NEURAL INFORMATION PROCESSING SYSTEMS

FAME: Forecasting Academic Impact via Continuous-Time Manifold Evolution

The Chinese University of Hong Kong

Jianrong Ding, Jianyuan Zhong, Zhengyan Shi, Qiang Xu

香港中文大學
The Chinese University of Hong Kong

Background: Scientific Evaluation

- LLMs are increasingly used to generate and evaluate research ideas.
- A fluent critique does not establish whether a manuscript will become influential.
- **FAME asks how a paper aligns with the evolving direction of its field.**

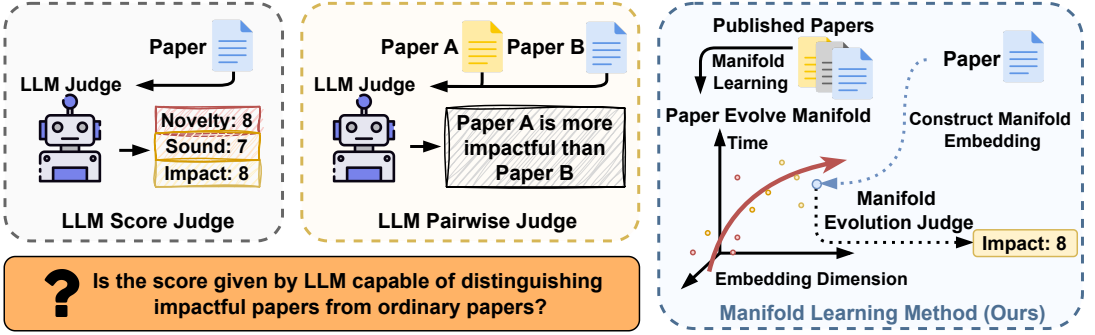



Task and Evaluation

- **Data:** 3,200 arXiv papers in Time Series Forecasting, Sparse Attention, and Image Generation.
- **Target:** a composite of four log-normalized signals: GitHub stars, citations, influential citations, and Altmetric scores.
- **Temporal split:** 18 monthly cutoffs from June 2024 to November 2025. Train on papers published by T and test on the next two months.
- **Comparison:** ML predictors, LLM judges with 30 historical paper–label examples, and SciJudge-4B. Report Spearman ρ_s over three runs.
- **Scope:** historical labels are finalized at T_{end}. This is a calibrated temporal split, not strict real-time forecasting.

Main Method: FAME

- Verify substantive inspiration between cited historical papers.
- Learn paper embeddings and continuous topic trajectories in a shared latent space.
- Forecast impact from a paper's alignment with topic momentum.

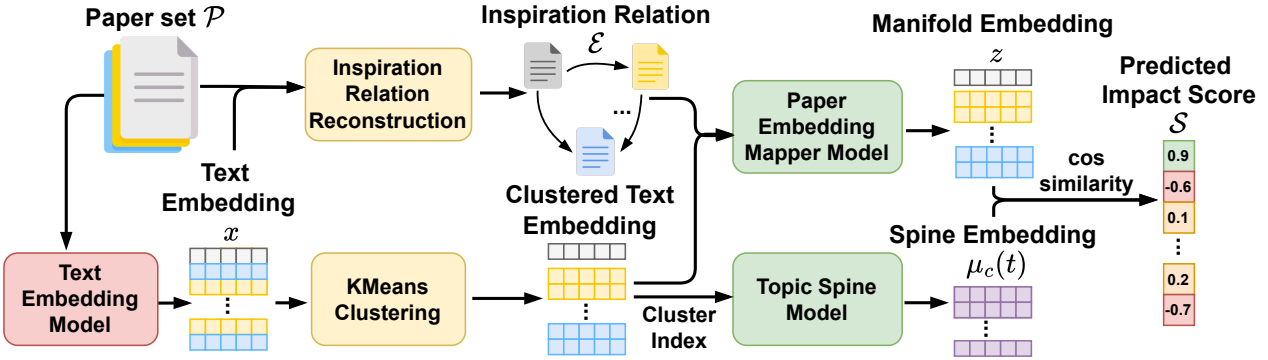

Learning the Manifold

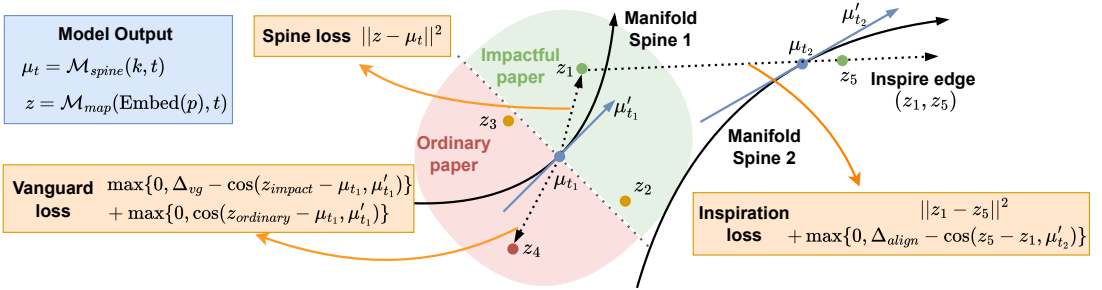

- **Spine binding**: anchor papers to their topic at publication time.
- **Inspiration**: align knowledge-flow edges with topic momentum.
- **Vanguard**: align locally high-impact papers with the forward direction.

$$\mathcal{L} = \alpha\mathcal{L}_{\text{spine}} + \beta\mathcal{L}_{\text{inspire}} + \gamma\mathcal{L}_{\text{vanguard}} \qquad (\alpha, \beta, \gamma) = (1, 0.1, 1).$$



Impact Inference

- Assign the unseen paper to a topic using the learned KMeans centroids.
- Map its content and publication time to the manifold.

$$\mathcal{S}(p) = \cos\left(\underbrace{z - \mu_c(t)}_{\text{paper residual}}, \quad \underbrace{\mu_c'(t)}_{\text{topic momentum}} \right)$$

- Higher angular alignment gives a higher predicted impact score.
- The topic tangent is estimated by a finite difference in continuous time.

Main Results: Impact Ranking

- FAME achieves mean Spearman correlation above 0.50 in all three periods.
- Both static predictors and text-based LLM judges have lower average correlations.

Evaluator	Jun–Nov 2024	Dec 2024–May 2025	Jun–Nov 2025
SVR	0.266	0.399	0.297
Transformer	0.278	0.321	0.294
TGCN	0.260	0.250	0.288
GPT-4.1-mini	0.205	0.126	0.145
DeepSeek-V3.2	–	0.210*	0.191
SciJudge-4B	–	0.223*	0.247
FAME	**0.535**	**0.513**	**0.533**

Means computed from the paper's monthly entries. *Five available months. Three runs per month.

Ablation: Geometric Supervision

- All three objectives contribute. Removing vanguard supervision causes the largest loss-ablation drop.

Variant	Jun–Nov 2024	Dec 2024–May 2025	Jun–Nov 2025
MLP + temporal embedding	0.277	0.332	0.346
Without spine loss	0.268	0.366	0.392
Without inspiration loss	0.210	0.335	0.318
Without vanguard loss	0.156	0.198	0.284
Full FAME	**0.535**	**0.513**	**0.533**

FAME as Context for LLM Evaluators

- Supplying manifold scores as explicit context improves every evaluated LLM.

LLM evaluator	Baseline ρ_s	+ FAME	Gain
GPT-4.1-mini	0.158	**0.317**	+0.159
GPT-4.1	0.070	**0.360**	+0.290
Gemini-3-Flash-Preview	0.190	**0.403**	+0.213
Gemini-3.1-Pro-Preview	0.169	**0.369**	+0.200
DeepSeek-V3.2	0.199	**0.382**	+0.183
GPT-5.2	0.179	**0.394**	+0.215

Reported means over each model's evaluated months. Compare the baseline and augmented version within each row.

Geometry of Impact

- Raw semantic embeddings entangle high-impact and ordinary papers.
- FAME organizes papers around evolving topic spines.

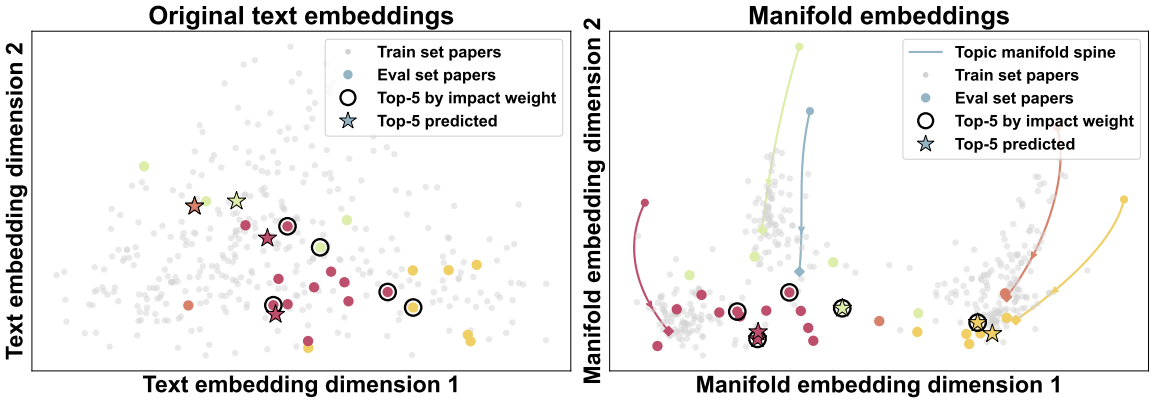

Stars: predicted top-5 papers. Circles: true top-5 papers. KernelPCA projections are illustrative rather than a quantitative distance test.

Conclusion and Scope

- **Model scientific evolution.** FAME combines verified inspiration with continuous topic trajectories.
- **Forecast with geometry.** Directional alignment improves manuscript impact rankings and supplies useful context to LLM evaluators.
- **Next steps.** Test time-truncated labels, early-stage hypotheses, and broader disciplinary settings.

Field trajectories are useful signals for scientific evaluation.



Thank you!

jrding25@cse.cuhk.edu.hk