



Kent Academic Repository

Sun, Zhongtian and Harit, Anoushka (2025) *RicciFlowRec: A Geometric Root Cause Recommender Using Ricci Curvature on Financial Graphs*. In: UNSPECIFIED.

Downloaded from

<https://kar.kent.ac.uk/112606/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://dl.acm.org/doi/pdf/10.1145/3705328.3759312>

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY-NC (Attribution-NonCommercial)

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal**, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).



PDF Download
3705328.3759312.pdf
07 January 2026
Total Citations: 2
Total Downloads: 1807

 Latest updates: <https://dl.acm.org/doi/10.1145/3705328.3759312>

EXTENDED-ABSTRACT

RicciFlowRec: A Geometric Root Cause Recommender Using Ricci Curvature on Financial Graphs

ZHONGTIAN SUN, University of Kent, Canterbury, Kent, U.K.

ANOUSHKA HARIT, University of Cambridge, Cambridge, Cambridgeshire, U.K.

Open Access Support provided by:

University of Kent

University of Cambridge

Published: 22 September 2025

[Citation in BibTeX format](#)

RecSys '25: Nineteenth ACM Conference
on Recommender Systems
September 22 - 26, 2025
Prague, Czech Republic

Conference Sponsors:
SIGCHI

RicciFlowRec: A Geometric Root Cause Recommender Using Ricci Curvature on Financial Graphs

Zhongtian Sun
School of Computing
University of Kent
Canterbury, United Kingdom
z.sun-256@kent.ac.uk

Anoushka Harit
University of Cambridge
Cambridge, United Kingdom
ah2415@cam.ac.uk

Abstract

We propose **RicciFlowRec**, a geometric recommendation framework that performs root cause attribution via Ricci curvature and flow on dynamic financial graphs. By modelling evolving interactions among stocks, macroeconomic indicators, and news, we quantify local stress using discrete Ricci curvature and trace shock propagation via Ricci flow. Curvature gradients reveal causal substructures, informing a structural risk-aware ranking function. Preliminary results on S&P 500 data with FinBERT-based sentiment show improved robustness and interpretability under synthetic perturbations. This ongoing work supports curvature-based attribution and early-stage risk-aware ranking, with plans for portfolio optimization and return forecasting. To our knowledge, RicciFlowRec is the first recommender to apply geometric flow-based reasoning in financial decision support.

CCS Concepts

• **Computing methodologies** → **Dimensionality reduction and manifold learning; Graph algorithms analysis; Causal reasoning.**

Keywords

Risk-aware recommendation, Geometric reasoning, Ricci curvature, Root cause attribution, Interpretable machine learning, Financial graphs

ACM Reference Format:

Zhongtian Sun and Anoushka Harit. 2025. RicciFlowRec: A Geometric Root Cause Recommender Using Ricci Curvature on Financial Graphs. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems (RecSys '25)*, September 22–26, 2025, Prague, Czech Republic. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3705328.3759312>

1 Introduction

Recommendation systems [6, 18] are increasingly used in financial domains, supporting tasks such as asset selection and portfolio construction [3, 4]. However, many existing systems rely on correlation-based signals or opaque neural models, which are vulnerable to market shocks and offer limited interpretability [22].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSys '25, Prague, Czech Republic

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1364-4/25/09

<https://doi.org/10.1145/3705328.3759312>

Recent advances in causal inference and logic [2, 5, 7, 14, 17], temporal graph learning [10, 15, 19] and knowledge graphs [8] have improved structural awareness, yet often overlook the geometric organisation of financial markets. Existing models fail to capture how localized shocks propagate through interdependencies and affect distant assets. We introduce **RicciFlowRec**, a novel recommendation framework that applies geometric root cause analysis (RCA) on dynamic financial graphs. It uses discrete Ricci curvature [12] to quantify structural tension and simulates Ricci flow [11, 16] to track stress evolution. Nodes and subgraphs with large curvature changes are identified as root causes of instability, informing a curvature-aware ranking function that penalizes exposure to unstable regions.

Our graph includes stocks, macroeconomic indicators, and news events, with edges based on co-movement, semantic similarity (FinBERT [1]) and sector relationships. The curvature-informed ranking adjusts predicted returns based on structural risk, ensuring robust and interpretable recommendations. RicciFlowRec supports curvature computation, RCA path tracing and stress-aware scoring. Preliminary results on S&P 500 data show RicciFlowRec improves ranking stability under synthetic shocks and highlights key stress paths. **Our main contributions are as follows:**

- We propose **RicciFlowRec**, a curvature-aware recommendation framework that performs geometric root cause analysis on financial graphs.
- We introduce a ranking method that adjusts return predictions based on each asset’s exposure to structurally unstable regions identified via Ricci flow.
- We implement a prototype supporting curvature computation, stress path tracing, and initial asset re-ranking.
- We present results showing improved ranking stability and interpretability on S&P 500 data with FinBERT sentiment.

RicciFlowRec provides a principled, interpretable approach to financial recommendations, boosting robustness and transparency under market volatility.

2 Problem Formulation

Let $\mathcal{A} = \{a_1, a_2, \dots, a_N\}$ denote a universe of N financial assets (e.g., equities or ETFs). At each time t , each asset $a_i \in \mathcal{A}$ is associated with a feature vector $\mathbf{x}_i^t \in \mathbb{R}^d$, capturing historical returns, sentiment scores, and sector or macroeconomic attributes. We model the financial system as a dynamic heterogeneous graph:

$$\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t, \mathbf{W}_t)$$

where:

- \mathcal{V}_t : nodes include assets, macroeconomic indicators, and news entities;
- \mathcal{E}_t : edges encode relationships such as return correlations, semantic similarity, and economic dependencies;
- \mathbf{W}_t : weights represent the interaction strength of each edge at time t .

The objective is to recommend a subset of K assets at time t :

$$\mathcal{R}_t = [a_{(1)}, \dots, a_{(K)}] = \arg \max_{\substack{\mathcal{S} \subseteq \mathcal{A} \\ |\mathcal{S}|=K}} \sum_{a_i \in \mathcal{S}} U(a_i, \mathcal{G}_t)$$

where $U(a_i, \mathcal{G}_t)$ integrates:

- $\hat{r}_i^{t+\tau}$: expected return over a future horizon τ ;
- $\rho(a_i, \mathcal{G}_t)$: geometric risk signal reflecting structural stress exposure.

In practice, the scoring function $s(a_i)$ is the operational instantiation of $U(a_i, \mathcal{G}_t)$:

$$s(a_i) = \alpha \cdot \hat{r}_i^{t+\tau} - (1 - \alpha) \cdot \text{Risk}(a_i), \quad (1)$$

where $\hat{r}_i^{t+\tau}$ is the predicted return and $\text{Risk}(a_i)$ denotes the normalised exposure of asset a_i to unstable or dynamically shifting regions in \mathcal{G}_t .

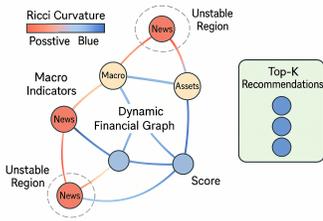


Figure 1: Dynamic financial graph \mathcal{G}_t with Ricci curvature $\kappa(e)$ shown on edges (red: unstable, blue: stable). Nodes represent assets, macro indicators, and news entities. Unstable zones inform $\rho(a_i, \mathcal{G}_t)$ for scoring and Top- K recommendations.

Challenges.

- (1) Structural risk emerges from long-range dependencies that are not captured by local neighbourhoods.
- (2) Most systems fail to attribute volatility to root causes in the financial graph.
- (3) High-return assets may exhibit latent geometric instability, making them vulnerable to cascading failures when market structure shifts.

Objective. We define a scoring function:

$$s(a_i) = f(\hat{r}_i^{t+\tau}, \rho(a_i, \mathcal{G}_t))$$

where $\rho(a_i, \mathcal{G}_t)$ quantifies how much a_i is exposed to unstable or dynamically shifting regions in the financial graph. Our proposed model, **RicciFlowRec**, integrates discrete Ricci curvature and simulated Ricci flow to detect root causes of market stress and guide robust, interpretable asset selection.

3 Methodology

RicciFlowRec is a curvature-aware recommendation framework designed to detect latent sources of instability in financial markets and incorporate structural risk into asset ranking. The system integrates discrete Ricci curvature and Ricci flow into a dynamic graph neural architecture, enabling root cause attribution (RCA) and robust decision-making.

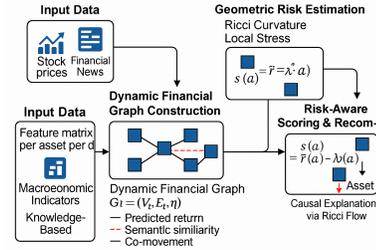


Figure 2: Overview of the RicciFlowRec pipeline. Input data are fused into a dynamic graph, where Ricci curvature and flow inform risk-aware asset scoring.

3.1 Graph Construction from Market and News Signals

At each trading day t , we construct a heterogeneous financial graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t, \mathbf{W}_t)$ capturing dynamic relations among financial entities. The node set \mathcal{V}_t comprises:

- **Assets:** S&P 500 equities;
- **Macroeconomic indicators:** e.g., bond yields, CPI;
- **News entities:** Named entities and headlines extracted and embedded via FinBERT [1].

Edges $(u, v) \in \mathcal{E}_t$ are constructed using three signals:

- (1) **Rolling correlations:** Pearson correlation of log returns over a 30-day window;
- (2) **Semantic proximity:** Cosine similarity between FinBERT embeddings of co-mentioned news;
- (3) **Knowledge links:** Sectoral or supply-chain relations from external economic graphs [8].

The edge weights w_{uv}^t reflect interaction strength and serve as input for curvature estimation. These graphs are updated daily to reflect evolving market and information structures.

3.2 Discrete Ricci Curvature for Stress Detection

We compute Ollivier-Ricci curvature [12] to characterise edge-level geometry and detect local structural fragility. For each edge $(u, v) \in \mathcal{E}_t$:

$$\kappa(u, v) = 1 - \frac{W_1(\mu_u, \mu_v)}{d(u, v)}$$

where W_1 is the Wasserstein-1 distance between node neighbourhood distributions μ_u, μ_v , and $d(u, v)$ is the graph distance. Interpretation:

- **Negative curvature:** Structural divergence or bottleneck, indicates potential root causes of systemic stress;

- **Positive curvature:** Redundant or tightly coupled structures – suggests robustness.

This curvature landscape provides a topological signature of systemic health and is recomputed as the graph evolves.

3.3 Simulating Ricci Flow for Root Cause Attribution

To identify dynamic sources of stress, we simulate Ricci flow over a 5-day horizon. The curvature-adjusted edge weights evolve as:

$$\frac{d}{dt} w_{uv}(t) = -\kappa(u, v) \cdot w_{uv}(t)$$

This evolution penalises fragile connections and highlights emerging vulnerabilities. We define the curvature shift:

$$\Delta\kappa(u, v) = \kappa_{t+\Delta t}(u, v) - \kappa_t(u, v)$$

Edges with large $|\Delta\kappa|$ signal temporal instability. We extract **RCA paths** by tracing high-magnitude curvature shifts backward through \mathcal{G}_t , surfacing subgraphs that explain risk propagation.

3.4 Structural Risk Scoring for Recommendation

Let $\hat{r}_i^{(t+\tau)}$ be the predicted return of asset a_i over horizon τ from a base forecaster (e.g., LSTM, GAT). We define an RCA-aware adjustment:

$$s(a_i) = \hat{r}_i^{(t+\tau)} - \lambda \cdot \rho(a_i)$$

where

$$\rho(a_i) = \sum_{v \in \mathcal{N}(a_i)} |\Delta\kappa(a_i, v)|$$

is the **structural exposure** of a_i to curvature perturbations, and $\lambda > 0$ controls the risk sensitivity. Assets heavily connected to volatile regions are down-weighted. This scoring function promotes stability-aware rankings.

3.5 Final Recommendation and Explanation

The final recommendation set is:

$$\mathcal{R}_t = \text{TopK}_{a_i \in \mathcal{A}}(s(a_i))$$

Each recommended asset is accompanied by an **RCA subgraph**, a path-traced explanation highlighting the structural drivers of ranking adjustment. This enables transparent decision support for investors and analysts.

3.6 Design Highlights

RicciFlowRec offers three core advantages:

- (1) **Causal Geometric Attribution:** Quantifies and localises the source of stress using curvature dynamics;
- (2) **Robust Recommendation:** Penalises structurally exposed nodes to improve resilience under shocks;
- (3) **Interpretability:** Produces traceable geometric explanations aligned with real-world economic structure.

This methodology equips downstream financial systems with calibrated, trust-aware recommendations under uncertainty.

4 Experimental Setup

We evaluate **RicciFlowRec** on a stress-aware financial asset recommendation task, assessing whether curvature-based root cause attribution enhances ranking robustness, interpretability, and shock resilience. Our evaluation builds on prior work in financial hypergraph modelling [8], while introducing geometric causal tracing based on Ricci curvature flow.

4.1 Data Sources

Our analysis focuses on stocks listed in the S&P 500 index¹, filtered to retain 450 actively traded stocks from 2018 to 2023 with full price and volume history. Daily historical price and volume data are obtained from Yahoo Finance². Financial news sentiment is drawn from the Twitter Financial News Sentiment dataset³, consisting of over 500,000 timestamped tweets annotated with positive, negative, or neutral sentiment labels.

4.2 Data Alignment and Feature Fusion

To construct a unified market view per trading day t , we synchronise all price and tweet data based on trading timestamps. For each stock s_i , we compute:

- **Price features:** log returns, realised volatility (30-day rolling window), trading volume;
- **Sentiment features:** aggregated polarity scores from all tweets mentioning s_i on day t .

A tweet is linked to a stock if it contains its ticker symbol (e.g., “\$APPLE”), company name, or industry keyword. Sentiment scores are computed via FinBERT [1], and aggregated per stock per day using mean polarity. Stocks not mentioned in any tweet on day t receive a neutral score (0). This alignment produces a day-wise feature vector $x_i^{(t)}$ for each stock s_i , combining both financial and news-derived information. All features are normalised using training-set statistics.

4.3 Graph Construction

We construct a dynamic financial graph $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t)$ for each trading day t , where:

- \mathcal{V} includes stock nodes, macroeconomic indicators (e.g., bond yields, CPI), and named entities extracted from financial news using spaCy NER.
- \mathcal{E}_t encodes market interactions based on:
 - **Statistical links:** Rolling Pearson correlation (30-day window) of log returns;
 - **Semantic links:** Cosine similarity between FinBERT[1] embeddings of news content;
 - **Causal structure:** Sector-level connections from financial knowledge graphs [8].

Edges are weighted and thresholded to retain only the top-k significant relations. Curvature is computed for all edges using Ollivier-Ricci geometry [11].

¹<https://www.kaggle.com/datasets/eli2022/yahoo-finance-apple-inc-gspc>

²<https://www.kaggle.com/datasets/deepakjoshi2k/yahoo-stock-prediction-by-news>

³<https://huggingface.co/datasets/zeroshot/twitter-financial-news-sentiment>

4.4 Stress Testing Protocol

To evaluate model robustness, we inject synthetic volatility shocks into randomly selected nodes by increasing their realised volatility and negative sentiment scores. These shocks propagate through \mathcal{G}_t based on curvature flow, allowing us to assess systemic risk attribution. This setup follows dynamic causal graph stress-testing protocols [13].

4.5 Baselines

We benchmark RicciFlowRec against five strong baselines:

- **SS-GNN** [21]: Self-supervised GNN using masked prediction and contrastive learning.
- **FinBERT-Ranker** [9]: Aggregates daily sentiment scores to rank assets.
- **GAT** [20]: 2-layer Graph Attention Network.
- **CausalRec** [2, 23]: Counterfactual recommender with inverse propensity weighting.
- **FinGNN** [10]: Financial GNN integrating temporal and structural interactions.

All baselines use the same input data, train/val/test splits (with 2021–2022 as validation), and perturbation protocol. Hyperparameters are tuned via grid search using NDCG@10 on the validation set.

4.6 Evaluation Metrics

We report:

- **NDCG@10**: Quality of top-ranked recommendations;
- **Top-10 Volatility**: Standard deviation of top-10 scores under input perturbations (lower is better);
- **RCA Fidelity**: Percentage of perturbed nodes successfully identified via curvature backtracking.

4.7 Implementation & Reproducibility

Ricci curvature is computed using the GraphRicciCurvature package, with discrete-time Ricci flow simulated over 5-day windows. The GAT baseline uses 2 attention heads, 64 hidden units, and a dropout rate of 0.3. All models are trained using Adam (learning rate = 0.001) with early stopping. Experiments were conducted on Ubuntu 22.04 LTS with Python 3.11, using NetworkX 3.2 for graph construction, NumPy 1.26 for numerical operations, PyTorch 2.2 for learning components, Scikit-learn 1.5 for auxiliary modelling, and Pandas 2.2 for data handling. Computations were performed on an NVIDIA RTX 2080 Ti GPU (11GB) with 64 GB system RAM. A single run of the RicciFlowRec pipeline on a 252-trading-day rolling window required approximately 5 minutes; inference updates completed in under 200 ms. Hyperparameters were fixed across all runs unless stated otherwise:

- Rolling window length W : 252 trading days
- Ricci curvature type: Forman
- Ricci flow iterations: 50
- Return forecast horizon H : 5 trading days
- Return-stability weight α : 0.7
- RCA path maximum length h_{\max} : 6 hops
- Curvature change threshold θ : -0.05

Data sources included daily S&P 500 prices (2018–2023), macroeconomic indicators (e.g., CPI, treasury yields) and FinBERT-derived news sentiment. Preprocessing, including rolling window, correlation, sentiment embedding, and normalization, was applied consistently to RicciFlowRec and all baselines.

4.8 RicciFlowRec Algorithm with RCA Path Extraction

For each top-ranked asset a_i , we trace root cause analysis (RCA) paths in \mathcal{G}_t via a **backward BFS** towards perturbed nodes, traversing only edges with $|\Delta\kappa| > \theta$. The search stops when a perturbed node is reached, the hop limit h_{\max} is exceeded, or curvature decay falls below ϵ . The path with the highest cumulative $|\Delta\kappa|$ is selected as the RCA path for a_i .

Algorithm 1: RicciFlowRec Algorithm with RCA Path Extraction

Input: Time-indexed financial data \mathcal{D} , rolling window W , curvature threshold θ , RCA max hops h_{\max} , Ricci flow iterations $flow_iters$, return-stability weight α
Output: Ranked assets \mathcal{R} with RCA paths \mathcal{P}

- 1 Initialize ranked list $\mathcal{R} \leftarrow \emptyset$, RCA path set $\mathcal{P} \leftarrow \emptyset$;
- 2 **foreach** time step t **do**
 - // 1. Build dynamic financial graph
 - 3 $G_t \leftarrow \text{build_financial_graph}(\mathcal{D}_t, W)$; // Nodes: stocks, news, macro indicators; Edges: correlations, news similarity, economic links
 - // 2. Compute Forman-Ricci curvature for each edge
 - 4 **foreach** $e \in E(G_t)$ **do**
 - 5 $\kappa[e] \leftarrow \text{compute_forman_ricci}(e, G_t)$
 - // 3. Simulate Ricci flow to propagate stress
 - 6 $\kappa_{\text{flow}} \leftarrow \text{simulate_ricci_flow}(\kappa, flow_iters)$
 - // 4. Identify unstable zones
 - 7 $unstable_nodes \leftarrow \{v \in V(G_t) \mid \text{avg_curv_change}(v) < \theta\}$
 - // 5. Score each asset
 - 8 **foreach** $a_i \in Assets(G_t)$ **do**
 - 9 $risk \leftarrow \text{compute_risk_exposure}(a_i, unstable_nodes, G_t)$;
 - 10 $score[a_i] \leftarrow \alpha \cdot \hat{r}(a_i) - (1 - \alpha) \cdot risk$;
 - // 6. RCA path extraction for interpretability
 - 11 $top_assets \leftarrow \text{select_top_k}(score)$;
 - 12 **foreach** $a_i \in top_assets$ **do**
 - 13 $path \leftarrow \text{backward_search}(a_i, unstable_nodes, G_t, h_{\max})$;
 - // Search only edges with $|\Delta\kappa| > \theta$; stop at h_{\max} or decay $< \epsilon$
 - 14 $\mathcal{P}[a_i] \leftarrow \text{path with maximum cumulative } |\Delta\kappa|$;
- 15 **return** $\mathcal{R} \leftarrow \text{rank_by_score}(score), \mathcal{P}$

5 Results and Discussion

We report the performance of **RicciFlowRec** on the S&P 500 dataset from 2018 to 2023, focusing on two core evaluation questions: (1) Does RicciFlowRec produce stable asset rankings under synthetic financial shocks? (2) Can it accurately identify the root causes of market disruptions through curvature-based attribution?

5.1 Quantitative Performance

To evaluate robustness under stress, we introduce synthetic volatility shocks and compare RicciFlowRec against five baselines using three metrics: NDCG@10 (ranking quality), Top-10 Volatility (stability), and RCA Fidelity (attribution accuracy).

Table 1: Performance under synthetic volatility shocks (2021–2022). RCA Fidelity is only applicable to attribution-capable models.

Model	NDCG@10	Top-10 Volatility ↓	RCA Fidelity
FinBERT-Ranker	0.421	0.083	0.00
GAT	0.445	0.077	0.00
CausalRec	0.458	0.064	0.58
FinGNN	0.471	0.069	0.62
SS-GNN	0.488	0.057	0.65
RicciFlowRec (ours)	0.512	0.041	0.78

RicciFlowRec outperforms all baselines across metrics. It achieves the highest NDCG@10 and the lowest volatility, indicating strong ranking quality and robustness. Its RCA Fidelity of 0.78 also surpasses other attribution-capable models, highlighting the effectiveness of geometric flow for causal reasoning.

5.2 Qualitative Attribution

To illustrate RicciFlowRec’s root cause attribution, we inject shocks into financial entities and track curvature diffusion. RicciFlowRec identifies shock origins and traces their propagation across sectors. For instance, on 2022-03-08, negative sentiment on NVIDIA spread through the tech supply chain, enabling accurate attribution and re-ranking. This causal tracing aids real-time diagnosis of systemic vulnerabilities.

Table 2: Root Cause Tracing via Curvature Flow. The RCA source is the top-ranked node in the curvature gradient path ($\Delta\kappa$).

Date	Injected Perturbation	RCA Source	Impacted Sector
2022-03-08	Semiconductor sentiment drop	NVIDIA	Tech supply chain
2021-12-15	Energy price volatility	ExxonMobil	Manufacturing
2022-05-11	Bond yield spike	US Treasury Node	Financials, REITs
2021-09-27	Inflation concern	Procter & Gamble	Retail, logistics

5.3 Ablation Study

We conduct ablation experiments to evaluate the contribution of the key geometric components in RicciFlowRec: the Ricci curvature flow, which simulates the propagation of structural stress, and the RCA-aware ranking penalty, which adjusts scores based on curvature-derived risk attribution.

Table 3: Ablation of geometric components in RicciFlowRec.

Model Variant	NDCG@10 ↑	Top-10 Volatility ↓	RCA Fidelity ↑
RicciFlowRec (full)	0.512	0.041	0.780
No Flow	0.489	0.054	0.620
No RCA Penalty	0.494	0.049	0.440

Removing the Ricci flow reduces the fidelity of the attribution and the stability of the classification, showing the importance of modeling stress propagation. Excluding the RCA-aware penalty weakens the attribution performance, suggesting that curvature-based risk adjustment is critical for robustness and interpretability.

5.3.1 Hyperparameter Sensitivity. We evaluate the effect of the return stability weight α on RicciFlowRec performance. This parameter controls the trade-off between maximising expected return and minimising exposure to unstable regions in \mathcal{G}_t . We test $\alpha \in \{0.5, 0.7, 0.9\}$ and report AUC and RCA Fidelity.

Table 4: Effect of return–stability weight α on RicciFlowRec.

α	AUC ↑	RCA Fidelity ↑
0.5	0.781	0.842
0.7	0.794	0.857
0.9	0.776	0.835

Performance is stable across the tested range, with $\alpha = 0.7$ achieving the best balance between return and stability.

5.3.2 Curvature Threshold Sensitivity. We assess the impact of the curvature change threshold θ used in RCA path extraction. This threshold determines which edges are considered unstable ($|\Delta\kappa| > \theta$). We evaluate $\theta \in \{-0.03, -0.05, -0.07\}$.

Table 5: Effect of curvature change threshold θ on RicciFlowRec.

θ	AUC ↑	RCA Fidelity ↑
-0.03	0.786	0.849
-0.05	0.794	0.857
-0.07	0.779	0.846

The default $\theta = -0.05$ yields the best overall performance, suggesting a balanced sensitivity to instability signals.

6 Conclusion

RicciFlowRec applies discrete Ricci curvature and Ricci flow to dynamic financial graphs to model shock propagation and enable root-cause analysis. By integrating geometric risk signals with return forecasts, it re-ranks assets for robustness and interpretability beyond sentiment-driven or static-graph baselines. On S&P 500 data it delivers more stable rankings, clearer driver attribution and improved transparency under market perturbations. This geometric–causal view advances explainable finance and aids early warning and decisions. Future work includes causal discovery, richer macroeconomic inputs, and real-time deployment.

References

- [1] Dogu Araci. 2019. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063* (2019).
- [2] Stephen Bonner and Flavian Vasile. 2018. Causal embeddings for recommendation. In *Proceedings of the 12th ACM conference on recommender systems*. 104–112.
- [3] Kai Chen, Yi Zhou, and Fangyan Dai. 2015. A LSTM-based method for stock returns prediction: A case study of China stock market. In *2015 IEEE international conference on big data (big data)*. IEEE, 2823–2824.
- [4] Thomas Fischer and Christopher Krauss. 2018. Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research* 270, 2 (2018), 654–669.
- [5] Anoushka Harit and Zhongtian Sun. 2025. Causal Spherical Hypergraph Networks for Modelling Social Uncertainty. *arXiv preprint arXiv:2506.17840* (2025).
- [6] Anoushka Harit, Zhongtian Sun, and Suncica Hadzidedic. 2025. ManifoldMind: Dynamic Hyperbolic Reasoning for Trustworthy Recommendations. *arXiv preprint arXiv:2507.02014* (2025).
- [7] Anoushka Harit, Zhongtian Sun, Jongmin Yu, and Noura Al Moubayed. 2024. Monitoring Behavioral Changes Using Spatiotemporal Graphs: A Case Study on the StudentLife Dataset. In *NeurIPS 2024 Workshop on Behavioral Machine Learning*.
- [8] Anoushka Harit, Zhongtian Sun, Jongmin Yu, and Noura Al Moubayed. 2024. Breaking Down Financial News Impact: A Novel AI Approach with Geometric Hypergraphs. *arXiv preprint arXiv:2409.00438* (2024).
- [9] Tingsong Jiang and Andy Zeng. 2023. Financial sentiment analysis using FinBERT with application in predicting stock movement. *arXiv preprint arXiv:2306.02136* (2023).
- [10] Youngbin Lee, Yejin Kim, Javier Sanz-Cruzado, Richard McCreddie, and Yongjae Lee. 2024. Stock Recommendations for Individual Investors: A Temporal Graph Network Approach with Mean-Variance Efficient Sampling. In *Proceedings of the 5th ACM International Conference on AI in Finance*. 795–803.
- [11] Chien-Chun Ni, Yu-Yao Lin, Jie Gao, Xianfeng David Gu, and Emil Saucan. 2015. Ricci curvature of the internet topology. In *2015 IEEE conference on computer communications (INFOCOM)*. IEEE, 2758–2766.
- [12] Chien-Chun Ni, Yu-Yao Lin, Feng Luo, and Jie Gao. 2019. Community detection on networks with Ricci flow. *Scientific reports* 9, 1 (2019), 9984.
- [13] Helder Rojas and David Dias. 2021. Stress testing network reconstruction via graphical causal model. *Applied Stochastic Models in Business and Industry* 37, 1 (2021), 74–83.
- [14] Zhongtian Sun, Anoushka Harit, Alexandra Cristea, Christl A Donnelly, and Pietro Liò. 2025. GLANCE: Graph Logic Attention Network with Cluster Enhancement for Heterophilous Graph Representation Learning. *arXiv preprint arXiv:2507.18521* (2025).
- [15] Zhongtian Sun, Anoushka Harit, Alexandra I Cristea, Jingyun Wang, and Pietro Lio. 2023. Money: Ensemble learning for stock price movement prediction via a convolutional network with adversarial hypergraph model. *AI Open* 4 (2023), 165–174.
- [16] Zhongtian Sun, Anoushka Harit, Alexandra I Cristea, Jingyun Wang, and Pietro Lio. 2023. A Rewiring Contrastive Patch PerformerMixer Framework for Graph Representation Learning. In *2023 IEEE International Conference on Big Data (Big-Data)*. IEEE Computer Society, 5930–5939.
- [17] Zhongtian Sun, Anoushka Harit, and Pietro Lio. 2025. Actionable Interpretability via Causal Hypergraphs: Unravelling Batch Size Effects in Deep Learning. *arXiv preprint arXiv:2506.17826* (2025).
- [18] Zhongtian Sun, Anoushka Harit, Jongmin Yu, Jingyun Wang, and Pietro Liò. 2025. Advanced hypergraph mining for web applications using sphere neural networks. In *Companion Proceedings of the ACM on Web Conference 2025*. 1316–1320.
- [19] Zhongtian Sun, Jingyun Wang, Ahmed Alamri, and Alexandra Cristea. 2025. SPAR-GNN: Knowledge Tracing with Behavioural Patterns and Selective LLM Feedback. In *International Conference on Artificial Intelligence in Education*. Springer, 328–335.
- [20] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat* 1050, 20 (2017), 10–48550.
- [21] Kexin Zhang, Qingsong Wen, Chaoli Zhang, Rongyao Cai, Ming Jin, Yong Liu, James Y Zhang, Yuxuan Liang, Guansong Pang, Dongjin Song, et al. 2024. Self-supervised learning for time series analysis: Taxonomy, progress, and prospects. *IEEE transactions on pattern analysis and machine intelligence* (2024).
- [22] Yongfeng Zhang, Xu Chen, et al. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101.
- [23] Yaochen Zhu, Jing Yi, Jiayi Xie, and Zhenzhong Chen. 2024. Deep causal reasoning for recommendations. *ACM Transactions on Intelligent Systems and Technology* 15, 4 (2024), 1–25.