em: https://www.b3.com.br/. Acesso em: 13 de agosto de 2024.
5. B3. (2025b). Índice Ibovespa (IBOVESPA) - Estatísticas Históricas. Disponível em: https://www.b3.com.br/pt_br/market-data-e-indices/indices/indices-amplos/indice-ibovespa-ibovespa-estatisticas-historicas.htm. Acesso em: 14 de outubro de 2025.
6. Ballings, M., Van den Poel, D., Hespeels, N., & Gryp, R. (2015). Evaluating multiple classifiers for stock price direction prediction. Expert systems with Applications, 42(20), 7046–7056.
7. Botache, D., Pérez, D., & Arbeláez, P. (2023). Unraveling the complexity of splitting sequential data: tackling challenges in video and time series analysis. arXiv preprint arXiv:2307.14294.
8. Brader, T. (2020). Campaigning for hearts and minds: How emotional appeals in political ads work. University of Chicago Press.
9. Campbell, S. D. (2005). A review of backtesting and backtesting procedures. Finance and Economics Discussion Series 2005-21, Board of Governors of the Federal Reserve System (U.S.).
10. Cooper, B. S. (2018). Interactive Planning and Sensing for Aircraft in Uncertain Environments with Spatiotemporally Evolving Threats [Tese de Doutorado, The Pennsylvania State University].
11. da Silva SANTOS, G. F., da Silva SANTOS, C. A., & MARQUES, N. L. (2021). A estratégia de buy and hold na composição da carteira de jovens investidores da baixada fluminense. *Caderno de Administração*, 29(2), 6–31.
12. Damodaran, A. (1997). *Avaliação de investimentos: ferramentas e técnicas para determinação do valor*. Qualitymark Editora Ltda.
13. Deboeck, G. J. (1994). Trading on the edge: neural, genetic, and fuzzy systems for chaotic financial markets. John Wiley & Sons.
14. El Naqa, I. & Murphy, M. J. (2015). What is machine learning?. Em Machine learning in radiation oncology: theory and applications (pp. 3–11). Springer International Publishing, Cham.

15. Haryanto, J. T. (2024). Short-Term Versus Long-Term Portfolio Management Strategies and the Selection of Securities. Advances in Management & Financial Reporting, 2(1), 11–23.
16. Investidor10. (2025). Investidor10 - Ajudando você a investir melhor. Disponível em: https://investidor10.com.br/. Acesso em: 13 de agosto de 2024.
17. Ma, Y., Han, R., & Wang, W. (2021). Portfolio optimization with return prediction using deep learning and machine learning. Expert Systems with Applications, 165, 113973.
18. Malek, N. H. A., Wan Yaacob, W. H. M., Mohammad, N. A., & Abdul-Ghaffar, A. (2022). Prediction of water quality classification of the Kelantan River Basin, Malaysia, using machine learning techniques. Water, 14(7), 1067.
19. Navlani, A., Fandango, A., & Idris, I. (2021). Python Data Analysis: Perform data collection, data processing, wrangling, visualization, and model building using Python. Packt Publishing Ltd.
20. Newman, J. L., Fuqua, D. R., Gray, E. A., & Simpson, D. B. (2006). Gender differences in the relationship of anger and depression in a clinical sample. Journal of Counseling & Development, 84(2), 157-162.
21. Oğuz, R. F., Eyercioğlu, Ö. K., & Avci, D. (2019). On the use of technical analysis indicators for stock market price movement direction prediction. Em 2019 27th Signal Processing and Communications Applications Conference (SIU) (pp. 1–4). IEEE.
22. Pixley, J. (2004). Emotions in finance: Distrust and uncertainty in global markets. Cambridge University Press.
23. Scikit-learn developers. (2024). sklearn.pipeline.Pipeline. Acessado em 29 de julho de 2025, de https://scikit-learn.org/stable/modules/generated/sklearn.pipeline.Pipeline.html
24. Status Invest. (2025). Status Invest - Análise de Ações, Fundos Imobiliários e muito mais. Disponível em: https://statusinvest.com.br/. Acesso em: 13 de agosto de 2024.
25. Shaik, A. B. & Srinivasan, S. (2018). A brief survey on random forest ensembles in classification model. Em International Conference on Innovative Computing and Communications: Proceedings of ICICC 2018, Volume 2 (pp. 253–260). Springer Singapore.
26. Silva, N. F., de Andrade, L. P., da Silva, W. S., de Melo, M. K., & Tonelli, A. O. (2024). Portfolio optimization based on the pre-selection of stocks by the Support Vector Machine model. Finance Research Letters, 61, 105014.
27. Zheng, J., Xin, D., Cheng, Q., Tian, M., & Yang, L. (2024). The random forest model for analyzing and forecasting the us stock market in the context of smart finance. arXiv preprint arXiv:2402.17194.