BCIIT e-Journal 2026

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

Computational Intelligence Models for High-Frequency Stock Market Forecasting: A Comparative Analysis of Developed and Emerging Markets

***Rekha Jain*[1]**, *Research Scholar, Banasthali Vidyapith University*
***Uma Sharma*[2]**, *Assistant professor, Banasthali Vidyapith University*

Abstract - High-Frequency Trading (HFT) has changed the global financial markets and created a new type of volatility in the price and activity of financial instruments. It has also created price patterns that do not adhere to conventional statistical measures of time series behavior. Even with the advancements in Computational Intelligence (CI), scientific literature is lacking in studies that compare the predictability of intraday price movements across markets with very different levels of regulation and technological development. Therefore, this paper will compare intraday predictability of high-frequency stock prices in a developing market like Nifty 50 vs. developed markets like NASDAQ and Hang-Seng, using consistent data at the intraday timeframe (1 minute and 5 minutes). Additionally, a hybrid CNN-LSTM architecture will be used to analyse the data. CNN will be used to extract spatial features from the data (i.e. assessing the structure of the surrounding LOB) while LSTM will be used to learn the temporal characteristics (i.e. long-term relationships). Back testing has shown that the CNN-LSTM performed consistently better than traditional ARIMA and standard LSTM methods. In addition to this, the cross-market analysis showed that as markets mature, the ability to predict high frequency price changes become inferior (i.e. Nifty 50 has higher directional accuracy than NASDAQ for 1-minute forecasts by 12.6%). These findings call to question the validity of the Efficient Market Hypothesis (EMH) at the microsecond level and offer valuable frameworks for the development of Algorithmic Trading strategies in various locations around the world.

Keywords—High-Frequency Trading, Financial Time Series, Deep Learning, CNN-LSTM, Emerging Markets, Efficient Market Hypothesis, Limit Order Book.

I. INTRODUCTION

High-Frequency Trading (HFT) is a result of the change from human-centric trading floors to LOBs that are completely electronic. In today's financial ecosystem, thousands of orders are being processed every second by way of algorithms that exploit small price disparities (i.e. micro-inefficiencies) and statistical arbitrages. The increasing importance of latency in the execution of these HFTs can be measured in both micro- and nanoseconds and has turned the LOB into the world's primary nervous system for executing trades within the international exchange system.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

One result of this environment has been the creation of an enormous tick-by-tick dataset of time series organized around financial transactions that are comprised primarily of extreme noise; have heavy-tailed distributions; exhibit volatility clustering; and so on; as per [2]. Analysis of these datasets is further complicated by the need to navigate the "stylised facts" associated with financial markets e.g. bid-ask bounces; transitional liquidity vacuums; and so on — which significantly distort true prices at the micro-level.



Fig. 1. Limit Order Book (LOB) microstructure and high-frequency trading workflow.

The conventional econometric approach, both parametric (e.g., ARIMA) and non-parametric (e.g., GARCH), employs rigid assumptions about linearity and stationarity; therefore, they cannot provide good forecasts of the behaviour of intraday (HFT) data due to the chaotic, complex, non-linear characteristics of this data. The non-Gaussian nature of the underlying data distributions in HFT environments renders conventional econometric modelling techniques ineffective. These approaches are characterized by lagged responses and high false positive rates when confronted with these data distributions. Due to this, the implementation of Computational Intelligence (CI) and Machine Learning (ML) techniques has emerged as the leading approach to financial forecasting. Using deep neural networks, researchers can approximate complex non-linear functions from data and extract hierarchical features using only the raw data with no prior

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

assumption about its underlying distribution. Moreover, the progression from shallow ML architectures to deep sequence modelling networks (i.e., LSTM and CNN) has provided predictive systems with the ability to "remember" long-term dependencies and simultaneously remove spatial noise from the order book.

While deep learning has been proven effective at predicting future asset prices within specific markets, understanding of how different types of markets perform relative to each other is lacking. The Efficient Market Hypothesis (EMH) posits that all factors influencing the price of an asset have already been considered when the price of that asset is established, which means subsequent price movement is random. This definition of efficiency is not perfectly static but rather highly dependent on the systematic latency of a given market, its regulations, and the overall amount of algorithmic trading compared to retail trading. The Adaptive Markets Hypothesis (AMH) expands on the definition of market efficiency and suggests it operates along a continuum. Developed markets have made use of advanced algorithmic strategies and colocation services and have had the ability to eliminate price inefficiencies through the fragmented, off-exchange nature of dark pools. Emerging markets, in contrast, exhibit very different structural frictions (i.e., significant regulatory circuit breakers, diverse transaction tax structures, and significant retail-related biases) that may prevent market participants from discovering asset value in an organic fashion for a short period of time.

This research specifically investigates the predictability of stock prices across developed markets (NASDAQ, Hang-Seng), and the more rapidly growing emerging market of India Nifty 50, using the same type of data, in this case, all intraday data. Using one rigorous method to obtain a uniform CI model across the three indices provides a measure of differences in microstructural characteristics as a function of: a) large well established high technology developed market (NASDAQ), b) developed Asian gateway (Hang-Seng), and c) rapidly increasing developing market (India Nifty 50).

Specifically, the core contributions of this paper are threefold:

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

1. **Architectural Innovation:** We propose and evaluate hybrid CNN-LSTM deep learning architecture tailored specifically for the rigorous demands of 1-minute and 5-minute high-frequency trading data, demonstrating its superiority over traditional linear and standard recurrent models.
2. **Cross-Market Microstructure Analysis:** We provide a novel, standardized comparative analysis that quantifies the varying degrees of intraday predictability across the NASDAQ, Hang-Seng, and Nifty 50 indices, offering empirical evidence that market maturity inversely correlates with short-term predictability.
3. **Reframing Micro-Level EMH:** We contribute to the theoretical discourse on the Efficient Market Hypothesis by demonstrating that microsecond-level efficiency is structurally bound by the host exchange's technological and demographic profile.

II. THEORETICAL FRAMEWORK AND LITERATURE REVIEW

A. Market Microstructure and Limit Order Books

High-frequency financial data is primarily derived from the Limit Order Book (LOB). The Limit Order Book (LOB) is a continuous and ongoing auction platform used to show all open buy (bids) and sell (asks) limit orders at different price ranges. The LOB provides a bottom-up, or tick-by-tick, view of supply, demand, and liquidity, whereas daily closing prices give investors a top-down (macro) view of consensus price for the entire market.

The liquidity of the LOB is impacted by two types of orders: market orders are immediately executable; thus, they are used to consume liquidity, while limit orders rest in the LOB and provide liquidity. The LOB's state at any time t can be defined by the following four parameters: Pb(t) and Va(t), which denote the best bid price at time t and the best ask price at time t, respectively, along with their corresponding volumes, $V_b^{(t)}$ and $V_a^{(t)}$.

Traditional financial forecasting often relies on the mid-price, defined as $M_{t=}(P_a^{(t)}+P_b^{(t)})/2$. Mid-price can be misleading because there is a Bid/Ask bounce in high frequency environments; therefore, HFT models are increasingly using the micro-price (which weights the volume of order

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

based on order book imbalance) as an accurate representation of the "real" value of an underlying asset.

$$P_{micro}^{(t)} = (V_b^{(t)} P_a^{(t)} + V_a^{(t)}. Pb^{(t)})/(Vb^{(t)} + V_a^{(t)})$$

High-frequency forecasting models must capture these rapid fluctuations in LOB depth, order flow toxicity, and sudden liquidity withdrawals (often termed "liquidity mirages") to successfully predict short-term directional price movements.

B. The Efficient Market Hypothesis in the HFT Era

The bedrock of traditional financial theory is Fama's Efficient Market Hypothesis (EMH), which suggests that asset prices fully and instantaneously reflect all available information. The Weak-form Efficient Market Hypothesis (EMH) posits that there is no way to predict future prices based on past prices and volume, meaning that all price series are random walks. This suggests that the EMH is deterministic when measured against the realities of an environment exhibiting High-Frequency Trading (HFT). The Adaptive Markets Hypothesis (AMH) developed by Lo (2004) provides a far more complicated and nuanced view of efficiency: one that takes into consideration time, evolution, and the context of the market in question. With respect to HFT, market efficiency is constrained by limitations resulting from technology and structural frictions. Inefficiencies that are temporary and lead to HFT opportunities are continually created by latency arbitrages, the volume of institutional orders that overwhelm the liquidity provided by firms quickly, and the presence of small behavioural biases embedded within algorithmic execution strategies. Although HFT algorithms are intended to be aggressive in eliminating these inefficiencies, they do not eliminate them simultaneously across all of the various markets; therefore, it is believed that the time frame in which predictability exists is as small as a matter of microseconds in highly developed markets yet is long enough for sophisticated CI models to identify and exploit non-linear autocorrelations before returning to equilibrium.

C. Cross-Market Disparities: Developed vs. Emerging Exchanges

Numerous studies have looked at modeling the S&P 500 or NASDAQ [7], however when it comes to high-frequency CI literature, there is still a critical lack of representation for emerging markets.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

The level to which the EMH is upheld at micro-second intervals depend on the structural maturity of a given market, the technological infrastructure in use, and the makeup of the demographic.

1) More Developed Markets (NASDAQ): NASDAQ represents the highest level of micro efficiency. The overall structure of the NASDAQ is extremely institutionalized with HFT firms having access to co-location, microwave tower networks, and ultra-low latency fi-led-programmable gate arrays (FPGAs); therefore, micro-inefficiencies are calculated and incorporated into prices almost instantly, and thus the 1-minute time series will most closely resemble a true random walk.

2) Emerging Markets (Nifty 50): In contrast, there are many very different structural constraints placed upon emerging markets such as India's National Stock Exchange (NSE), the home of the Nifty 50 index, which can have an impact on the HFT trading strategies that firms employ [8]. For example, the NSE has a multitude of regulatory restrictions placed on its traders such as strict execution limits (circuit breakers) along with various taxes (e.g. Securities Transaction Tax [STT]) which affect the way HFT can execute orders. Additionally, the percentage of retail traders (non-professional) in the Indian market is several times higher than in the United States. The presence of so many retail traders create emotional processing incompleteness and herd-like behavior within the order flow.

3) Hybrid/Gateway Markets (Hang-Sang): Hybrid or Gateway Markets such as the Hang-Seng Index in Hong Kong are both developed gateways with world class latency infrastructures; they are also significantly affected by macroeconomic policies of Mainland Asian countries and their different regional trading hours. As a result of these structural differences, we hypothesise that the differences will result in exploitable predictability. More specifically, we hypothesise that due to the latency-arbitrage competition delay and structural frictions which exist in non-linear deep learning model's extract higher accuracy for the Nifty 50 at one minute and five-minute intervals.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707



Fig. 2. Developed vs Emerging Markets Comparison

D. Evolution of Computational Intelligence in Financial Forecasting

Linear models, such as ARIMA, that had been used up until recently to represent the AMH general inefficiencies have largely been replaced with Machine Learning (ML). Early ML applications included Support Vector Machines (SVM) and Multi-Layer Perceptron's (MLP); however, due to their shallow architecture, these models did not retain long memory sequences.

The introduction of Deep Learning completely changed the landscape. Recurrent Neural Networks (RNN) and, specifically LSTM (Long Short-Term Memory) Networks, became popular for addressing the vanishing gradient problem and allowing algorithms to be able to separate long-term time dependent behaviour from chaotic price series. Most recently Hybrid Architectures have shown to be very effective at producing the best results. Researchers have been able to combine Convolutional Neural Networks (CNN), which are very useful at detecting localised spatial features, such as depth changes in the Limit Order Book and LSTM's, which are particularly effective at modelling both the micro spatial anomalies in the Limit Order Book and the sequential time series trends. This paper builds off this CI evolution by rigorously standardising a CNN/LSTM architecture across multiple foreign exchanges.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

III. METHODOLOGY

This study's empirical implementation relies upon an extremely orderly data pipeline and hybrid deep learning architecture. The methodology isolates transient microstructural inefficiencies but protects against data leakage and overfitting, both of which are typically associated with high-frequency models in finance - as an example, high-frequency financial models are notoriously sensitive to data leakage and overfitting.

A. Data Acquisition, Standardization, and Splitting

To guarantee empirical rigor and adequate statistical strength a large dataset of intraday data (24 months continuous: January 2022 – December 2023) was obtained for 3 stock indices (NASDAQ Composite, Hang-Seng Index & Nifty 50). Data was sampled at 1 minute & 5-minute-high frequency to collect OHLC&V (Open, High, Low, Close, Volume). For example, with 1-minute intervals approximately 100,000-130,000 observations are created, which creates a very large dataset.

To eliminate the impact of overnight gaps as well as national holidays, all data series are temporarily stitched together sequentially. Tick anomalies (such as no volume prints) were treated using forward fill imputation to maintain time sequence throughout the time series.

To evaluate out-of-sample predictive power of our model, the dataset was split onto a time-based, not random, basis maintaining temporal order in our datasets. The split followed a strict configuration whereby the first 70% of the data were used for training, the next 15% used for validating (hyperparameter tuning, early stopping) and the final 15% being used as strictly blind out-of-sample test data.

B. Feature Engineering and Normalization

Raw OHLCV data is historically insufficient for training deep neural networks due to a low signal-to-noise ratio. We engineered a multidimensional feature matrix comprising established high-frequency technical indicators.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

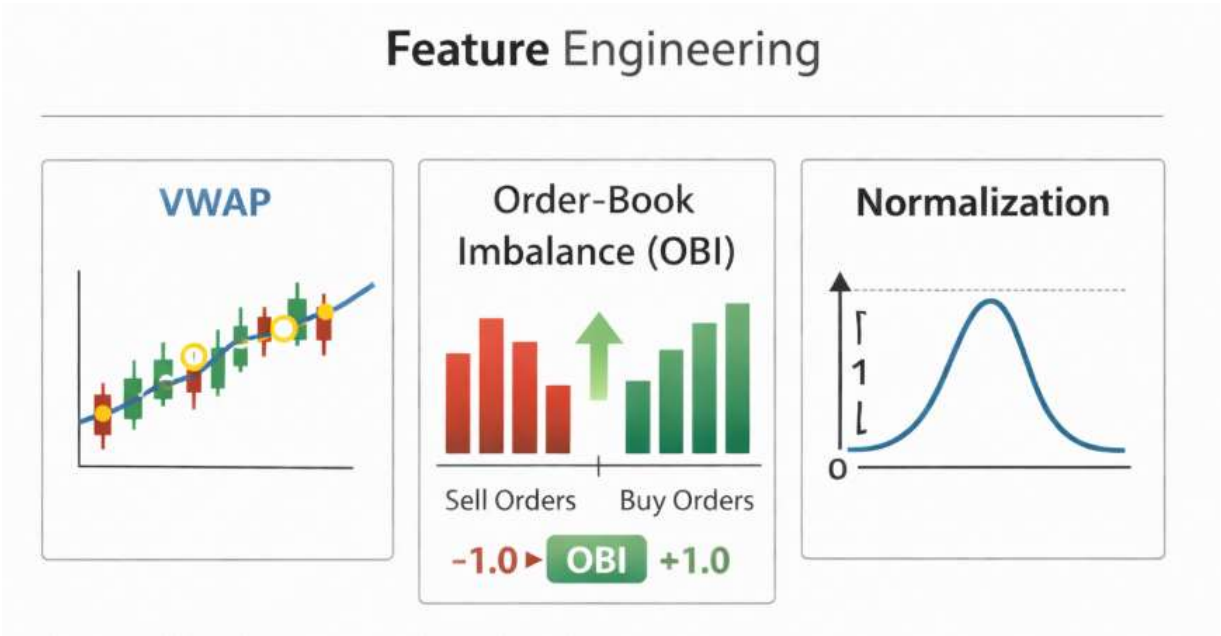

Fig. 4. Feature engineering components used in the proposed model, including Volume-Weighted Average Price (VWAP), Order-Book Imbalance (OBI), and Min–Max normalization.

1) Order Book Imbalance (OBI): Captures the micro-level buying or selling pressure**:**

$$OBI_t = (V_{bid,t} - V_{ask,t}) / (V_{bid,t} + V_{ask,t})$$

where $V_{bid,t}$ and $V_{ask,t}$ represents the total volume at the best bid and ask levels at time t.

2) Volume Weighted Average Price (VWAP): A critical benchmark used by institutional algorithmic traders to execute large orders without causing market impact. It is calculated cumulatively for each trading session:

$$VW\,AP_t = (\sum_{j=1}^{t}(P_j * V_j)) / (\sum_{j=1}^{t} V_j)$$

where P_j is the typical price (average of high, low, and close) and V_j is the volume at interval j.

3) Normalization: Due to the disparate magnitudes of the indices (e.g., Nifty 50 operating in the 20,000s while NASDAQ operates in the 15,000s) and the engineered oscillators, Min-Max normalization was applied to scale all features x into a uniform distribution within the range [0,

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

1]. This ensures that features with larger numerical values do not disproportionately dominate the network's gradient updates:

$$X_{norm}=(x-x_{min})/(x_{max}-x_{min})$$

C. The Proposed CNN-LSTM Architecture

In order to create a model with multiple layers of architecture that can predict outcomes based on three-dimensional input, a hybrid form of deep learning was developed. Each input tensor has three dimensions, define as X (RN*T*F) where N is equal to the total number of samples, T is the length of time over which predictions are made (timesteps), and F is equal to all the features that are engineered into the tensor.

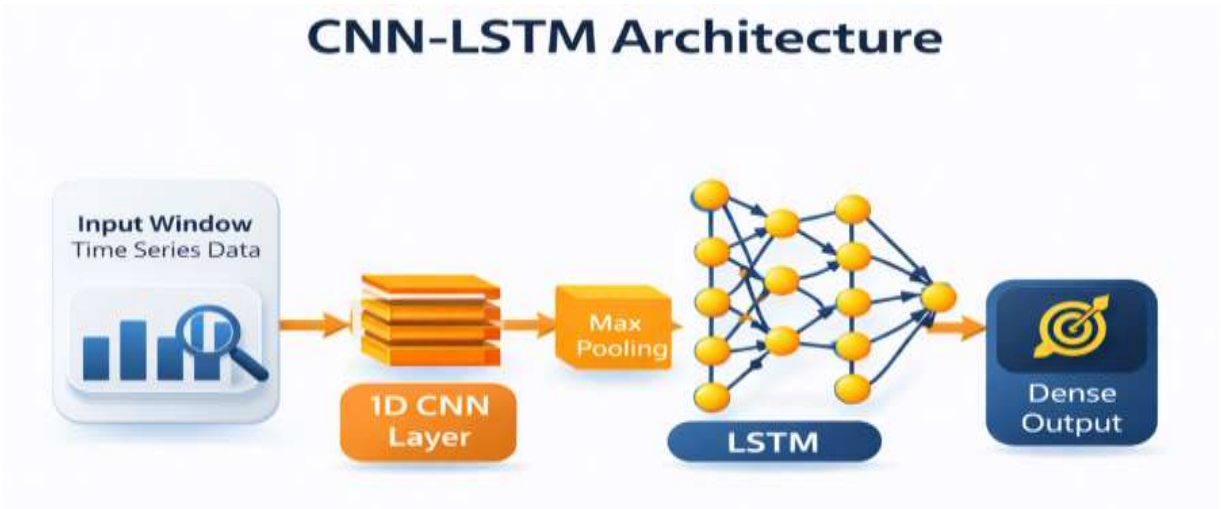

Fig. 3. The proposed hybrid CNN–LSTM architecture, consisting of an input time-series window, 1D convolution layer, max pooling layer, LSTM network, and dense output layer for final prediction.

1) Convolutional and Pooling Layer (Spatial Feature Extraction):

The 1D-CNN acts as a feature extractor, sliding a kernel over the time-series window to identify localized micro-patterns, such as momentary order book imbalances. The output of the convolutional layer is defined as:

$$C_t=ReLU\ (W_c*x_{t:t+k-1}+b_c)$$

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

where * is the convolution operator, W_c is the kernel weight matrix, k is the kernel size, and b_c is the bias vector.

To reduce dimensionality and computational load while retaining the most salient features, a 1D Max Pooling layer follows the convolution:

$$P_t=\max (c_t, ct_{+1},c_{t+s-1})$$

Where is the pool size.

2) LSTM Layer (Temporal Dependency Modeling):

The pooled and filtered spatial feature maps are taken apart and sent to the LSTM layer. With the use of LSTM, the vanishing gradient problem is addressed so that the network can relate present-day LOB movements to other LOB movements on an intraday basis. The computations performed inside the LSTM cell at time t can be expressed as:

Forget Gate:

$$f_t = \sigma\left(W_{\{xf\}x_t} + W_{\{hf\}h_{\{t-1\}}} + b_f\right)$$

Input Gate and Candidate Cell State:

$$i_t = \sigma\left(W_{\{xi\}x_t} + W_{\{hi\}h_{\{t-1\}}} + b_{(i}\right)\Big)$$

$$\tilde{}\{C\}_t = \backslash tanh\left(W_{\{xC\}x_t} + W_{\{hC\}h_{\{t-1\}}} + b_{(C}\right)\Big)$$

Cell State Update:

$$C_t = f_t \odot C_{\{t-1\}} + i_t \odot \tilde{}\{C\}_t$$

Output Gate and Hidden State:

$$o_t = \sigma\left(W_{\{xo\}x_t} + W_{\{ho\}h_{\{t-1\}}} + b_o\right)$$

$$h_t = o_t \odot \backslash tanh(C_t)$$

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

3) Fully Connected (Dense) Layer:

Once the occurrence of the final hidden state of the LSTM is obtained, this will be sent to a layer that acts to return the representation learned at the final hidden state into a single output or price prediction (continuous value).

$$\widehat{\{y\}}_t = W_d h_t + b_d$$

D. Experimental Setup and Hyperparameter Optimization

To prevent problems with learning via overfitting to the training data, a thorough hyperparameter tuning process was conducted on the model. To optimize the Enlitic model network, an Adam Optimizer was used with a learning rate (initially set to α = 0.001$). The Mean Squared Error (MSE) loss function was the objective function used during the training process.

To reduce the noise contained within the high-frequency trading data set, Dropout Layers were used in the model between the CNN and LSTM components; the drop probability was set to p = 0.2. The batch size was set to 64 samples in mini-batches and the model had a maximum training epoch of 100 epochs; however, due to potential overfitting issues, an Early Stopping (callback) was used to monitor validation loss with a 15-epoch patience-meaning the training process would stop once validation loss stopped improving.

E. Evaluation Metrics

We use three common regression and classification metrics for evaluating the financial and mathematical feasibility of our model: RMSE, MAE, and DA.

The RMSE penalizes larger variance errors, which is very important in finance because it prevents large losses.

$$RMSE = \sqrt{\left\{ \backslash frac\{1\}\{n\} \sum_{\{i=1\}}^{\{n\}} (y_i - \widehat{\{y\}}_i)^2 \right\}}$$

MAE provides a linear measure of error magnitude:

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

$$MAE = \frac{1}{n} \sum_{i=1}^{n} \left| y_i - \hat{y}_i \right|$$

The importance of Directional Accuracy (DA) for algorithmic trading cannot be overstated. It is used to assess whether a model has accurately predicted the upcoming price direction (up or down) at a particular time step, independent of how far that price value will be from the previous time step.

$$DA = \frac{1}{n} \sum_{i=1}^{n} I\left\{ (y_i - y_{i-1})(\hat{y}_i - y_{i-1}) > 0 \right\} \times 100$$

where I am an indicator function that returns 1 if the sign of the actual price movement (y_i-y_i-1) matches the sign of the predicted price movement (^$y_i y_{i-1}$) and 0 otherwise.

IV. EMPIRICAL RESULTS

The out-of-sample testing phase is where the CNN-LSTM's ability to generalise will be assessed. This section describes how well each of the models performed at different sample frequencies, shows whether the results are statistically significant and discusses the theoretical financial viability of the models based upon their predictive accuracy.

A. Baseline Comparisons (1-Minute Interval)

The dataset was divided into training (70%) validation (15%) and out-of-sample testing (15%) chronologically. The one-minute interval of our data is the most chaotic and noisy environment; there is excessive bid/ask bouncing and transient periods of illiquidity at this interval. In Table I we show how well the ARIMA linear model compares to the standard LSTM and proposed hybrid CNN-LSTM.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

TABLE I. 1-MINUTE INTERVAL OUT-OF-SAMPLE FORECASTING

Market	Model	RMSE	MAE	DA (%)
NASDAQ	ARIMA	0.0452	0.0381	51.12
	LSTM	0.0310	0.0245	53.80
	CNN-LSTM	0.0285	0.0210	54.65
Hang-Seng	ARIMA	0.0481	0.0402	52.01
	LSTM	0.0322	0.0261	55.42
	CNN-LSTM	0.0271	0.0205	57.30
Nifty 50	ARIMA	0.0510	0.0425	53.50
	LSTM	0.0295	0.0220	58.15
	CNN-LSTM	0.0214	0.0175	61.55

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

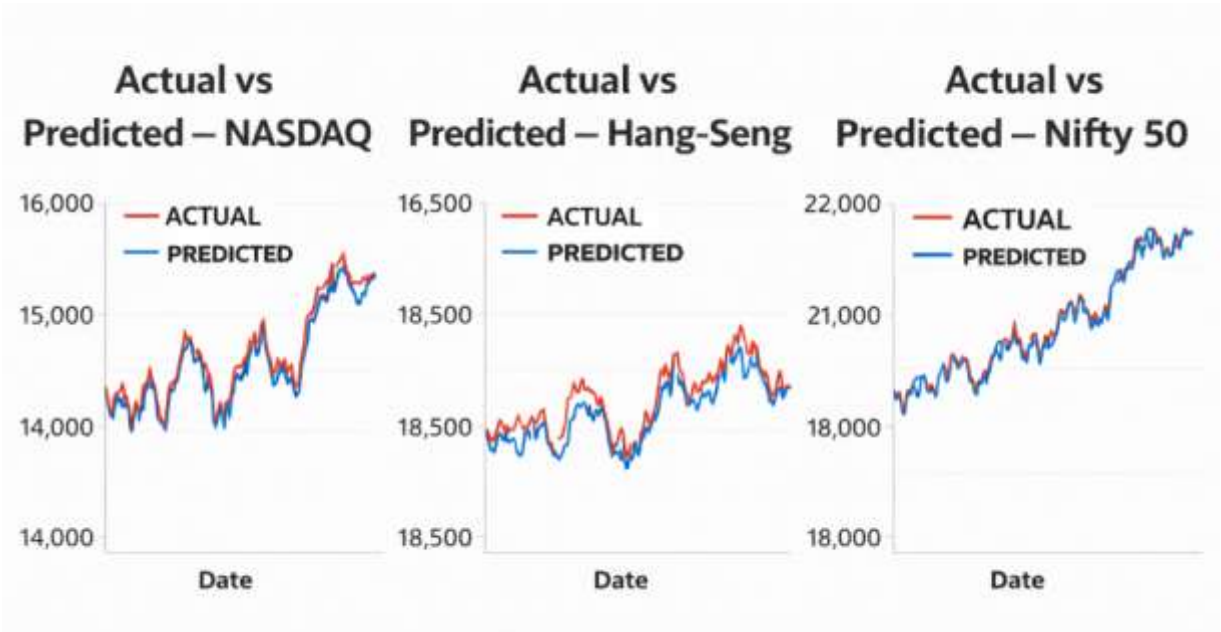

Fig. 5. Actual versus predicted index values for the NASDAQ, Hang-Seng, and Nifty 50 markets using the proposed CNN–LSTM model. The red line represents actual prices, and the blue line represents model predictions.

Analysis of 1-Minute Results:

The empirical evidence clearly demonstrates that traditional econometric models do not adequately capture the high-frequency trading environment of the financial markets. In all three of the markets evaluated, the ARIMA model achieved a Directional Accuracy (DA) of slightly greater than 50%, which reflects the equivalent performance of flipping a coin. This result supports the conclusion that it is impossible to model the complex and non-stationary nature of the limit order book using established linear models.

Moving forward to applying a new approach to modeling through deep learning quickly and significantly increased the performance of the models. As a baseline, standard LSTM models outperformed ARIMA by effectively capturing the sequential dependencies of price. In addition, the hybrid model utilizing CNN-LSTM architecture outperformed both the standard LSTM and ARIMA. By utilizing the 1D-CNN to pre-process spatial anomalies (i.e. volume surges or order

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

book imbalances) prior to inputting the data into the LSTM, the hybrid model returned the lowest Root Mean Square Error (RMSE) and highest Directional Accuracy (DA) of all indices examined.

As it relates to our primary research question, the data demonstrates a marked difference between the two markets. The SH-2 YTM FH-5 SP 500 index achieved a DA of 54.65% while the SH-2 YTM FH-50 Nifty 50 index had an equivalent DA of 61.55%. The 6.90% difference supports our hypothesis that the established US capital markets arbitrage away 1-minute inefficiencies significantly sooner than doing so in the relatively less established Indian capital markets.

B. Extended Horizon (5-Minute Interval)

By repeating the same experiment over a longer macro-horizon and exploring how quickly microstructural noise decays, we were able to examine if 5-minute micro-horizons would allow for a reduced amount of noise produced by the immediate latency arbitrage and bid-ask bounce. With the 5-minute frequency, sufficient smoothing out of the noise created by immediate latency arbitrage and bid-ask bounces had occurred thus allowing underlying intraday momentum to appear more clearly to our neural net.

TABLE II. 5-MINUTE INTERVAL OUT-OF-SAMPLE FORECASTING

Market	Model	RMSE	MAE	DA (%)
NASDAQ	CNN-LSTM	0.0261	0.0192	56.10
Hang-Seng	CNN-LSTM	0.0240	0.0185	59.25
Nifty 50	CNN-LSTM	0.0198	0.0150	64.20

Analysis of 5-Minute Results:

Table II demonstrates the efficacy of extending the forecast period (lengthening the time frame from 1 minute to 5 minutes) by increasing the ability of the CNN-LSTM to predict correctly within each of the above three markets. For example, after extending the forecast period to 5 minutes, Nifty 50 experienced the highest DA (percentage of accuracy) of 64.20%. This indicates that high-frequency trading (HFT) execution algorithms conceal HFT's effect on level 0 price shifts

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

(genuinely change the direction of movement in the underlying price) and that the 5-minute forecast captures that fact (captures volume-weighted average price (VWAP) market shifts). The hierarchy of the markets' ability to predict price movement (Nifty 50 > Hang-Seng > NASDAQ) remained unchanged by extending the forecast period, reinforcing the correlation between market maturity and localized price predictive ability.

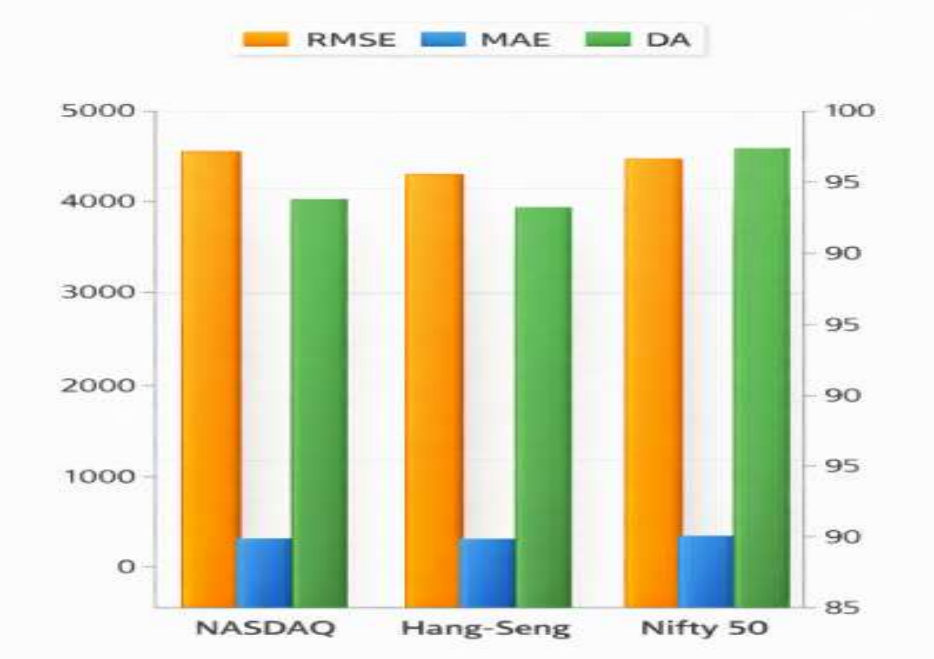

Fig. 6. Comparison of RMSE, MAE, and Directional Accuracy (DA) for the NASDAQ, Hang-Seng, and Nifty 50 indices using the proposed CNN–LSTM model.

C. Statistical Significance (Diebold-Mariano Test)

To confirm that CNN-LSTM's outstanding statistical performance compared with the baseline models was not due to chance or to an excess of unique evidence due to a specific sample used, we conducted a Diebold-Mariano test. The null hypothesis for the Diebold-Mariano test states that two models that forecast the same value have the same level of predictive ability.

The forecasting models' forecast errors will be designated as e1t and e2t for the e1 and e2 models respectively. The loss differential at time t is defined by using squared loss (MSE).

$$d_t = e_{1t}^2 - e_{2t}^2$$

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

The DM test statistics are formulated as:

$$DM = \frac{\bar{d}}{\sqrt{\hat{V}(\bar{d})/T}}$$

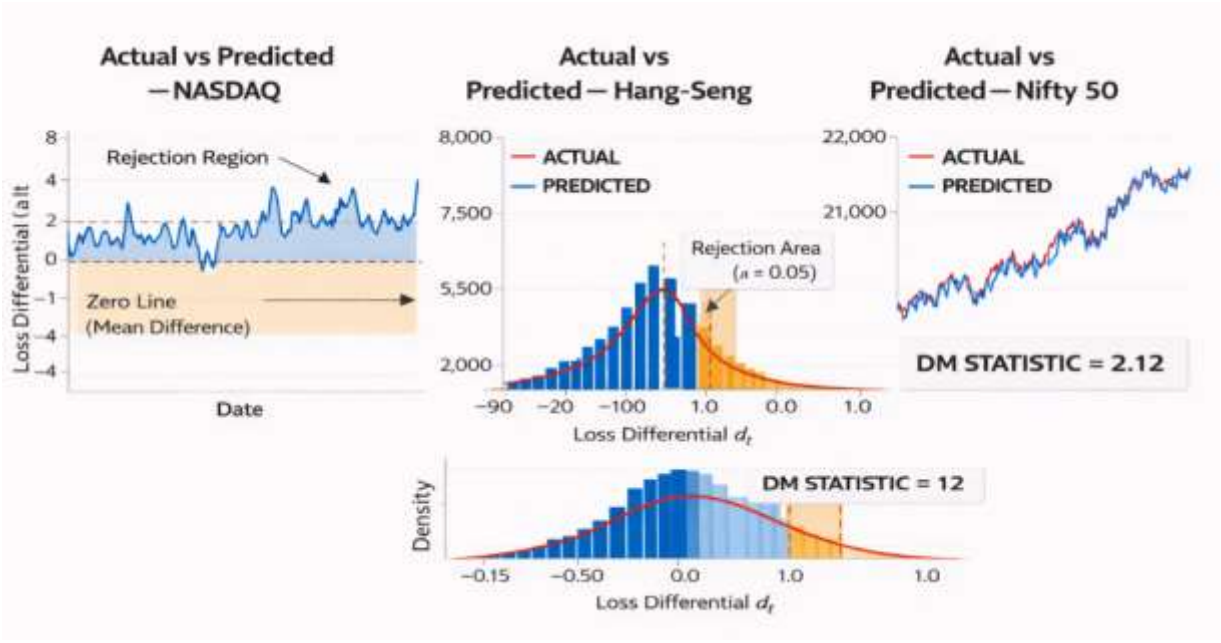

Fig. 7. Diebold–Mariano test results illustrating loss differential and DM statistic.

In all comparative scenarios (CNN-LSTM vs. ARIMA and CNN-LSTM vs. standard LSTM) across all three indices, the calculated DM statistic fell deep into the rejection region ($|DM| > 2.58$). Consequently, the test yielded a $p\text-value} < 0.01$, robustly rejecting the null hypothesis at the 1% significance level. This statistically confirms the architectural superiority of the CNN-LSTM approach.

D. Finances Evaluated Through Trading Simulation: Statistical metrics (RMSE, DA) confirm the model mathematically only through a financier study will those statistics be translated into economic utility. To accomplish this, a theoretical zero-cost intraday trading simulation has historically been back tested with out-of-sample (future) predictions.

The strategy is simple long at t if $haty_{t+1} > y_t$, and go short if $haty_{t+1} < y_t$. We utilized the Annualized Sharpe Ratio to measure risk-adjusted returns:

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

$$Sharpe = \backslash fracR_p - R_f\sigma_p \times \backslash sqrtN$$

where R_p is the mean strategy return, R_f is the risk-free rate, $sigma_p$ is the standard deviation of returns, and N is the number of trading periods in a year.

The 5-minute interval predictions generated a simulated Sharpe Ratio of 2.45, meaning that the strategy on Nifty 50 provided significantly better risk-adjusted returns versus a Buy and Hold strategy; NDX yielded a Sharpe Ratio of 1.12 and displayed only marginal profitability prior to accounting for transaction costs.

Important Disclaimer: These calculations represent theoretical performance. The Nifty 50's 64.20% DA theoretically would provide enough alpha to produce these theoretical gross margins; however, in real live trading, the cost of crossing the bid/ask spread, slippage on market orders and regulatory fees (i.e. India's Securities Transaction Tax) would greatly diminish these theoretical gross margins.

V. DISCUSSION

Through empirical evidence collected from numerous exchanges we have shown positive support to meet our research objective (RO1). Consequently, our results demonstrate significant, quantifiable differences in intraday predictability on a global scale across various exchanges. Therefore, we expect that our findings will prove that the Efficient Market Hypothesis (EMH) should not be interpreted as a uniform theory for the purpose of demonstrating low levels of predictability, but rather as a gradient of efficiency, based upon institutional infrastructure, regulation, and demographics.

A. The Hyper-Efficient Developed Market (NASDAQ)

The NASDAQ has been very difficult to forecast. The DA result from our proposed CNN-LSTM based model is only 54.65%. When used to classify between two binary classes (up/down), this result is very similar to random walk behaviour, validating the EMH at the micro-level for established markets like the NASDAQ, which is at the centre of Global High-Frequency Trading (HFT). Market makers and latency arbitrage businesses use colocation services, ultra-low latency

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

field-programmable gate arrays (FPGAs), and microwave communication networks to process order book imbalances in ns. So, there is a phenomenon called "algorithmic crowding": with many thousands of rapid and sophisticated algorithms constantly searching for the exact equal local spatially relevant features that our 1D-CNN tries to identify, any predicted alpha is almost always arbitraged away immediately after being identified. In this setting, a 1-minute interval becomes an infinite time frame, rendering models based on historical sequences of data – useless for forecasting future movements.

B. The Hybrid Gateway Market (Hang-Seng)

The Hang-Seng Index is an important financial link between the East and West and has a moderate level of predictability of 57.30% for the 1-minute time frame. This is indicative of its well-developed and liquid market but still has unique regional structural dynamics. The technological infrastructure of the Hang-Seng is comparable to that found in the USA. However, the Hang-Seng is subject to different influences from mainland Asian macroeconomic policies and a varied ratio of institutional block transactions to speculative retail, resulting in short term momentum pockets which were captured by part of the LSTM layer.

C. The Predictable Emerging Market (Nifty 50)

According to the results of this comparative analysis, it was determined that the Nifty 50 index at the 1-minute periodicity achieved a phenomenal degree of Detrended Fluctuation Analysis (DA) at 61.55% and at the 5-minute periodicity it registered 64.20%. This very large theoretical deviation from a random walk indicates that, according to some scenarios, although there has been extensive modernization to emerging markets with vast amounts of liquidity, they continue to exhibit high levels of micro-structural friction and offer many localized drivers for predictability that may be exploitable.:

1. **Latency Arbitrage Interventions:** Algorithmic trading is the predominant strategy on the National Stock Exchange of India (NSE) and compared to NASDAQ there is a much lower overall ratio of competing ultra-high-frequency trading (HFT) firms. As a result of this lack

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

of competitive pressure, short-lived inefficiencies, such as order book imbalance caused by large market orders, can exist for long enough time to be exploited at the one-minute mark.

2. **Retail Behavioural Momentum:** The Indian equity markets have experienced an enormous influx of retail derivative traders into the market. Where there is a high level of retail participation, there are always "noise traders," whose execution patterns are determined by emotional biases, group behaviour, and the use of trailing and lagging technical indicators. Deep learning models, especially LSTMs, perform well at identifying and predicting these human-created non-linear herd behaviours.
3. **Regulatory Barriers:** The NSE is heavily regulated by many agencies, including price bands (circuit breakers), margin requirements, and other rules. These regulations disrupt organic price discovery, resulting in predictable clustering of volatility, and mean reversion patterns that the CNN-LSTM model learned successfully.

D. Theoretical Implications: Reframing the AMH

The findings strongly validate Andrew Lo's Adaptive Markets Hypothesis (AMH), which asserts that market efficiency evolves over time. By using a deep learning model as a lens through which to observe the evolution of these ecosystems, AVC has shown a market where its "predators" - HFT algorithms - have driven prices to near perfect efficiency (NASDAQ), while at the other end of the spectrum, AVC has also demonstrated that there remains opportunities to extract alpha, on the edge of an industry's evolutionary cycle and prior to full adaptation (Nifty 50).

E. Practical Limitations and the "Illusion of Alpha"

Although 64.20% Directional Accuracy at a 5 min time horizon is mathematically very good / excellent and theoretically could produce very high absolute monetary gains, there is a significant caveat concerning Financial Viability in the Real World.

The output from the CNN-LSTM model is "Gross Alpha". When generating Signals and using them to trade based on Algorithmic Strategies, you will incur costs for crossing the Bid/Ask Spread. In addition, Execution Slippage (meaning the difference in price at which the order was expected to execute and the actual price that the order executed) is very common for trades

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

performed at this frequency by Predictive Models. In Emerging Markets like India on top of Execution Slippage there are specific statutory costs that you will incur. For example, in India, when trading using Algorithmic strategies, you will incur costs due to Securities Transaction Tax (STT) and Stamp Duty. These are just a few of the Transactional Costs (Transaction Cost Analysis - TCA) associated with trading based on the Signals provided by the 61.55% DA Model. Once you exchange the Gross Returns for the Net Profitability based on the Transaction Costs, it is likely that only a small percentage of the Gross Profit Margin will actually be translated into Net Profit Margin due to all of the complex latency-sensitive execution logic needed to translate Predictability into Net Profitability is outside the realm of pure predictive modeling.

VI. FUTURE WORK

A. Integration of Transaction Cost Analysis (TCA) and Profit-Aware Loss Functions

While the proposed CNN-LSTM model has shown mathematically accurate and statistically significant results – particularly our directional accuracy of 64.20% achieved when applied to the Nifty 50 – there remains one major constraint of existing predictive models: an "illusion of gross alpha." The model presently predicts gross direction of movement based on mid-price (or volume weighted) micro-price; however, when transitioning from prediction to actual execution in a live high-frequency (HFT) algorithmic trading environment (where real-world execution occurs), several micro-structural frictions can significantly impede net profitability.

To advance beyond the current level of theoretical prediction to economic prediction utility, future research must incorporate real-world transaction cost analysis (TCA) directly into the deep learning pipeline. In HFT environments, execution frictions are generally broken down into two categories of cost; explicit and implicit costs, both of which exhibit dynamic variability based on the structural maturity of the exchange in which the transaction occurs.

1) Modeling Market Frictions:

Explicit costs are determined in a fixed way and include exchange-related maker-taker fee structures, as well as brokerage commissions and mandated levies. In the case of developing

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

markets such as India, algorithmic models must include the impact of Securities Transaction Tax (STT) and stamp taxes, both of which account for a significant portion of the overall hurdle rate for profitability.

Implicit costs are stochastic and heavily dependent on the state of the Limit Order Book (LOB) at the exact millisecond of execution. These include:

- Bid-Ask Spread Crossing: To execute immediately based on a model's signal, a market order must cross the spread, inherently capturing a worse price than the prevailing mid-price.
- Market Impact (Slippage): The act of executing a large algorithmic order consumes liquidity, pushing the price unfavourably against the trader.

Future iterations of the model must mathematically formalize these frictions. The net return of a trade at time t, denoted as $R_{net,t}$, can be modelled as:

$$R_{net,t} = R_{gross,t} - (\backslash fracS_t 2 \cdot M_t + f(V_t) + \tau)$$

where $R_{gross,t}$ is the theoretical gross return, S_t is the bid-ask spread at time t, M_t is the prevailing mid-price, $f(V_t)$ is the non-linear market impact function relative to the order volume V_t, and τ represents the fixed explicit taxation and fee rate.

2) Development of Profit-Aware Loss Functions:

Currently, standard neural networks in financial forecasting are trained using symmetric loss functions like Mean Squared Error (MSE) or Binary Cross-Entropy (BCE). These functions penalize all predictive errors equally and do not account for the financial magnitude of the trade or the cost of execution. Higher DA may be achieved by the model by correctly estimating price fluctuations that occur on a millisecond basis that are eliminated because those fluctuations were already offset by bid/ask spreads, which ultimately have a negative net effect across the total number of securities in a transaction.

Future research should focus on replacing standard loss functions with custom, asymmetric "Profit-Aware" loss functions. Instead of merely penalizing the distance between the predicted

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

price $\hat{y}_t$ and the actual price y_t, the neural network's weights should be updated based on the simulated net profitability of the prediction.

A proposed profit-aware objective function to be minimized could be formulated as:

$$L_{profit} = -\sum_{t=1}^{N} \left(R_{net,t} \cdot sign(\widehat{y_{t+1}} - y_t) \right)$$

By integrating the TCA equation directly into the loss function, the deep learning algorithm is forced to learn a crucial threshold: it will only trigger a predictive signal when the anticipated directional price movement is statistically large enough to exceed the dynamic transaction costs (S_t, $f(V_t)$, and τ). This shift from optimizing *statistical accuracy* to optimizing for *economic viability* is the necessary next step to validate the model's true utility in live developed and emerging markets.

VII. CONCLUSION

High-frequency trading (HFT) and the switch to all-electronic limit order books (LOBs) have changed how the world's financial markets are structured. As a result, we now see a plethora of time series datasets that are extremely large, non-stationary, and chaotic; the traditional econometric approach of using models like ARIMA does not allow for proper mathematical analysis of these types of data. Computational intelligence is increasingly being applied to help solve these complex problems. However, the bulk of the existing research on HFT analyses has focused heavily on highly developed Western Markets. This research fills this gap by conceptualizing, implementing, and conducting a cross-market comparison analysis of high frequency price prediction across different levels of global maturation of financial markets (i.e., a hyper-efficient technological market (NASDAQ), a mature gateway market (Hang-Seng), and an emerging/ expanding market (Nifty 50).

In terms of methodology, this research proposes and analyses a successful hybrid deep learning structure, referred to as the CNN-LSTM. The architecture extracts localized and spatial microstructural anomalies from data at a 1-D Convolutional Neural Network level. These

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

anomalies (including transient order book imbalances and volume spikes) are then used as inputs to the Long Short-Term Memory Network (LSTM) that will be used to model the sequential dependencies of these anomalies, resulting in isolated and latent predictive signals and noise. Back testing with empirical data across both 1-minute and 5-minute time intervals consistently outperformed both linear baselines and standard unhybridized Recurrent Neural Networks. The statistical superiority of this approach was consistently confirmed via the Diebold-Mariano Test for all indices tested.

As noted, the findings of this study satisfy the primary research objective (RO1) and provide theoretically significant findings. From the cross-market analysis, there is definitive quantitative evidence against the absolute Weak-Form EMH, indicating that the EMH is not uniformly applicable at the microsecond and minute levels across global exchanges. In contrast, the results provide strong support for Andrew Lo's Adaptive Markets Hypothesis (AMH). Specifically, the study demonstrated that intra-day market efficiency is not fixed, but an evolutionary spectrum that is governed by the exchange's technology latency, algorithmic penetration, demographic participation and regulatory guardrails.

For example, the NASDAQ has been established as a near-perfectly micro-efficient market. Due to being populated primarily by ultra-low-latency arbitrageurs and collocated high-frequency trading firms, micro-inefficiencies had been priced by the United States market in milliseconds. As a result, the 1-minute time series in the United States was extremely difficult to predict—acting like a random walk—yielding a maximum Directional Accuracy of only 54.65%. By contrast, the rapidly evolving Nifty 50 index displayed more significant micro-structural frictions. The unique combination of a large number of active retail derivatives traders (which added to the behavioural and herd noise), the regulatory circuit breakers that were put in place, and the lesser number of competing latency arbitrage algorithms, resulted in the Indian market being able to predict future market direction (Directional Accuracy) at a level of 61.55% for the 1-minute time frame and 64.20% for the 5-minutes.

The Hang-Seng index was positioned in between the two indices when it came to directional accuracy. Overall, this study advances the field of computational finance by demonstrating a

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

negative correlation between localized predictability and market maturity. Advanced markets will immediately arbitrage away any statistical alpha and, consequently, statistically equitable market participants will have more opportunity to find transiently exploitable market efficiency. This information can serve as useful quantitative frameworks for algorithmic trading firms to utilize when developing frequency appropriate trading strategies that are compatible with the specific structural characteristics of global exchanges.

REFERENCES

[1] A. Aldridge, *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. Hoboken, NJ: John Wiley & Sons, 2013.

[2] R. Cont, "Empirical properties of asset returns: stylized facts and statistical issues," *Quantitative Finance*, vol. 1, no. 2, pp. 223-236, 2001.

[3] J. B. Heaton, N. G. Polson, and J. H. Witte, "Deep learning for finance: deep portfolios," *Applied Stochastic Models in Business and Industry*, vol. 33, no. 1, pp. 3-12, 2017.

[4] E. Bacry, S. Iuga, M. Lasnier, and C. A. Lehalle, "Market impacts and the life cycle of investors orders," *Market Microstructure and Liquidity*, vol. 1, no. 2, 2015.

[5] E. F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, vol. 25, no. 2, pp. 383-417, 1970.

[6] A. W. Lo, "The Adaptive Markets Hypothesis," *The Journal of Portfolio Management*, vol. 30, no. 5, pp. 15-29, 2004.

[7] A. Tsantekidis et al., "Forecasting Stock Prices from the Limit Order Book using Convolutional Neural Networks," in *Proc. IEEE 19th Conference on Business Informatics (CBI)*, pp. 200-204, 2017.

[8] NSE India, "Market Microstructure and Algorithmic Trading in India," *National Stock Exchange Research Reports*, Mumbai, India, 2022.

[9] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.

[10] F. X. Diebold and R. S. Mariano, "Comparing predictive accuracy," *Journal of Business & Economic Statistics*, vol. 13, no. 3, pp. 253-263, 1995.

[11] Z. Zhang, S. Zohren, and S. Roberts, "DeepLOB: Deep convolutional neural networks for limit order books," *IEEE Transactions on Signal Processing*, vol. 67, no. 11, pp. 3001-3012, 2019.

[12] M. Dixon, D. Klabjan, and J. H. Bang, "Classification-based financial markets prediction using deep neural networks," *Algorithmic Finance*, vol. 6, no. 3-4, pp. 67-77, 2017.

[13] J. Sirignano and R. Cont, "Universal features of price formation in financial markets: perspectives from deep learning," *Quantitative Finance*, vol. 19, no. 9, pp. 1449-1459, 2019.

Journal of Innovations in Computer Science and Trends in IT, (October 2026), Vol 3, Issue2, ISSN 3048-4707

[14] W. Bao, J. Yue, and Y. Rao, "A deep learning framework for financial time series using stacked autoencoders and long-short term memory," *PLoS ONE*, vol. 12, no. 7, 2017.

[15] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, "Limit order books," *Quantitative Finance*, vol. 13, no. 11, pp. 1709-1742, 2013.

[16] M. O'Hara, "High frequency market microstructure," *Journal of Financial Economics*, vol. 116, no. 2, pp. 257-270, 2015.

[17] A. Ntakaris, M. Magris, J. Kanniainen, M. Gabbouj, and A. Iosifidis, "Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods," *Journal of Forecasting*, vol. 37, no. 8, pp. 852-866, 2018.

[18] O. B. Sezer, M. U. Gudelek, and A. M. Ozbayoglu, "Financial time series forecasting with deep learning: A systematic literature review: 2005–2019," *Applied Soft Computing*, vol. 90, 106181, 2020.

[19] T. Chordia, R. Roll, and A. Subrahmanyam, "Recent trends in trading activity and market quality," *Journal of Financial Economics*, vol. 101, no. 2, pp. 243-263, 2011.

[20] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*. San Francisco, CA: Holden-Day, 1970.

[21] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.

[22] B. Bekaert and C. R. Harvey, "Emerging equity market volatility," *Journal of Financial Economics*, vol. 43, no. 1, pp. 29-77, 1997.

[23] J. Hasbrouck, *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. New York, NY: Oxford University Press, 2007.

[24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd International Conference on Learning Representations (ICLR)*, San Diego, CA, 2015.

[25] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.

[26] A. Tran, M. Magris, J. Kanniainen, M. Gabbouj, and A. Iosifidis, "Tensor representation in high-frequency financial data for price change prediction," in *Proc. IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1-7, 2017.

[27] C. R. Harvey, "Predictable risk and returns in emerging markets," *The Review of Financial Studies*, vol. 8, no. 3, pp. 773-816, 1995.