Behavioral Finance Model: Google Trends and LSTM in Forecasting Retail Assets on the Brazilian Stock Exchange

Lucas D. M. Mendonça[1]
Diego L. Cardoso[1]

[1]Laboratório de Pesquisa Operacional – Universidade Federal do Pará (UFPA)
Belém – PA – Brasil

lucas.mendonca@castanhal.ufpa.br

diego@ufpa.br

Abstract. *Social data mining, such as Google searches, has been increasingly explored in financial time series forecasting. However, the literature frequently ignores the semantic noise inherent to retail search queries. This paper investigates the impact of different Google Trends syntactic categories on neural network-based predictive models for Brazilian Stock Exchange (B3) retail assets, aiming to mitigate this noise. Analyzing error and volatility metrics across nine assets, our results suggest that the predictive value of online attention is not universal but instead depends on the type of query used. In 77.8% of the analyzed assets, the incorporation of Google Trends data outperformed the pure OHLC model, most notably for AMER3, where searching for the commercial name reduced the MAPE from 643.18% to 99.57%. Based on these observations, we propose an exploratory taxonomy of four categories and argue that financial predictive architectures benefit from semantic parameterization adapted to each asset's characteristics.*

1. Introduction

The stock market plays a fundamental role in the modern economy, facilitating the raising of funds for companies and offering profitability opportunities for investors. Traditionally, forecasting the prices of financial investments is considered an inherently complex task due to high volatility, non-linearity, and the influence of multiple exogenous factors, this premises are supported by the Efficient Market Hypothesis (EMH) [Fama 1970]. In contrast, the Behavioral Finance literature challenges the EMH proposal, arguing that investors do not act in a purely rational way, but are subject to cognitive biases and emotional reactions that distort prices. This perspective suggests the existence of informational inefficiencies and timing anomalies in markets.

It is in this efficiency gap that Deep Learning-based models find applicability, given their ability to identify subtle non-linear patterns and long-term dependencies that challenge the assumption of fully random prices [Shiller 2003]. With the current abundance of data, multivariate architectures, which integrate heterogeneous sources, have achieved superior performance. Such models are essential to map not only temporal patterns, but also the cause-and-effect relationship (or strong correlation) between different explanatory variables [Gopali et al. 2024].

In the Brazilian context, the retail sector exhibits particular volatility and been sensitive to economic cycles and consumer sentiment. Although traditional statistical

methods, such as Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroscedasticity (GARCH), have been widely used in financial modeling, they often fail to capture long-term dependencies and complex nonlinear dynamics, especially during times of acute uncertainty. In this scenario, Artificial Neural Networks (ANNs), specifically LSTM architectures, have been reported to perform well [Churi et al. 2023].

Aligned with this framework, this work proposes a predictive model focused on the retail sector of the B3 (Brazilian Stock Exchange). The main contribution and differentiating factor of this research lies in the integration of traditional financial time series OHLC (Open, High, Low, Close) with alternative investor attention data, extracted from Google Trends, composing a multivariate architecture. The central objective is to investigate whether the incorporation of online search volume increases the predictive accuracy of Deep Learning models and, fundamentally, to analyze how the query syntax (search by brand name, financial ticker, or aggregated terms) interacts with the asset's volatility regimes to separate the purely directional signal from statistical noise.

2. Related Works

Recent literature in Information Systems reinforces the effectiveness of hybrid models that incorporate social media analytics and digital footprints. [Bantis et al. 2023] Studies have shown that integrating Google Trends data into dynamic factor models significantly improves forecasting accuracy, demonstrating notable performance gains in emerging economies such as Brazil. Simultaneously [Leippold et al. 2022], using neural networks, they identified that the strong presence of retail investors in the market increases the predictability of short-term returns in the Chinese market, while [Wang 2022] achieved success by using deep learning focused exclusively on predicting the performance of the retail sector. These studies clearly demonstrate the efficiency of machine learning models in the predictability of various aspects from the financial sector.

At the intersection of behavioral finance and machine learning [Andrei and Hasler 2014] demonstrate, through equilibrium modeling, that the variance of stock returns grows along with uncertainty and attention, attracting capital flows, and it is also established that investor attention and market uncertainty are covariant. This makes it clear that to perform the analysis and prediction of financial assets, it is necessary to consider these factors and improve methods for capturing investors' attention.

[Hamid and Heiden 2015] empirically validate this correlation, proving that investor attention, along with the accuracy of predictive models based on Google Trends, increases notably during "panic states"and periods of high volatility. Inversely, the authors warn that periods of low uncertainty turn search data into white noise. Furthermore, [Petropoulos et al. 2022] demonstrate that deep learning architectures powered by online queries are especially effective at detecting impending turbulence.

The superiority of hybrid and multivariate predictive architectures over purely univariate approaches transcends the business domain. Demonstrating this capability [Mesquita et al. 2023] demonstrate the integrated use of multiple algorithms to significantly optimize the forecasting in physical non-linear systems, like the flood hydrology.

Currently, hybrid models using market data and Google Trends have high predic-

tive potential, proving capable of capturing investors' attention when analyzing import data, as demonstrated in [Andrade et al. 2025]. Proven in [Barber and Odean 2007] the effect of attention is even more noticeable in the share market, where individual investors become buyers of stocks that attract attention. Presented in [Taleb 2010] the concept of events with low probability and high impact when they materialize, generating systemic attentional shocks, referred to as *Black Swan*, attract the attention of investors and have great media appeal, generating potential interest and research. However, as already explained by [Cebrián and Domenech 2023] when using Google Trends metrics, all results obtained from any approach based on this data are subject to the context in which they are used. This necessitates a methodical curation of the values entered for the searches. For this article, we used a clear and objective taxonomic analysis to evaluate the ideas when exposed to the Brazilian retail scenario.

3. Methodology

The methodology of this work follows a quantitative experimental approach, structured in the stages of Data Collection, Pre-processing, Modeling, and Evaluation.

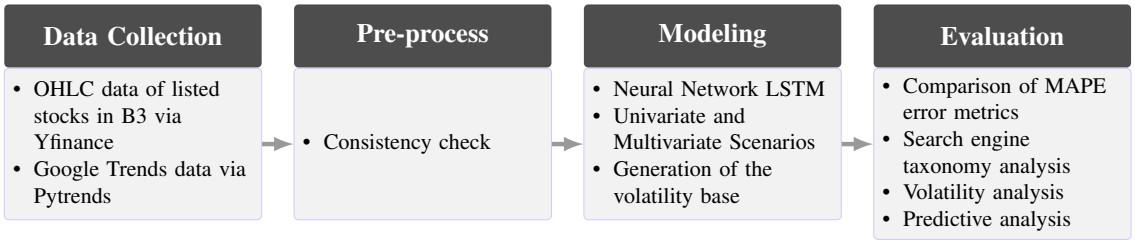


Figura 1. Methodological flow of the research: processing and analysis stages.

3.1. Collection and Data Source

The empirical sample for this study consists of nine assets listed on the B3: AMER3, ASAI3, BHIA3, CEAB3, GMAT3, LREN3, MGLU3, PCAR3, and PETZ3. The dataset covered a time window of 200 weeks for training and 10 weeks for technical validation totaling 210 weeks (January 2021 to January 2025).

In contrast to massive quantitative sampling, a methodological approach was chosen: **Selective heterogeneity sampling**, following the concept of [Bolfarine and de Oliveira Bussab 2005]. The literature on financial data mining demonstrates that neural network-based models respond differently across market structures. Therefore, the selection of this specific group was designed to act as a microcosm of the Brazilian retail sector, ensuring the representation of four distinct commercial subsegments:

- E-commerce and Home Appliances (MGLU3, BHIA3, AMER3);
- Clothing and Fashion (LREN3, CEAB3);
- Food and Wholesale (PCAR3, GMAT3, ASAI3);
- Specialized Retail (PETZ3).

In addition to sectoral diversification, this asset basket ensures a fundamental variance for the test: the algorithm's exposure to multiple stress regimes. The sample ranges from historically stable stocks (such as GMAT3) to assets subject to extreme speculation

and systemic shocks (such as AMER3 and BHIA3). This variability is broader to assess the predictive effectiveness of the Google Trends and Concept Drift.

1. Market Data (B3): Collected via python open-source package Yfinance[1], which provides daily financial market data retrieved from the Yahoo Finance Public API (OHLC).
2. Trend Data (Google Trends): Collected via the open-source Python package Pytrends.[2], which provides a framework for automating the download of Google Trends reports.

3.2. Semantic Taxonomy in Google Trends

In collecting data from Google Trends, it was necessary to use the Python tool PyTrends, which provides reliability and assists in automated data collection, as demonstrated in [Fahrudin et al. 2022], it proceed to collect the listed companys through two different forms of identification: (i) o *ticker*, trading code that identifies the company on the stock exchange; and (ii) the trade name, used in non-financial queries. This distinction allows isolating the attentional variable and verifying possible divergences between the two representations, as well as aggregating the variables to verify possible correlational equality and noise reduction.

3.3. Preprocessing and Feature Engineering

For the analysis and insertion of the collected data into the model, a sample field was selected that corresponds to the start date (01-03-2021) and end date (01-01-2025) Since the data have different periodicities (stocks are daily, Google Trends is weekly), a temporal alignment process was performed. The stock data were grouped to understand the weekly movement on the stock exchange, allowing the LSTM model to process the inputs simultaneously. To mitigate the discrepancy in magnitude and variance between the different variables of the time series, specifically the OHLC financial quotation data and the search volume indices of *Google Trends*, the standard scaler (Z-score) standardization technique was applied due to his proven efficiency in [Suthiponpisal and Tancharoen 2024] to reduce noise in the stock market. This process transforms each feature from the dataset so that it starts operating with a zero average. ($\mu = 0$) and unit standard deviation ($\sigma = 1$).

The transformation is mathematically defined by the following equation:

$$z = \frac{x - \mu}{\sigma} \tag{1}$$

Where:

- z represents the standardized value of the observation to be entered into the model;
- x corresponds to the original value of the variable in the time series (asset price or search index);
- μ is the sample mean of the respective characteristic calculated over the training set;
- σ is the sample standard deviation of the same characteristic.

[1] `https://pypi.org/project/yfinance/`
[2] `https://pypi.org/project/pytrends/`

The application of this technique, implemented computationally through the library *Scikit-Learn* in the Python environment, constitutes a fundamental preprocessing step. By geometrically aligning the scales of financial and social data, it prevents variables of greater numerical magnitude from dominating the model's cost function. In this way, standardization ensures faster, more stable, and efficient convergence of synaptic weights in gradient descent-optimized neural networks.

3.4. Predictive Architecture Modeling

To isolate the attentional variable, four neural network configurations were designed:

- Base Model (OHLC): purely technical Inputs (Open, High, Low e Close).
- Trademark Model (OHLC + Company Name): Addition of the Google Trends metric of searches at the Company Name (eg., "Lojas Renner").
- Financial Model (OHLC + Ticker): Insertion of the Google Trends metric of searches estrictely at the Ticker code (eg., "LREN3").
- Sectoral/Aggregated Model (OHLC Agregate): Insertion of both metrics used in the Google Trends search engine for retail companies.

The artifact developed in Python uses the Keras/TensorFlow library. The implemented architecture consists of a sequential network composed of:

- **Input Gate:** Sliding window configured in the format $(N = 4, F = 4)$ for the base model and $(N = 4, F = 5)$ for the Google Trends models, Processed by a 64-unit LSTM layer with sequence return to capture complex temporal patterns.
- **Hidden Layer:** A second 64-unit LSTM layer (without sequence return) to condense the temporal information into a feature vector, followed by a dense layer of 64 neurons with ReLU activation function for nonlinear interpretation of features.
- **Dropout:** One 20 Dropout layer positioned before the output gate to mitigate overfitting, randomly deactivating neurons during training.
- **Output Gate:** a dense layer with a single neuron and linear activation, responsible for projecting the predicted final value for the action.

The training was conducted over 200 epochs with a batch size of 32, using the *Adaptive Moment Estimation* optimizer with learning rate of 0,001 and MAE (*Mean Absolute Error*) loss function. A (20%) dropout layer It was applied between the intermediate dense layer and the output gate for regularization, parameters used based on the comparative study of [Dhanawansa et al. 2022].

3.5. Volatility Calculation and Regime Change Analysis

To measure the degree of uncertainty and stochastic fluctuations of each asset, historical volatility was calculated based on weekly windows. Initially, the daily price data collected between January 2021 and January 2025 were divided into 3 groups: Complete (training + validation) from 01-03-2021 to 01-05-2025(March 1, 2021, to May 1, 2025), Train from 01-03-2021 to 11-03-2024(March 1, 2021, to March 11, 2024) and Validation from 11-04-2024 to 01-13-2025(November 4, 2024, to January 13, 2025) and then aggregated on a weekly basis, extracting the respective opening, high, low, and closing values for each period. From the closing price series (P_t), continuous logarithmic returns were calculated, as given by the equation:

$$R_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \tag{2}$$

Weekly volatility (σ_{weekly}) of each asset was defined as the sample standard deviation of these returns (R_t). For standardization purposes and alignment with financial literature, the metric was annualized assuming 52 operational weeks, through the following conversion:

$$\sigma_{annual} = \sigma_{weekly} \times \sqrt{52} \times 100 \tag{3}$$

A crucial methodological aspect of this research was the structural segregation of the volatility analysis between the training period (representing the first 95% of the time series) and the validation period (the last 5%, equivalent to approximately 10 weeks of prediction). This chronological partitioning allowed the geometric quantification of Concept Drift.

Through the calculation of the volatility ratio between the windows. ($\sigma_{\text{validation}}/\sigma_{\text{train}}$) and the measuring of the maximum percentage price range during the test period, it was possible to dynamically classify the assets. This analytical approach proved fundamental for mapping the algorithm's resilience when transitioning from chronic stress regimes (in training) to directional calm phases (in validation), as well as for isolating anomalous systemic shock events (Black Swans, eg,. AMER3) described by [Taleb 2010], theoretically justifying the gains and losses in accuracy of the neural network conditional on the market's state of attention.

4. Results and Discussion

Using the database created from the available tools, it was possible to generate a heat map of the weekly correlation between market data, specifically the closing value as it is the output variable of the model, and Google trend data (TradeMark, Ticker, and Aggregate), resulting in the Figure shown. 2

The figure 2 assists in the decision-making process and shows potential as a tool for a multi-criteria decision-making approach. It is possible with the Figure 2 understand the predictive power of trend variables in predicting stock closing prices based on the correlation between these assets and their respective trend metrics in the 3 aspects addressed, showing where the data have the greatest similarity in their temporal variations.

The constructed LSTM model comprises a predictive architecture capable of ingesting historical data from retail companies and cross-referencing it with online attention tracking. To isolate the effect of semantic noise inherent in users' internet search behavior, and to understand its correlation with market stress, the model's performance was evaluated in four distinct scenarios:

1. **Baseline (OHLC Only):** Univariate model using only price history as baseline.
2. **Multivariate (Trademark):** Adding searches by the company's trade name (e.g., "Lojas AmericanasC&A").
3. **Multivariate (Ticker):** Addition of searches strictly by trading code on B3 (e.g., "AMER3", "CEAB3").

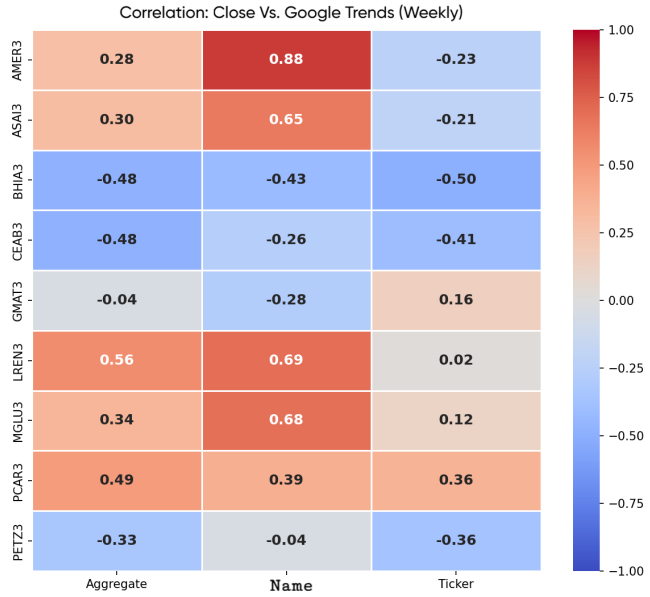


Figura 2. Heat map of the weekly correlation between closing prices and trend metrics from Google Trends by asset.

4. **Multivariate (Agregated):** Addition of searches in Google Trends by ticker symbol on the B3 stock exchange and company trade name.

The figure 3 presents a comparison between the traditional ”OHLC Only”model and models with added Google Trends data. The generated image visually demonstrates the MAP of the traditional models compared to the best of the 3 proposed Trends models (Aggregate, Brand, and Ticker) arranged against the asset's annualized volatility axis.

The Table 1 presents the evolution of the error metric (MAPE) for the validation set under the four training scenarios, ordered in descending order by the annualized volatility of the asset during the validation period.

The Table 1 analisys demonstrates the hypothesis of multiple attentional variables beneficial to the predictive model. The performance gain was predominantly found in models using Google Trends data, especially ”Ticker”and ”Brand.”The low performance of the ”Pure OHLC”and ”Aggregated”models when compared to the others is evident, highlighting the importance of using models with social sector data and keyword taxonomy.

Tabela 1. Error Comparison (MAPE %) vs. Volatility: Base Model and Google Trends Semantics Variation

Asset	Vol. (Validation)	OHLC Only	Trends (Aggregate)	Trends (Trademark)	Trends (Ticker)	Best Approach
AMER3	280.23%	643.18%	275.07%	**99.57%**	564.18%	Trademark
CEAB3	51.33%	**6.18%**	9.12%	7.55%	7.30%	OHLC
MGLU3	45.11%	16.57%	**13.66%**	17.25%	21.79%	Agregado
PCAR3	36.15%	27.86%	22.29%	**20.90%**	44.30%	Marca
BHIA3	31.25%	112.38%	69.93%	117.56%	**67.92%**	Ticker
PETZ3	28.23%	5.45%	5.82%	7.57%	**4.30%**	Ticker
ASAI3	27.83%	8.84%	7.41%	**7.27%**	9.63%	Marca
LREN3	24.99%	2.82%	4.17%	4.75%	**2.20%**	Ticker
GMAT3	22.71%	**2.85%**	3.52%	4.71%	4.10%	OHLC

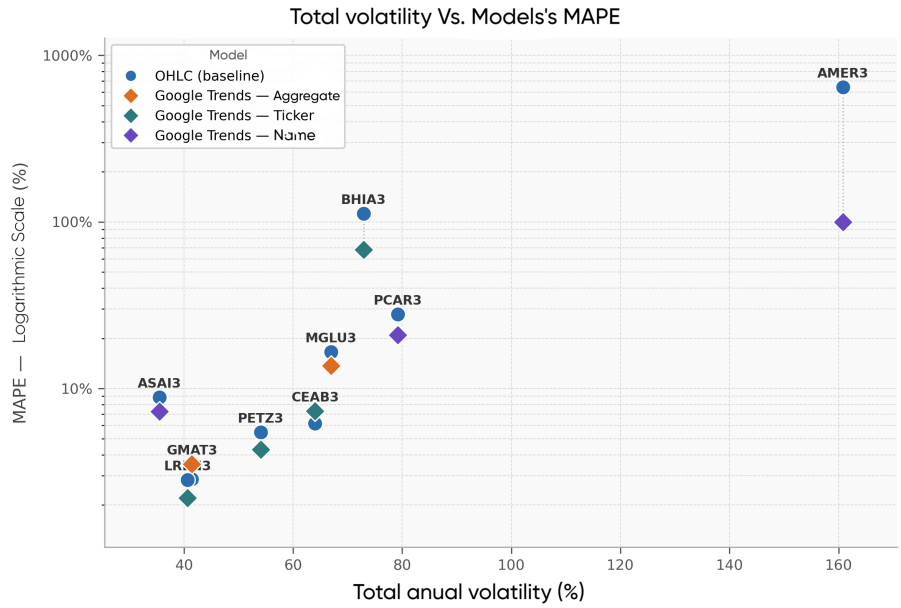


Figura 3. Relationship between total annualized volatility and MAPE of LSTM models by asset. Circles represent the "OHLC Only"baseline model and diamonds represent the best model among those from Google Trends (Aggregate, Ticker or Trademark). Both axes are on a logarithmic scale. AMER3 and BHIA3 stand out as assets with high volatility and the highest predictive error in the baseline model.

5. Discussions

Given the artifacts generated based on the models used in the design of systems based on Deep Learning and social mining for a better understanding of asset behavior and its susceptibility to the models used in this article, we propose the classification of assets through analysis into four taxonomic categories based on established theories of behavioral finance and results from the neural network models used:

Taxonomic Category I — Technical Predominance: Only in 22,2% of the cases analyzed (GMAT3 and CEAB3), the purely financial "Pure OHLC"model outperformed approaches using social data. This result is consistent with the Efficient Market Hypothesis [Fama 1970] and demonstrating through the Table 1 where online search data for these assets does not improve the models and becomes white noise, lacking incremental predictive power over price history.

Taxonomic Category II — Organic Investor Attention: In 33.3% of cases (PETZ3, LREN3, and BHIA3), the ticker-search model outperformed the other approaches. This phenomenon is explained by the investor attention theory of [Andrei and Hasler 2014] and finds an example of it in the Table 1 where the assets in question show improvement in the Trends(Ticker) model and worsening in the Trends(Trademark) model, it is possible to infer that queries by trading code capture organic financial intent, filtering out the noise of retail consumers who seek the commercial brand without interest in the assets. The exclusive use of the ticker, therefore, acts as a semantic isolator from diffuse institutional attention.

Taxonomic Category III — Institutional Attention in Extreme Events: In ca-

ses AMER3, PCAR3, and ASAI3, searching by trademark yielded the best predictive results. This regime corresponds to the phenomenon of attention-grabbing stocks described by [Barber and Odean 2007] and is equally consistent with the concept of the Black Swan [Taleb 2010], where low-probability, high-impact events such as corporate collapses and systemic restructurings, upon materializing, generate attentional shocks; attention extends beyond the specialized financial audience and reaches society in general, which consumes news about the trademark. This behavior was captured by the neural network via search volume for the trademark and can be seen in the Table 1 where, in stark contrast to the previous category (II), the assets show improvement in the Trends (Brand) model and decline in the Trends (Ticker) model. Assets that best support this thesis and corroborate the Taxonomic classification employed are:

- **AMER3 (The Peak of Signal):** Amidst the corporate collapse[3], the search for the trademark "Lojas Americanas"captured the panic of society and reduced the neural network error (MAPE) from a dramatic 643.18% to 99.57%.
- **PCAR3 (Institutional Attention):** Undergoing significant restructuring[4], using the group trademark significantly reduced the predictive error of 27.86% from the strictly financial model to 20.90%.

Taxonomic Category IV — Macro Sentiment: The outlier MGLU3 is the only asset that exhibits multidimensional characteristics; none of the attentional dimensions, financial or commercial, individually capture the asset's movement with maximum efficiency. The improvement in its predictability therefore occurs diffusely across the two semantic domains (Trademark and Ticker) and can be verified by the superiority of the Aggregate model as shown in the image 2, consistent with the argument of [Andrei and Hasler 2014] where the covariance between attention and market uncertainty can be multidimensional, making semantic aggregation the optimal strategy for capturing attentional signals.

6. Conclusion

The experiments conducted demonstrate the superiority in different taxonomic categories of the model with social data extracted from Google Trends compared to the "Pure OHLC"LSTM model, and identify its susceptible variation in the taxonomy employed in collecting Google trend data in accordance with market momentum characteristics. More relevant than the simple addition of these trend variables is how the query syntax interacts with the volatility regime of each asset. With this correlation and analysis of the assets, a behavioral taxonomy in four categories was proposed.

Assets with low uncertainty and efficient pricing dynamics, such as GMAT3 and CEAB3, did not show predictive gain with search data, behaving consistently with the Efficient Market Hypothesis. In contrast, assets under organic financial attention, such as PETZ3, LREN3, and BHIA3, responded positively to the ticker search, which filters

[3]January 11, 2023, Americanas S.A. disclosed a relevant fact regarding accounting inconsistencies of approximately R$20 billion, later revised to over R$40 billion, culminating in the filing for judicial reorganization on January 19, 2023. AMERICANAS S.A. *Comunicados e Fatos Relevantes*. Available at: ri.americanas.io. Accessed in: March 2026.

[4]Grupo Pão de Açúcar (GPA) conducted a corporate restructuring process between 2020 and 2024, including the separation of Assaí's operations in 2021. Available at: gpari.com.br. Accessed in: March 2026.

out noise from consumers uninterested in the asset. In cases of extreme events, such as the collapse of AMER3, whose volatility reached 280% During the validation period, searching by trade name proved to be the strongest signal, reducing the MAPE by 643% for 99.57%, showing that systemic shocks go beyond the specialized financial public and mobilize broad social attention. Finally, MGLU3 stood out as a case of attentional multidimensionality, where only the aggregation of the two semantic dimensions satisfactorily captured the predictive signal.

These findings reinforce that the effectiveness of Google Trends as a predictive variable is not an intrinsic property of the tool, but a function of the taxonomy adopted in the collection and the current market regime. Predictive financial architectures, therefore, require semantic parameterization sensitive to the unique characteristics of each asset, making any generic or aggregated approach that disregards this conditioning inappropriate and points to the success of using a multi-criteria approach guided by market and social indicators.

6.1. Future Works

As a natural outcome of this research, it is suggested to expand the methodological framework through the implementation of different architectural families for the purpose of comparative benchmarking. Future works must include the application of traditional parametric statistical models from the Box-Jenkins family, specifically ARIMA and SARIMA. Being able to evaluate whether the ability to purify the exogenous signal of Google Trendsis based on linear predictive approaches, or whether the extraction of this attentional value is exclusive to the non-linearity of neural networks.

Additionally, within the scope of Deep Learning, it is recommended to explore one-dimensional Convolutional Neural Networks (1D-CNN) or hybrid architectures (CNN-LSTM). This algorithmic evolution aims to optimize the automatic extraction of local features in time series, potentially being more effective in capturing, isolating and weighting very short-term ”stochastic spikes”generated by investor attention and social spillover in online searches. The analysis carried out is still capable of great expansion when macroeconomic data is added, in addition to the possible creation of a decision support model based on methods with a multi-criteria approach.

6.2. acknowledgments.

Thanks to PPGEE and Laboratório de Pesquisa Operacional who provided the tolls who made possible this research and this work and to CNPq CAPES for supporting the researchers involved.

Referências

Andrade, C. F. O., Herrera, P. A. L., and Bonilla, D. A. G. (2025). Google trends as an early signal in international flower trade. *Ornamental Horticulture*, 31:e312959.

Andrei, D. and Hasler, M. (2014). Investor attention and stock market volatility. *The Review of Financial Studies*, 28(1):33–72.

Bantis, E., Clements, M. P., and Urquhart, A. (2023). Forecasting gdp growth rates in the united states and brazil using google trends. *International Journal of Forecasting*, 39(4):1909–1924.

Barber, B. M. and Odean, T. (2007). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial Studies*, 21(2):785–818.

Bolfarine, H. and de Oliveira Bussab, W. (2005). *Elementos de amostragem*. Editora Blucher.

Cebrián, E. and Domenech, J. (2023). Is google trends a quality data source? *Applied Economics Letters*, 30(6):811–815.

Churi, A., Chakraborty, D., Khatwani, R., Pinto, G., Shah, P., and Sekhar, R. (2023). Stock price prediction using deep learning and sentiment analysis. In *2023 2nd International Conference on Futuristic Technologies (INCOFT)*, pages 1–6. IEEE.

Dhanawansa, V., Samarasinghe, P., Yogarajah, P., Gardiner, B., and Karunasena, A. (2022). Comparative study of deep learning parameter selection for multi-output regression on head pose estimation. In *2022 IEEE International Conference on Industrial Technology (ICIT)*, pages 1–6.

Fahrudin, T., Asniar, and Faizul Ula, M. (2022). Comparison the new computer sciences study programs in indonesia using pytrends. In *2022 International Conference on Informatics, Multimedia, Cyber and Information System (ICIMCIS)*, pages 379–383.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2):383–417.

Gopali, S., Siami-Namini, S., Abri, F., and Namin, A. S. (2024). A comparative multivariate analysis of var and deep learning-based models for forecasting volatile time series data. *IEEE Access*, 12:155423–155436.

Hamid, A. and Heiden, M. (2015). Forecasting volatility with empirical similarity and google trends. *Journal of Economic Behavior & Organization*, 117:62–81.

Leippold, M., Wang, Q., and Zhou, W. (2022). Machine learning in the chinese stock market. *Journal of Financial Economics*, 145(2):64–82.

Mesquita, L. A., Ranieri, C. M., and Ueyama, J. (2023). Artificial neural networks applied to time series for flood prediction. In *Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, pages 712–725. SBC.

Petropoulos, A., Siakoulis, V., Stavroulakis, E., Lazaris, P., and Vlachogiannakis, N. (2022). Employing google trends and deep learning in forecasting financial market turbulence. *Journal of Behavioral Finance*, 23(3):353–365.

Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of Economic Perspectives*, 17(1):83–104.

Suthiponpisal, V. and Tancharoen, D. (2024). A comparative evaluation of noise reduction versus data normalization techniques in stock market prediction using transformer models. In *2024 8th International Conference on Information Technology (InCIT)*, pages 775–780.

Taleb, N. N. (2010). *The Black Swan: The Impact of the Highly Improbable, with a New Section on Robustness and Fragility*, volume 2. Random house trade paperbacks.

Wang, C.-H. (2022). Considering economic indicators and dynamic channel interactions to conduct sales forecasting for retail sectors. *Computers & Industrial Engineering*, 165:107965.