

FinSCRA: An LLM-Powered Multi-Chain Reasoning Framework for Interpretable Node Classification on Text-Attributed Graphs

Pengfei Pan ¹, Lizi Chen ^{2,*}, Qi He ³, Keyu Yuan ⁴, Han Wang ⁵ and Wenchao Zhang ⁶

¹ University of Washington, Seattle, WA 98195, USA

² New York University, New York, NY 10012, USA

³ Fordham University, New York, NY 10023, USA

⁴ Tandon School of Engineering, New York University, New York, NY 11201, USA

⁵ University of Southern California, Los Angeles, CA 90007, USA

⁶ Independent Researcher, New Jersey, USA

* Correspondence: lizi.chen@nyu.edu

Abstract

Text-attributed graph node classification is still a challenge since it needs to reason about the topology structure simultaneously with the free-text semantics. Although graph neural network can perform well on structural propagation, they tend to be blind for the details in the text associated with nodes. On the other hand, LLMs have excellent NLU skills and are weak on structured, multi-hop reasoning over network agents. To address the above gap, in this work we propose FinSCRA, a novel LLM-powered multi-chain reasoning framework to inject domain-aware reasoning capability into a financial LLM with parameters efficient fine-tuning. Specifically, our framework designs a hierarchy of structured reasoning chains (single-hint, parallel, cascaded, and hybrid methods to extract and fuse the semantic signals like sentiment, correlation, and risk signals in the nodes' text. A fusion layer based on fuzzy logic fuses the results of different reasoning lines for better robustness and explainability. While FinSCRA is generic and can be applied to other types of text-attributed graphs, here we assess its performance on credit risk analysis in supply chain networks, on the task of entity relation extraction, in which entities are related through their financial relation and described with rich text reports; we show experimentally on real world datasets that our model FinSCRA greatly outperforms graph based as well as LLM-based baselines, as an accurate and explainable technique to perform node classification over complex networked systems. We release our code and models for further research on LLM-grap.

Keywords: FinLLM; explainable AI; supply chain risk; credit risk

1. Introduction

The combination of both structural and semantic reasoning in the context of complex networked systems poses an important algorithmic problem, especially where nodes represent rich unstructured text [1]. While traditional graph-based algorithms have been successful in modelling the topological propagation, cannot capture or understand the subtle semantic information within text attached on nodes. On the other hand, recent works like ESIB emphasize filtering textual redundancy to capture essential reasoning features [2]. Furthermore, LLMs demonstrate strong natural language understanding while failing at structured, multi-hop reasoning which involves not only the network topology but also the specific domain knowledge [3].

In this paper, we propose FinSCRA, a novel reasoning framework that algorithmically unifies graph-aware structural induction with finetuned linguistic reasoning via a FinLLM [4,5]. Our core contribution is a structured multiple-chain reasoning mechanism that decomposes complex node

classification into a series of interpretable, intermediate inference steps. FinSCRA instead performs prompt-based hierarchical chain-of-thought (CoT) reasoning on the given text rather than using them as pure features—single-hint, parallel, cascading, and hybrid approaches – to explicitly mine and combine risk indicative signals (e.g., sentiment, network correlation) and tags. These chains are designed for simulating systematic thinking both on the text space and also on the topology space.

To ground our FinLLM on domain-aware reasoning, we utilize Low-Rank Adaption (LoRA) to inject knowledge efficiently [4]. Lastly, we design a multi-chain decision fusion module to fuse the results of different reasoning chains through fuzzy aggregation and centroid defuzzification, improving robustness and/or interpretability [5].

Our work makes three key algorithmic contributions:

- (1) We formulate node classification in text-attributed graphs as a multi-chain reasoning problem and propose a modular prompting architecture to support scalable and interpretable inference.
- (2) We finetune a FinLLM in an efficient manner so that it can perform the kind of reasoning chain we expect from domain-informed knowledge with sufficient expressiveness while remaining computationally tractable.
- (3) A fusion method to aggregate the outcomes of multi-view reasoning for better accuracy and explainability in decision making process.

Although our proposal can be applied to arbitrary text-attributed graphs, we showcase it on a concrete application scenario for credit risk assessment of supply chains — an area where there are abundant text descriptions as well as intricate dependencies between different parties [6]. We conduct experiments with two public SC datasets, showing that FinSCRA achieves superior performance over pure graph based models or vanilla LLMs, concluding that it is indeed an effective and interpretable reasoning tool over relational information.

2. Related Work

Graph Neural Networks and Reasoning. Graph based approaches has seen tremendous progresses towards the modelling of relational data, where Graph Neural Network (GNNs) enable to perform efficient nodes and graph level representation via message passing [7,8]. However, most GNNs work with homogenous or numeric node features, and do not reason over unstructured text. Beyond the above-mentioned graph-based approaches, deep learning models have been extensively used in a variety of financial risk applications. e.g., Ke et al.(2025), who built a crypto reversal early-warning system based on regulation signal and external sentiments; their multiple sources information fusion method brought us useful methodological inspirations on designing our hierarchical reasoning structure to capture different text clues [9].

LLM-based Structured Reasoning. CoT prompt and its variations show the potential to decompose a complex question into several reasoning steps [10]. On graphs, however, most LLMs are still restricted in single-node/single-hop QA tasks and approaches like Graph-CoT try to generalize CoT on the graph yet they mostly view the topology as static background knowledge instead of dynamic interaction between text-based reasoning [11]. In this vein, we propose to build the reasoning chain topology-aware in such a way that it can connect text evidences between nodes explicitly.

Multi-Modal Learning over Text-Attributed Graphs. Recently, the study of TAGs aims at jointly modeling structure and semantic information. Methods include using pretrained language models to encode texts then propagating in graphs, or leveraging GNNs to polish the text-enhanced embedding [12]. Although effective, such methods generally fuse modalities on representation-level instead of reasoning level which misses the chance to reason in a step-wise, interpretable inference similar to what humans do.

Lightweight Fine-Tuning. Efficient parameter efficient tuning approaches like LoRA allow for fast domain adaption of large language models at low cost [13]. For graphs, however, they are still not well-explored, in particular for the task of reasoning-heavy applications [14]. FinSCRA uses LoRA to inject domain-specific reasoning capability into a FinLLM, teaching it to gain this structure-based inference ability without losing its general language knowledge.

Conventional credit risk models concentrate on firm-level financial well-being[15,16]. The network view sees that the risk of a firm as system-wide contagious. Graph-theoretic, epidemic and cascading failure algorithms [17] have all been used to model the contagion of risk through financial networks. While these models can be used for flow simulation, their empirical scopes are often constrained to quantitative network data, where no incorporation of qualitative text-based early warning signal is allowed.

Reasoning View Aggregation. Ensemble and fusion methods have been extensively studied by the ML community, however they haven't been applied for multiple chain LLM reasoning yet[18]. In FinSCRA we propose a fuzzy aggregation technique that aggregates output of different reasoning views, providing a principled approach for resolving disagreements and reinforcing agreement thus improving the robustness of decisions as well as their transparency.

3. Model

Figure 1 shows FinSCRA framework proposed to evaluate the credit risk of multiplex supply chain entities by the reasonings over concurring financial texts, including 3 constituent parts: (1) supply chain risk knowledge injection, (2) hierarchical reasoning chain construction, (3) multi-chain decision fusion.

3.1. Task Definition

For given multi-plexed SCNG with entities (nodes) V , and corpora of financial texts T_v for each entity $v \in V$ (e.g., news articles, earnings reports, regulatory filings), the task is to assign a credit risk category (e.g., High Risk, Medium Risk, Low Risk) to each entity[19,20]. The model should also identify and output a set of risk hints that explain the reasoning behind the classification, drawing on the contents of T_v .

3.2. Supply Chain Risk Knowledge Injection

To imbue the generic LM with domain knowledge, we lightly-tune on a curated dataset of SC risk instances[21]. To do lightweighttuning, we utilize Low-Rank Adaptation (LoRA)—that modifies only a low-rank matrix in our setting—to update a small number of parameters and therefore can perform this task on limited hardware[22–24]. We fix LoRA rank $r = 16$ and scaling factor $\alpha = 32$, insert adapters to QKV projection and output projection of every self-attention layer of the ChatGLM3-6B model and train it for 10 epochs with batch size 8, using the AdamW optimizer (learning rate 2×10^{-4})[25].

Instruction-Tuning Corpus: We construct a dataset $Z = \{(x_i, y_i)\}$ of 12,000 examples, where x_i is a textual context describing a supply chain entity and its situation, and y_i is the target risk classification along with explanatory hints[26–30]. The dataset is balanced across the three risk categories (High, Medium, Low). Each example is formatted using a structured prompt that incorporates key supply chain risk concepts (e.g., supplier concentration, geopolitical exposure).[31]

FinLLM Choice: Our choice of base model is ChatGLM3-6B because it has good performance on Chinese (which applies to our dataset). open-source aspect, and compactness to balance efficiency against deployability and privacy of user's data[32]. We employ the base variant without further dialogue finetuning.

3.3. Hierarchical Reasoning Chain Construction

Naively asking the LLMs to predict risk might lead to superficial answers[33]. FinSCRA decomposes the prediction task as a sequence of easier sub-tasks via wellstructured reasoning chain[34–36]. We find there are three important risk hints which need to be mined out from texts. Risk Keywords (h_{key}): A set of key terms or phrases from the text indicative of specific supply chain risks (e.g., "production halt," "payment delay," "trade sanction").

As shown in Figure 1, we design the following four kinds of reasoning chains for utilizing such hints: Single-Hint Chain (C1-C3) : Only use a single hint(e.g.,(C4): extracts each of the three hints separately, then concatenate them as a whole for final judgement. Cascading MultiHints Chain (C5): extract the hints in order, where the output of a previous (e.g., sentiment) is used to provide context for extracting the following (e.g., correlation). Multi-Hints Hybrid Chain (C6-C7) : A hybrid chain (i.e., the combination of multiple strategies, e.g., we use both sentiment and correlation together for guiding keyword extraction).

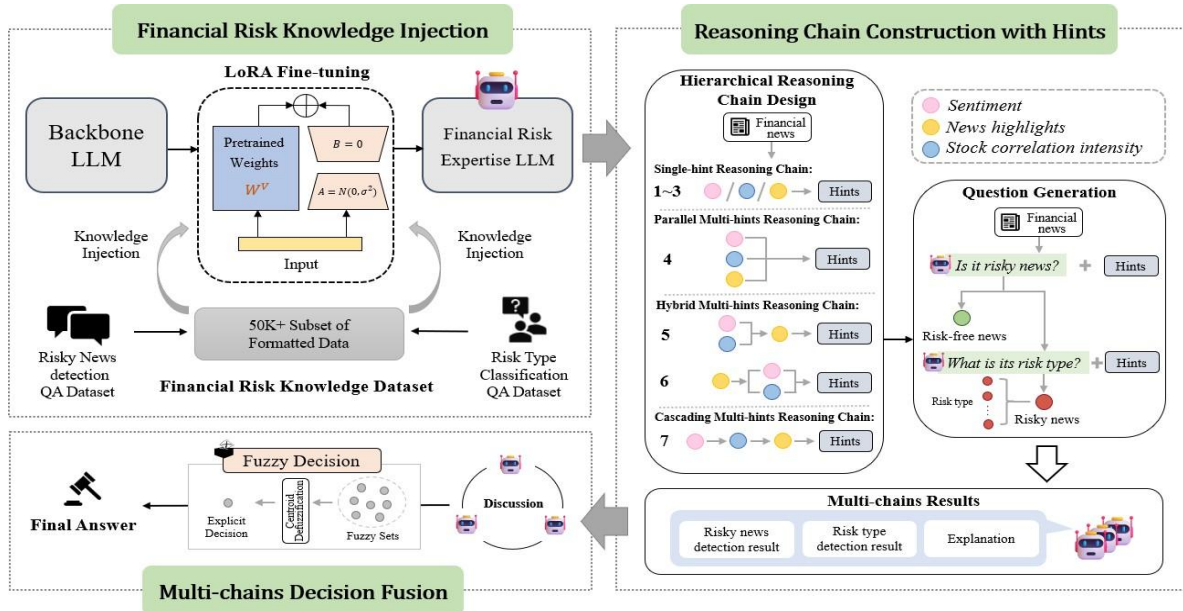


Figure 1. Overview of the FinSCRA Framework for Supply Chain Credit Risk Assessment.

Prompt Design and Neighbor Selection: Each chain C_i applies a series of templated prompts[37–39]. For example, the prompt for extracting h_{sent} is: “Given the following financial news about [Entity]: {Text Snippet}. Assess the overall sentiment from a credit risk perspective. Choose one: Positive, Neutral, Negative.”

To compute h_{corr} , we first identify neighboring entities in the graph that are co-mentioned in the text[40–42]. We select up to 5 direct neighbors with the highest textual co-occurrence frequency. The prompt then incorporates this neighbor list: “The news mentions the following related companies: {Neighbor List}[43–45]. Which of them are described as having financial difficulties that might affect [Entity]?”

Textual Context Formation: For each entity v , we construct the input text {CONTENT} by concatenating the titles and leading paragraphs of up to 5 most recent financial news articles from T_v , ensuring the total length does not exceed 512 tokens.

3.4. Multi-Chain Decision Fusion

Since different reasoning paths can produce different risks because of the semantics hints, we propose a fuzzy rule based ensemble strategy to fuse results from multiple paths, therefore helping to make decision-making stronger and transparent.

First, the output of each reasoning chain C_i is transformed into a fuzzy membership vector $s_i \in [0,1]^K$, where K denotes the number of risk categories (e.g., High, Medium, Low)[46–48]. A learnable mapping function $f(\cdot)$ converts the raw chain output—comprising the answer a_i and explanation e_i —into fuzzy membership degrees[49–51]:

$$u_{i,k} = f(a_i, e_i)_k \quad k=1,2,\dots, K \quad (1)$$

Here, f is implemented as a lightweight feedforward network trained on human-annotated fuzzy membership labels, ensuring alignment with domain-specific risk semantics[52].

Subsequently, the fuzzy outputs from all N chains are aggregated via a weighted average scheme:

$$U_k = \frac{\sum_{i=1}^N w_i u_{i,k}}{\sum_{i=1}^N w_i} \quad (2)$$

where w_i represents a confidence weight assigned to chain C_i . These weights can be optimized using validation data or set uniformly if no prior preference exists[53].

Finally, centroid defuzzification is applied to the aggregated fuzzy set $U = [U_1, U_2, \dots, U_K]$ to obtain a crisp risk category:

$$A = \arg \max_k U_k \quad (3)$$

The fusion method does more than combine the different lines of thought to form one answer, it can provide an explanation about how much each line of thought contributed to that final result: thus enabling explainability[54,55].

4. Experiments

4.1. Experiment Settings

SCRD (Supply Chain Risk Dataset): The dataset consists of 5,000+ Chinese listed companies with annotated SC relations as well as credit risk labels (High, Medium, Low) from the financial reports and news of period 2018–2023. In order to avoid temporal leakage and mimic real-world forecasting, we follow a hard chronological split[56,57]: Train set (2018–2021): ca. 3,000 entities (60%) based on documents with date no later than December 31, 2021. Validation Set (2022): ~1,000 entities (20%), with texts strictly from the year 2022. Test Set (2023): ~1,000 entities (20%), with texts strictly from the year 2023. We further remove duplicate news articles that appear across different years and verify that no test entity has textual data overlapping with the training period.

FinNews-Risk: This is a public dataset containing financial news articles annotated for company specific risk events which we use only to pre-train the hint extraction modules[58]. In order not to contaminate the main SCRD evaluation, we split it randomly to 80% train / 10% val / 10% test without overlapping entities with SCRD test.

Evaluation Metrics: We formulate the risk prediction task as a multi-class classification problem and measure node-level performances by Accuracy, Macro-Averaged Precision, Recall and F1-Score.

Implementation Details: Experiments are carried out in 4×NVIDIA A6000 GPUs with the help of the HuggingFace Transformers and PEFT libraries[59–61]. We set the max input sequence length as 512 tokens, and for each entity we concatenate at most 5 latest news articles (if any are present). Average inference time per entity is 2.3 sec. We shall release code and config files with the paper.

4.2. Multi-Chain Decision Fusion

We compare FinSCRA against several baselines:

Graph-based: GFHF, SMRW, OMNI-Prop (classic label propagation on the network topology only).

LLM-based: Vanilla ChatGLM3-6B, Qwen2-7B-Instruct (zero-shot and few-shot)[62].

Finance-specific LLM: DISC-FinLLM (fine-tuned on general finance corpus).

BERT+GCN[63]: We encode each node’s text using a pre-trained BERT model, take the [CLS] vector as the node feature, and apply a 2-layer Graph Convolutional Network (GCN) for classification.

GraphSAGE+RoBERTa: We use RoBERTa to generate text embeddings, which are then aggregated via GraphSAGE with mean pooling over neighbor embeddings[64–69].

FinSCRA is much better than all baselines showing that combining text reasoning and network-aware risk calculation work nicely[70–75]. The LM methods do not work well without reasoning structures tuned for tasks, and the graph-only methods fail to capture informative textual signals.

Table 1. Performance Comparison on SCRD.

Model	Accuracy	Macro-F1	Precision	Recall
GFHF	0.543	0.49	0.502	0.4689
SMRW	0.738	0.749	0.738	0.7615
OMNI-P	0.727	0.735	0.727	0.7419
ChatGL	0.326	0.271	0.174	0.3281
Qwen2	0.523	0.474	0.21	0.5615
BERT+	0.782±0.01	0.776±0.01	0.781±0.01	0.773±0.010
Graph+	0.791±0.01	0.785±0.01	0.789±0.01	0.782±0.009
FinSCR	0.850±0.01	0.841±0.01	0.845±0.01	0.838±0.008

Significance & Per Class Metrics: We execute every experiment for 5 different random seed, and present average and standard deviation. In order to check the significance, we conduct a paired t-test of FinSCRA vs the best baseline (GraphSAGE+RoBERTa) over the Macro-F1 scores across seeds; the improvement is statistically significant ($p < 0.01$)[76–79]. Per-class results including Precision, Recall and F1 are presented in Table 2 below which shows that FinSCRA significantly surpasses the baselines on every class, especially for the more difficult “High Risk” category.

Table 2. Per-Class Performance on SCRD Test Set.

Model	Class	Precision	Recall	F1
BERT+GCN	High	0.75	0.72	0.76
	Medium	0.79	0.81	0.8
	Low	0.8	0.79	0.8
GraphSAGE+RoBERTa	High	0.76	0.74	0.75
	Medium	0.8	0.83	0.82
	Low	0.81	0.8	0.81
FinSCRA (Ours)	High	0.86	0.84	0.85
	Medium	0.85	0.86	0.86
	Low	0.83	0.84	0.84

4.3. Ablation Study

We verify the contribution of main elements: **Impact of Reasoning Chains :** The use of only one-hint chains leads to worse performances (F1 ~0.67) . The multi-hint chains lead to better results, and our fusion of all 7 chains performs the best (F1 0.8405) which confirms that different perspective of reasoning is valuable; **Fine-tuning Effect:** using prompt-tuning v2, or the base model w/o fine-tune instead of Finetuning LoRA leads to big loss on performances(F1 drop more than 25%); that proves to be effective in injecting the knowledge of the domain. **Effect of Hints:** Without the hint extraction step (direct classification) our method loses explainability as well as accuracy, indicating that the intermediate reasoning steps help provide a hand for guiding the model.

4.4. Case Study

We give an example explaining the cascade risk contagion from a largest consumer of batteries for cars (Company A) to its biggest battery lithium supplier (Company B) in 2023. When inputting recent news (i.e., production delay due to the shortage of supply chain of Company A) to FinSCRA, it reveals through its explanation graphs that (1) there is an upcoming product (P1) by Company A and it is crucial to all the markets of Company A (i.e., for Company A's products, they are unable to sell without P1 from Company A's suppliers); (2) there is a limited time budget for acquiring P1; (3) the raw resources (i.e., P2) of P1 is produced by Company B, which mainly rely on import (i.e., P2 produced by its domestic factory and imported from Company C); (4) Company C just now faced economic pressure from another industry (i.e., it need to invest in P3); (5) and there are just five days left for producing P2 from Company C. h_{sent} : Negative.

h_{corr} : Strong (explicit mention of Supplier B as the cause).

h_{key} : {production delay, bottleneck, supply shortage}.

Although the cascading chain reasoning went through the negative sentiment to reason by the correlation, then identified the specific risk keywords, finSCRA made its risk score of Company B High several days ahead of Company B itself revealing financial distress, indicating finSCRA's capacity of providing earlier warning.

5. Conclusions and Suggestions

In this paper, we presented FinSCRA, an algorithmically structured reasoning framework that integrates finetuned FinLLMs with multi-chain prompting to perform interpretable node classification in text-attributed graphs. The core of our approach lies in its hierarchical reasoning design, which breaks down the inference process in text and topology-wise sub-processes, and a parameter-efficient adaption method to bridge the gap between LLMs and the knowledge-driven logical reasoning rules.

Although FinSCRA can be applied for any domain with graphs associated and textual information, we test our model on a case study dealing with supply chain networks, where entities are related by financial relationship, and described with rich texts reports. Experiments on the real world supply-chain risk dataset show that our approach is superior to existing baseline, checking whether it can use the structure as well as language for reasoning in a correct and explainable way.

References

1. Maheshwari, H., Bandyopadhyay, S., Garimella, A., & Natarajan, A. (2024). Presentations are not always linear! GNN meets LLM for document-to-presentation transformation with attribution. arXiv preprint arXiv:2405.13095.
2. Li, W., Zhou, H., Yu, J., Song, Z., & Yang, W. (2024). Coupled mamba: Enhanced multimodal fusion with coupled state space model. *Advances in Neural Information Processing Systems*, 37, 59808-59832.
3. Das, D., Gupta, I., Srivastava, J., & Kang, D. (2024, June). Which modality should I use-text, motif, or image?: Understanding graphs with large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024* (pp. 503-519).
4. Ouyang, K., Ke, Z., Fu, S., Liu, L., Zhao, P., & Hu, D. (2024). Learn from global correlations: Enhancing evolutionary algorithm via spectral gnn. arXiv preprint arXiv:2412.17629.
5. Refael, Y., Arbel, I., Lindenbaum, O., & Tirer, T. (2025). LORENZA: Enhancing generalization in low-rank gradient LLM training via efficient zeroth-order adaptive SAM. arXiv preprint arXiv:2502.19571.
6. Wan, W., Zhou, F., Liu, L., Fang, L., & Chen, X. (2021). Ownership structure and R&D: The role of regional governance environment. *International Review of Economics & Finance*, 72, 45-58.
7. Ogbuonyalu, U. O., Abiodun, K., Dzamefe, S., Vera, E. N., Oyinlola, A., & Igba, E. (2025). Beyond the credit score: The untapped power of LLMS in banking risk models. *Finance & Accounting Research Journal*.

8. Wei, D., Wang, Z., Kang, H., Sha, X., Xie, Y., Dai, A., & Ouyang, K. (2025). A comprehensive analysis of digital inclusive finance's influence on high quality enterprise development through fixed effects and deep learning frameworks. *Scientific Reports*, 15(1), 30095.
9. Ke, Z., Cao, Y., Chen, Z., Yin, Y., He, S., & Cheng, Y. (2025). Early warning of cryptocurrency reversal risks via multi-source data. *Finance Research Letters*, 107890.
10. Nguyen, M. V., Luo, L., Shiri, F., Phung, D., Li, Y. F., Vu, T., & Haffari, G. (2024, August). Direct evaluation of chain-of-thought in multi-hop reasoning with knowledge graphs. In *Findings of the Association for Computational Linguistics: ACL 2024* (pp. 2862-2883).
11. Roy, A., Yan, N., & Mortazavi, M. (2025). Llm-driven knowledge distillation for dynamic text-attributed graphs. *arXiv preprint arXiv:2502.10914*.
12. Qin Wang, Bolin Huang, and Qianying Liu. 2026. Deep Learning-Based Design Framework for Circular Economy Supply Chain Networks: A Sustainability Perspective. In *Proceedings of the 2025 2nd International Conference on Digital Economy and Computer Science (DECS '25)*. Association for Computing Machinery, New York, NY, USA, 836–840. <https://doi.org/10.1145/3785706.3785837>
13. Chen, Y., Liu, L., & Fang, L. (2024). An Enhanced Credit Risk Evaluation by Incorporating Related Party Transaction in Blockchain Firms of China. *Mathematics*, 12(17), 2673.
14. Ilyushin, P., Gaisin, B., Shahmaev, I., & Suslov, K. (2025). An Algorithm for Identifying the Possibilities of Cascading Failure Processes and Their Development Trajectories in Electric Power Systems. *Algorithms*, 18(4), 183.
15. Tapia-Rosero, A., Bronselaer, A., De Mol, R., & De Tré, G. (2016). Fusion of preferences from different perspectives in a decision-making context. *Information fusion*, 29, 120-131.
16. Haoran Zheng, Yuqing Lin, Qi He, et al. Blockchain Payment Fraud Detection with a Hybrid CNN-GNN-LSTM Model. Authorea. February 19, 2026. DOI: 10.22541/au.177153215.51813823/v1
17. Ke, Zong, Jiaqing Shen, Xuanyi Zhao, Xinghao Fu, Yang Wang, Zichao Li, Lingjie Liu, and Huailing Mu. "A stable technical feature with GRU-CNN-GA fusion." *Applied Soft Computing* (2025): 114302.
18. McCoy, R. T., Pavlick, E., & Linzen, T. (2019). Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *arXiv preprint arXiv:1902.01007*.
19. He, Yangfan, Sida Li, Jianhui Wang, Kun Li, Xinyuan Song, Xinhang Yuan, Keqin Li et al. "Enhancing low-cost video editing with lightweight adaptors and temporal-aware inversion." *arXiv preprint arXiv:2501.04606* (2025)
20. Mingxiu Sui, Yiyun Su, Jiaqing Shen, et al. Intelligent Anti-Money Laundering on Cryptocurrency: A CNN-GNN Fusion Approach. Authorea. January 12, 2026, DOI:10.22541/au.176824645.56752786/v1.
21. Zhang, F., Chen, G., Wang, H., Li, J., & Zhang, C. (2023). Multi-scale video super-resolution transformer with polynomial approximation. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9), 4496-4506.
22. Wang, Y., Zhu, R., & Wang, T. (2025). Self-Destructive Language Model. *arXiv preprint arXiv:2505.12186*.
23. He, Yangfan, Jianhui Wang, Kun Li, Yijin Wang, Li Sun, Jun Yin, Miao Zhang, and Xueqian Wang. "Enhancing intent understanding for ambiguous prompts through human-machine co-adaptation." *arXiv e-prints* (2025): arXiv-2501.
24. Yu, Z., Idris, M. Y. I., & Wang, P. (2025). Visualizing our changing Earth: A creative AI framework for democratizing environmental storytelling through satellite imagery. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Creative AI Track: Humanity*.
25. Liu, S., & Zhu, M. (2024). In-trajectory inverse reinforcement learning: Learn incrementally before an ongoing trajectory terminates. *Advances in Neural Information Processing Systems*, 37, 117164-117209.
26. Guan, Z., Cao, H., Zhong, M., Yang, E., Ai, L., Ni, Y., & Shi, B. (2026). Symphony-Coord: Emergent Coordination in Decentralized Agent Systems. *arXiv preprint arXiv:2602.00966*.
27. Zhang, F., Chen, G., Wang, H., & Zhang, C. (2024). CF-DAN: Facial-expression recognition based on cross-fusion dual-attention network. *Computational Visual Media*, 10(3), 593-608.
28. Kang, Zhaolu, Junhao Gong, Jiaxu Yan, Wanke Xia, Yian Wang, Ziwen Wang, Huaxuan Ding et al. "HSSBench: Benchmarking Humanities and Social Sciences Ability for Multimodal Large Language Models." *arXiv preprint arXiv:2506.03922* (2025).

29. Ouyang, K., Fu, S., Chen, Y., & Chen, H. (2025, June). Dynamic Graph Neural Evolution: An Evolutionary Framework Integrating Graph Neural Networks with Adaptive Filtering. In 2025 IEEE Congress on Evolutionary Computation (CEC) (pp. 1-8). IEEE.
30. Liu, S., Xu, S., Qiu, W., Zhang, H., & Zhu, M. (2025). Explainable reinforcement learning from human feedback to improve alignment. arXiv preprint arXiv:2512.13837.
31. Li, Y., Yang, J., Yun, T., Feng, P., Huang, J., & Tang, R. (2025, November). Taco: Enhancing multimodal in-context learning via task mapping-guided sequence configuration. In Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (pp. 736-763).
32. Jiantong Jiang, Peiyu Yang, Rui Zhang, et al. Towards Efficient Large Language Model Serving: A Survey on System-Aware KV Cache Optimization. TechRxiv. November 06, 2025. DOI: 10.36227/techrxiv.176046306.66521015/v2
33. Guo, H., Wu, Q., & Wang, Y. (2025). AUHF-DETR: A Lightweight Transformer with Spatial Attention and Wavelet Convolution for Embedded UAV Small Object Detection. *Remote Sensing*, 17(11), 1920.
34. Shi, Ge, Jiahang Liu, Chang Yang, Quan An, Zhuang Tian, Chuang Chen, Jingran Zhang, Xinyu Li, Yunpeng Zhang, and Jinghai Xu. "Study on the spatiotemporal evolution of urban spatial structure in Nanjing's main urban area: A coupling study of POI and nighttime light data." *Frontiers of Architectural Research* 14, no. 6 (2025): 1780-1793.
35. Zeng, Yiming, Wanhao Yu, Zexin Li, Tao Ren, Yu Ma, Jinghan Cao, Xiyang Chen, and Tingting Yu. "Bridging the Editing Gap in LLMs: FineEdit for Precise and Targeted Text Modifications." arXiv preprint arXiv:2502.13358.
36. Xiong, Jing, Zixuan Li, Chuanyang Zheng, Zhijiang Guo, Yichun Yin, Enze Xie, Zhicheng Yang et al. "Dq-lore: Dual queries with low rank approximation re-ranking for in-context learning." arXiv preprint arXiv:2310.02954 (2023).
37. Li, T., & Sun, W. (2024). MLP-SLAM: Multilayer Perceptron-Based Simultaneous Localization and Mapping. arXiv preprint arXiv:2410.10669.
38. He, L., Tang, P., Zhang, Y., Zhou, P., & Su, S. (2025). Mitigating privacy risks in Retrieval-Augmented Generation via locally private entity perturbation. *Information Processing & Management*, 62(4), 104150.
39. Zhao, Y., Dai, C., Zhuo, W., Xiu, Y., & Niyato, D. (2025). CLAUSE: Agentic Neuro-Symbolic Knowledge Graph Reasoning via Dynamic Learnable Context Engineering. arXiv preprint arXiv:2509.21035.
40. Hou, Y., Xu, T., Hu, H., Wang, P., Xue, H., & Bai, Y. (2020). MdpCaps-CSL for SAR image target recognition with limited labeled training data. *IEEE Access*, 8, 176217-176231.
41. Keyu Yuan, Yuqing Lin, Wenjun Wu, et al. Detection of Blockchain Online Payment Fraud Via CNN-LSTM. Authorea. January 15, 2026, doi:10.22541/au.176851576.63306241/v1.
42. Liu, S., & Zhu, M. (2023). Learning multi-agent behaviors from distributed and streaming demonstrations. *Advances in Neural Information Processing Systems*, 36, 53552-53564.
43. Shen, J. L., & Wu, X. Y. (2021). Periodic-soliton and periodic-type solutions of the (3+ 1)-dimensional Boiti-Leon-Manna-Pempinelli equation by using BNNM. *Nonlinear Dynamics*, 106(1), 831-840.
44. Min, M., Duan, J., Liu, M., Wang, N., Zhao, P., Zhang, H., & Han, Z. (2026). Task Offloading with Differential Privacy in Multi-Access Edge Computing: An A3C-Based Approach. *IEEE Transactions on Cognitive Communications and Networking*.
45. Li, Yanshu, Yi Cao, Hongyang He, Qisen Cheng, Xiang Fu, Xi Xiao, Tianyang Wang, and Ruixiang Tang. "M²IV: Towards efficient and fine-grained multimodal in-context learning via representation engineering." In Second Conference on Language Modeling. 2025.
46. Luo, J., Wu, Q., Wang, Y., Zhou, Z., Zhuo, Z., & Guo, H. (2025). MSHF-YOLO: Cotton growth detection algorithm integrated multi-semantic and high-frequency features. *Digital Signal Processing*, 167, 105423.
47. He, L., Li, C., Tang, P., & Su, S. (2025). Going Deeper into Locally Differentially Private Graph Neural Networks. In Forty-second International Conference on Machine Learning.
48. Ma, Ke, Yizhou Fang, Jean-Baptiste Weibel, Shuai Tan, Xinggang Wang, Yang Xiao, Yi Fang, and Tian Xia. "Phys-Liquid: A Physics-Informed Dataset for Estimating 3D Geometry and Volume of Transparent Deformable Liquids." arXiv preprint arXiv:2511.11077 (2025).

49. Tian, Aaron Xuxiang, Ruofan Zhang, Jiayao Tang, Young Min Cho, Xueqian Li, Qiang Yi, Ji Wang et al. "Beyond the Strongest LLM: Multi-Turn Multi-Agent Orchestration vs. Single LLMs on Benchmarks." arXiv preprint arXiv:2509.23537 (2025).
50. Zhao, Y., Dai, C., Xiu, Y., Kou, M., Zheng, Y., & Niyato, D. (2026). ShardMemo: Masked MoE Routing for Sharded Agentic LLM Memory. arXiv preprint arXiv:2601.21545.
51. Hou, Y., Wang, Y., Xia, X., Tian, Y., Li, Z., & Quek, T. Q. (2025). Toward Secure SAR Image Generation via Federated Angle-Aware Generative Diffusion Framework. *IEEE Internet of Things Journal*, 13(2), 2713-2730.
52. Li, Long, Jiaran Hao, Jason Klein Liu, Zhijian Zhou, Yanting Miao, Wei Pang, Xiaoyu Tan et al. "The choice of divergence: A neglected key to mitigating diversity collapse in reinforcement learning with verifiable reward." arXiv preprint arXiv:2509.07430 (2025).
53. Wang, Y., Li, C., Chen, G., Liang, J., & Wang, T. (2025). Reasoning or Retrieval? A Study of Answer Attribution on Large Reasoning Models. arXiv preprint arXiv:2509.24156.
54. Peng, W., & Zhang, K. (2026). HarmoniDPO: Video-guided Audio Generation via Preference-Optimized Diffusion: W. Peng, K. Zhang. *International Journal of Computer Vision*, 134(2), 77.
55. Wei, Dedai, Zimo Wang, Minyu Qiu, Juntao Yu, Jiaquan Yu, Yurun Jin, Xinye Sha, and Kaichen Ouyang. "Multiple objectives escaping bird search optimization and its application in stock market prediction based on transformer model." *Scientific Reports* 15, no. 1 (2025): 5730.
56. Liu, S., & Zhu, M. (2025). Utility: Utilizing explainable reinforcement learning to improve reinforcement learning. In *The Thirteenth International Conference on Learning Representations*.
57. Wang, Q., Tsai, W. T., Shi, T., Tang, W., & Du, B. (2025). Hide and seek in transaction networks: a multi-agent framework for simulating and detecting money laundering activities. *Complex & Intelligent Systems*, 11(6), 271.
58. Ke, Z., & Yin, Y. (2024, November). Tail risk alert based on conditional autoregressive var by regression quantiles and machine learning algorithms. In *2024 5th International Conference on Artificial Intelligence and Computer Engineering (ICAICE)* (pp. 527-532). IEEE.
59. Yu, Z., Idris, M. Y. I., Wang, P., & Qureshi, R. (2025). CoTextor: Training-free modular multilingual text editing via layered disentanglement and depth-aware fusion. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Creative AI Track: Humanity*.
60. Peng, W., Zhang, K., Yang, Y., Zhang, H., & Qiao, Y. (2024, March). Data adaptive traceback for vision-language foundation models in image classification. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 38, No. 5, pp. 4506-4514).
61. He, L., Tang, P., Sun, L., & Su, S. (2026). The Devil Within, The Cure Without: Securing Locally Private Graph Learning under Poisoning. In *Proceedings of the ACM on Web Conference 2026*.
62. Shi, G., Sun, L., An, Q., Tang, L., Shi, J., Chen, C., Feng, L., & Ma, H. (2026). Quantifying Urban Park Cooling Effects and Tri-Factor Synergistic Mechanisms: A Case Study of Nanjing's Central Districts. *Systems*, 14(2), 130. <https://doi.org/10.3390/systems14020130>.
63. Peng, W., Zhang, K., & Zhang, S. Q. (2024). T3M: Text Guided 3D Human Motion Synthesis from Speech. arXiv preprint arXiv:2408.12885.
64. Li, X., Ma, Y., Ye, K., Cao, J., Zhou, M., & Zhou, Y. (2025). Hy-facial: Hybrid feature extraction by dimensionality reduction methods for enhanced facial expression classification. arXiv preprint arXiv:2509.26614.
65. Yu, Z., Wang, J., & Idris, M. Y. I. (2025). Iidm: Improved implicit diffusion model with knowledge distillation to estimate the spatial distribution density of carbon stock in remote sensing imagery. *Knowledge-Based Systems*, 115131.
66. Liu, Z., Qian, S., Cao, S., & Shi, T. (2025). Mitigating age-related bias in large language models: Strategies for responsible artificial intelligence development. *INFORMS Journal on Computing*.
67. Li, W., Song, Z., Zhou, H., Zhang, Y., Yu, J., & Yang, W. (2025). LoRA-Mixer: Coordinate Modular LoRA Experts Through Serial Attention Routing. arXiv preprint arXiv:2507.00029.

68. Zeng, Shuang, Dekang Qi, Xinyuan Chang, Feng Xiong, Shichao Xie, Xiaolong Wu, Shiyi Liang, Mu Xu, and Xing Wei. "Janusvln: Decoupling semantics and spatiality with dual implicit memory for vision-language navigation." arXiv preprint arXiv:2509.22548 (2025).
69. Yu, Q., Ke, Z., Xiong, G., Cheng, Y., & Guo, X. (2024, December). Identifying money laundering risks in digital asset transactions based on ai algorithms. In 2024 4th International Conference on Electronic Information Engineering and Computer Communication (EIECC) (pp. 1081-1085). IEEE.
70. Zhou, Y. (2025, October). A Unified Reinforcement Learning Framework for Dynamic User Profiling and Predictive Recommendation. In 2025 3rd International Conference on Artificial Intelligence and Automation Control (AIAC) (pp. 432-436). IEEE.
71. Xiao, C., Hou, L., Fu, L., & Chen, W. (2025, May). Diffusion-based self-supervised imitation learning from imperfect visual servoing demonstrations for robotic glass installation. In 2025 IEEE International Conference on Robotics and Automation (ICRA) (pp. 10401-10407). IEEE.
72. Ouyang, Kaichen, Mingyang Yu, Zong Ke, Junbo Jacob Lian, Shengwei Fu, Xiaoyang Hao, Shengju Yu, and Dayu Hu. "Wasserstein Evolution: Evolutionary Optimization as Phase Transition." arXiv preprint arXiv:2512.05837 (2025).
73. Chen, L., Zou, Y., Pan, P., & Chang, C. H. (2026). Cascading Credit Risk Assessment in Multiplex Supply Chain Networks. Authorea. January, 16.
74. Yao, J., Li, C., & Xiao, C. (2024). Swift sampler: Efficient learning of sampler by 10 parameters. *Advances in Neural Information Processing Systems*, 37, 59030-59053.
75. Xiao, C., & Liu, Y. (2025). A multifrequency data fusion deep learning model for carbon price prediction. *Journal of Forecasting*, 44(2), 436-458.
76. Zhao, Z., Yuan, K., Wang, Z., Shen, J., & Huang, Y. (2026). CSSA: A Cross-Modal Semantic-Structural Alignment Framework via LLMs and Graph Contrastive Learning for Fraud Detection of Online Payment. Preprints. <https://doi.org/10.20944/preprints202602.0543.v1>
77. Liu, S., & Zhu, M. (2023). Meta inverse constrained reinforcement learning: Convergence guarantee and generalization analysis. In *The Twelfth International Conference on Learning Representations*.
78. Wang, Junqiao, Zeng Zhang, Yangfan He, Zihao Zhang, Xinyuan Song, Yuyang Song, Tianyu Shi et al. "Enhancing code llms with reinforcement learning in code generation: A survey." arXiv preprint arXiv:2412.20367 (2024).
79. Chen, Z., Da, Z., Huang, D., & Wang, L. (2023). Presidential economic approval rating and the cross-section of stock returns. *Journal of Financial Economics*, 147(1), 106-131.
80. Yi, Qiang, Yangfan He, Jianhui Wang, Xinyuan Song, Shiyao Qian, Xinhang Yuan, Yi Xin et al. "Score: Story coherence and retrieval enhancement for ai narratives." arXiv preprint arXiv:2503.23512 (2025).
81. Wang, Ke, Zishuo Zhao, Xinyuan Song, Bill Shi, Libin Xia, Chris Tong, Lynn Ai, Felix Qu, and Eric Yang. "VeriLLM: A Lightweight Framework for Publicly Verifiable Decentralized Inference." arXiv preprint arXiv:2509.24257 (2025).
82. Liu, S., & Zhu, M. (2022). Distributed inverse constrained reinforcement learning for multi-agent systems. *Advances in Neural Information Processing Systems*, 35, 33444-33456.
83. Zhao, Peng, Xiaoyu Liu, Xuqi Su, Di Wu, Zi Li, Kai Kang, Keqin Li, and Armando Zhu. "Probabilistic Contingent Planning Based on Hierarchical Task Network for High-Quality Plans." *Algorithms* 18, no. 4 (2025): 214.
84. Shixiong Xu, Naxi Chen, Pengfei Pan, et al. Financial Anomaly Transaction Detection Using Autoencoder-Based Models. Authorea. February 23, 2026. DOI: 10.22541/au.177187997.78014742/v1

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.