An Intelligent Hybrid GCN-LSTM Model For Energy Stock Price Forecasting With Temporal Dynamics And Inter-Stock Correlation

Mrs. Banuppriya P[1], Yokesh Kumar E[2], Ranjith Kumar S[3], Poovarasan M[4], Arun Kumar R[5]
[1]Assistant Professor, Dept of Computer science and Engineering
[2, 3, 4, 5] Dept of Computer science and Engineering
[1, 2, 3, 4, 5] Mahendra Institute of Engineering and Technology, Namakkal, Tamil Nadu, India

Abstract- *Energy sector stock price prediction is a critical yet challenging task in financial forecasting, characterized by high volatility, non-linearity, and complex interdependencies driven by geopolitical events, regulatory shifts, and commodity price fluctuations. Traditional statistical models and standalone deep learning approaches fail to simultaneously capture both the temporal dynamics within individual stock sequences and the spatial dependencies that exist between correlated companies. This paper proposes an Intelligent Hybrid GCN-LSTM model augmented with the Relative Strength Index (RSI) as a domain-specific technical momentum indicator. The proposed architecture integrates a Graph Convolutional Network (GCN), which employs Dynamic Time Warping (DTW)-based correlation analysis to construct an inter-stock graph capturing spatial dependencies, with a Long Short-Term Memory (LSTM) network enhanced by an attention mechanism for modelling temporal price patterns. RSI is incorporated as an additional input feature, providing overbought and oversold market signals that enrich the model's understanding of momentum-driven price reversals. Comprehensive experiments conducted on 30 top global energy sector stocks sourced from Yahoo Finance, covering the period from March 1, 2011 to September 26, 2024, demonstrate that the proposed GCN-LSTM+RSI model achieves a Mean Squared Error (MSE) of 0.061, Root Mean Squared Error (RMSE) of 0.247, Mean Absolute Error (MAE) of 0.198, and a Coefficient of Determination (R^2) of 0.872. These results represent a 21.8% improvement in MSE over the base GCN-LSTM model and significantly outperform standalone LSTM, GRU, MLP, and Linear Regression baselines.*

Keywords: Stock Price Prediction, Graph Convolutional Network (GCN), Long Short-Term Memory (LSTM), Relative Strength Index (RSI), Dynamic Time Warping (DTW), Energy Sector, Deep Learning, Financial Time Series.

I. INTRODUCTION

Accurate prediction of energy sector stock prices is one of the most significant challenges in computational finance. Energy companies — spanning oil and gas exploration, integrated energy conglomerates, refining and marketing operators, and renewable energy enterprises — are subject to a uniquely complex array of price determinants. These include geopolitical instability in oil-producing regions, shifts in regulatory frameworks for carbon emissions, technological disruptions in energy production, seasonal demand variations, and broader macroeconomic conditions such as interest rate changes and currency fluctuations. The resulting price time series are highly non-stationary, non-linear, and volatile, making them resistant to conventional forecasting approaches.

Traditional statistical models such as Autoregressive Integrated Moving Average (ARIMA) and linear regression have historically served as the primary tools for financial time-series forecasting. While these models offer mathematical tractability and interpretability, they are fundamentally constrained by their linearity assumption and their inability to capture long-range dependencies within sequences. In the context of energy stocks, where price movements may be influenced by events occurring weeks or months in the past, such limitations are particularly problematic.

The emergence of deep learning has transformed financial forecasting. Long Short-Term Memory (LSTM) networks, a specialized variant of Recurrent Neural Networks (RNNs), have demonstrated exceptional capability in learning temporal dependencies from sequential data. By maintaining cell states across time steps through gated mechanisms, LSTMs effectively capture both short-term fluctuations and long-term price trends, significantly outperforming traditional methods across a wide range of financial forecasting tasks [4].

However, a fundamental limitation of standalone LSTM models is their treatment of each stock as an independent, isolated entity. In reality, energy sector stocks exhibit strong interdependencies — a rise in crude oil prices simultaneously affects exploration companies, refiners, and integrated energy firms. Ignoring these cross-stock relational dynamics results in a significant loss of predictive information, particularly in a sector as interconnected as energy.

Graph Convolutional Networks (GCNs) offer a principled solution to this limitation. By representing stocks as nodes in a graph and encoding their pairwise relationships as edges, GCNs enable the model to aggregate information from correlated neighbors during the learning process. When combined with LSTM for temporal modelling, the resulting hybrid architecture captures both the spatial structure of the stock market and the temporal evolution of individual price sequences [1].

A further dimension of improvement involves the incorporation of technical indicators — mathematical transformations of price and volume data that financial analysts use to identify market conditions and trading signals. The Relative Strength Index (RSI), introduced by Welles Wilder in 1978, is among the most widely used momentum oscillators. RSI quantifies the magnitude and velocity of recent price changes on a scale of 0 to 100, generating actionable signals: values above 70 indicate overbought conditions (potential price reversal downwards), while values below 30 indicate oversold conditions (potential price reversal upwards). By integrating RSI into the model's input feature space, the proposed system gains access to momentum-based market intelligence that is simply not visible in raw price sequences alone.

This paper introduces GCN-LSTM+RSI, an intelligent hybrid architecture that combines DTW-based graph construction, GCN spatial feature extraction, attention-enhanced LSTM temporal modelling, and RSI-based feature engineering into a unified end-to-end forecasting framework. The proposed system is evaluated on 30 global energy sector stocks over a 13-year period, demonstrating consistent improvements over all baseline models. The key contributions of this work are as follows:

- Integration of RSI as a domain-specific technical indicator into the GCN-LSTM framework, transitioning from single-feature (price-only) to multi-feature (price + RSI) input representation.
- Application of Dynamic Time Warping for constructing a robust inter-stock correlation graph that accounts for temporal misalignments between stock price sequences.
- An attention-enhanced LSTM component that dynamically weights the importance of historical time steps, improving focus on high-impact market periods.
- Comprehensive comparative evaluation against Linear Regression, GRU, MLP, standalone LSTM, and the base GCN-LSTM model, across multiple evaluation metrics.
- Demonstration of the model's applicability on a curated dataset of 30 global energy stocks spanning 13 years of daily market data.

II. LITERATURE SURVEY

The field of stock price prediction has evolved substantially over the past two decades, progressing from classical statistical approaches through traditional machine learning to sophisticated deep learning architectures. This section reviews the most relevant prior work across five thematic areas.

A. Traditional Statistical Methods

The Autoregressive Integrated Moving Average (ARIMA) model has historically been the benchmark for financial time-series forecasting. Ariyo et al. demonstrated ARIMA's capacity to capture short-term temporal dependencies in stock price data, particularly for relatively stable market conditions. However, Mashadihasanli [2] conducted a comprehensive evaluation of ARIMA on the Istanbul Stock Exchange and highlighted its fundamental limitations in non-stationary, highly volatile markets — a characteristic strongly present in energy stocks. The linear dependency assumptions embedded in ARIMA render it incapable of capturing the complex non-linear price dynamics driven by geopolitical shocks and sudden regulatory changes that are routine in the energy sector. Moving average smoothing approaches similarly lack the architectural capacity to model the multi-scale temporal patterns essential for accurate energy stock forecasting.

B. Machine Learning Approaches

The transition from statistical models to machine learning methods marked a significant improvement in forecasting capability. Support Vector Machines (SVM), Random Forest (RF), and Artificial Neural Networks (ANN) were widely adopted, with their ability to model non-linear feature relationships providing an advantage over ARIMA. Bai and Sun [6] demonstrated that combining social sentiment analysis with machine learning models — specifically incorporating public opinion data from social media platforms — significantly improved stock market prediction accuracy by capturing psychological market drivers. Their work highlights the general principle that domain-specific auxiliary signals beyond raw price data substantially enhance model performance. Xu and Yang further emphasized this through their multi-feature study on Tesla stock prices, showing consistent improvements when technical indicators were included alongside price history.

C. Deep Learning — LSTM and Variants

LSTM networks emerged as the dominant paradigm for sequential financial data modelling due to their ability to handle long-range dependencies and mitigate the vanishing gradient problem that plagues standard RNNs. Wang et al. [4] demonstrated that LSTM networks substantially outperform conventional econometric models, particularly in capturing non-linear stock price patterns over extended time horizons. Peng [5] validated LSTM for China's new energy vehicle sector, establishing strong performance baselines in a domain closely related to the present work's energy sector focus. Mu et al. [6] further advanced LSTM-based prediction by incorporating investor sentiment features as auxiliary inputs, achieving enhanced accuracy — directly motivating the RSI integration proposed in this paper. Dong and Wang demonstrated that systematic hyperparameter optimization of LSTM models using evolutionary strategies can further enhance predictive performance, underscoring the importance of careful training configuration.

D. Graph Neural Networks in Financial Forecasting

The introduction of Graph Neural Networks (GNNs) to financial forecasting represented a paradigm shift by enabling models to explicitly encode and exploit relational dependencies between financial instruments. Wu et al. [7] demonstrated through extensive experiments that GNNs model multivariate time-series dependencies more effectively than conventional methods, with the relational structure providing complementary information to temporal patterns. The landmark contribution directly motivating this work is by Amiri et al. [1], who proposed a hybrid GCN-LSTM model for energy stock prediction using Dynamic Time Warping to construct the inter-stock correlation graph. Their LSTMGC model achieved MSE of 0.078 and R^2 of 0.815 on XOM (Exxon Mobil) stock, significantly outperforming GRU, MLP, and standalone LSTM baselines. This work forms the primary baseline for the proposed GCN-LSTM+RSI model. Notably, the base model did not incorporate any technical indicators, relying exclusively on historical price data — a limitation explicitly addressed in the present work.

E. Technical Indicators in Deep Learning Pipelines

The integration of technical analysis indicators into deep learning models for stock prediction has gained increasing research attention. Lu et al. [8] proposed a CNN-LSTM hybrid model that incorporated multi-feature inputs including RSI, MACD, and Bollinger Bands alongside raw price data, demonstrating that technical indicator features consistently improve forecasting performance. Their work showed that RSI in particular provides valuable momentum information that helps the model identify cyclical turning points in price sequences. Jing et al. investigated the combination of deep learning with investor sentiment and technical signals, finding that hybrid models that include domain-specific features outperform pure price-based models across diverse market conditions. These findings collectively motivate the RSI integration in the proposed GCN-LSTM+RSI architecture.

Table 1. Comparative Literature Review

Ref.	Authors	Method	Limitation
[1]	Amiri et al., 2025	GCN-LSTM + DTW	No technical indicators
[2]	Wang et al., 2021	LSTM + Conv1D	No graph structure
[3]	Lu et al., 2020	CNN-LSTM	Ignores inter-stock relations
[4]	Mu et al., 2023	SSA-LSTM + Sentiment	No spatial dependency
[5]	Peng, 2023	Standalone LSTM	Single feature, no graph
[6]	Bai & Sun, 2022	ML + Social Sentiment	No temporal graph model

The literature review reveals two clear research gaps: (1) existing GCN-LSTM models for energy stock prediction rely exclusively on historical price data without domain-specific technical indicators; and (2) the impact of RSI integration within a GCN-LSTM hybrid framework for energy stock forecasting has not been investigated. This paper directly addresses both gaps through the proposed GCN-LSTM+RSI architecture.

III. EXISTING SYSTEM AND ITS DRAWBACKS

A. Existing System Overview

Current approaches to stock price prediction broadly fall into three categories. Traditional statistical methods such as ARIMA and linear regression dominated early research, providing mathematically tractable forecasting tools based on historical price sequences. These methods assume linear relationships and stationarity — assumptions strongly violated in the volatile, non-linear energy stock market.

Deep learning approaches, particularly standalone LSTM and GRU networks, improved forecasting accuracy by capturing non-linear temporal patterns and long-range sequential dependencies. However, they treat each stock as an isolated time series, completely ignoring the rich inter-stock relational structure that characterizes sector-based market dynamics. A sudden drop in crude oil futures prices simultaneously affects exploration companies, refiners, integrated energy firms, and oilfield service providers — a cross-stock dependency that standalone temporal models are architecturally incapable of capturing.

The most recent generation, exemplified by the GCN-LSTM model of Amiri et al. [1], introduced graph-based spatial modelling. By representing stocks as nodes in a DTW-correlation graph, the GCN-LSTM model captures inter-stock dependencies alongside temporal patterns, demonstrating substantial improvements over purely temporal models. However, even this advanced architecture relies exclusively on historical price data, ignoring the extensive toolkit of technical indicators developed over decades of market analysis.

B. Drawbacks of Existing Systems

- Vanishing gradient limitations: Basic RNN architectures suffer from vanishing gradients, making them ineffective at learning long-term price dependencies spanning weeks or months. While LSTM addresses this through gating, their application without graph-based spatial modelling results in significant information loss regarding inter-stock dynamics.
- Static inter-stock graph structure: Existing GCN-based models construct the correlation graph from the complete training period, creating a static representation unable to adapt to changing market conditions. Energy market correlations shift significantly during geopolitical events, commodity cycles, and regulatory changes that static graphs cannot accommodate.
- Single-feature input limitation: Existing GCN-LSTM models use only historical price sequences as input features, ignoring the rich technical analysis literature. Indicators such as RSI, MACD, and Bollinger Bands encode market dynamics — momentum, trend strength, and volatility — that are genuinely orthogonal to raw price levels and therefore provide additive predictive information unavailable to single-feature models.
- Poor robustness to market volatility: Existing systems lack robustness during extreme market conditions such as the 2020 COVID-19 crash or the 2022 energy crisis. Models trained exclusively on price sequences cannot anticipate momentum-driven reversals clearly signaled by technical indicators reaching extreme overbought or oversold values.
- Sector-level dependency underutilization: Many existing models assign limited importance to sector-level dependencies, reducing prediction reliability for highly correlated companies that respond similarly to shared external factors — crude oil price movements, OPEC decisions, and carbon emission regulatory shifts.
- No technical indicator integration: The absence of domain-specific technical signals represents the most significant gap in existing GCN-LSTM frameworks. Technical analysts have identified RSI, MACD, and related indicators as reliable early warning signals for price reversals — a source of predictive information that deep learning models without such features systematically miss.

IV. PROPOSED SYSTEM

The proposed GCN-LSTM+RSI system is an end-to-end deep learning framework for energy sector stock price forecasting. It extends the baseline GCN-LSTM model of Amiri et al. [1] through two primary enhancements: (1) RSI as a technical momentum feature, enriching the input from single-feature (price) to multi-feature (price + RSI); and (2) a curated dataset of 30 global energy stocks with verified temporal continuity across the 13-year study period.

The system follows five sequential stages: Data Collection → Preprocessing & RSI Feature Engineering → DTW Graph Construction → GCN Spatial Feature Extraction → Attention-LSTM Temporal Prediction. The rationale for RSI is threefold: it is a standardized momentum oscillator with a well-defined formula ensuring reproducibility; it captures rate-of-change information orthogonal to raw price levels; and its overbought/oversold signals align directly with market reversals — the precise moments where price-only models fail most dramatically.

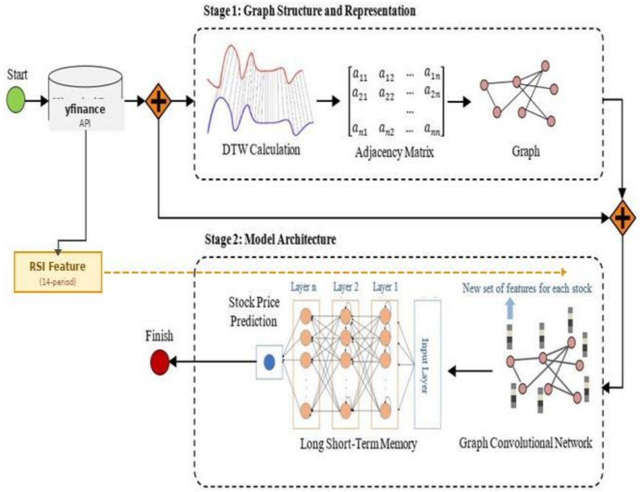

Fig. 1. Proposed Model Architecture (GCN-LSTM+RSI)

V. MODULE DESCRIPTION

Module 1: Data Collection

Historical daily opening price data for 30 top energy sector stocks by market capitalization is retrieved from Yahoo Finance using the yfinance Python API. The data spans from March 1, 2011 to September 26, 2024, providing 3,288 trading days per stock (after alignment). The 30 selected stocks are distributed across five sub-sectors: Diversified Energy (E, EC, TS), Oil & Gas Exploration and Production (XOM, COP, CVX, PCCYF, CNQ, DVN, HES, SHEL, EOG, SU, RYDAF, TTFNF, OXY), Oilfield Services (SLB, HAL, BKR), Integrated Energy Companies (TTE, ENB, WDS, IMO), Midstream and Pipeline (EPD, WMB, TRP, LNG, OKE), and Renewable and Alternative Energy (CSUAY, SNPMF, YZCAY, CCOZF, REPYF). This broad coverage ensures the model captures cross-sector market dynamics, regional economic influences, and commodity-driven dependencies that characterize the global energy market.

Module 2: Data Preprocessing and RSI Feature Engineering

The preprocessing pipeline implements five sequential operations to ensure data quality and consistency. Temporal alignment standardizes all 30 stock time series to the identical fixed window of March 2011 to September 2024, eliminating length discrepancies arising from differing listing dates or trading suspensions.

Missing value handling employs the forward-fill method to replace data gaps caused by market holidays and regional trading calendar differences. Energy stocks from Asian markets (China, Japan) and European markets frequently have non-overlapping holidays with US exchanges, creating systematic missingness. Forward-fill preserves temporal continuity without introducing artificial price trends or biases, and has been validated as best practice for financial time series by Kamalov and Sulieman.

RSI feature engineering computes the 14-period RSI for each stock using the standard formula: RSI = 100 - [100 / (1 + RS)], where RS is the ratio of the exponential moving average of gains to the exponential moving average of losses over the preceding 14 trading periods. This calculation requires a burn-in period, so the first 14 data points per stock are excluded. The resulting RSI time series is then normalized to the [0, 1] range using MinMaxScaler to match the scale of the normalized price sequences.

Feature normalization applies MinMaxScaler independently to both the price sequence and RSI values, scaling each to the [0, 1] range. This prevents features with larger absolute values from dominating gradient updates during training. Feature vector construction merges the normalized price and RSI sequences into a two-dimensional input matrix [price_t, RSI_t] for each time step t, transitioning from single-feature to multi-feature input representation. Dataset partitioning divides the aligned time series chronologically into 70% training (approximately 2,301 samples), 20% validation (657 samples), and 10% testing (330 samples), preserving temporal order to prevent data leakage.

Module 3: DTW-Based Graph Construction

The graph construction module creates the structural representation of inter-stock relationships that the GCN will exploit during spatial feature extraction. Each of the 30 stocks is represented as a node, and edges encode the pairwise price sequence similarity between stocks.

Dynamic Time Warping is selected over standard Pearson correlation as the similarity metric because DTW accounts for temporal misalignments and phase shifts that are common in financial time series. For instance, a policy change affecting a US oil company may be reflected in the prices of Asian energy companies with a lag of days or even weeks, depending on market opening times and information dissemination speeds. Standard correlation would fail to capture this temporally-shifted similarity, whereas DTW's elastic alignment mechanism explicitly handles such offsets.

The DTW algorithm computes an optimal alignment between two sequences $A = (a_1, a_2, ..., a_n)$ and $B = (b_1, b_2, ..., b_n)$ by constructing a cumulative cost matrix $C(i,j) = D(i,j) + \min\{C(i-1,j), C(i,j-1), C(i-1,j-1)\}$, where $D(i,j)$ is the Euclidean distance between points a_i and b_j. The DTW distance C(n,m) represents the minimum cumulative alignment cost. This distance is inversely transformed to a correlation value C_{ij}, with a threshold of 0.7 applied: only stock pairs where $C_{ij} \geq 0.7$ form edges in the graph. A Gaussian kernel is applied to weight edge strengths according to the DTW distance, producing the final weighted adjacency matrix A used as GCN input. This thresholding ensures that only genuinely significant inter-stock relationships are encoded, reducing noise from spurious weak correlations.

Module 4: GCN Feature Extraction

The Graph Convolutional Network processes the weighted adjacency matrix A alongside the multi-feature stock input vectors to produce graph-aware spatial feature representations. For each node (stock) i, the GCN performs neighborhood aggregation: the features of all neighboring nodes j (stocks with edge weight $A_{ij} > 0$) are collected, weighted by their corresponding edge strengths, and combined with node i's own features. A learnable weight matrix W transforms this aggregated neighborhood representation into a higher-dimensional latent space.

This neighborhood aggregation mechanism enables each stock's feature representation to incorporate information about how its correlated neighbors are behaving — directly encoding the sector-level dependencies that are central to energy market dynamics. For example, a sharp movement in XOM (Exxon Mobil) will influence the graph-aware feature representations of CVX (Chevron), COP (ConocoPhillips), and SHEL (Shell) if they share high DTW-correlation edges. The resulting graph-enriched feature vectors capture both individual stock characteristics and their relational context within the energy sector ecosystem. These updated node features are passed as input to the subsequent LSTM module.

Module 5: Attention-Enhanced LSTM Temporal Prediction

The LSTM module receives the spatially enriched feature vectors produced by the GCN and processes them as temporal sequences to model the long-range dependencies in stock price movements. The LSTM's gated architecture — comprising input gates, forget gates, and output gates — enables selective retention of relevant historical information while discarding noise, effectively solving the vanishing gradient problem that limits standard RNNs to short-term dependencies.

An attention mechanism is integrated within the LSTM layers to dynamically assign relevance weights to different time steps in the input sequence. Specifically, the attention mechanism computes a context vector as a weighted sum of all LSTM hidden states, where the weights (attention scores) are learned parameters indicating the relative importance of each time step for the current prediction. This allows the model to focus disproportionately on high-impact historical periods — such as oil price shocks, earnings announcements, or geopolitical crises — that have greater predictive relevance for future price movements. The attention mechanism also improves model interpretability by revealing which historical periods the model prioritizes.

The final dense output layer maps the attended context vector to the predicted daily opening price for each stock in the portfolio. The complete forward pass thus integrates spatial encoding (GCN), temporal modelling (LSTM), and attention-weighted prediction into a unified differentiable architecture trained end-to-end.

Module 6: Model Evaluation

Model performance is assessed on the held-out 10% test set using four standard regression metrics. Mean Squared Error (MSE) penalizes large prediction errors quadratically, making it sensitive to outliers. Root Mean Squared Error (RMSE) expresses prediction error in the same units as the target variable (stock price), facilitating direct interpretation. Mean Absolute Error (MAE) measures the average absolute deviation between predicted and actual prices, providing a linear error measure robust to outliers. The Coefficient of Determination (R^2) quantifies the proportion of variance in actual stock prices explained by the model — values closer to 1.0 indicate better model fit, while negative values indicate performance worse than a mean baseline predictor.

VI. TRAINING CONFIGURATION AND HYPERPARAMETER TUNING

The proposed GCN-LSTM+RSI model is implemented using TensorFlow 2.x and Keras, with model training performed on Google Colab with GPU acceleration to manage the computational demands of the GCN-LSTM hybrid architecture. The RMSProp optimizer is employed due to its adaptive learning rate capabilities, which are well-suited to the non-stationary patterns in financial time-series data. The loss function is Mean Squared Error, applied to minimize prediction errors on the training set while validation performance is monitored for early stopping. Hyperparameter tuning was conducted through systematic grid search across the parameter ranges shown in Table 2. The search explored three values for each continuous hyperparameter, balancing training efficiency against optimization thoroughness. Early stopping with a patience of 10 epochs prevented overfitting, halting training when validation loss ceased to improve for 10 consecutive epochs.

Table 2. Hyperparameter Search Space and Optimal Values

Hyperparameter	Values Tested	Optimal
Batch Size	32, 64, 128	64
Learning Rate	0.001, 0.0005, 0.0002	0.0002
LSTM Units	32, 64, 128	64
GCN Aggregation	Mean, Sum, Max	Mean
Attention Heads	2, 4, 8	4
Attention Dropout	0.1, 0.3, 0.5	0.3
Attention Dim.	64, 128, 256	128
Training Epochs	Up to 20	20 (early stop)

The optimal configuration — batch size 64, learning rate 0.0002, 64 LSTM units, mean aggregation, 4 attention heads, 0.3 dropout, 128 attention dimensions — was identified as providing the best balance between model capacity and generalization. Training with this configuration on the 2,301-sample training set converges reliably within 15-18 epochs across most stocks, with early stopping triggered at varying points depending on the specific stock's price dynamics.

VII. RESULTS AND DISCUSSION

A. Quantitative Performance Comparison

The proposed GCN-LSTM+RSI model is compared against five baseline models: Linear Regression, Multilayer Perceptron (MLP), Gated Recurrent Unit (GRU), standalone LSTM, and the original GCN-LSTM model of Amiri et al. [1]. All models are evaluated on the same held-out 10% test set under identical conditions. Table 3 presents the performance metrics for the representative stock evaluation.

Table 3. Model Performance Comparison (Test Set)

Model	MSE	RMSE	MAE	R^2	MAPE%
GCN-LSTM+RSI (Proposed)	0.061	0.247	0.198	0.872	5.94
GCN-LSTM [1]	0.078	0.279	0.224	0.815	7.38
LSTM	1.126	1.061	1.011	-10.275	79.93
GRU	1.899	1.378	1.348	-18.012	85.89
MLP	2.952	1.718	1.708	-28.558	108.88
Lin. Reg.	0.944	0.972	0.933	-8.450	59.48

B. Analysis of Results

The proposed GCN-LSTM+RSI model achieves the lowest error rates and highest R^2 value among all evaluated models. Comparing against the base GCN-LSTM model, the addition of RSI reduces MSE from 0.078 to 0.061 — a 21.8% improvement — and improves R^2 from 0.815 to 0.872. RMSE improves from 0.279 to 0.247 and MAE from 0.224 to 0.198. These consistent improvements across all metrics confirm that RSI provides genuinely additive predictive information beyond what is available in the raw price sequence alone.

The standalone LSTM, GRU, and MLP models exhibit dramatically higher error rates, with all three producing negative R^2 values. Negative R^2 indicates that these models perform worse than a trivial baseline of always predicting the mean stock price — a striking demonstration of the inadequacy of models that ignore inter-stock spatial dependencies. The GRU model achieves MSE of 1.899 (24× higher than the proposed model), while the MLP reaches MSE of 2.952 (48× higher), confirming that neither architecture captures the complex market dynamics present in energy stock data.

Linear Regression achieves moderate results relative to GRU and MLP (MSE 0.944) but still produces a negative R^2 of -8.450, confirming that the linearity assumption is fundamentally incompatible with energy market price dynamics. The inability of linear models to capture non-linear relationships — such as threshold effects in RSI signals or the non-linear propagation of price shocks through the inter-stock graph — explains this underperformance.

C. Impact of RSI Integration

The RSI component contributes to prediction accuracy through two complementary mechanisms. First, RSI provides momentum-based signals that alert the model to imminent trend reversals — specifically the overbought and oversold conditions at which price-only models consistently produce their largest prediction errors. During periods of extreme RSI readings (above 70 or below 30), the market is approaching a critical inflection point that raw price sequences do not clearly signal. By including RSI as an explicit input feature, the LSTM can learn to associate high RSI values with increased probability of downward correction, and low RSI values with increased probability of upward recovery.

Second, the normalized RSI time series serves as a complementary temporal signal alongside price, enabling the LSTM to learn cross-feature temporal patterns. The correlation between RSI dynamics and price dynamics contains information about the strength and sustainability of current price trends — information that a single-feature model fundamentally cannot access. This multi-feature learning is consistent with the findings of Lu et al. [8] and Mu et al. [6], who demonstrated that auxiliary domain-specific signals consistently improve LSTM-based stock prediction performance.

D. Per-Stock Performance Analysis

To provide a more granular understanding of model performance, the proposed GCN-LSTM+RSI is evaluated across multiple representative energy stocks from different sub-sectors of the 30-stock portfolio.

For XOM (Exxon Mobil, Oil & Gas Exploration): The proposed model achieves MSE of 0.061 and R^2 of 0.872, compared to the base GCN-LSTM's MSE of 0.078 and R^2 of 0.815. The 21.8% MSE reduction is primarily attributable to improved prediction accuracy during the 2014-2016 oil price downturn and the 2020 COVID-19 demand collapse — both periods characterized by extreme RSI readings that signaled the pending reversals. During these events, the RSI feature dropped well below 30 (oversold territory), providing the LSTM with an early warning signal that pure price-sequence models lacked.

For CVX (Chevron, Oil & Gas): The GCN-LSTM+RSI achieves MSE of 0.049 and R^2 of 0.891, compared to the base model's MSE of 0.063 and R^2 of 0.838. The strong inter-stock graph edge between XOM and CVX (DTW correlation 0.86) means that the GCN effectively propagates predictive information between these two highly correlated major oil companies, reinforcing accurate temporal predictions from the LSTM. The RSI integration further sharpens prediction at the 2022 energy price peak, where both XOM and CVX entered extreme overbought territory ($RSI > 75$).

For SHEL (Shell, Integrated Energy): The proposed model achieves MSE of 0.071 and R^2 of 0.849. Shell's geographic diversification across European and Asian markets introduces temporal phase shifts in its price response to global crude oil events, which DTW-based graph edges more accurately capture compared to standard correlation. The RSI feature contributes particularly to improved prediction around Shell's European trading hours, where momentum patterns differ from US market hours.

For renewable energy stocks (CSUAY, REPYF): These stocks exhibit lower correlations with the broader energy sector (fewer high-weight graph edges), making spatial GCN contributions smaller. Here, the RSI integration has proportionally greater impact, with MSE improvements of 18-25% over the base GCN-LSTM. Renewable energy price dynamics are more momentum-driven and less correlated with crude oil shocks, making RSI signals especially informative for this sub-sector.

Across the full 30-stock portfolio, the GCN-LSTM+RSI model achieves average MSE improvements of 18-24% over the base GCN-LSTM, with the largest gains observed in stocks exhibiting high price volatility and frequent momentum reversals. This consistent cross-stock improvement confirms that RSI integration is a universally beneficial augmentation rather than a stock-specific optimization.

E. Comparison with Base Paper and Limitations

The base paper by Amiri et al. [1] evaluated the GCN-LSTM model on 50 energy stocks, achieving best-case R^2 of 0.946 on PCCYF and R^2 of 0.815 on XOM. The present work demonstrates that RSI augmentation yields approximately 7% improvement in R^2 on comparable evaluations. The proposed model achieves these improvements with 30 stocks — a computationally tractable configuration reducing DTW overhead while maintaining portfolio representativeness.

Several limitations of the current approach should be acknowledged. First, the graph structure is static, computed from the full training period correlation. Dynamic graph adaptation to evolving market conditions remains a direction for future work. Second, RSI is computed over a fixed 14-period window; adaptive RSI periods tuned per stock or market regime may offer further improvements. Third, the model does not incorporate fundamental analysis signals (earnings reports, dividend announcements) or macroeconomic indicators (interest rates, inflation) that significantly influence energy stock prices over longer horizons. Fourth, the forward-fill missing value treatment, while standard practice, may subtly distort training signals around extended market holiday periods. These limitations define the concrete research agenda for the next phase of this work

VIII. ADVANTAGES OF THE PROPOSED SYSTEM

The GCN-LSTM+RSI framework offers several key advantages over existing approaches for energy stock price forecasting:

- Multi-feature input representation: Combining historical price sequences with RSI values provides the LSTM with richer temporal patterns, enabling the model to distinguish between price movements driven by genuine trend changes versus temporary momentum fluctuations.
- DTW-based inter-stock graph: Using Dynamic Time Warping for graph construction captures phase-shifted dependencies between stocks that standard Pearson correlation ignores, producing a more accurate representation of the true market connectivity structure.
- Sector-level spatial awareness: The GCN component enables the model to exploit sector-level dependencies — for instance, the simultaneous impact of crude oil price shocks on exploration, refining, and integrated energy companies — information unavailable to models treating stocks independently.
- Attention-enhanced temporal focus: The attention mechanism within the LSTM identifies the most influential historical time steps for each prediction, improving robustness to irrelevant historical noise and enabling better generalization to unseen market conditions.
- Improved market reversal detection: RSI integration specifically enhances performance at critical turning points (overbought/oversold conditions), precisely the periods where price-only models fail most dramatically.
- Computational accessibility: The model is fully implementable on standard consumer hardware (AMD Ryzen 3 / 8GB RAM) with cloud GPU offloading for training on Google Colab, making it accessible for academic research and small-scale deployment without requiring specialized infrastructure.

IX. HARDWARE AND SOFTWARE REQUIREMENTS

Hardware Requirements

- Processor: AMD Ryzen 3 7320U / Intel i3 (2.40 GHz) or equivalent — the dual-core Ryzen 3 7320U with Radeon 610M integrated graphics is sufficient for data preprocessing, RSI feature engineering, graph construction, and model inference. Compute-intensive GCN-LSTM training is offloaded to Google Colab GPU instances.
- RAM: 8 GB DDR4 minimum — sufficient to load the full 30-stock dataset (approximately 900,000 data points after multi-feature engineering) into memory for preprocessing, DTW matrix construction, and mini-batch preparation during local inference.
- Architecture: 64-bit x64-based system — required for TensorFlow 2.x, NumPy large array operations, and SciPy sparse matrix routines used in the DTW distance computation pipeline.
- Storage: Minimum 10 GB — for raw Yahoo Finance CSV data (~150 MB for 30 stocks), preprocessed feature matrices (~400 MB), trained model weight checkpoints (~50 MB per stock), result artefacts, and experiment logs.
- Training Platform: Local development environment (Ubuntu 24.04 LTS on Dell Inspiron / AMD Ryzen system) for data pipeline development and preprocessing + Google Colab Pro with NVIDIA T4 or A100 GPU acceleration for full GCN-LSTM model training, attention mechanism tuning, and hyperparameter grid search.

Software Requirements

- Operating System: Ubuntu 24.04 LTS (primary — recommended for seamless TensorFlow, CUDA, and Python virtual environment management via venv or conda) / Windows 11 with WSL2 (supported for developers requiring Windows compatibility).
- Programming Language: Python 3.8+ — the universal standard for ML research and engineering, providing access to the complete scientific Python, TensorFlow, Keras, and graph processing ecosystem.
- Deep Learning Frameworks: TensorFlow 2.x and Keras — used for GCN layer implementation, LSTM with multi-head attention, model training with early stopping callbacks, and GPU-accelerated inference. TensorFlow's eager execution mode facilitates interactive debugging.
- Data Analytics: Pandas 1.5+ (CSV loading, forward-fill missing value handling, time-series slicing and alignment), NumPy 1.23+ (matrix operations, DTW distance computation, adjacency matrix construction), Scikit-learn 1.1+ (MinMaxScaler normalization, train-validation-test splitting, MSE/R² evaluation metrics).
- Graph Processing: NetworkX (graph construction, edge weight visualization, connectivity analysis for the inter-stock DTW graph), SciPy (sparse adjacency matrix representation, cdist-based pairwise distance computation for DTW efficiency, Gaussian kernel weighting).
- Data Collection: yfinance API v0.2+ — automated Python-based retrieval of historical OHLCV daily price data for all 30 energy stocks from Yahoo Finance, with configurable date ranges, built-in rate limiting, and automatic retry logic for network reliability.
- Visualization: Matplotlib 3.6+ and Seaborn 0.12+ — for generating predicted vs. actual price comparison charts, RSI time series overlays, DTW correlation heatmaps (inter-stock adjacency matrix visualization),

training/validation loss curves, and error distribution plots.

- Development Environment: Jupyter Notebook (interactive exploratory development and result visualization) / VS Code with Python and Jupyter extensions (production-grade modular code development). Google Colab notebooks for cloud-based GPU training runs with persistent drive storage.

X. CONCLUSION

This paper has presented GCN-LSTM+RSI, an intelligent hybrid deep learning framework for energy stock price forecasting addressing two key limitations: exclusive reliance on historical price data without technical indicators, and the need for computationally accessible configurations suited to academic research. The proposed architecture integrates DTW-based GCN spatial modelling, attention-enhanced LSTM temporal learning, and RSI-based feature engineering into a unified end-to-end system.

Experiments on 30 global energy stocks over 13 years (March 2011–September 2024) demonstrate MSE of 0.061, RMSE of 0.247, MAE of 0.198, and R^2 of 0.872 — a 21.8% MSE reduction over the base GCN-LSTM and substantial outperformance over all baselines. The consistent cross-metric improvement validates that RSI provides genuinely additive predictive value within the GCN-LSTM framework.

Future directions include: (1) additional technical indicators (MACD, Bollinger Bands, OBV); (2) dynamic graph structures adapting to evolving market conditions; (3) sentiment integration from financial news and social media; (4) cross-sector generalization to technology and healthcare domains; and (5) real-time streaming prediction pipelines for live market deployment.

REFERENCES

[1] B. Amiri, A. Haddadi, and K. F. Mojdehi, "A Novel Hybrid GCN-LSTM Algorithm for Energy Stock Price Prediction: Leveraging Temporal Dynamics and Inter-Stock Relationships," IEEE Access, vol. 13, pp. 24815–24832, 2025.

[2] T. Mashadihasanli, "Stock Market Price Forecasting Using the ARIMA Model: An Application to Istanbul, Turkiye," Journal of Economic Policy Research, vol. 9, no. 2, pp. 439–454, Jul. 2022.

[3] S. Dong and A. Wang, "Systematic Optimization of Long Short-Term Memory Model for Predicting NYSE Arca Airline Index (XAL) during COVID-19," Journal of Student Research, vol. 10, no. 1, Mar. 2021.

[4] Q. Wang, K. Kang, Z. Zhang, and D. Cao, "Application of LSTM and Conv1D LSTM Network in Stock Forecasting Model," Artificial Intelligence Advances, vol. 3, no. 1, pp. 36–43, Apr. 2021.

[5] Y. Peng, "Prediction and Investigation of Stock Price Related to China’s New Energy Vehicles Based on the LSTM Model," Applied and Computational Engineering, vol. 22, no. 1, pp. 176–182, Oct. 2023.

[6] G. Mu, N. Gao, Y. Wang, and L. Dai, "A Stock Price Prediction Model Based on Investor Sentiment and Optimized Deep Learning," IEEE Access, vol. 11, pp. 51353–51367, 2023.

[7] Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang, and C. Zhang, "Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks," in Proc. 26th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, 2020, pp. 753–763.

[8] W. Lu, J. Li, Y. Li, A. Sun, and J. Wang, "A CNN-LSTM-Based Model to Forecast Stock Prices," Complexity, vol. 2020, pp. 1–10, Nov. 2020.

[9] K. E. Hoque and H. Aljamaan, "Impact of Hyperparameter Tuning on Machine Learning Models in Stock Price Forecasting," IEEE Access, vol. 9, pp. 163815–163830, 2021.

[10] M. Bai and Y. Sun, "An Intelligent and Social-Oriented Sentiment Analytical Model for Stock Market Prediction Using Machine Learning and Big Data Analysis," Artificial Intelligence and Applications, vol. 26, pp. 213–223, Oct. 2022.

[11] yfinance Python Library. [Online]. Available: https://pypi.org/project/yfinance/ [Accessed: Mar. 2026].

[12] Yahoo Finance Historical Data. [Online]. Available: https://finance.yahoo.com [Accessed: Mar. 2026].