Stocks in Sync: Cluster-Aware Deep Learning for Multi-stock Forecasting in Financial Market

Soumyadeep Basak[1], Shirsendu Roy[1], Sarnaavho Pal[1], Shubham Sahu[1], Deepsubhra Guha Roy[2], and Piyali Datta[3](✉)

[1] Department of CSE(AIML), Institute of Engineering and Management, University of Engineering and Management, Kolkata, India

[2] IEM Centre of Excellence for Cloud Computing and IoT, Department of CSE(AIML), Institute of Engineering and Management, University of Engineering and Management, Kolkata, India

[3] IEM-IIT Mandi Centre for Joint Research on Human Computer Interaction, Department of CSE(AIML), Institute of Engineering and Management, University of Engineering and Management, Kolkata, India

datta.piyali.in@gmail.com

Abstract. Forecasting stock price movements is a complex task due to inherent market volatility, nonlinear dependencies, and inter-stock correlations. This paper proposes a cluster-informed deep learning framework for the simultaneous prediction of the next-day opening price direction across 46 constituents of the NIFTY 50 index. Unlike traditional single-stock models, our approach leverages shared sectoral patterns by integrating unsupervised similarity-based clustering with a unified multi-output architecture. We first identify co-movement relationships among stocks using a combination of correlation and cosine similarity metrics, followed by hierarchical clustering. These relationships inform the design of a hybrid neural model: a Time-Distributed dense encoder for dimensionality reduction, parallel LSTM layers for temporal representation learning, and a shared LSTM block leading to stock-specific output heads. Extensive experiments demonstrate that the proposed model consistently outperforms statistical and rule-based baselines in directional accuracy and the F1-score. The results highlight the efficacy of incorporating sector-level dependencies and structured temporal learning in robust financial forecasting.

Keywords: Stock Forecasting · Deep Learning · Time Series Prediction · NIFTY 50 · LSTM · Clustering · Multi-Output Learning · Sectoral Modeling · Feature Engineering

1 Introduction

Predicting the directional movement of stock prices remains a core challenge in quantitative finance due to the noisy, nonlinear, and highly interdependent

C. Zaroliagis et al. (Eds.): ICAA 2026, LNCS 16423, pp. 207–219, 2026.
https://doi.org/10.1007/978-3-032-15621-1_17

nature of financial time series [1]. Traditional univariate models often treat each stock in isolation, overlooking broader market dynamics such as sectoral trends and stock-to-stock co-movements—factors especially prominent in indices like the NIFTY 50 [2]. Recent advances in deep learning have driven sequential modeling by the use of recurrent models like Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU) [3]. However, the majority of models treat stocks individually or employ graph-based models, which is associated with scalability, interpretability, and generalization issues. In this paper, we introduce a novel cluster-informed multi-stock forecasting framework that explicitly models inter-stock dependencies beyond graph neural networks. We construct stock clusters using co-movement similarity measures and correlations or cosine-based metrics followed by hierarchical clustering to identify latent sectoral structures. These clusters guide the design of Cluster-Aware Dual-LSTM(CA-DLSTM), a shared neural architecture comprising a TimeDistributed encoder, parallel and shared LSTM layers of varying sequence lengths, a shared LSTM block, and stock-relevant dense heads that output binary predictions for directional movement. We propose a Ticker-Embedded LSTM that embeds stock identity into the learning process by augmenting time series data with a ticker column. A learnable ticker embedding is initialized via Multidimensional Scaling on the hierarchical clustering distance matrix, then concatenated with LSTM outputs before prediction. Although it captures cross-stock temporal dependencies, it underperforms the cluster-informed model (mean F1 ~ 0.64). Incorporating hierarchical clustering and structural embeddings enables modeling of both sector-level and stock-specific behavior, and experiments on NIFTY 50 show our models outperform statistical and rule-based baselines in accuracy and F1 score. This work makes the following key contributions:

- A cluster-informed deep learning framework has been developed that uses co-movement-based clustering to uncover latent sectoral structure among NIFTY 50 stocks.
- A scalable cluster aware dual LSTM architecture with multi-output has been designed with shared encoders and stock-specific dense heads for simultaneous prediction across 46 stocks.
- A Ticker-Embedded LSTM with an MDS-initialized stock embedding that incorporates structural priors into interleaved sequence modeling.
- We empirically demonstrate that both models outperform traditional baselines and standalone LSTM approaches in directional accuracy and F1 score.

2 Related Work

Stock Prediction with Deep Learning. Deep learning has become popular in financial forecasting for modeling nonlinear and time-dependent market data. Recurrent models like LSTM and GRU are widely used for single-stock prediction with technical or price features [3,4], but they remain stock-specific and lack scalability for indices such as NIFTY 50. Multi-stock models [4] aim to address this issue but often neglect inter-stock relations. GNNs capture such ties via correlations [7], though they add complexity and depend on manual graph design. LSTM and GRU perform better under volatility [3,5]. Recent work explores ticker-aware models with trainable embeddings; we initialize them using Multidimensional Scaling on a co-movement distance matrix from hierarchical clustering [1]. Our cluster-informed model improves generalization, showing the value of embedding sectoral structure.

Clustering and Co-Movement in Finance. Clustering has long been used in finance for pattern discovery, asset allocation, and reducing modeling complexity [2]. Methods like hierarchical clustering and k-means with correlation or cosine distances identify economically meaningful groups of co-moving stocks [1,6], reflecting macroeconomic effects or industry dynamics. Prior work often applies clustering as a pre-processing step for diversification, strategy design [2], or anomaly detection [4], but rarely integrates it into model design. Our approach(Cluster-Aware Dual-LSTM and Ticker-Embedded LSTM) instead embeds sector-driven co-movement patterns, obtained via hierarchical clustering on correlation and directional similarity, directly into the modeling pipeline. This enables the model to capture inter-stock dependencies and generalize better across sectors.

Multi-output Learning and Evolutionary Approaches. Multi-task and multi-output learning have gained traction in financial forecasting as they capture shared temporal and structural patterns across assets [7,8,13], improving generalization in indices like NIFTY 50 with strong intersectoral dependencies [2]. Genetic Algorithms (GA) have also been applied to evolve interpretable rule-based trading systems, though they struggle with nonlinear temporal dynamics in noisy markets. Our approach combines these directions in a cluster-conscious multi-output hybrid model. Using correlation- and cosine-based similarity measures, hierarchical clustering separates NIFTY 50 stocks [1,6], guiding a TimeDistributed encoder and dual LSTM backbone with multi-scale support [3,5]. Structural awareness is further introduced via trainable ticker embeddings initialized with Multidimensional Scaling on the clustering distance matrix. These embeddings, concatenated with recurrent outputs, feed stock-specific dense heads for prediction. While TA-LSTM improves over baselines, our CA-DLSTM achieves superior directional accuracy and F1 scores, demonstrating the value of integrating clustering, temporal modeling, and multi-output learning.

3 Model Architecture and Methodological Approach

Here, within this section, we outline the end-to-end process that we use to predict direction of movement of NIFTY 50 stocks. Our pipeline is composed of a clustering-based preprocessing step followed by a variety of neural and statistical models that we obtain from these structured inputs. The deep learning models, ticker-aware models, and statistical baselines are evaluated using accuracy and also the F1 score.

3.1 Data Preparation

3.1.1 Data Preprocessing We begin by collecting historical daily stock price data for the NIFTY 50 constituents over a fixed time window. Missing values arising from trading holidays or corporate actions are forward-filled using the last available valid entry to maintain time series alignment across all stocks.

For the LSTM-based models, we generate a binary *movement label* column to represent the prediction target. The label is computed as:

$$y_t = \begin{cases} 1, & \text{if } (O_{t+1} - C_t) > 0 \\ 0, & \text{otherwise} \end{cases}$$

where C_t is the closing price on day t and O_{t+1} is the opening price on day $t+1$. This labeling scheme enables the models to learn directional movement prediction rather than raw price forecasting. All feature series(including the technical indicators used by all the models: RSI, EMA, Bollinger Bands and MACD) are standardized to zero mean and unit variance prior to similarity computation and model training. Additionally, rolling windows are applied to structure inputs for temporal models, with each sample consisting of a fixed-length sequence of past observations. This preprocessing ensures the dataset is clean, normalized, and temporally organized for both clustering and predictive modeling tasks.

3.1.2 Clustering via Co-Movement Similarity The first step involves grouping stocks based on historical co-movement patterns to reveal latent sectoral or behavioral structures. Given daily log returns $\mathbf{r}_i = [r_i^1, r_i^2, ..., r_i^T]$ for each stock i, a similarity matrix $S \in \mathbb{R}^{N\times N}$ is constructed using multiple time-series similarity metrics from prior literature. In contrast to standard approaches based on Pearson or Spearman correlation, we emphasize directional co-movement—specifically, how frequently two stocks move in the same direction on the same day. This offers a more discrete and behaviorally-aligned similarity measure and better captures structural alignment between financial instruments.

First, we define a movement direction vector for each stock:

$$m_i^t = \begin{cases} +1 & \text{if } r_i^t > 0 \\ -1 & \text{if } r_i^t < 0 \\ 0 & \text{if } r_i^t = 0 \end{cases} \quad \text{(movement direction for stock } i \text{ on day } t)$$

Then, the exponentially-weighted co-movement similarity between stocks i and j is computed as:

$$S_{ij} = \sum_{t=1}^{T} \lambda^{T-t} \cdot \delta(m_i^t = m_j^t) \quad \text{where } \delta(a = b) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{otherwise} \end{cases}$$

Here, $\lambda \in (0, 1)$ is a decay parameter controlling how much recent alignment is emphasized. This formulation captures the number of days on which stocks exhibit the same movement direction (i.e., both positive or both negative), while giving more weight to recent occurrences to avoid overfitting to outdated or noisy alignment. The similarity was normalized as:

$$\tilde{S}_{ij} = \frac{S_{ij}}{\sum_{t=1}^{T} \lambda^{T-t}} \quad \text{(normalized co-movement score between } i \text{ and } j\text{)}$$

To identify clusters without predefined sizes, agglomerative hierarchical clustering is employed. This method allows flexibility in grouping based on a dendrogram threshold rather than requiring a fixed number of clusters. Empirically, many resulting clusters aligned with sectoral divisions—e.g., banking, automotive, pharmaceutical, and energy—while some minor clusters showed cross-sectoral or noisy groupings. These clusters are validated using a secondary similarity metric from independent literature, confirmed their robustness, and this clustering process informs subsequent model architectures, enabling sector-aware forecasting through shared layers or structural embeddings (Fig. 1).

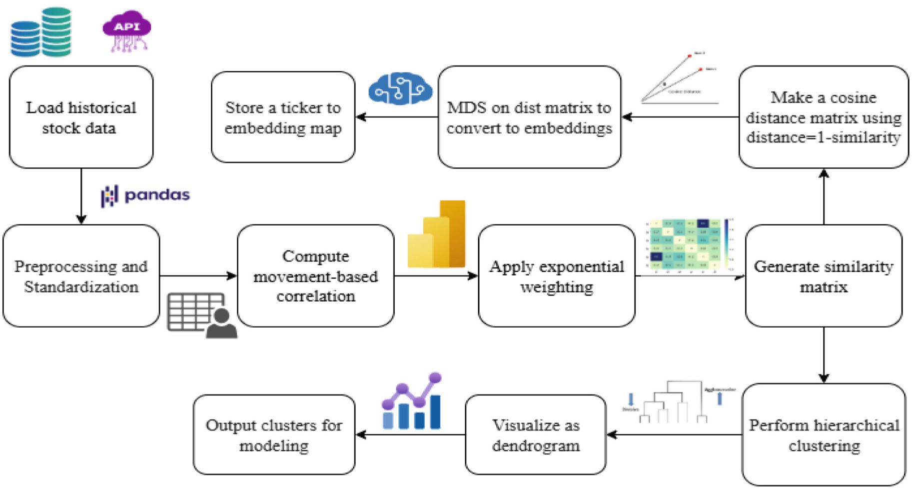

Fig. 1. Clustering via co-movement similarity.

3.2 Model Architecture

We evaluate five distinct models spanning deep learning, ticker-aware mechanisms, and statistical baselines. Each model is described below.

3.2.1 Cluster-Aware Dual LSTM (CA-DLSTM) This is our primary deep learning architecture for multi-stock forecasting. Stock features (basic and technical indicators) are extracted for all 46 stocks and concatenated row-wise:

$$[\mathrm{Open}_A, \mathrm{Open}_B, ..., \mathrm{Close}_A, \mathrm{Close}_B, ..., \mathrm{RSI}_A, \mathrm{RSI}_B, ...]$$

After sequence generation and scaling, this high-dimensional ($8 \times$ feature-set) time-series data are passed through a *TimeDistributed encoder layer* to reduce dimensionality while preserving temporal context (which performs better than PCA for time-series inputs due to temporal continuity).

The core temporal modeling is handled by two parallel LSTM encoders operating on different horizons: a short-term window ($T_1 = 30$) and a long-term window ($T_2 = 60$). Each sequence branch outputs a hidden representation:

$$\mathbf{h}_{\text{short}} = \mathrm{LSTM}_{T_1}(\mathbf{z}_{1:T_1}), \quad \mathbf{h}_{\text{long}} = \mathrm{LSTM}_{T_2}(\mathbf{z}_{1:T_2})$$

These representations are concatenated and passed through a shared dense fusion layer:

$$\mathbf{h}_{\text{global}} = \phi(\mathbf{W} \cdot [\mathbf{h}_{\text{short}} \| \mathbf{h}_{\text{long}}] + \mathbf{b})$$

Finally, individual dense heads are used for per-stock predictions:

$$\hat{y}_i = \sigma(\mathbf{W}_i \cdot \mathbf{h}_{\text{global}} + b_i)$$

Here, $\hat{y}_i$ is the predicted probability of upward movement for the i^{th} stock in the cluster, and $\phi(\cdot)$ is a non-linear activation (e.g., ReLU or tanh). The model is trained using a balanced binary cross-entropy loss across all heads. It achieves a strong average F1 score of 0.768, with a maximum of 0.87, outperforming several baseline approaches (Fig. 2).

Algorithm 1: Cluster-Aware Dual LSTM (CA-DLSTM)

Input: Historical price series of N stocks
Output: Upward movement probabilities $\hat{y}_i$ for each stock
Step 1: Preprocessing
Compute movement labels; handle missing data; standardize features
Step 2: Clustering
Construct directional co-movement similarity matrix S with exponential decay Apply agglomerative clustering to obtain stock groups $\{C_1, C_2, \dots\}$
foreach *cluster* C_k **do**
    Step 3: Sequence Modeling
    Generate sequences with technical indicators Encode to lower dimension with Time-Distributed Dense layer Extract temporal features via two LSTM branches (short-term, long-term) Fuse representations into a shared global layer
    Step 4: Prediction
    foreach *stock* $i \in C_k$ **do**
        Predict $\hat{y}_i$ using a stock-specific dense head
Step 5: Training
Optimize using balanced binary cross-entropy loss across all heads

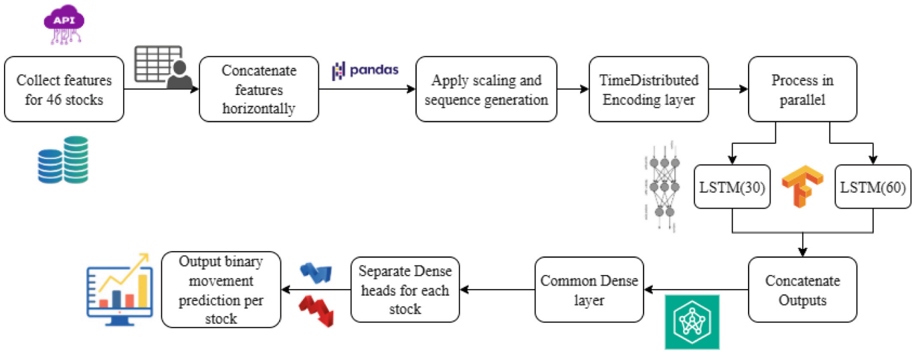

Fig. 2. Cluster-Aware Dual LSTM (CA-DLSTM)

3.2.2 Ticker-Embedded LSTM (TE-LSTM) This model introduces a *ticker-aware embedding layer* that provides the network with stock-specific identity information. The data is structured by concatenating sequences from all stocks vertically, producing a larger dataset where each row includes a *ticker identifier*.

Embedding Initialization: Each ticker is mapped to a trainable embedding vector, initialized using *Multidimensional Scaling (MDS)* applied to a distance matrix derived from co-movement-based similarity:

$$\mathrm{sim}(\mathbf{r}_i, \mathbf{r}_j) = \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\left(\mathrm{sign}(r_i^t) = \mathrm{sign}(r_j^t)\right), \quad D_{ij} = 1 - \mathrm{sim}(\mathbf{r}_i, \mathbf{r}_j)$$

$$\mathbf{E} = \mathrm{MDS}(D), \quad \mathbf{e}_i^{(0)} \in \mathbb{R}^d$$

This structurally-informed initialization ensures that each ticker embedding starts with a spatial layout consistent with sectoral co-movement. The temporal features (excluding ticker) for each sample are processed using an LSTM:

$$\mathbf{h}_i = \mathrm{LSTM}(\mathbf{x}_{1:T}^{(i)}), \quad \mathbf{h}_i \in \mathbb{R}^h$$

The temporal output is concatenated with the corresponding embedding and passed through fully connected layers:

$$\hat{y}_i = \sigma\left(\mathbf{W} \cdot [\mathbf{h}_i \,\|\, \mathbf{e}_i] + b\right)$$

Temporal Modeling: Despite its conceptual advantage, this model achieves a modest F1 score (∼0.72), suggesting that while stock identity is important, temporal modeling must also be guided by inter-stock relationships beyond embeddings alone (Fig. 3).

Algorithm 2: Ticker-Embedded LSTM (TE-LSTM)

Input: Historical sequences $\{\mathbf{x}_{1:T}^{(i)}\}_{i=1}^{N}$, ticker identifiers $\{id_i\}_{i=1}^{N}$
Output: Predicted movement probabilities $\{\hat{y}_i\}_{i=1}^{N}$

Embedding Initialization:
Construct similarity matrix S from directional agreement Compute distance matrix $D = 1 - S$ Obtain initial embeddings $\mathbf{e}_i^{(0)}$ using MDS on D

Temporal Encoding:
foreach *stock* $i = 1 \ldots N$ **do**
Encode $\mathbf{x}_{1:T}^{(i)}$ using LSTM to obtain hidden state $\mathbf{h}_i$

Fusion and Prediction:
foreach *stock* $i = 1 \ldots N$ **do**
Concatenate $\mathbf{h}_i$ with embedding $\mathbf{e}_i$ Compute output $\hat{y}_i = \sigma(\mathbf{W} \cdot [\mathbf{h}_i \,\|\, \mathbf{e}_i] + b)$

Training:
Optimize parameters using binary cross-entropy loss across all stocks

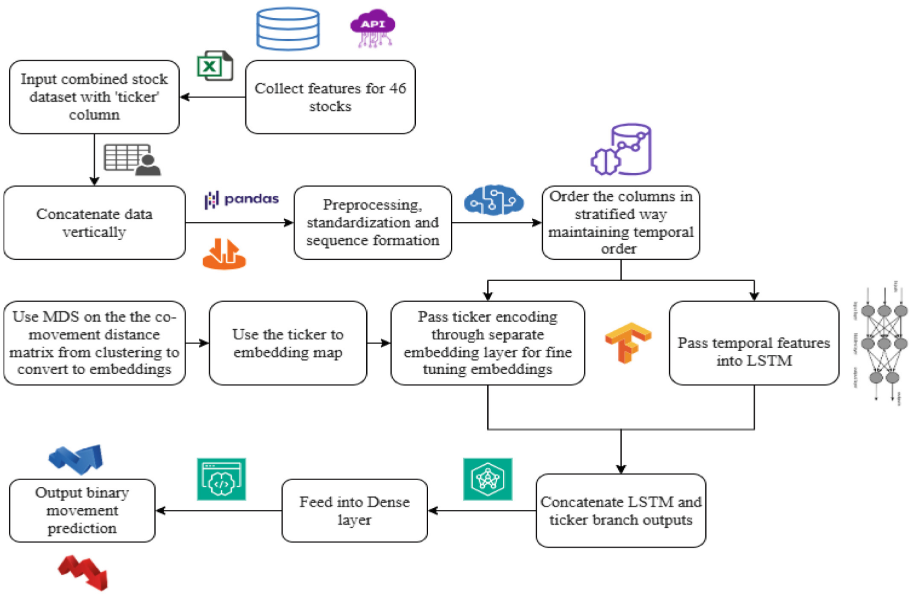

Fig. 3. Ticker-Embedded LSTM(TE-LSTM)

3.2.3 Genetic Programming-Based Rule Learning (Baseline) To construct an interpretable rule-based baseline, we apply *Genetic Programming (GP)* to evolve technical indicator combinations that best predict directional stock movements. The rules are learned from a population of expressions combining indicators such as RSI, MACD, and Bollinger Bands, with arithmetic and logical operators. Two strategies are evaluated by manually defined rule sets from literature and optimized rules via GP tuned specifically to the NIFTY 50 stock data. An example evolved rule takes the form: if RSI $> 70 \wedge$ MACD $< 0 \Rightarrow$ Sell. While rules evolved via GP show marginal improvement over static rule sets, both approaches yield accuracy and F1 scores in the 0.63–0.65 range. These results are competitive for non-deep models, but fall short of deep learning counterparts.

3.2.4 Vector AutoRegression (VAR) Baseline As a classical statistical model, *Vector AutoRegression (VAR)* is employed as a benchmark. Stocks from similar sectors (same clusters) are grouped, and VAR models are trained on *percentage returns* (preferred over log returns for empirical stability).

Rather than predicting raw prices, the model forecasts percentage returns:

$$\mathbf{r}_{t+1} = \mathbf{A}_1\mathbf{r}_t + \mathbf{A}_2\mathbf{r}_{t-1} + \cdots + \mathbf{A}_p\mathbf{r}_{t-p} + \boldsymbol{\varepsilon}_{t+1}$$

These are then converted to binary movement labels:

$$\hat{y}_{t+1} = \mathbb{1}(r_{t+1} > 0)$$

This formulation aligns with classification objectives and removes non-stationarity. Despite its simplicity, the VAR model performs comparably to rule-based baselines, achieving F1 scores between 0.64–0.66. Its consistency makes it a valid benchmark for evaluating modern deep learning models.

4 Experimental Results

To evaluate the effectiveness of our proposed model, we compare it against both internal baselines and recent external benchmarks from the literature. Specifically, we consider: (1) a cluster-based Vector Autoregression (VAR) model, (2) a cluster-level Genetic Algorithm (GA)-optimized rule-based system, (3) state-of-the-art multi-stock deep learning models such as DGDNN [14], Graph-WaveNet [18], DTML [15].

Each model is assessed on next-day stock movement prediction using two key metrics: **Accuracy** and **F1-score**. These metrics are standard in financial time series classification tasks [3,4].

Evaluation Metrics. Accuracy measures the overall proportion of correctly predicted labels:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \tag{1}$$

where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives respectively. The F1-score is the harmonic mean of precision and recall, capturing a balance between false positives and false negatives:

$$\text{F1-score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}, \tag{2}$$

with Precision $= \frac{TP}{TP+FP}$ and Recall $= \frac{TP}{TP+FN}$.

Table 1. Comparative Performance of Multi-Stock Prediction Models (Next-Day Directional Classification).

Model	Mean F1	Mean Acc.
Cluster-Aware Dual-LSTM	0.77	0.72
Ticker-Embedded LSTM	0.72	0.70
VAR Baseline	0.64	0.62
GA Rule Baseline	0.63	0.62
DGDNN [14]	0.63	0.65
GraphWaveNet [18]	0.60	0.59
DTML (Transformer) [15]	0.58	0.58
HMG-TF (Transformer) [16]	0.59	0.57
DGRCL (Dyn. Graph + CL) [17]	0.66	0.53

4.1 Comparative Performance

It is worth noting that several of the external benchmarks such as HMG-TF and DGRCL are evaluated on very large universes (e.g., NASDAQ-1000+ stocks), where the task is substantially more challenging due to higher volatility and the need to capture diverse inter-stock dependencies. By contrast, our framework focuses on 6–10 coherent clusters covering the 46 NIFTY 50 stocks. Despite operating on a smaller but sectorally diverse set, our model achieves consistently higher mean F1-scores and accuracy, highlighting the advantage of combining cluster-aware structure with multi-output prediction. In addition to the tabular summary, we provide comparative plots to illustrate the performance distribution across models:

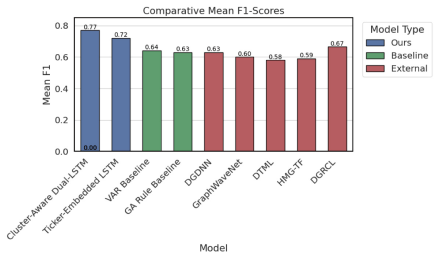

Fig. 4. Mean F1-scores of proposed models compared against statistical baselines and external benchmarks.

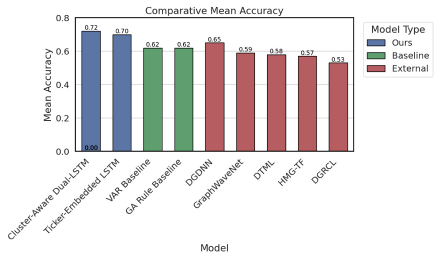

Fig. 5. Mean Accuracy of proposed models compared against statistical baselines and external benchmarks.

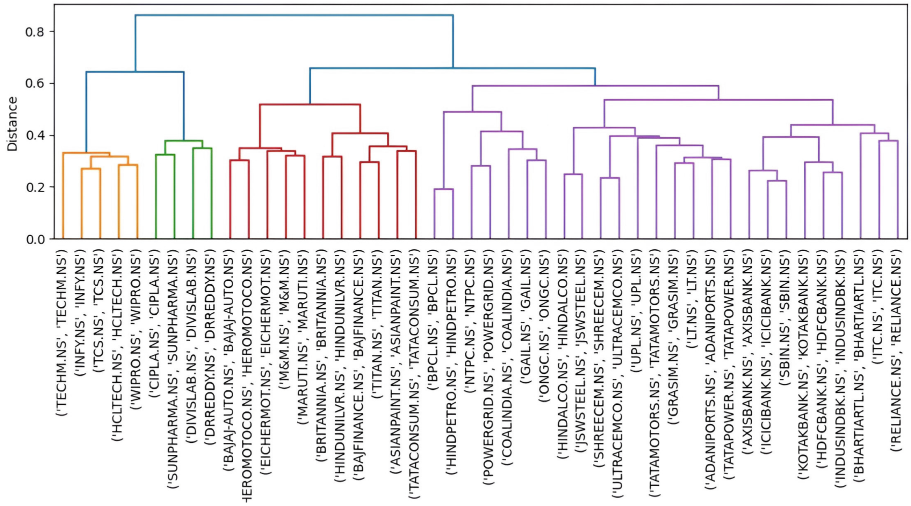

Fig. 6. Dendrogram of Nifty 50 stocks using movement direction similarity (correlation and co-movement based).

4.2 Cluster-Aware Insights

The clustering analysis validates that many stocks with sectoral ties (e.g., banking, automotive, pharmaceutical, energy) are grouped together, while a few small clusters emerge as cross-sectoral or noisy. This supports the motivation for designing a single model with shared representations and sector-aware heads. As shown in Table 1 and Figs. 4–5, the proposed CA-DLSTM and TE-LSTM architectures outperform both classical baselines (VAR, GA) and external deep learning benchmarks (e.g., GraphWaveNet, DTML, DA-RNN). The performance margin is more pronounced for clusters with strong intra-sector co-movement, demonstrating the benefit of leveraging sector-level dependencies. The hybrid architecture's ability to capture both global (sectoral) and local (stock-specific) temporal dynamics significantly improves robustness under volatile conditions.

Compared to single-stock models, the unified framework reduces overfitting and benefits from cross-asset representation sharing [3,7,8] (Fig. 6).

5 Conclusion

In this study, we presented a cluster-driven deep learning framework for synchronous, large-scale prediction of the direction of the next-day stock price movement within the NIFTY 50 index. Our method successfully captures asset-specific temporal dynamics as well as sectoral co-movement patterns by combining unsupervised clustering with sophisticated feature engineering and a hybrid recurrent architecture (CA-DLSTM) and (TE-LSTM). This design outperforms conventional baselines and allows for scalable multi-output forecasting. Strong generalizability is provided by the model's data-driven sector representation and modular architecture, which enable it to be applied to different asset universes and financial markets. Additionally, the framework provides opportunities for future improvements such as adding attention mechanisms, incorporating macroeconomic indicators, or expanding to include intraday high-frequency data. Overall, our results highlight the critical role of sector-aware multi-task learning and deep sequence modeling in advancing robust and interpretable financial forecasting systems.

Acknowledgements. The authors would like to express their sincere gratitude to the IEM Centre of Excellence for Cloud Computing and IoT for providing financial and infrastructural support through research grants IEMT(S)/2024/02-G19 and IEMT(S)/2023/02-G05 and IEDC-CSE at the Institute of Engineering and Management, Kolkata, for offering the technical resources and collaborative environment.

References

1. Dolphin, R., Smyth, B., Xu, Y., Dong, R.: Measuring financial time series similarity with a view to identifying profitable stock market opportunities. arXiv preprint arXiv:2107.03926 (2021)
2. Bjärkby, S.: A Cluster Analysis of Stocks to Define an Investment Strategy. Master's thesis, Lund University (2019)
3. Gao, Y., Wang, R., Zhou, E.: Stock prediction based on optimized LSTM and GRU models. J. Comput. Sci. Appl. (2021)
4. Kim, Y., Choi, H., et al.: Financial time-series forecasting using deep neural networks. Expert Syst. Appl. **173**, 114620 (2021)
5. Hema, K., Mounika, B., Mounesh, A., Santhosh Reddy, C., Mohan Babu, C.: Advanced stock market prediction using hybrid GRU-LSTM techniques. Int. J. Adv. Res. Innov. Ideas Educ. (IJARIIE) **11**(1), 1236–1242 (2025)
6. Li, Z., Wang, Y., et al.: A similarity measurement for time series and its application to the stock market. Knowl.-Based Syst. **227**, 107213 (2021)
7. Park, H.J., et al.: Stock market forecasting using a multi-task approach based on external events during the COVID-19 pandemic. Appl. Soft Comput. **113**, 107839 (2022)

8. Yuan, X., Wu, Q., Xu, W.: Stock index forecasting during COVID-19 pandemic using multi-task learning. J. Comput. Sci. **52**, 101344 (2021)
9. Sims, C.A.: Macroeconomics and reality. Econometrica **48**(1), 1–48 (1980)
10. Feng, X., Li, Y., Wang, J.: Improving S&P stock prediction with time series stock similarity. In: Proceedings of ICMLA (2019)
11. Zhao, Z., Wang, J., et al.: Stock market co-movement assessment using a three-phase clustering method. Appl. Soft Comput. **94**, 106450 (2020)
12. Chen, S.-H., Yeh, C.-H.: Using genetic programming to model volatility in financial time series. Appl. Finan. Econ. **12**(8), 543–552 (2002)
13. Toshniwal, S., Bansal, A., et al.: Multi-task learning for financial forecasting. arXiv preprint arXiv:1809.10336 (2018)
14. Liu, Z., Zhu, K., Wang, W., Zhou, J., Sun, J.: Decoupled graph diffusion neural network for stock movement prediction. arXiv preprint arXiv:2307.03062 (2023)
15. Wang, Z., Tang, J., Chang, Y.: DTML: a deep transformer-based multi-level contextual learning framework for stock movement prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence (2021)
16. Wu, X., Wang, J., Chen, H.: HMG-TF: A hierarchical multi-graph transformer for stock movement prediction. In: Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence, pp. 5311–5318 (2021)
17. Huang, Z., Zhou, Y., Lin, H.: Dynamic graph contrastive learning for stock movement prediction. Expert Syst. Appl. **260**, 125849 (2025)
18. Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., Zhang, C.: Graph WaveNet for deep spatial-temporal graph modeling. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 1907–1913 (2019)