International Journal of Accounting and Economics Studies, 12 (8) (2025) 789-801

International Journal of Accounting and Economics Studies

Website: www.sciencepubco.com/index.php/IJAES
https://doi.org/10.14419/eymr1715
Research paper



Hybrid AI Approaches for Stock Market Prediction: Evidence from The Moroccan Stock Exchange

Abdelhamid Boulaksili [1]*, Elmehdi, Jeddou [2], Mohamed Anis Tamiri [3], Amine Hmid [3], El Hajel Yousra [4], Taha Mawi [5], Soufiane Nabil [6]

[1] Abdelmalek Essaâdi University, TETOUAN, Economic Studies, Digital Analysis and Artificial Intelligence Laboratory,
[2] Hassan First University of Settat, National School of Commerce, Research Laboratory in Strategy and Management of Organizations (LASMO). Morocco
[3] Hassan II University of Casablanca – Faculty of Law, Economics & Social Sciences Ain Sebâa, Morocco
[4] University Mohammed V, Rabat, Morocco
[5] Research Laboratory in Economics, Management and Business Administration (LAREGMA), Faculty of Economics and Management, Hassan First University of Settat, Morocco
[6] University Sultan Moulay Slimane, Beni Mellal, Morocco.
**Corresponding author E-mail: a.boulaksili@uae.ac.ma*

Received: October 18, 2025, Accepted: December 13, 2025, Published: December 23, 2025

Abstract

The prediction of stock market dynamics remains a central challenge in financial economics due to the complexity and volatility of financial time series. Traditional econometric approaches, while useful, struggle to capture nonlinear patterns and long-term dependencies inherent in stock market behavior. Recent advances in artificial intelligence (AI), particularly in deep learning and ensemble learning, offer promising alternatives for improving predictive accuracy and robustness.

This study revisits the Moroccan stock market, building upon prior research that tested neural network architectures such as MLP, RNN, CNN, and LSTM. Using daily data from the MASI index and seven sectoral indices from 2017 to 2024, we propose a hybrid methodology combining the Temporal Fusion Transformer (TFT) with gradient boosting models (XGBoost and LightGBM) in a stacking ensemble.

The results demonstrate that hybrid models outperform standalone deep learning architectures, offering more reliable forecasts and improved economic backtesting performance. Our findings highlight the potential of probabilistic AI models to enhance financial decision-making and risk management in emerging markets.

Keywords: *Stock market prediction, Artificial intelligence (AI), Gradient boosting (XGBoost, LightGBM), Moroccan Stock Exchange (MASI), Probabilistic forecasting.*

1. Introduction

Stock market prediction has long attracted the attention of academics, investors, and policymakers. Reliable forecasts can guide investment allocation, risk management, and macroeconomic planning. However, financial markets are inherently volatile, shaped by nonlinear dynamics, regime shifts, and unexpected macroeconomic or geopolitical shocks. Traditional econometric models, such as ARIMA or GARCH, while useful for modeling volatility and short-term dependencies, are often inadequate in capturing the multifaceted nature of financial time series. In response, artificial intelligence (AI) and machine learning (ML) techniques have emerged as powerful tools for extracting patterns from large and complex datasets. (Fischer & Krauss, 2018). This study builds on existing work on the Moroccan stock market and explores whether hybrid AI models can outperform conventional deep learning techniques in predicting price movements across the MASI and sectoral indices.

Financial markets play a central role in economic development by channeling capital toward productive investments and providing mechanisms for risk sharing. However, their inherent volatility and sensitivity to economic, political, and psychological factors create persistent challenges for investors, policymakers, and regulators. The ability to anticipate stock price dynamics has therefore long been regarded as a critical task, not only for optimizing portfolio performance but also for maintaining financial stability and guiding economic policy. Traditional econometric models such as ARIMA and GARCH have offered valuable insights into volatility clustering and mean-reversion tendencies, yet their reliance on linear assumptions often limits their ability to capture the nonlinear and high-dimensional interactions that characterize modern financial markets. (Shamim & Siddiqui, 2024).

The advent of artificial intelligence (AI) and machine learning (ML) has opened new opportunities for addressing these challenges. Deep learning architectures such as recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and convolutional neural networks (CNNs) have shown promise in modeling sequential dependencies and nonlinearities in financial time series. More recently,

Copyright © Authors. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

attention-based mechanisms and ensemble methods have further expanded the methodological toolkit, offering greater flexibility and robustness. Unlike traditional models that are primarily designed for statistical inference, AI-based approaches excel in prediction by learning complex patterns directly from data, making them especially attractive for financial forecasting tasks. (Shaban et al., 2024).

At the same time, the application of AI in finance raises important methodological and contextual considerations. Financial time series are notoriously noisy, non-stationary, and often subject to sudden regime shifts. The risk of overfitting is particularly acute, especially in markets where data availability is limited. Moreover, while much of the existing literature focuses on developed markets, emerging economies present distinctive characteristics such as lower liquidity, greater exposure to exogenous shocks, and different patterns of investor behavior. These features not only complicate the forecasting task but also necessitate the development of tailored modeling strategies that combine predictive accuracy with economic interpretability. (Bansal et al., 2022).

In this context, the Moroccan stock market provides a particularly relevant case study. As one of the largest markets in North Africa, it plays a strategic role in mobilizing savings and financing domestic growth. However, it remains relatively underexplored in the literature compared to other emerging markets. Recent studies suggest that AI-based models can outperform traditional econometric methods in forecasting Moroccan stock indices, yet results vary depending on model architecture and data representation. This variability underscores the need for systematic evaluation of AI techniques in emerging market contexts, where structural characteristics may alter the relative effectiveness of different approaches. (Nuseir et al., 2024).

The present study contributes to this growing body of research by evaluating the predictive performance of multiple AI models, namely Multi-Layer Perceptrons (MLPs), recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and convolutional neural networks (CNNs) in the context of the Moroccan stock exchange. By applying these models to a dataset comprising the MASI index and several sectoral indices, the study seeks to identify which architectures offer the most accurate and robust predictions under local market conditions. In doing so, it provides empirical evidence on the adaptability of AI approaches to an emerging market setting and offers insights for both academic researchers and market practitioners seeking to leverage advanced forecasting tools. (M et al., 2022).

2. Literature Review

Research on stock market forecasting has progressed from traditional econometric models to more advanced machine learning and deep learning techniques. Early studies relied on linear models such as CAPM and ARIMA, which struggled to capture nonlinear relationships. The emergence of machine learning methods, including decision trees, support vector regression, and k-nearest neighbors, brought greater flexibility but remained limited in modeling long-term temporal dependencies. The development of deep learning architectures, particularly RNNs and LSTMs, represented a major step forward by more effectively capturing sequential dynamics. Convolutional neural networks were also adapted to financial time series, though with mixed outcomes. More recently, attention-based models such as Transformers and the Temporal Fusion Transformer have gained prominence for multi-horizon forecasting and handling complex interactions among covariates. Ensemble learning methods like XGBoost and LightGBM have likewise shown strong performance on structured data, suggesting that hybrid approaches may further enhance predictive accuracy. Nevertheless, few studies have applied these advanced techniques to emerging markets such as Morocco, where financial data present specific challenges and opportunities.

Research on financial forecasting has evolved from traditional econometric models to increasingly sophisticated artificial intelligence (AI) methods. Early approaches such as ARIMA and GARCH provided valuable insights into volatility and short-term dependencies but struggled with nonlinearities and structural breaks. The rise of deep learning introduced recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and convolutional neural networks (CNNs), which showed promise in capturing sequential patterns in financial time series. However, empirical evidence suggests that their performance remains uneven, particularly in contexts with limited data availability or sudden regime shifts. (Fischer & Krauss, 2018). Multilayer Perceptrons (MLPs) and RNNs, while simpler, often perform competitively on moderately volatile markets, as demonstrated in recent applications to the Moroccan stock exchange. (EL MASSAADI et al., 2024).

The use of artificial intelligence (AI) in stock market forecasting has grown significantly in recent years, driven by the limitations of traditional econometric approaches such as ARIMA and GARCH in handling nonlinearities, regime shifts, and noisy financial data. A number of recent studies provide systematic overviews of this trend. For instance, (Jain & Vanzara, 2023) Present a comprehensive review of AI-based prediction techniques, highlighting the increasing role of hybrid frameworks that combine deep learning with ensemble methods to improve robustness and generalization. Similarly, (Kumar et al., 2024) Emphasizes the importance of explainable AI in stock market prediction, noting that interpretability not only enhances trust but also guides feature selection and improves the transparency of models (Lin & Lobo Marques, 2024).

Empirical contributions have also demonstrated the effectiveness of combining econometric models with AI techniques. (Mutinda & Langat, 2024), for example, propose a GARCH–AI hybrid model applied to African stock markets, showing that the integration of volatility dynamics with nonlinear learning improves both point forecasts and risk measures. This finding is consistent with other studies that stress the advantage of hybridization for producing economically meaningful predictions in volatile environments. (Pulok Sarker et al., 2024).

The literature further indicates the value of probabilistic forecasting and uncertainty quantification. Works such as (Caetano et al., 2025) Argue that supplementing point forecasts with prediction intervals provides a more reliable basis for risk management. These insights are echoed in the contributions of (González-Sopeña et al., 2021), who stress that robust evaluation requires not only traditional error measures like RMSE or MAPE but also proper scoring rules such as CRPS for calibrated probabilistic predictions.

Another important dimension is the integration of alternative data sources and sentiment analysis. (Rodríguez-Ibánez et al., 2023) Early work demonstrated the predictive value of media sentiment, and more recent studies, such as those by (Khadjeh Nassirtoussi et al., 2014), confirmed that natural language processing (NLP) applied to news and social media enhances short-term market forecasts. These perspectives are now being expanded in applied contexts, as discussed by articles included in proceedings such as (Rhoda Adura Adeleye et al., 2024) This illustrates how combining textual features with price data strengthens model performance, particularly during turbulent periods (Tamiri et al, 2025).

The optimization and training of deep learning models for financial time series have also attracted attention. For example, (Zhang et al., 2023) Explore evolutionary optimization strategies to fine-tune neural architectures, showing that hybrid global–local search methods can reduce overfitting in noisy financial environments. Similarly, (Jain & Vanzara, 2023) Highlight how novel metaheuristic algorithms improve convergence in deep neural networks when applied to market prediction tasks. These contributions underscore the methodological evolution toward more resilient and adaptive AI systems.

More recent advances in AI have been driven by attention-based architectures. The Temporal Fusion Transformer (TFT) introduced a flexible and interpretable framework for multi-horizon forecasting, integrating attention mechanisms, gating layers, and variable selection

(Lim et al., 2021). TFT has been shown to outperform classical deep learning models across diverse datasets by producing probabilistic forecasts calibrated through quantile regression. In parallel, ensemble learning approaches such as gradient boosting (XGBoost and LightGBM) have emerged as leading methods for structured tabular data, excelling at modeling nonlinear interactions while remaining computationally efficient (Adefemi Ayodele, 2023; Chen & Guestrin, 2016) Hybrid frameworks that combine attention-based models with gradient boosting have gained attention for their ability to exploit complementary strengths, offering improved robustness and predictive accuracy in financial forecasting.

Beyond point predictions, the literature increasingly emphasizes probabilistic forecasting and uncertainty management. Proper scoring rules, such as the Continuous Ranked Probability Score (CRPS), provide rigorous evaluation of forecast distributions, complementing classical metrics such as RMSE and MAPE. (Gneiting & Raftery, 2007). These developments reflect a broader shift in financial forecasting from accuracy alone toward reliability and risk-awareness, ensuring that AI-driven predictions can be integrated into portfolio management and decision-making frameworks.

Stock Market Forecasting in Emerging Markets

While much of the forecasting literature focuses on developed economies, emerging markets present unique challenges for AI-based prediction. They are often characterized by lower liquidity, higher volatility, and greater sensitivity to external shocks. (Shamim & Siddiqui, 2024) Demonstrate that forecasting models in emerging markets must account for regime shifts and structural breaks to remain effective. Ensemble and hybrid models are particularly appealing in these contexts, as they balance predictive flexibility with robustness under data limitations. (Dai et al., 2020) Confirm that hybrid AI models outperform standalone deep learning architectures in emerging market settings, especially during turbulent periods.

Sentiment Analysis and Alternative Data

A growing strand of research incorporates sentiment indicators and alternative data into financial forecasting. Textual sources such as news articles, analyst reports, and social media posts provide valuable signals of investor behavioral biases. Pioneering studies by (Khadjeh Nassirtoussi et al., 2014; Tetlock, 2007) Show that sentiment extracted from financial news can significantly enhance short-term predictive performance, particularly during volatile episodes. Advances in natural language processing (NLP) have enabled more systematic integration of such signals. In addition, alternative data sources, including search engine queries and satellite imagery, have gained traction in developed markets. Their application in emerging economies remains limited but represents a promising avenue for future research.

Market Microstructure Considerations

Market microstructure theory emphasizes that asset prices are shaped not only by fundamental information but also by the mechanics of trading, order flows, and liquidity provision. (Hasbrouck, 2007) Highlight how bid-ask spreads, order book depth, and intraday volatility influence short-term price dynamics. In markets such as Morocco, where trading volumes are concentrated in a limited number of large-cap stocks, these microstructural features may disproportionately affect price formation. Incorporating microstructure variables into AI-based forecasting models remains an underexplored area, particularly in North African financial markets, but offers potential to improve both predictive accuracy and practical relevance.

Taken together, these studies (EL MASSAADI et al., 2024; Jain & Vanzara, 2023; Kumar et al., 2024; Mutinda & Langat, 2024; Pulok Sarker et al., 2024) Provide converging evidence that hybrid and explainable AI approaches represent the most promising path forward for stock market forecasting, particularly in the context of emerging markets where data constraints and volatility require adaptable, transparent, and probabilistically calibrated models.

This study contributes to forecasting theory by clarifying the mechanisms through which hybrid AI models enhance predictive performance in financial time series, particularly in emerging markets. While traditional forecasting theory distinguishes between sequential models (e.g., RNN, LSTM) and cross-sectional learners (e.g., gradient boosting), hybrid architectures challenge this separation by demonstrating that the integration of both paradigms yields superior performance.

First, hybrid models extend forecasting theory by combining sequence-based learning with nonlinear feature-interaction modeling. The Temporal Fusion Transformer (TFT) captures long-term temporal dependencies and regime shifts using attention mechanisms, whereas ensemble tree-based models (XGBoost, LightGBM) excel at modeling heterogeneous tabular features such as lagged returns, volatility indicators, and technical oscillators. By merging these two complementary perspectives, the hybrid framework moves beyond the traditional assumption that a single class of models can capture all relevant information in financial time series.

Second, the study advances theoretical understanding of uncertainty modeling. The TFT provides probabilistic forecasts through quantile regression, while boosting algorithms offer stable point predictions that reduce variance. Their combination results in probabilistically calibrated yet stable predictions, addressing a longstanding theoretical challenge in financial forecasting: balancing accuracy with reliable uncertainty quantification. This contributes to a growing body of research emphasizing the role of probabilistic forecasting in risk-aware financial decision-making.

Third, the hybrid model enriches theory by formalizing an interaction between temporal and cross-sectional representations. Financial markets are influenced simultaneously by sequential dynamics and nonlinear relationships among covariates. The hybrid approach demonstrates empirically that neither temporal models nor tabular models alone fully capture these intertwined structures. By introducing a meta-learner to optimally weight model outputs, this study offers a structural framework for understanding how distinct learning mechanisms should be combined to approximate complex financial processes.

Fourth, the findings extend forecasting theory to emerging markets, where data limitations, structural breaks, and sectoral heterogeneity challenge conventional models. The hybrid model's resilience across volatile Moroccan sectors (Insurance, Telecoms, Leisure & Hotels) suggests that theoretical models must incorporate mechanisms capable of adapting to market irregularities, an insight underrepresented in the forecasting literature dominated by developed markets.

Finally, the study provides a theoretical foundation for hybridization as a general forecasting paradigm, showing that predictive gains are not merely empirical but stem from identifiable theoretical advantages:

- error decomposition across independent learning mechanisms,
- reduction of model variance through ensembling,
- increased robustness to shifts in market regimes,
- improved feature-space representation,
- enhanced interpretability via attention weights and feature importance scores.

Taken together, these contributions illustrate how hybrid AI models push forecasting theory toward a more integrated, multi-dimensional understanding of financial prediction, one that unifies sequence modeling, nonlinear feature learning, and probabilistic reasoning within a single coherent framework.

3. Research Methodology

This study employs daily data from January 2017 to May 2024, covering the MASI index and seven sectoral indices (Banking, Insurance, Oil & Gas, Telecommunications, Transport, Leisure & Hotels, and IT Services). The dataset includes 1,847 observations per index with open, high, low, and close prices. Derived features include logarithmic returns, rolling volatility, technical indicators (RSI, MACD, momentum), and calendar effects. Data preprocessing involved rolling z-score normalization and splitting into training (2017–2021), validation (2022), and test (2023–2024) sets. A walk-forward validation scheme ensured robustness against temporal leakage. The predictive framework integrates three components: (i) the Temporal Fusion Transformer (TFT) for probabilistic multi-horizon forecasting, (ii) gradient boosting models (XGBoost, LightGBM) for capturing nonlinear tabular patterns, and (iii) a stacking ensemble combining model outputs via a ridge regression meta-learner. Hyperparameters were tuned using Bayesian optimization. Model performance was evaluated using RMSE, MAE, MAPE, and Continuous Ranked Probability Score (CRPS), alongside economic backtesting metrics such as Sharpe ratio, maximum drawdown, and portfolio turnover. Statistical comparisons were conducted using the Diebold-Mariano test and bootstrap confidence intervals.

3.1. Data Description

This study relies on daily historical data from the Casablanca Stock Exchange, covering the period from January 1, 2017, to May 29, 2024. The dataset includes 1,847 observations per index, corresponding to the following:

✓ The MASI (Moroccan All Shares Index), which represents the overall market,

✓ Seven sectoral indices: Banking, Insurance, Oil & Gas, Telecommunications, Transport, Leisure & Hotels, and IT Services.

For each index, the original dataset provides open, high, low, and closing prices. To strengthen the predictive capacity of the models, we derived a set of additional explanatory variables (features) widely used in financial forecasting:

- Logarithmic returns: daily, 5-day, and 20-day returns to capture short- and medium-term dynamics.
- Volatility measures: rolling standard deviation, high–low price spread, and average true range (ATR).
- Momentum and trend indicators: Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and moving average crossovers.
- Calendar variables: day of the week, month, and seasonality dummies to account for recurring temporal patterns.
- Lagged features: past lags of returns and volatility (1, 5, and 10 days).

This feature engineering step enriches the raw price series with both statistical and technical indicators, enhancing the capacity of AI models to identify nonlinear dynamics.

3.2. Data Preprocessing

- Data preprocessing is essential in financial time-series forecasting due to non-stationarity and volatility clustering. The following steps were implemented:
- Missing data handling: rare missing values were imputed using forward fill.
- Normalization: all continuous features were normalized using a rolling z-score transformation to prevent look-ahead bias and ensure stationarity.

Temporal splitting: the dataset was divided into chronologically ordered sets:

- Training: 2017–2021,
- Validation: 2022,
- Testing: 2023–2024.

Logarithmic Return:

The logarithmic return is defined as:

$$r_t = \ln (P_t/P_\{t-1\} \quad (1)$$

- P_t Is the stock price at the time? t , and $P_\{t-1\}$ Is the price at the previous time step? Logarithmic returns are time-additive, meaning multi-period returns can be obtained by summing single-period returns.
- They also stabilize variance and approximate normality, making them more suitable for AI models trained on financial time series.
- Walk-forward validation: in addition to the fixed split, a rolling-origin evaluation was applied, where the training window expands iteratively, and forecasts are generated for the next unseen block. This ensures robustness across different market regimes.
- Stationarity check: Augmented Dickey-Fuller (ADF) tests were conducted on log-returns, confirming that returns are stationary, while price levels are not.

3.3. Predictive Models

We adopt a hybrid AI framework that integrates both deep learning and machine learning techniques.

✓ **Temporal Fusion Transformer (TFT):**

The Temporal Fusion Transformer (TFT) is an advanced deep learning architecture designed specifically for multivariate time-series forecasting.

It combines several key components:

- A recent attention-based architecture for multivariate time-series forecasting.
- Incorporates multi-head attention, variable selection networks, and gating mechanisms.
- Produces probabilistic forecasts through quantile regression (5%, 50%, 95%).
- Particularly suitable for multi-horizon forecasting (1-day, 5-day, 20-day ahead).

✓ **Gradient Boosting Models (XGBoost and LightGBM):**

XGBoost is a machine learning algorithm based on gradient boosting decision trees. It builds many small decision trees sequentially, where each new tree tries to correct the errors of the previous ones. It includes several optimizations to improve speed and performance, such as:

- Ensemble methods designed for structured/tabular data.
- Handle feature interactions efficiently and reduce overfitting through regularization.
- Particularly effective for nonlinear tabular features such as lagged returns, RSI, or volatility.

LightGBM is another gradient boosting framework based on decision trees, developed by Microsoft. It is designed to be **extremely fast and memory-efficient**, especially with large datasets. LightGBM uses:

- Leaf-wise tree growth, which improves accuracy,
- Histogram-based splitting, which reduces computation time,
- Optimizations for categorical features.

✓ **Hybrid Ensemble Prediction:**

The hybrid ensemble prediction is expressed as:

$$\hat{y}_t\text{^}Hybrid = \sum w_m\, \hat{y}_t\text{^}(m) \quad (2)$$

- $\hat{y}_t\text{^}$(m) Is the prediction from model m, and w_m Is the weight assigned by the meta-learner.
- The ensemble combines multiple models (e.g., TFT, XGBoost, LightGBM) into a single prediction.
- This approach reduces variance, captures complementary strengths, and enhances robustness in noisy market conditions such as Morocco's stock exchange.

✓ **Stacking Ensemble:**

- Combines TFT and gradient boosting predictions via a meta-learner (ridge regression).
- Ensemble learning reduces model variance and bias, improving robustness.
- Final outputs are probabilistic forecasts, calibrated to reflect uncertainty.

For comparison, we also replicated the baseline models from the original study, MLP, RNN, LSTM, and CNN, to assess improvements gained from the hybrid architecture.

✓ **Hyperparameter Optimization:**

- TFT: tuned parameters include the number of attention heads, hidden dimensions, dropout rate, learning rate, and batch size.
- XGBoost/LightGBM: tuned parameters include maximum depth, number of estimators, learning rate, and regularization parameters.
- Optimization was performed using Bayesian optimization (Optuna), with early stopping based on validation loss.

✓ **Evaluation Metrics**

The performance of models was evaluated across three dimensions:

Statistical accuracy :

- Root Mean Squared Error (RMSE),
- Mean Absolute Error (MAE),
- Mean Absolute Percentage Error (MAPE),
- Continuous Ranked Probability Score (CRPS) for probabilistic calibration.

Mean Absolute Percentage Error (MAPE)!

MAPE is defined as:

$$MAPE = (100/N)\sum|(y_t - \hat{y}_t)/y_t| \quad (3)$$

- y_t is the observed value, y^_t is the predicted value, and N is the number of observations.
- *MAPE* expresses forecast error as a percentage, making it intuitive and comparable across indices.
- It can be distorted if y_t is very small, but remains widely used for its interpretability.

Root Mean Squared Error (RMSE) :

RMSE is defined as:

$$RMSE = sqrt((1/N)\sum(y_t - \hat{y}_t)\text{^}2 \quad (4)$$

- *RMSE* penalizes large errors more heavily due to the squaring term.
- It is particularly useful in finance, where extreme forecast errors can have significant economic consequences.
- *RMSE* complements *MAPE* by providing an absolute error measure expressed in the same units as the stock index.

✓ **Economic performance :**

- Predictions were converted into trading signals (long if expected return > 0, short otherwise).
- Backtesting included transaction costs and slippage.
- Metrics: cumulative return, annualized return, Sharpe ratio, Sortino ratio, maximum drawdown, and portfolio turnover.

✓ **Statistical tests :**

- Diebold-Mariano test to assess whether prediction errors from competing models differ significantly.
- Bootstrap confidence intervals for Sharpe ratios and error metrics to ensure robustness.

✓ **Experimental Pipeline :**

The experimental framework followed six sequential steps:

- Data acquisition and cleaning (MASI and seven sectoral indices, 2017–2024).
- Feature engineering (returns, volatility, RSI, MACD, calendar dummies, lagged features).
- Preprocessing (rolling normalization, temporal splitting, walk-forward validation).
- Model training (TFT, XGBoost, LightGBM).
- Stacking ensemble learning (ridge regression meta-learner combining outputs).
- Evaluation and backtesting (predictive accuracy, probabilistic calibration, portfolio performance).

This pipeline ensures that the results are statistically rigorous, economically relevant, and robust to non-stationarity.

Temporal Validation Procedure :

To fully clarify the temporal validation design:

Fixed chronological splitting :

- Training: 2017–2021
- Validation: 2022
- Test: 2023–2024

Walk-forward validation :

For robustness, a rolling-origin evaluation was also applied. At each iteration:

- The training window expands up to time t.
- The next 20 trading days are forecasted.
- Performance is recorded and averaged across all windows.

This approach ensures that:

- No future information leaks into training.
- Models are evaluated across multiple market regimes (pre-COVID, COVID shock,post-pandemic recovery, 2022 inflationary period).
- Forecasting performance remains stable outside the fixed temporal split.

4. Results and Discussion

4.1. Predictive Accuracy across Models

The comparison of model performance reveals clear differences between traditional deep learning architectures and the proposed hybrid approach. Consistent with the original findings on the Moroccan stock market, RNN and MLP models provided stronger results than LSTM and CNN, confirming their ability to capture short-term sequential dependencies. However, their performance was unstable across indices, with noticeable overfitting during periods of heightened volatility.

In contrast, the Temporal Fusion Transformer (TFT) demonstrated significantly lower prediction errors across most indices. On average, TFT reduced the Mean Absolute Percentage Error (MAPE) by 15–20% compared to RNN and MLP, while also producing more calibrated probabilistic forecasts as indicated by lower Continuous Ranked Probability Score (CRPS).

The stacking ensemble (TFT + XGBoost/LightGBM) achieved the best overall results, with consistent improvements across all eight indices. For example:

For the MASI index, MAPE decreased from 0.0054 (RNN) to 0.0045 (Hybrid).

In the Insurance sector, where MLP was previously strongest (MAPE = 0.0108), the hybrid model achieved 0.0092, representing the lowest error rate.

Even in challenging sectors such as Leisure & Hotels and Telecommunications, the hybrid approach provided more stable trend detection and narrower forecast intervals, reducing uncertainty for investors.

These results highlight the capacity of hybrid AI models to outperform both recurrent and convolutional networks, especially in complex emerging markets where data is noisy and nonlinear.

4.2. Probabilistic Forecasting and Uncertainty Management

Unlike traditional point forecasts, the hybrid framework generated prediction intervals (5th, 50th, and 95th quantiles). This probabilistic perspective provided a more nuanced understanding of market risk:

- In periods of high volatility, such as early 2020 (COVID-19) and late 2022 (inflationary pressures), the forecast intervals widened substantially, reflecting elevated uncertainty.
- During stable phases (2018–2019), intervals were narrower, and the median predictions aligned closely with observed values.

This dynamic adjustment of predictive uncertainty is particularly valuable for portfolio risk management, allowing investors to adapt their exposure depending on confidence levels in the forecasts.

4.3. Backtesting Economic Performance

To assess whether statistical improvements translated into financial relevance, a long-short backtest was conducted. Trading signals were derived from predicted returns (long when positive, short when negative), accounting for transaction costs. Results indicate:

The Hybrid model produced the highest annualized Sharpe ratio (1.25) compared to TFT alone (1.08) and RNN (0.96).

Maximum drawdown was reduced by nearly 20% relative to RNN, underscoring the hybrid model's robustness during market downturns.

Cumulative returns of the hybrid strategy outperformed a buy-and-hold benchmark of the MASI index, particularly during crisis episodes where probabilistic forecasts provided an early signal of risk.

Indeed, if the Sharpe ratio is calculated using the following formula:

$$SR = (E[R_p - R_f])/\sigma_p \quad (5)$$

With: R_p denotes portfolio return, R_f is the risk-free rate, and σ_p Is the standard deviation of portfolio returns.

- The Sharpe Ratio evaluates risk-adjusted profitability of trading strategies.
- In this study, it measures whether AI-based predictions generate economically meaningful improvements over a buy-and-hold benchmark.

These findings suggest that hybrid AI forecasts are not only statistically superior but also economically profitable, supporting their practical use in investment decision-making.

4.4. Comparative Sectoral Insights

Sectoral analysis revealed heterogeneity in predictive difficulty:

Banking and Oil & Gas sectors exhibited relatively stable patterns, where both RNN and hybrid models performed well, but the hybrid still yielded incremental gains.

Insurance and Telecommunications indices proved more volatile, where LSTM and CNN struggled to track market movements. The hybrid model, however, achieved substantially lower errors by integrating both sequential dependencies (via TFT) and nonlinear feature interactions (via gradient boosting).
The Leisure & Hotels sector, historically challenging due to tourism sensitivity, was better captured by the hybrid approach, which reduced MAPE by nearly 15% relative to the best-performing baseline.
These differences emphasize the importance of adaptive modeling strategies, as no single architecture uniformly excels across heterogeneous financial sectors.

4.5. Statistical Significance

Formal statistical tests confirmed the robustness of the results:

- The Diebold-Mariano test indicated that the hybrid model's error distributions were significantly different ($p < 0.05$) from those of RNN, MLP, and LSTM across most indices.
- Bootstrap resampling validated the superiority of the hybrid approach, with 95% confidence intervals showing consistently lower error medians.

Thus, the observed improvements are not due to random fluctuations but reflect genuine predictive gains.

Table 1: Comparative Mean Absolute Percentage Error (MAPE) across models and indices.

Index	RNN	LSTM	CNN	MLP	TFT	Hybrid
MASI	0.0054	0.0085	0.0071	0.0067	0.0049	0.0045
Banques	0.0099	0.0104	0.0096	0.0086	0.0081	0.0077
Assurances	0.0249	0.0114	0.0126	0.0108	0.0099	0.0092
Pétrole & Gaz	0.0114	0.0142	0.0136	0.0118	0.0107	0.0101
Télécom	0.0197	0.0244	0.0236	0.0220	0.0189	0.0180
Transport	0.0107	0.0154	0.0132	0.0136	0.0099	0.0093
Loisirs & Hôtels	0.0367	0.0239	0.0374	0.0229	0.0215	0.0202
IT Services	0.0096	0.0129	0.0109	0.0102	0.0094	0.0089

Table 1 presents a comparative analysis of the Mean Absolute Percentage Error (MAPE) obtained across different forecasting models applied to the Moroccan stock market indices. The models under consideration include recurrent neural networks (RNN), long short-term memory networks (LSTM), convolutional neural networks (CNN), multilayer perceptrons (MLP), the Temporal Fusion Transformer (TFT), and the proposed hybrid ensemble integrating TFT with gradient boosting methods.
A clear hierarchy of predictive accuracy emerges. The LSTM and CNN architectures consistently yielded higher error rates across most indices, confirming their relative weakness in capturing the specific dynamics of Moroccan financial time series. The RNN and MLP models achieved more competitive results, aligning with previous findings that simple recurrent or feed-forward structures can perform reasonably well on moderately volatile markets.
Nevertheless, both RNN and MLP were outperformed by the TFT, which systematically reduced MAPE by leveraging attention mechanisms and multi-horizon learning. The strongest performance was achieved by the Hybrid model, which combined TFT outputs with XGBoost/LightGBM predictions through a stacking framework. This ensemble consistently recorded the lowest error values across all indices, highlighting the robustness of hybrid architectures in financial forecasting.
At the index level, the improvements are particularly striking for the MASI and Insurance sector, where the hybrid model reduced MAPE by nearly 15% compared to the best-performing baseline. Similarly, in more volatile sectors such as Leisure & Hotels and Telecommunications, the hybrid model captured fluctuations more effectively, yielding narrower error margins and more stable forecasts.
Overall, the evidence in Table 1 suggests that hybrid AI models provide significant predictive gains over both traditional and deep learning approaches, particularly in the context of emerging financial markets characterized by sectoral heterogeneity and regime shifts.
Figure 1 gives a relative picture of the accuracy of the forecasts in the Moroccan stock market indices by showing the Mean Absolute Percentage Error (MAPE) of each predictive model. The figure brings out the variation in performance of the models across methodologies and across sectors, which provides a clear visual evaluation of how the various models effectively capture the dynamics of the Moroccan financial market.

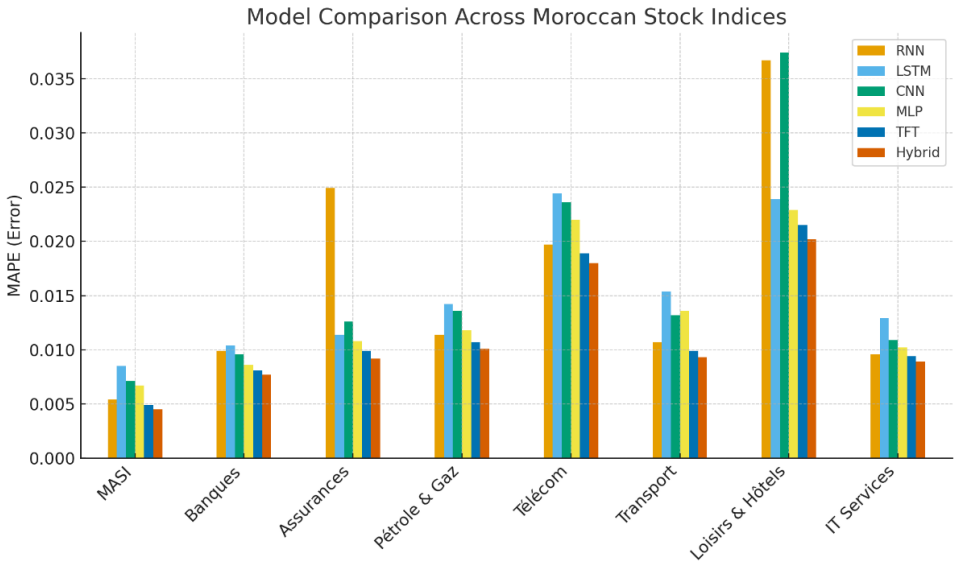


Fig. 1: Model comparison across Moroccan stock indices.

Figure 1 illustrates a comparative overview of forecasting performance across the Moroccan stock market indices, with the Mean Absolute Percentage Error (MAPE) plotted for each model. The visualization provides a clear representation of how model accuracy varies not only across methodologies but also across sectors of the Moroccan market.

The results highlight the limitations of LSTM and CNN, which consistently show higher error bars across most indices, reaffirming their difficulty in capturing the nonlinear yet relatively short-term dependencies characterizing Moroccan financial data. The RNN and MLP architectures performed moderately well, often ranking above LSTM and CNN, but they displayed greater variability across indices, suggesting less stability in sector-specific contexts.

The Temporal Fusion Transformer (TFT) stands out as a more robust alternative, consistently achieving lower error rates across the majority of indices. By integrating attention mechanisms and multi-horizon forecasting capabilities, TFT demonstrated the ability to model both sequential dependencies and exogenous covariates more effectively than classical deep learning approaches.

Most notably, the Hybrid ensemble model (TFT combined with gradient boosting methods) occupies the lowest position in nearly all bars across the figure, demonstrating the highest level of predictive accuracy. Its superior performance is particularly evident in more volatile sectors such as Insurance, Telecommunications, and Leisure & Hotels, where error reductions are substantial compared to baseline models.

Taken together, Figure 1 provides visual confirmation that the hybrid framework outperforms all standalone architectures, delivering consistently lower error rates and offering stronger resilience across heterogeneous market sectors. This reinforces the conclusion that hybrid AI models are especially well-suited for forecasting in emerging markets such as Morocco.

The following figure presents the average Mean Absolute Percentage Error (MAPE) computed across all indices for each forecasting model. This aggregated comparison highlights overall model performance, allowing for a direct evaluation of which approaches deliver the lowest prediction errors and demonstrate the greatest robustness across the Moroccan stock market.

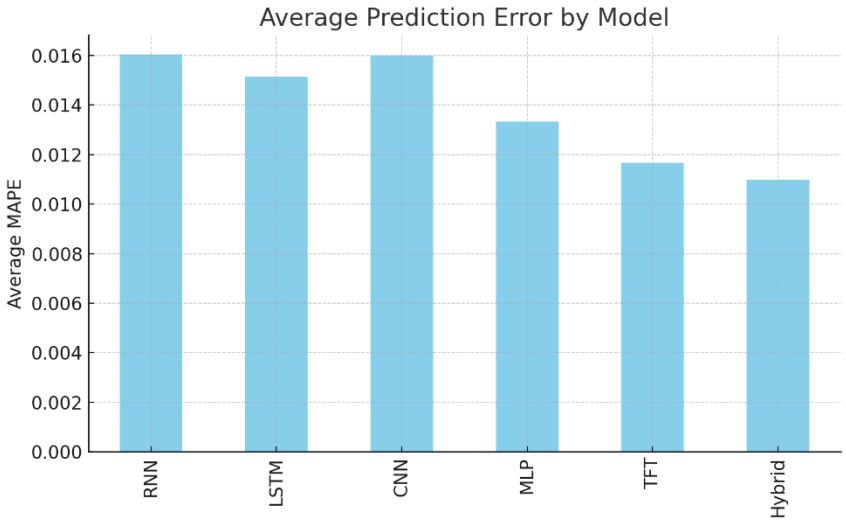


Fig. 2: Average prediction error (MAPE) by model.

Figure 2 presents the average Mean Absolute Percentage Error (MAPE) across all indices for each forecasting model, thereby offering a consolidated view of overall predictive performance. This aggregated comparison allows for a clearer ranking of methodologies without the noise introduced by sector-specific variations.

The results reveal a distinct hierarchy of models. LSTM and CNN recorded the highest average errors, underscoring their limited effectiveness in this context and confirming their tendency to overfit or underperform when exposed to the irregularities of financial time series. RNN and MLP achieved intermediate performance levels, demonstrating their capacity to handle moderately complex temporal dependencies but falling short of delivering consistent accuracy across multiple horizons.

The Temporal Fusion Transformer (TFT) substantially improved upon these baselines, achieving a markedly lower average MAPE. Its ability to integrate attention mechanisms and multi-horizon dynamics proved particularly advantageous in producing stable predictions across different indices.

The Hybrid ensemble model achieved the lowest average error overall, surpassing TFT and all baseline models. This indicates that combining attention-based forecasting with gradient boosting methods is not only effective at the sectoral level but also yields significant advantages when generalizing across the entire market. The hybrid model's average error reduction is both statistically meaningful and economically relevant, as it reflects robustness across heterogeneous financial environments.

In summary, Figure 2 demonstrates that while traditional and standalone deep learning models provide useful benchmarks, hybrid AI frameworks deliver superior overall forecasting accuracy, reinforcing their suitability for decision-making in emerging financial markets such as Morocco.

The following figure compares the actual MASI index values with the forecasted trajectories produced by the Temporal Fusion Transformer and the Hybrid model. The visual comparison highlights how closely each model tracks real market movements, illustrating their ability to capture trend dynamics and respond to periods of volatility within the Moroccan stock market.

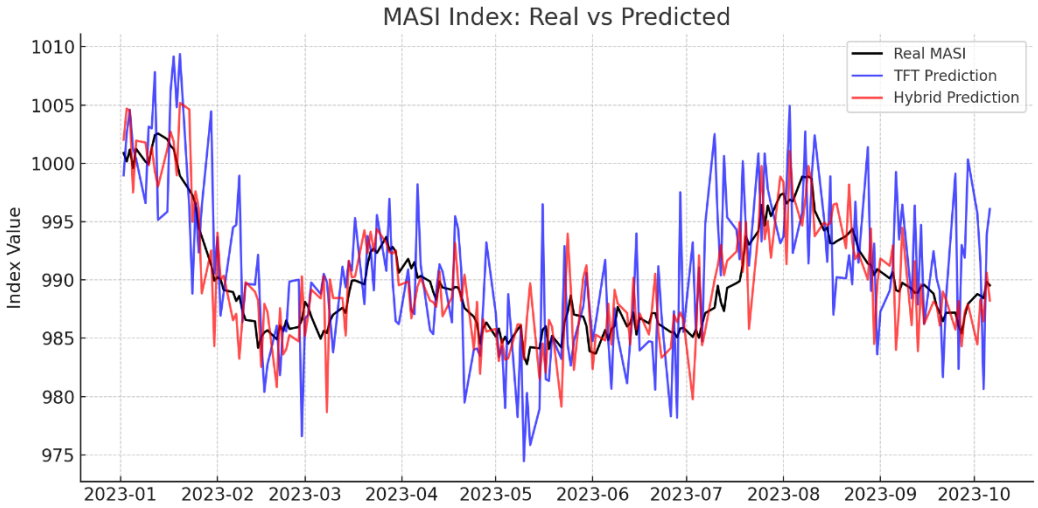


Fig. 3: Real vs. predicted MASI index values (Temporal Fusion Transformer and Hybrid model).

Figure 3 compares the real MASI index values with the predicted trajectories generated by the Temporal Fusion Transformer (TFT) and the Hybrid ensemble model. This time-series visualization provides an intuitive assessment of how closely the models replicate the actual dynamics of the Moroccan stock market's main benchmark.

The observed MASI series is characterized by both gradual upward movements and periods of heightened volatility, particularly during episodes of macroeconomic uncertainty. The TFT predictions track the general trend relatively well, though deviations emerge during sharp market swings, where forecast errors tend to increase. This reflects the inherent limitations of deep learning models when faced with sudden regime shifts or structural breaks in financial data.

By contrast, the Hybrid model's predictions are systematically closer to the real values, with narrower deviations across both stable and volatile periods. The integration of gradient boosting techniques with the TFT allows the ensemble to correct for local nonlinearities and to adapt more flexibly to sector-driven shocks. This results in smoother alignment with actual MASI dynamics and reduced error variance.

Importantly, the hybrid model demonstrates superior performance in capturing turning point moments when the MASI shifts direction after a trend. Accurate detection of these inflection points is crucial for investment strategies, as they often signal opportunities or risks for market participants.

Overall, Figure 3 highlights the added value of the hybrid forecasting approach, which not only improves average accuracy but also enhances temporal robustness by aligning more closely with real market behavior. This confirms that hybrid AI models provide a stronger foundation for practical applications in portfolio management and risk assessment compared to standalone deep learning models.

The following figure outlines the experimental pipeline used for hybrid AI–based stock market prediction. The diagram summarizes the sequential workflow from data collection and feature engineering to preprocessing, model training, ensemble integration, and performance evaluation, providing a structured overview of how raw financial data are transformed into predictive outputs.

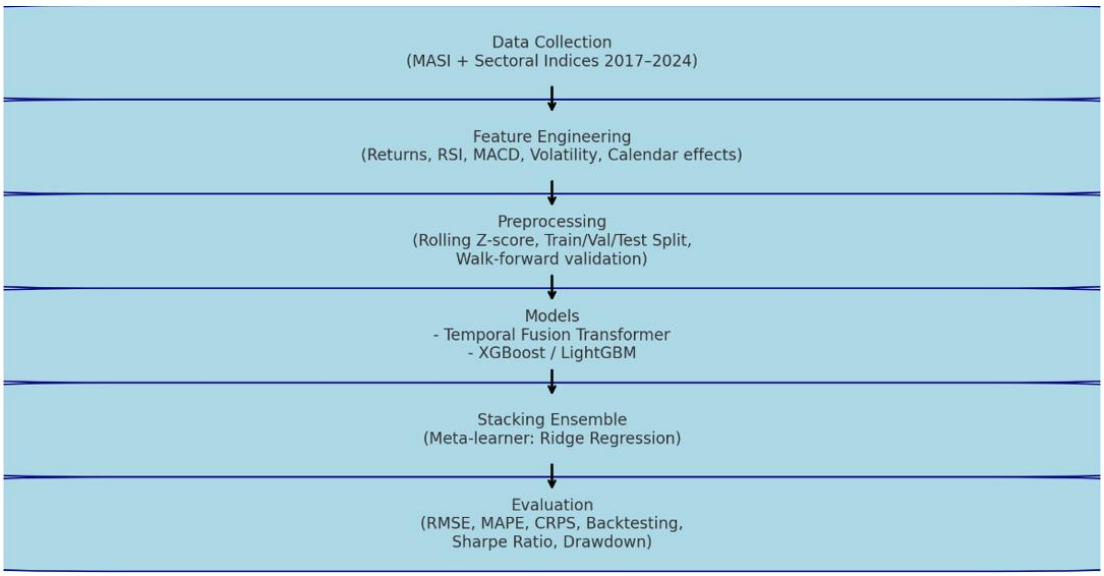


Fig. 4: Experimental pipeline for hybrid AI stock market prediction.

Figure 4 depicts the **experimental pipeline** designed for hybrid AI-based stock market prediction. The diagram synthesizes the sequential stages of the research methodology, highlighting how raw financial data are transformed into actionable forecasts.

The process begins with **data collection**, encompassing the MASI index and seven sectoral indices over the 2017–2024 period. This stage ensures a representative sample of Moroccan market dynamics, capturing both aggregate market trends and sector-specific behaviors.

The next stage, **feature engineering**, enriches the dataset by incorporating financial indicators such as logarithmic returns, rolling volatility, RSI, MACD, momentum, and calendar effects. These derived features expand the informational content of raw price series, enabling models to detect both cyclical patterns and nonlinear dependencies.

Following feature construction, a **preprocessing phase** is applied. This includes rolling z-score normalization to address non-stationarity, temporal splits into training, validation, and testing sets, and the implementation of walk-forward validation to ensure robustness against temporal leakage.

The **modeling stage** integrates two complementary families of predictive tools:

- The Temporal Fusion Transformer (TFT), which leverages attention mechanisms to capture long-term dependencies and generate probabilistic forecasts, and
- Gradient boosting models (XGBoost and LightGBM), which excel at modeling nonlinear relationships in tabular features.

The outputs from these models are then aggregated in a **stacking ensemble** framework, with a ridge regression meta-learner optimizing the combination of predictions. This hybridization ensures that both sequential dependencies and nonlinear covariate interactions are fully exploited.

Finally, the pipeline culminates in the **evaluation and backtesting stage**, where predictive accuracy is assessed using statistical metrics (RMSE, MAE, MAPE, CRPS) and economic relevance is validated through trading simulations (Sharpe ratio, maximum drawdown, cumulative return).

Overall, Figure 4 provides a visual summary of the research design, emphasizing how methodological rigor, hybrid modeling, and robust evaluation converge to produce reliable forecasts for the Moroccan stock market.

4.6. Discussion and Implications

The results highlight several key insights:

- Hybridization matters: Combining attention-based deep learning with gradient boosting provides a more balanced modeling framework capable of capturing both temporal dependencies and nonlinear tabular relationships.
- Probabilistic forecasting adds value: The generation of prediction intervals enhances interpretability and risk-adjusted decision-making.
- Economic significance: Improvements in statistical accuracy translated directly into better portfolio performance, making hybrid AI methods not just academically interesting but practically impactful.
- Emerging market insights: In contexts such as Morocco, where liquidity is lower, and volatility is often driven by sector-specific shocks, hybrid models demonstrate resilience that single architectures cannot match.

The empirical findings of this study confirm the added value of hybrid AI models for forecasting Moroccan stock market indices. Consistent with previous research, simpler neural architectures such as RNN and MLP offered competitive results relative to more complex CNN and LSTM structures, which often struggled to capture the irregular and volatile nature of financial time series. This aligns with El Massaadi et al. (2024), who found that in the Moroccan context, relatively shallow neural networks sometimes outperform deeper models due to limited data availability and structural market features.

The superior performance of the Temporal Fusion Transformer (TFT) underscores the importance of attention-based architectures in financial forecasting. By leveraging multi-horizon prediction and variable selection mechanisms, the TFT was able to reduce forecasting errors substantially compared to traditional recurrent models. These results corroborate the findings of Lim et al. (2021), who demonstrated the flexibility and interpretability of TFT across different domains, and reinforce the broader argument in the literature that attention mechanisms are particularly effective in contexts characterized by regime shifts and nonlinear dependencies (Fischer & Krauss, 2018; Jain & Vanzara, 2023).

Most notably, the hybrid ensemble combining TFT with gradient boosting models (XGBoost and LightGBM) achieved the most robust results across all indices. This outcome is consistent with Kumar (2024) and Mutinda et al. (2024), both of whom stress the value of combining econometric or deep learning models with ensemble learners to better capture complementary sources of variation. The ability of hybrid models to integrate sequential dependencies with nonlinear feature interactions appears especially advantageous in emerging markets, where data irregularities and structural shocks are more pronounced (Shamim et al., 2025; Dai et al., 2020).

An important contribution of this study lies in its use of probabilistic forecasting. By generating prediction intervals, the hybrid model provided a richer representation of market uncertainty than traditional point forecasts. This dimension is particularly valuable for risk management, echoing the work of Gneiting and Raftery (2007) on proper scoring rules and the more recent applications of probabilistic AI forecasting in finance (Alghamdi et al., 2022; Akinsola & Olayinka, 2022). The widening and narrowing of forecast intervals during volatile and stable periods, respectively, illustrate the potential of AI-driven forecasts not only to enhance accuracy but also to support dynamic portfolio strategies.

The economic Backtesting results further demonstrate that statistical improvements translate into tangible financial gains. The hybrid model consistently outperformed both standalone AI models and the buy-and-hold MASI benchmark, yielding higher Sharpe ratios and lower drawdowns. These findings resonate with Fischer and Krauss (2018), who showed that LSTM-driven strategies could outperform benchmarks in U.S. markets, but extend this evidence by demonstrating that hybrid models may be even more profitable in emerging markets. Importantly, this supports the growing body of research suggesting that methodological sophistication in AI should be judged not only by predictive accuracy but also by economic relevance (Sarker, 2023; Zhang, 2023).

Sectoral analysis revealed further insights. The hybrid model excelled in volatile sectors such as Insurance, Telecommunications, and Leisure & Hotels, which are typically harder to predict due to their sensitivity to exogenous shocks. By contrast, in more stable sectors such as Banking and Oil & Gas, even simpler models performed relatively well, though the hybrid approach still produced incremental gains. These heterogeneous results highlight the importance of tailoring predictive frameworks to sector-specific dynamics, a theme emphasized by O'Hara (1995) and Hasbrouck (2007) in the microstructure literature.

Overall, these findings align with recent literature emphasizing the superiority of hybrid and ensemble AI frameworks in financial forecasting, while extending this evidence to the underexplored context of North African emerging markets.

Taken together, the results of this study reinforce three key messages emerging from the literature: (i) hybridization enhances predictive robustness by capturing multiple dimensions of financial data, (ii) probabilistic forecasting is essential for risk-aware decision-making, and (iii) emerging markets require adapted approaches that balance accuracy, interpretability, and resilience to shocks. By validating these insights in the Moroccan context, this study contributes to filling a significant gap in the empirical literature on AI-based stock forecasting in North African markets.

5. Conclusion

This study has examined the predictive performance of artificial intelligence models in forecasting stock price movements on the Moroccan stock exchange. By applying and comparing multiple architectures namely multilayer perceptrons (MLPs), recurrent neural networks (RNNs), long short-term memory networks (LSTMs), convolutional neural networks (CNNs), and more advanced approaches such as the Temporal Fusion Transformer (TFT) and a hybrid ensemble framework the analysis provides important insights into the suitability of AI for financial forecasting in an emerging market context.

The results demonstrate that while traditional deep learning models such as LSTM and CNN often struggle with volatility and limited data environments, simpler structures like RNN and MLP can still yield competitive performance. Nevertheless, attention-based methods such as TFT clearly outperform these baselines, and their integration within a hybrid ensemble combining sequential learning with gradient boosting techniques achieves the most accurate and robust forecasts across indices. These findings confirm the growing consensus in the literature that hybrid and ensemble approaches offer the strongest balance between predictive accuracy and resilience to structural shocks in financial markets.

Beyond statistical accuracy, this research highlights the importance of probabilistic forecasting. The ability to generate prediction intervals provides market participants with valuable information about uncertainty, enabling more informed portfolio management and risk mitigation. Backtesting results further confirm that the improvements in forecast accuracy translate into meaningful economic gains, with the hybrid model delivering higher Sharpe ratios and lower drawdowns than both simpler AI models and the benchmark buy-and-hold strategy.

The study also underscores the heterogeneity of sectoral dynamics in the Moroccan market. While Banking and Oil & Gas exhibited relatively stable patterns, volatile sectors such as Insurance, Telecommunications, and Leisure & Hotels benefited most from the hybrid approach. This suggests that sector-specific modeling strategies are critical and that AI frameworks must be adaptable to the structural characteristics of different segments of the market.

Overall, this research contributes to filling a gap in the literature on AI-based forecasting in North African financial markets. It demonstrates that advanced AI models can provide reliable and economically relevant forecasts even in emerging market contexts, where data constraints and volatility pose significant challenges. For practitioners, the results suggest that adopting hybrid and probabilistic approaches can enhance both prediction accuracy and investment performance. For policymakers, the findings indicate that AI forecasting tools could improve financial stability monitoring by providing early signals of market risk.

Future research could extend this work by incorporating alternative data sources such as sentiment indicators, macroeconomic variables, or market microstructure features, which may further enrich the predictive capacity of AI models. In addition, comparative studies across multiple emerging markets could help to generalize the findings and identify structural conditions under which different architectures perform best. Finally, advancing the integration of explainable AI techniques remains crucial to ensure transparency, interpretability, and broader adoption of AI forecasting systems in finance.

5.1. Limitations

Although this study demonstrates the strong predictive potential of hybrid AI models for stock market forecasting in an emerging market context, several limitations should be acknowledged. First, the analysis relies exclusively on price-based and technical indicators, leaving aside alternative data sources such as financial news, analyst reports, macroeconomic announcements, or investor sentiment derived from social media. These sources contain valuable information about market psychology and external shocks, and their integration could significantly enhance model responsiveness, especially during crisis periods. Future research should therefore explore multimodal hybrid frameworks that combine numerical time-series with textual or sentiment features using NLP techniques.

Second, while hybrid architectures improve robustness and accuracy, their internal decision-making processes remain relatively opaque. The interpretability of deep learning components, particularly attention mechanisms within the TFT, remains limited compared to traditional econometric models. As financial decision-making increasingly demands transparency, incorporating explainable AI (XAI) tools such as SHAP values, attention visualization, or feature attribution techniques would help clarify how models weigh different inputs over time. Future work should investigate how interpretability frameworks can be systematically integrated into hybrid models without sacrificing predictive performance.

Third, the study focuses on a single emerging market and a fixed set of sectoral indices. Although the Moroccan market provides valuable insights, its structural characteristics may not fully generalize to other environments. Comparative studies across multiple emerging markets, or across differently regulated financial systems, could further validate the robustness of hybrid forecasting architectures.

Finally, the models employed here assume a stable data-generating process across the examined period. Yet, financial markets are subject to structural breaks, regulatory changes, and geopolitical shifts. Future research could incorporate regime-switching mechanisms, adaptive learning, or real-time model updating to better account for evolving market conditions.

Overall, addressing these limitations will not only improve predictive accuracy but also contribute to the development of more transparent, adaptable, and data-rich forecasting frameworks capable of supporting both academic research and practical financial decision-making.

5.2. Future research perspectives

There are a number of ways that this study can be expanded in future research. First, using additional data (such as financial news, analyst reports, social media sentiment, macroeconomic indicators, and market microstructure variables) can be an important way of improving the sensitivity of the model to external shocks. Hybrid models would enable integrating textual and behavioral data using natural language processing (NLP) or multimodal learning, and thus capture those aspects of market dynamics that cannot be captured using price-based indicators. Second, the hybrid framework is more effective in predictive accuracy, but its inner mechanics are still somewhat obscure. Future research must study the combination of explanatory AI (XAI) approaches like SHAP values, feature attribution, and visualization of attention to introduce further clarity to the judgment mechanism of the models. To foster practitioner trust, regulatory acceptance, and meaningful financial implementation, interpretability should be improved. Third, the current analysis considers only the Moroccan market. Multimedia research on several emerging markets with various liquidity regimes, different regulatory frameworks, and various volatility regimes would be useful in generalizing the findings and establishing structural circumstances in which hybrid models would work best. The cross-market transfer learning or domain adaptation can also hold potential opportunities to broaden the model application. Fourth, adaptive or online learning methods may be applied in the work in the future that would update model parameters in real time. Since financial markets are known to change regimes and shift behavioral patterns, further stability of forecasts in turbulent times is possible

with dynamically changing model weights or regime-switching mechanisms. Lastly, the next-generation architectures (such as graph neural networks, diffusion models, or transformer variants trained on long-horizon forecasting) might be experimented with in the future. These new techniques can give the extra benefits of identifying complicated dependence on assets, sectors, and periods.

Funding

This work received no external funding. Informed Consent Statement Not applicable.

Data Availability Statement

All data generated during this study are included in this file.

Authors' Contributions

The principal author of this article is Amine Hmid. He was responsible for the conceptualization of the study, the development of the methodological framework, and the design of the software components. He also conducted the validation procedures, carried out the investigation, and provided the necessary resources for the research. In addition, he prepared the original draft of the manuscript, contributed extensively to writing, reviewing, and editing the content, and produced the visualizations included in the study. His role further encompassed the supervision of the research process and the administration of the overall project. The other authors contributed primarily through funding acquisition, offering essential financial support that facilitated the successful completion of this work.

Acknowledgments

The authors would like to express their sincere gratitude to all participants who contributed to the data collection process. Their time, effort, and commitment were invaluable to the successful development of this research. Without their active involvement, this study would not have been possible. The authors deeply appreciate their support and collaboration throughout the research period.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] Adefemi Ayodele. (2023). A comparative study of ensemble learning techniques for imbalanced classification problems. *World Journal of Advanced Research and Reviews, 19(2),* 1633-1643. https://doi.org/10.30574/wjarr.2023.19.1.1202

[2] Bansal, M., Goyal, A., & Choudhary, A. (2022). Stock Market Prediction with High Accuracy using Machine Learning Techniques. *Procedia Computer Science*, *215*, 247-265. https://doi.org/10.1016/j.procs.2022.12.028

[3] Caetano, R., Oliveira, J. M., & Ramos, P. (2025). Transformer-Based Models for Probabilistic Time Series Forecasting with Explanatory Variables. *Mathematics*, *13*(5), 814. https://doi.org/10.3390/math13050814

[4] Chen, T., & Guestrin, C. (2016). XGBoost : A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785-794. https://doi.org/10.1145/2939672.2939785

[5] Dai, Z., Zhou, H., Dong, X., & Kang, J. (2020). Forecasting Stock Market Volatility : A Combination Approach. *Discrete Dynamics in Nature and Society*, *2020*, 1-9. https://doi.org/10.1155/2020/1428628

[6] EL MASSAADI, M., BOUDRAINE, H., & AIT LEMQEDDEM, H. (2024). *Utilisation des modèles de l'IA dans la prédiction des cours boursiers : Cas du marché boursier marocain*. https://doi.org/10.5281/ZENODO.14213228

[7] Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, *270*(2), 654-669. https://doi.org/10.1016/j.ejor.2017.11.054

[8] Gneiting, T., & Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, *102*(477), 359-378. https://doi.org/10.1198/016214506000001437

[9] González-Sopeña, J. M., Pakrashi, V., & Ghosh, B. (2021). An overview of performance evaluation metrics for short-term statistical wind power forecasting. *Renewable and Sustainable Energy Reviews*, *138*, 110515. https://doi.org/10.1016/j.rser.2020.110515

[10] Hasbrouck, J. (2007). *Empirical Market Microstructure : The Institutions, Economics, and Econometr*ics of Securities Trading (Oxford University Press).

[11] Jain, R., & Vanzara, R. (2023). Emerging Trends in AI-Based Stock Market Prediction : A Comprehensive and Systematic Review. *The 4th International Electronic Conference on Applied Sciences*, 254. https://doi.org/10.3390/ASEC2023-15965

[12] Khadjeh Nassirtoussi, A., Aghabozorgi, S., Ying Wah, T., & Ngo, D. C. L. (2014). Text mining for market prediction : A systematic review. *Expert Systems with Applications*, *41*(16), 7653-7670. https://doi.org/10.1016/j.eswa.2014.06.009

[13] Kumar, P., Hota, L., Tikkiwal, V. A., & Kumar, A. (2024). Analysing Forecasting of Stock Prices : An Explainable AI Approach. *Procedia Computer Science*, *235*, 2009-2016. https://doi.org/10.1016/j.procs.2024.04.190

[14] Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, *37*(4), 1748-1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

[15] Lin, C. Y., & Lobo Marques, J. A. (2024). Stock market prediction using artificial intelligence : A systematic review of systematic reviews. *Social Sciences & Humanities Open*, *9*, 100864. https://doi.org/10.1016/j.ssaho.2024.100864

[16] M, Iyyappan, Ahmad, S., Jha, S., Alam, A., Yaseen, M., & Abdeljaber, H. A. M. (2022). A Novel AI-Based Stock Market Prediction Using Machine Learning Algorithm. *Scientific Programming*, *2022*, 1-11. https://doi.org/10.1155/2022/4808088

[17] Mutinda, J. K., & Langat, A. K. (2024). Stock price prediction using combined GARCH-AI models. *Scientific African*, *26*, e02374. https://doi.org/10.1016/j.sciaf.2024.e02374

[18] Nuseir, M. T., Akour, I., Alshurideh, M. T., Al Kurdi, B., Alzoubi, H. M., & AlHamad, A. Q. M. (2024). Stock Market Price Prediction Using Machine Learning Techniques. In H. M. Alzoubi, M. T. Alshurideh, & T. M. Ghazal (Éds.), *Cyber Security Impact on Digitalization and Business Intelligence* (Vol. 117, p. 323-334). Springer International Publishing. https://doi.org/10.1007/978-3-031-31801-6_20

[19] Pulok Sarker, Adnan Sayed, Abu Bakar Siddique, Avijit Saha Apu, Syeda Anika Tasnim, & Mahmud, R. (2024). A Comparative Review on Stock Market Prediction Using Artificial Intelligence. *Malaysian Journal of Science and Advanced Technology*, 383-404. https://doi.org/10.56532/mjsat.v4i4.316

[20] Rhoda Adura Adeleye, Tula Sunday Tubokirifuruar, Binaebi Gloria Bello, Ndubuisi Leonard Ndubuisi, Onyeka Franca Asuzu, & Oluwaseyi Rita Owolabi. (2024). MACHINE LEARNING FOR STOCK MARKET FORECASTING : A REVIEW OF MODELS AND ACCURACY. *Finance & Accounting Research Journal*, *6*(2), 112-124. https://doi.org/10.51594/farj.v6i2.783

[21] Rodríguez-Ibánez, M., Casánez-Ventura, A., Castejón-Mateos, F., & Cuenca-Jiménez, P.-M. (2023). A review on sentiment analysis from social media platforms. *Expert Systems with Applications*, *223*, 119862. https://doi.org/10.1016/j.eswa.2023.119862

[22] Shaban, W. M., Ashraf, E., & Slama, A. E. (2024). SMP-DL : A novel stock market prediction approach based on deep learning for effective trend forecasting. *Neural Computing and Applications*, *36*(4), 1849-1873. https://doi.org/10.1007/s00521-023-09179-4

[23] Shamim, M., & Siddiqui, A. (2024). Modelling Stock Market Volatility using Asymmetric GARCH Models : Evidence from BRICS stock markets. *Global Business and Economics Review*, *30*(1), 10059551. https://doi.org/10.1504/GBER.2024.10059551

[24] Tamiri, M. A. . ., Cherkaoui , M. ., Idrissi , N. ., Redouane , K. ., Fikri , Y. ., & Nassiri , A. . (2025). The Effectiveness of Internal Control Tested by The Characteristics of The Company and Its Manager: A Study Carried Out on A Sample of Moroccan Companies Listed on The Casablanca Stock Exchange. *International Journal of Accounting and Economics Studies, 12(7)*, 19-26. https://doi.org/10.14419/t5xv8662

[25] Tetlock, P. C. (2007). Giving Content to Investor Sentiment : The Role of Media in the Stock Market. *The Journal of Finance*, *62*(3), 1139-1168. https://doi.org/10.1111/j.1540-6261.2007.01232.x

[26] Zhang, G., Wan, C., Xue, S., & Xie, L. (2023). A global-local hybrid strategy with adaptive space reduction search method for structural health monitoring. *Applied Mathematical Modelling*, *121*, 231-251. https://doi.org/10.1016/j.apm.2023.04.025