Liyuan CHEN, Gaoguo JIA, Dongsheng GU, Jiangpeng YAN, Yuhang JIANG, Xiu LI, Xiaojun ZENG, 2025. MENTOR: a multi-agent framework for event and narrative trend prediction with optimized reasoning. *Frontiers of Information Technology & Electronic Engineering*, 26(10):1847-1861. https://doi.org/10.1631/FITEE.2500608

MENTOR: a multi-agent framework for event and narrative trend prediction with optimized reasoning

Key words: Narrative economics; Multi-agent; Event detection; Event forecasting

Corresponding author: Xiaojun ZENG
E-mail: x.zeng@manchester.ac.uk
ORCID: https://orcid.org/0000-0002-2320-2495

Motivation

- Financial markets are profoundly influenced by evolving narratives that shape investor sentiment and market behavior. However, predicting future events and quantifying their economic impacts from these narratives remain challenging.
- Bridging narrative economics with large language model reasoning is crucial to achieve more structured, interpretable, and data-driven forecasting. It is of paramount importance to design a framework that can integrate multi-agent collaboration and iterative reasoning to uncover the link between narrative dynamics and financial outcomes.

Main idea

- To operationalize narrative economics and provide intelligent assistance for financial forecasting, we propose MENTOR, a multi-agent framework for event and narrative trend prediction.
- In MENTOR, we introduce key tasks including trending event detection, future event forecasting, and industry index ranking, optimized through a teacher–student iterative reasoning mechanism.
- MENTOR effectively links narrative dynamics to measurable market outcomes, enhancing forecasting accuracy and interpretability across diverse financial datasets.

Framework

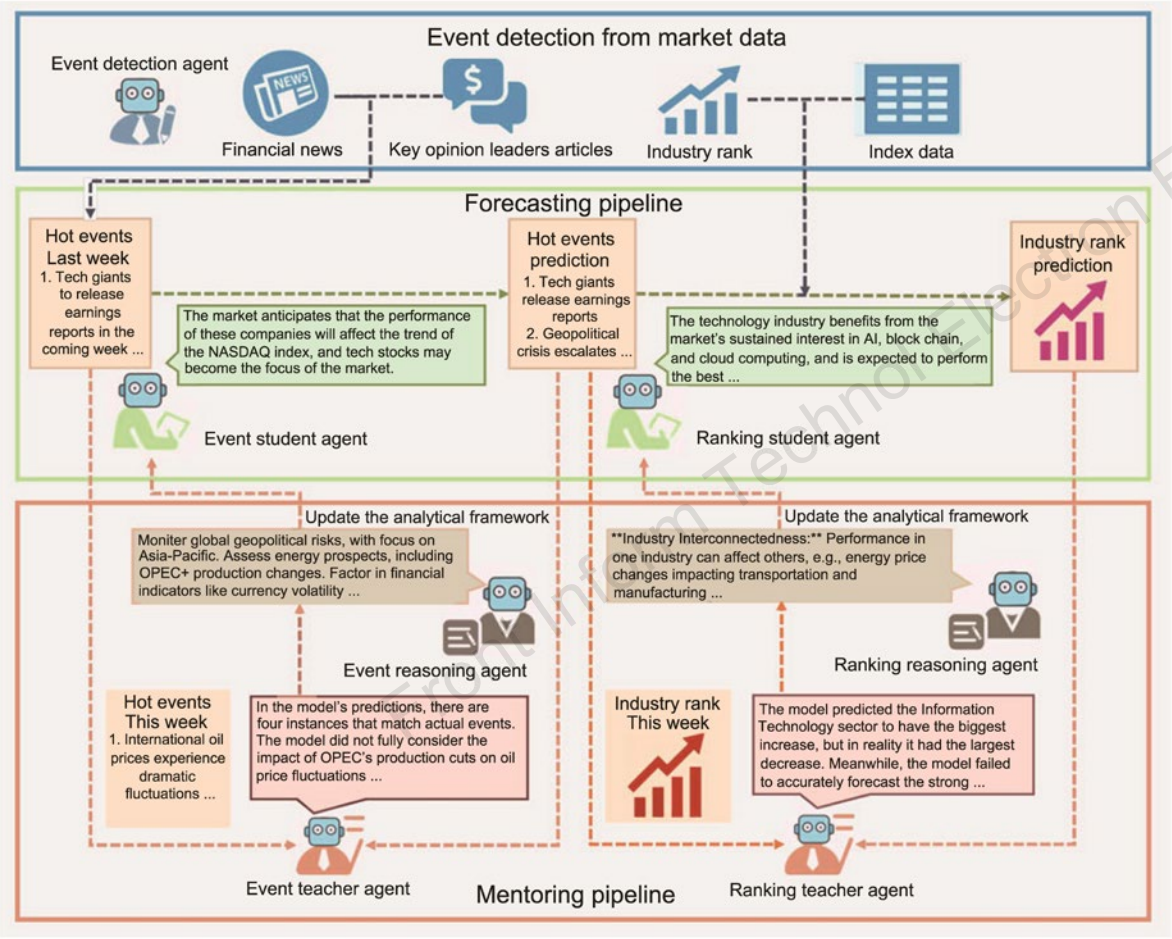

MENTOR applies narrative economics via structured reasoning. It decomposes financial forecasting into three sequential subtasks: trending event detection, future event prediction, and industry index ranking.

Overview of the iterative process in the MENTOR framework

Method

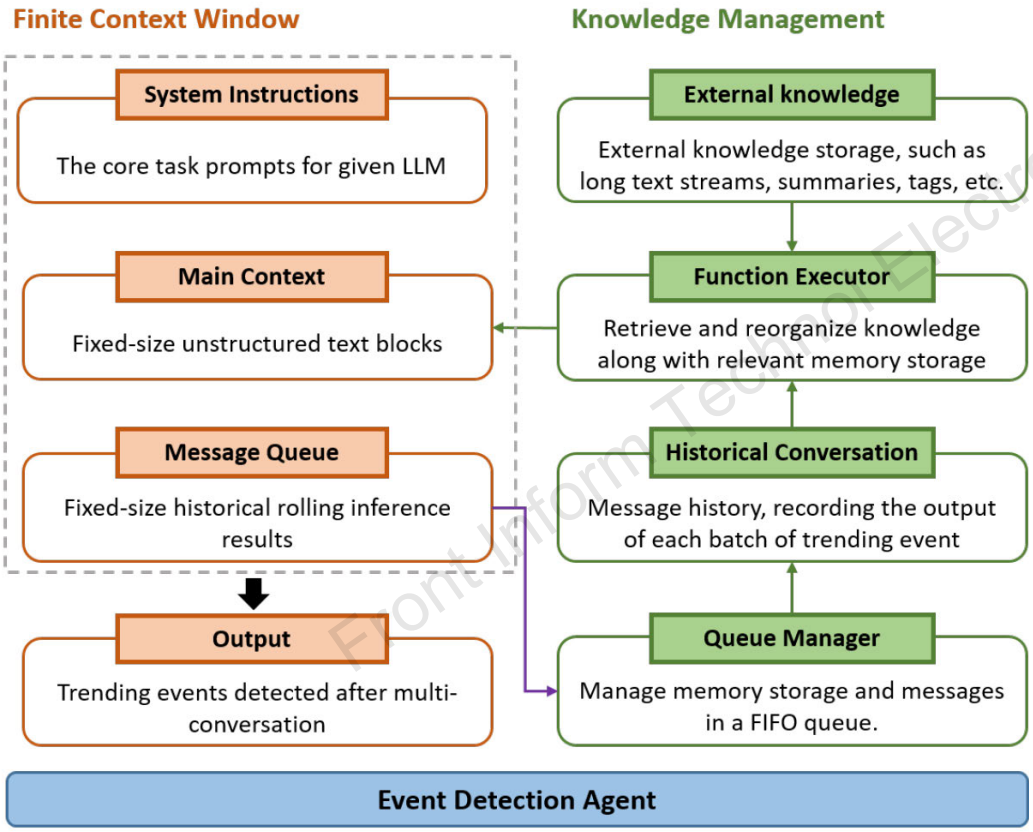

Workflow of the event detection agent

1. Trending event detection: Extract and rank trending events from large-scale narrative data

Method

2. Future event prediction and industry index ranking: Building on detected trending events, our framework forecasts future events likely to emerge from current narrative constellations in financial markets. Rather than treating prediction as simple classification, we extrapolate narrative trajectories based on the principle that new narratives typically evolve from existing ones through association or amplification. These predicted events are then mapped to sector-level financial impacts, generating ranked lists of industries expected to benefit or suffer in the next time period. This two-step approach creates a reliable bridge between narrative dynamics and measurable market outcomes.

Method

Algorithm 1 MENTOR: rolling strategy evolution

Require: weekly news $\{\mathcal{N}^{(n)}\}$ and weekly returns $\{\boldsymbol{R}^{(n)}\}$ for $n = 1, 2, \ldots, N$

Ensure: predictions $\hat{\mathcal{H}}^{(n)}, \hat{\boldsymbol{R}}^{(n)}$ for each weekly time step n

$\mathcal{S}_{\text{event}} \leftarrow \mathcal{S}_{\text{event}}^{(0)}$

$\mathcal{S}_{\text{ranking}} \leftarrow \mathcal{S}_{\text{ranking}}^{(0)}$

for each time step $n = 1, 2, \ldots, N$ **do**

  $\mathcal{H}_{\text{last}} \leftarrow \delta(\mathcal{N}^{(n-1)})$

  $\boldsymbol{R}_{\text{last}} \leftarrow \boldsymbol{R}^{(n-1)}$

  $\hat{\mathcal{H}}^{(n)} \leftarrow \mathcal{S}_{\text{event}}(\mathcal{H}_{\text{last}})$

  $\hat{\boldsymbol{R}}^{(n)} \leftarrow \mathcal{S}_{\text{ranking}}(\hat{\mathcal{H}}^{(n)}, \boldsymbol{R}_{\text{last}})$

  Observe ground truth $\mathcal{H}^{(n)}$, $\boldsymbol{R}^{(n)}$

  $\Delta\mathcal{S}_{\text{ranking}} \leftarrow \mathcal{T}_{\text{ranking}}(\hat{\mathcal{H}}^{(n)}, \mathcal{H}^{(n)}, \hat{\boldsymbol{R}}^{(n)}, \boldsymbol{R}^{(n)}, \mathcal{S}_{\text{ranking}})$

  $\Delta\mathcal{S}_{\text{event}} \leftarrow \mathcal{T}_{\text{event}}(\hat{\mathcal{H}}^{(n)}, \mathcal{H}^{(n)}, \mathcal{S}_{\text{event}}, \Delta\mathcal{S}_{\text{ranking}})$

  $\mathcal{S}_{\text{event}} \leftarrow \mathcal{G}_{\text{event}}(\mathcal{S}_{\text{event}}, \Delta\mathcal{S}_{\text{event}})$

  $\mathcal{S}_{\text{ranking}} \leftarrow \mathcal{G}_{\text{ranking}}(\mathcal{S}_{\text{ranking}}, \Delta\mathcal{S}_{\text{ranking}})$

  Output $(\hat{\mathcal{H}}^{(n)}, \hat{\boldsymbol{R}}^{(n)})$

end for

3. TextGrad optimization loop: Student agents generate predictions while teacher agents provide textual feedback as "gradients" to iteratively refine reasoning strategies. This rolling weekly optimization adapts dynamically to market changes without model retraining.

Major results

Table 3 Performance comparison of different methods on industry ranking tasks

Category	Method	S&P 500 (11 industries)			A-share (9 industries)		
		Spearman	Kendall_τ	Hit_rate_3	Spearman	Kendall_τ	Hit_rate_3
Rule-based	Momentum	0.159	0.120	0.364	0.100	0.065	0.444
Time-series	ARIMA	0.142	0.131	0.373	0.138	0.089	0.444
RL-based	SAC	0.089	0.067	0.303	0.092	0.071	0.389
	A2C	0.082	0.061	0.288	0.088	0.068	0.378
	PPO	0.095	0.073	0.318	0.097	0.075	0.401
LLM-based	StkFEP_DeepSeek	0.050	0.049	0.333	0.103	0.060	0.444
	StkFEP_O1	0.072	0.068	0.303	0.101	0.102	0.389
	SEP_DeepSeek	0.078	0.042	0.208	0.118	0.069	0.472
	SEP_O1	0.121	0.098	0.333	0.134	0.098	0.408
	MENTOR_DeepSeek (ours)	0.220	0.170	0.364	0.141	**0.128**	**0.488**
	MENTOR_O1 (ours)	**0.221**	**0.180**	**0.394**	**0.173**	0.125	0.439

Best results are in bold. RL: reinforcement learning; SAC: soft actor–critic; A2C: advantage actor–critic; PPO: proximal policy optimization

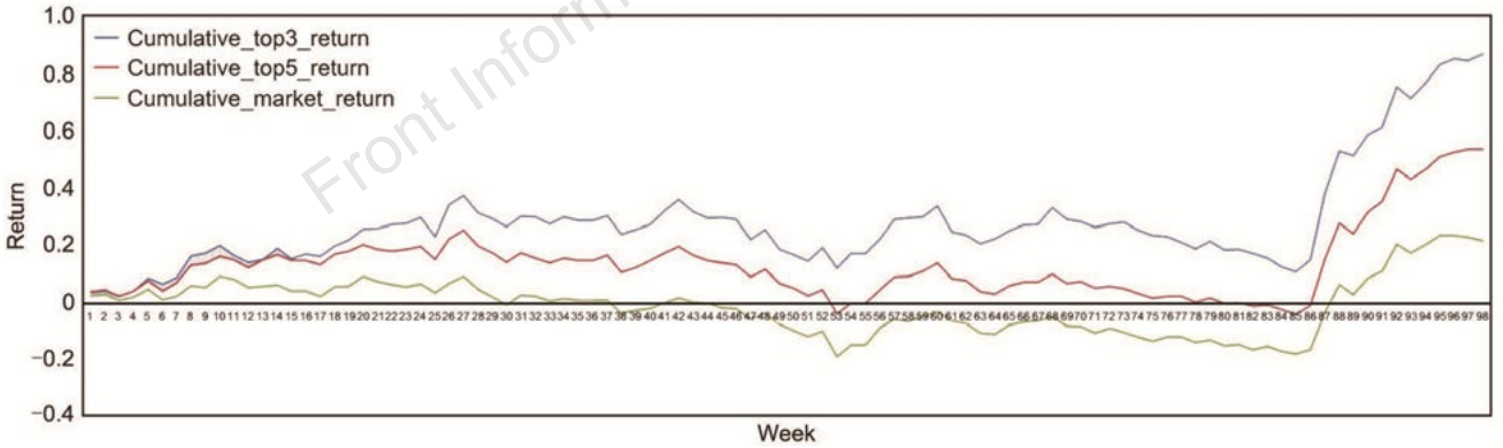

Fig. 4 Cumulative return of the Top K portfolios. The figure illustrates the backtesting results for the DeepSeek-R1 teacher and DeepSeek-V3 student pair; more results are reported in Section 5.6 (References to color refer to the online version of this figure)

Conclusions

- This work introduces MENTOR, a multi-agent reasoning framework that bridges narrative economics and financial forecasting. By decomposing complex predictive tasks into structured subtasks—event detection, future event prediction, and industry index ranking—MENTOR achieves systematic and interpretable forecasting.
- Through its teacher–student iterative reasoning mechanism, MENTOR enhances reasoning quality, connects evolving narratives with measurable market outcomes, and demonstrates strong predictive performance across both Chinese and global financial datasets. This framework provides a new pathway toward explainable and narrative-driven financial intelligence systems.