Forecasting Market Volatility Through Dynamic Financial Networks

Pulikandala Nithish Kumar[1] Nneka Umeorah[1] (Presenter) Alex Alochukwu[2]

[1] Cardiff University, UK
[2] Albany State University, USA

Insurance Data Science Conference
June 9–10, 2026
Hannover, Germany

Volatility Clustering in the VIX

Key Observation:
Large volatility spikes are typically followed by further periods of elevated volatility, illustrating the phenomenon of volatility clustering.

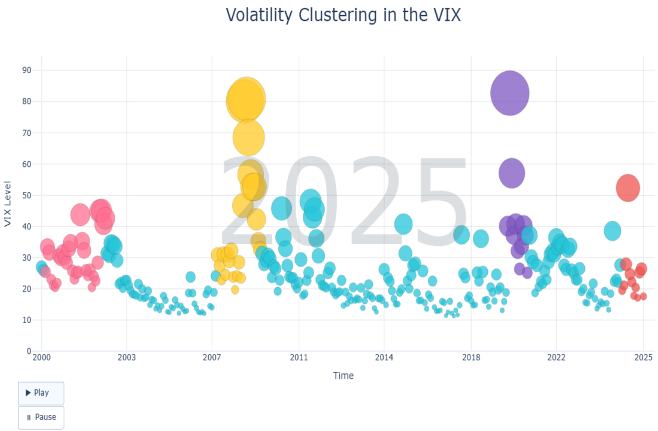

Volatility Clustering and Contagion

Why we need a temporal and graph-based framework

Empirical Evidence: Volatility Clustering

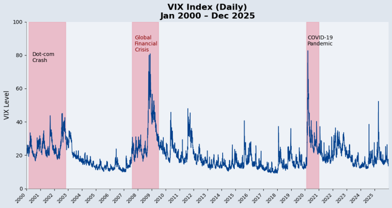

Volatility clusters over time and propagates across interconnected markets.

Implications

1. Volatility is Persistent
Large market movements are followed by further large movements, while calm periods tend to remain calm.

2. Markets are Interconnected
Financial shocks rapidly spill across markets.

3. Dependence Structures Evolve Over Time
Cross-market relationships strengthen during crises and weaken during stable periods.

Aim and Objectives

To develop a forecasting framework that jointly captures **volatility persistence** and **cross-market spillover effects** within a dynamic financial network.

Objectives

1. Construct dynamic financial networks using volatility spillover relationships among global equity markets.
2. Develop a temporal graph-based forecasting framework that integrates historical volatility dynamics with network information.
3. Evaluate forecasting performance across multiple forecasting horizons and market regimes.
4. Compare the proposed framework against classical econometric, deep learning, and graph-based benchmark models.

Existing Research

- ARCH/GARCH models effectively capture volatility clustering and conditional heteroskedasticity (Engle, 1982; Bollerslev, 1986).
- Volatility spillover frameworks measure directional shock transmission across financial markets (Diebold & Yilmaz, 2012, 2014).
- Machine learning approaches improve nonlinear volatility forecasting performance (Kim, 2003; Andersen et al., 2003).
- GNNs have recently been applied to model financial interconnectedness and spillovers (Kipf & Welling, 2017; Veličković et al., 2018).
- Temporal graph models have shown promise for modelling evolving financial dependencies (Son et al., 2023).

Our Contribution: We propose a graph-based deep learning framework for modelling market volatility using dynamic financial networks.

The framework jointly captures:

- volatility persistence over time,
- cross-market interconnectedness,
- and directional spillover transmission.

Market-Time Graph Representation

Dataset Structure

Component	Description
Markets	8 global equity indices [GSPC, GDAXI, FCHI, FTSE, NSEI, N225, KS11, HSI]
Period	November 2007 – June 2022 [Frequency » Daily]
Volatility Proxy	Squared daily log returns
Training Set	Nov. 2007 – Aug. 2014 (50%)
Validation Set	Sep. 2014 – Dec. 2017 (20%)
Test Set	Jan. 2018 – Jun. 2022 (30%)

Market-Time Structure: Each observation corresponds to one market index observed on one trading day:

$$(i, t) = \text{index } i \text{ at day } t.$$

For each index, rolling windows of past volatility proxies are used to predict future volatility.

Graph Structure: At each modelling period, the eight indices are treated as graph nodes, while edges are constructed using either correlation or Diebold–Yilmaz volatility spillovers. The resulting adjacency matrix is used by the TemporalGAT model to learn cross-market dependence.

Formulation

The Pearson correlation between markets i and j is defined as

$$\rho_{ij} = \frac{\mathrm{Cov}(X_i, X_j)}{\sigma_i \sigma_j}$$

where X_i and X_j denote the volatility proxy series for markets i and j, respectively.
The correlation adjacency matrix is then constructed as

$$A^{(\mathrm{corr})} = [\rho_{ij}]_{i,j=1}^{N},$$

where each entry $A_{ij}^{(\mathrm{corr})}$ is the correlation between i and j.

The correlation network captures symmetric market dependence and volatility co-movement across financial markets.

Correlation Graph

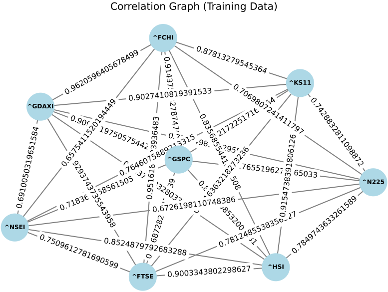

Graph Representation

- Nodes represent global stock indices.
- Edge weights correspond to pairwise correlations.
- Stronger edges indicate higher market co-movement.

Volatility Spillover Network

Diebold–Yilmaz Spillover Measure

Using a VAR model, generalized forecast error variance decomposition measures how much of market i's forecast uncertainty is explained by shocks from market j:

$$\theta_{ij}^{g}(H) = \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} \left(e_i^\top A_h \Sigma e_j \right)^2}{\sum_{h=0}^{H-1} \left(e_i^\top A_h \Sigma A_h^\top e_i \right)}$$

where H is the forecast horizon, A_h are VAR coefficient matrices, Σ is the covariance matrix of errors, and e_i, e_j are selection vectors.

The spillover adjacency matrix is constructed as

$$A^{(\mathrm{spill})} = \left[\widetilde{\theta}_{ij}^{g}(H) \right]_{i,j=1}^{N},$$

where each entry represents the volatility transmitted from market j to i.

A larger value of $A_{ij}^{(\mathrm{spill})}$ indicates stronger volatility transmission from market j to market i.

Volatility Spillover Network

Training Spillover Graph

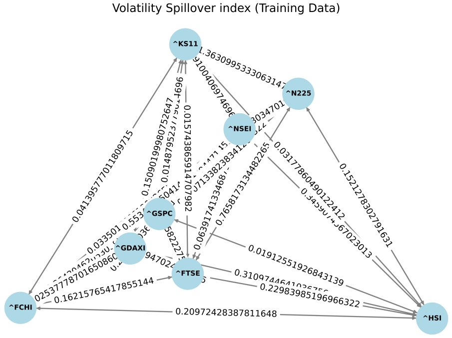

Interpretation

- Edge weights measure spillover intensity.
- Network captures directional volatility transmission and cross-market feedback effects.
- Edge directions reveal how volatility shocks propagate through the global financial system.

Highlight: Unlike correlation networks, spillover networks capture the direction of volatility transmission between markets. This directed network forms the primary graph structure used within the TemporalGAT framework.

TemporalGAT: Model Architecture

Architecture Flow

Volatility Sequences
Rolling volatility windows

↓

Temporal LSTM Module
Learns volatility persistence

↓

Temporal Node Embeddings
Hidden market representations

↓

Spillover Adjacency Matrix
Directional market links A_t

↓

GCN + GAT Layers
Propagate embeddings through graph

↓

Multi-Horizon Forecasts
$t+1,\ t+5,\ t+15,\ t+21$

Core Mathematical Components

1. Temporal Encoding

$$h_t = \mathrm{LSTM}(X_t, h_{t-1})$$

2. Graph Convolution

$$H^{(l+1)} = \sigma\left(\hat{A}H^{(l)}W^{(l)}\right)$$

3. Graph Attention

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k\in\mathcal{N}_i}\exp(e_{ik})}$$

4. Multi-Horizon Forecasting

$$\widehat{\mathbf{y}}_{i,t} = (\hat{y}_{i,t+1}, \hat{y}_{i,t+5}, \hat{y}_{i,t+15}, \hat{y}_{i,t+21})$$

Key Idea: The LSTM first summarizes each market's volatility history. The resulting node embeddings are then propagated through the spillover graph using GCN and GAT layers.

Model Hyperparameter Configuration

Grid Search and Model Selection

Search Space

- Hidden dim:[32, 64, 128]
- GAT attention heads: [4, 8]
- Learning rates: [0.0001, 0.001, 0.01]
- Total configurations: 18
- Training epochs: 70

Best Model Configuration

Hyperparameter	Value
Hidden Dimensions	64
Number of Heads	4
Learning Rate	0.001

Experimental Setup

- GPU training on Google Colab
- Data split: 50% / 20% / 30%
- Evaluation metrics: MAFE, MSE, MAPE, R^2
- Forecast horizons: $h = 1, 5, 15, 21$

Comparative Forecasting Benchmarks

Classical Econometric Benchmarks

GARCH(1,1): Linear volatility dynamics

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

- Models time-varying volatility using past shocks and past variance
- α: reaction to new information; β: volatility persistence

EGARCH(1,1): Asymmetric volatility dynamics

$$\log(\sigma_t^2) = \omega + \beta\log(\sigma_{t-1}^2) + \alpha\frac{|\varepsilon_{t-1}|}{\sigma_{t-1}} + \gamma\frac{\varepsilon_{t-1}}{\sigma_{t-1}}$$

- Models log-volatility (ensures positivity); γ: asymmetry parameter
- Captures leverage effects (negative shocks increase volatility more)

HAR-RV: Multi-scale volatility model

$$RV_{t+1} = \beta_0 + \beta_d RV_t + \beta_w RV_t^{(w)} + \beta_m RV_t^{(m)} + \varepsilon_t$$

- Uses realised volatility at daily, weekly, and monthly horizons
- Captures heterogeneous market behaviour across time scales

Comparative Forecasting Benchmarks

Benchmark Models - contd.

Category	**Model**	**Purpose**
Deep Learning	BM LSTM	Fully-connected neural network Temporal volatility forecasting
Graph-Based	SGNN-GATM DGNN-GATM C-TGATM GARCH-TGATM	Static graph representation Deep graph representation Correlation-based TemporalGAT GARCH-driven TemporalGAT
Proposed	**TGATM**	**Temporal encoding + dynamic spillover networks**

Key Difference

Unlike existing benchmarks that rely on **fixed graph structures**, TGATM learns **time-varying spillover networks** and combines them with **temporal attention** to capture evolving market dependencies for forecasting.

TGATM vs Classical Econometric Models

Indices	TGATM	GARCH(1,1)	EGARCH(1,1)	HAR-RV
MSE				
GSPC	0.02214	0.111975	0.030199	0.055457
GDAXI	0.00897	0.023510	0.032299	0.027697
FCHI	0.01093	0.033062	0.037769	0.038240
FTSE	0.01781	0.026167	0.021205	0.034271
NSEI	0.01740	0.032884	0.027441	0.036480
N225	0.00037	0.013279	0.012935	0.015714
KS11	0.01028	0.016869	0.016905	0.018984
HSI	0.00528	0.034410	0.031404	0.045763
Average	0.011647	0.036520	0.026270	0.034076
MAFE				
GSPC	4.90×10^{-4}	1.25×10^{-2}	9.12×10^{-4}	3.08×10^{-3}
GDAXI	8.05×10^{-5}	5.53×10^{-4}	1.04×10^{-3}	7.67×10^{-4}
FCHI	1.19×10^{-4}	1.09×10^{-3}	1.43×10^{-3}	1.46×10^{-3}
FTSE	3.17×10^{-4}	6.85×10^{-4}	4.50×10^{-4}	1.18×10^{-3}
NSEI	3.03×10^{-4}	1.08×10^{-3}	7.53×10^{-4}	1.33×10^{-3}
N225	1.39×10^{-7}	1.76×10^{-4}	1.67×10^{-4}	2.47×10^{-4}
KS11	1.06×10^{-4}	2.85×10^{-4}	2.86×10^{-4}	3.60×10^{-4}
HSI	2.79×10^{-5}	1.18×10^{-3}	9.86×10^{-4}	2.09×10^{-3}
Average	1.81×10^{-4}	2.20×10^{-3}	7.53×10^{-4}	1.31×10^{-3}

Key Result: TGATM achieves the lowest average MSE and MAFE, and consistently outperforms the classical econometric benchmarks.

Average Prediction Errors

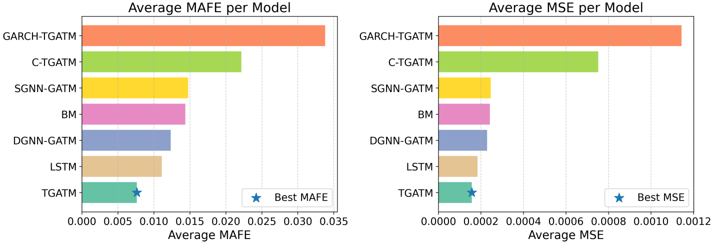

Key Result:
TGATM achieves the lowest average MAFE and MSE across competing graph and deep learning benchmarks, demonstrating the effectiveness of combining temporal encoding with spillover-aware graph attention learning.

TGATM Performance Across Market Regimes

Regime-Dependent Forecasting Performance

Average MAFE Across Regimes

Regime	1	5	15	21
COVID Crisis (Feb–Dec 2020)	0.0445	0.0470	0.0462	0.0320
Inflation Shock (Jan–Jun 2022)	0.0327	0.0360	0.0381	0.0227
Post-COVID Recovery (Jan–Dec 2021)	0.0143	0.0160	0.0154	0.0115
Pre-COVID Calm (Jan 2017–Jan 2019)	0.0100	0.0106	0.0108	0.0160

Columns report 1-, 5-, 15-, and 21-day ahead forecasts.

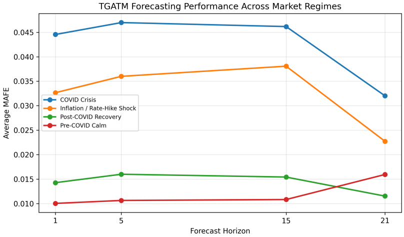

- Forecasting errors rise substantially during crisis regimes.
- COVID-19 exhibits the strongest forecasting difficulty.
- Calm market periods produce significantly lower forecast errors.
- Long-horizon volatility dynamics become more stable during stressed periods.

Robustness: Fixed vs Dynamic Spillover Graphs

Expanding-Window Forecasting Framework

Graph Specifications

- **Dynamic Graph:** Spillover network is re-estimated at each forecast origin using an expanding window.
- **Fixed Graph:** A single spillover network is estimated once and held fixed during expanding-window forecasting.
- **Original TGATM Setup:** Main experimental design used in the paper without expanding-window re-estimation.

Graph Type	Horizon	MAFE	MSE	RMSE	MAPE (%)
Dynamic Graph	1	0.022398	0.001277	0.024755	42.9318
Dynamic Graph	5	0.023415	0.001319	0.025755	45.9362
Dynamic Graph	15	0.024036	0.001393	0.026295	47.4470
Dynamic Graph	21	0.023514	0.001338	0.025778	45.3593
Fixed Graph	1	0.022390	0.001275	0.024749	42.9295
Fixed Graph	5	0.023435	0.001321	0.025773	45.9643
Fixed Graph	15	0.024089	0.001394	0.026349	47.5777
Fixed Graph	21	0.023508	0.001336	0.025772	45.3399
Baseline Graph	1	0.006305	0.000055	0.007442	21.4795
Baseline Graph	5	0.006034	0.000050	0.007104	27.8136
Baseline Graph	15	0.011800	0.000184	0.013566	46.3775
Baseline Graph	21	0.017903	0.000372	0.019276	69.4633

Fixed and dynamic graphs exhibit nearly identical forecasting performance, showing limited gains from continuously re-estimating the spillover network in the expanding-window setting.

Robustness: Diebold–Mariano Analysis

Statistical Comparison of Forecasting Accuracy

The Diebold–Mariano (DM) test evaluates whether forecasting errors from two competing models are statistically different.

DM Statistics (MSE Loss)

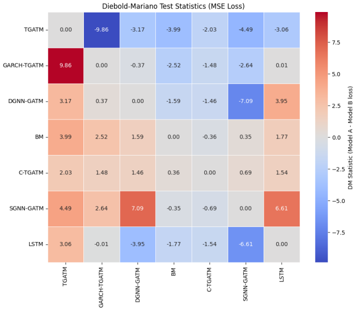

DM p-values (MSE Loss)

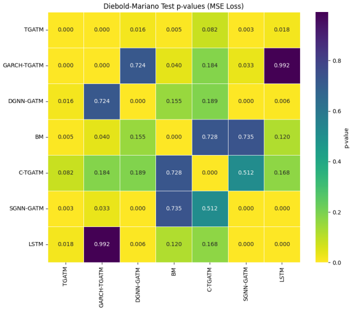

Large negative DM statistics and low p-values indicate that TGATM achieves significantly lower forecasting errors than competing graph and deep learning architectures.

Conclusion and Future Directions

Conclusion

- Proposed a TGATM for multi-horizon volatility forecasting.
- Integrated temporal sequence learning, graph convolution, graph attention, volatility spillover networks.
- TGATM outperformed classical econometric models, deep learning and static graph benchmarks.
- Results demonstrate strong regime-dependent forecasting behaviour.

Future Work

- Incorporating heterogeneous financial networks.
- Extending to intraday and high-frequency volatility forecasting.
- Integrating macroeconomic and sentiment-based spillover channels.

Temporal graph learning provides a powerful framework for modelling evolving financial contagion and volatility transmission.

References

Engle, R. F. (1982).
Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation.
Econometrica, 50(4), 987–1007.

Bollerslev, T. (1986).
Generalized Autoregressive Conditional Heteroskedasticity.
Journal of Econometrics, 31(3), 307–327.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003).
Modeling and Forecasting Realized Volatility.
Econometrica, 71(2), 579–625.

Kim, K.-j. (2003).
Financial Time Series Forecasting Using Support Vector Machines.
Neurocomputing, 55(1–2), 307–319.

Kipf, T. N., & Welling, M. (2017).
Semi-Supervised Classification with Graph Convolutional Networks.
In *International Conference on Learning Representations (ICLR)*.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018).
Graph Attention Networks.
In *International Conference on Learning Representations (ICLR)*.

Diebold, F. X., & Yilmaz, K. (2012).
Better to Give than to Receive: Predictive Directional Measurement of Volatility Spillovers.
International Journal of Forecasting, 28(1), 57–66.

Diebold, F. X., & Yilmaz, K. (2014).
On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms.
Journal of Econometrics, 182(1), 119–134.

Son, H., et al. (2023).
Spatial-Temporal Graph Neural Networks for Volatility Forecasting.
Expert Systems with Applications.

Thank You

Questions & Discussion