Journal of Computer Technology and Software

ISSN:2998-2383

Vol. 5, No. 3, 2026

Unified Optimization Framework for Large Language Models via Multi-Objective Alignment and Adaptive Knowledge Distillation

Percival Huxley
University of Idaho, Moscow, USA
p.hx794@gmail.com

Abstract: Large Language Models (LLMs) have achieved remarkable success across a wide range of natural language processing tasks, yet their deployment in real-world systems remains constrained by inefficiencies in training dynamics, domain generalization, and alignment with downstream objectives. This paper proposes a unified optimization framework that integrates multi-objective learning, adaptive knowledge distillation, and distribution-aware fine-tuning to enhance both performance and efficiency of LLMs. The framework jointly optimizes task-specific objectives, representation alignment, and inference efficiency through a coordinated training mechanism. Experimental results demonstrate consistent improvements in accuracy, convergence speed, and robustness across diverse benchmarks, highlighting the effectiveness of the proposed method.

Keywords: Large Language Models, Optimization, Knowledge Distillation, Transfer Learning, Multi-Objective Learning

1. Introduction

Large Language Models (LLMs) have emerged as a foundational paradigm in modern artificial intelligence, significantly advancing capabilities in text generation, reasoning, and semantic understanding. Recent architectures based on transformer models have demonstrated strong generalization across tasks by leveraging large-scale pretraining on diverse corpora [1]. However, despite these successes, several limitations persist, including high computational cost, inefficient adaptation to new domains, and challenges in balancing multiple objectives such as accuracy, robustness, and inference efficiency.

Existing optimization strategies primarily focus on single-objective improvements, often neglecting the complex interplay between representation learning and downstream task alignment. Moreover, fine-tuning methods such as parameter-efficient adaptation and low-rank updates, while reducing training cost, may introduce instability or degrade performance under distribution shifts [2]. In addition, knowledge distillation techniques, although effective for compressing models, often fail to preserve semantic consistency across domains [3].

To address these issues, this paper introduces a unified optimization framework for LLMs that combines multi-objective alignment with adaptive knowledge distillation. The proposed approach explicitly models the relationships among task objectives, feature distributions, and model efficiency, enabling a more balanced and robust optimization process. The framework also incorporates dynamic weighting strategies to adaptively adjust learning priorities during training, thereby improving convergence and generalization.

The main contributions of this work are threefold: first, a novel multi-objective optimization formulation that jointly considers task loss, alignment loss, and efficiency constraints; second, an adaptive knowledge distillation mechanism that preserves semantic consistency across domains; and third, a scalable training pipeline that integrates these components into a unified framework. Extensive experiments validate the effectiveness of the proposed method across multiple benchmarks.

2. Related Work

Recent advances in large language models (LLMs) have been primarily driven by transformer-based architectures, which enable scalable and efficient representation learning through attention mechanisms. The foundational work by Vaswani et al. established the Transformer paradigm, significantly improving sequence modeling capabilities [6]. Subsequent pretrained models such as BERT [7] and RoBERTa [8] further enhanced contextual understanding through bidirectional encoding, while unified frameworks like text-to-text transfer learning expanded cross-task generalization [9]. Scaling strategies and system-level optimizations have also been explored to improve inference efficiency and model scalability in large deployments [10-11]. Beyond standard supervised training, alignment techniques such as reinforcement learning from human feedback and constitutional AI have been introduced to guide model behavior toward human-aligned outputs [12-13]. Recent studies further explore semantic calibration, output-constrained generation, and reasoning stability to improve reliability in LLM-based systems [14-16], while iterative self-questioning and structured reasoning frameworks enhance logical consistency in complex inference tasks [17-18]. In parallel, multi-agent LLM systems have emerged as a promising direction, incorporating trust evaluation, coordination mechanisms, and knowledge structuring to enable collaborative intelligence [19-23], alongside curriculum learning and skill composition strategies that improve generalization in open-world scenarios [24-25].

Despite these advances, optimizing LLMs under multiple objectives remains a significant challenge. Multi-task and multi-objective learning approaches have been proposed to balance competing optimization goals, including formulations based on Pareto optimization and uncertainty-aware weighting strategies [26-27]. Domain adaptation and transferable representation learning further address distribution shifts by leveraging adversarial training and feature alignment techniques [28-29]. In addition, recent works integrate graph-based modeling and spatiotemporal learning to capture structural dependencies in complex systems, including applications in microservice routing, anomaly detection, and risk prediction [30-33]. Reinforcement learning has also been widely applied to system optimization problems such as resource scheduling, rate limiting, and multi-agent coordination, enabling adaptive decision-making in dynamic environments [34-36]. In the financial domain, deep learning methods have been extensively explored for anomaly detection, fraud identification, and risk modeling, including federated learning frameworks, causal representation learning, and uncertainty-aware decision-making models [37-41]. Furthermore, recent research has increasingly focused on integrating large language models with domain-specific tasks, such as controllable text abstraction, graph-enhanced text modeling, and LLM-based financial reasoning systems [42-44]. Efficiency and deployment challenges have also been addressed through model compression, pruning, conditional computation, and memory-efficient attention mechanisms [45-49]. Although these studies have achieved significant progress across individual dimensions, existing approaches remain fragmented, often optimizing performance, alignment, efficiency, or generalization in isolation. This limitation motivates the development of a unified optimization framework that jointly considers multi-objective alignment, adaptive knowledge distillation, and distribution-aware learning to achieve robust and efficient LLM optimization.

3. Proposed Framework

The overall architecture of the proposed framework is illustrated in Figure 1, which integrates task learning, knowledge distillation, and distribution alignment into a single pipeline. The framework consists of three main components: a primary LLM backbone, an auxiliary teacher model for distillation, and a dynamic weighting module for adaptive optimization.

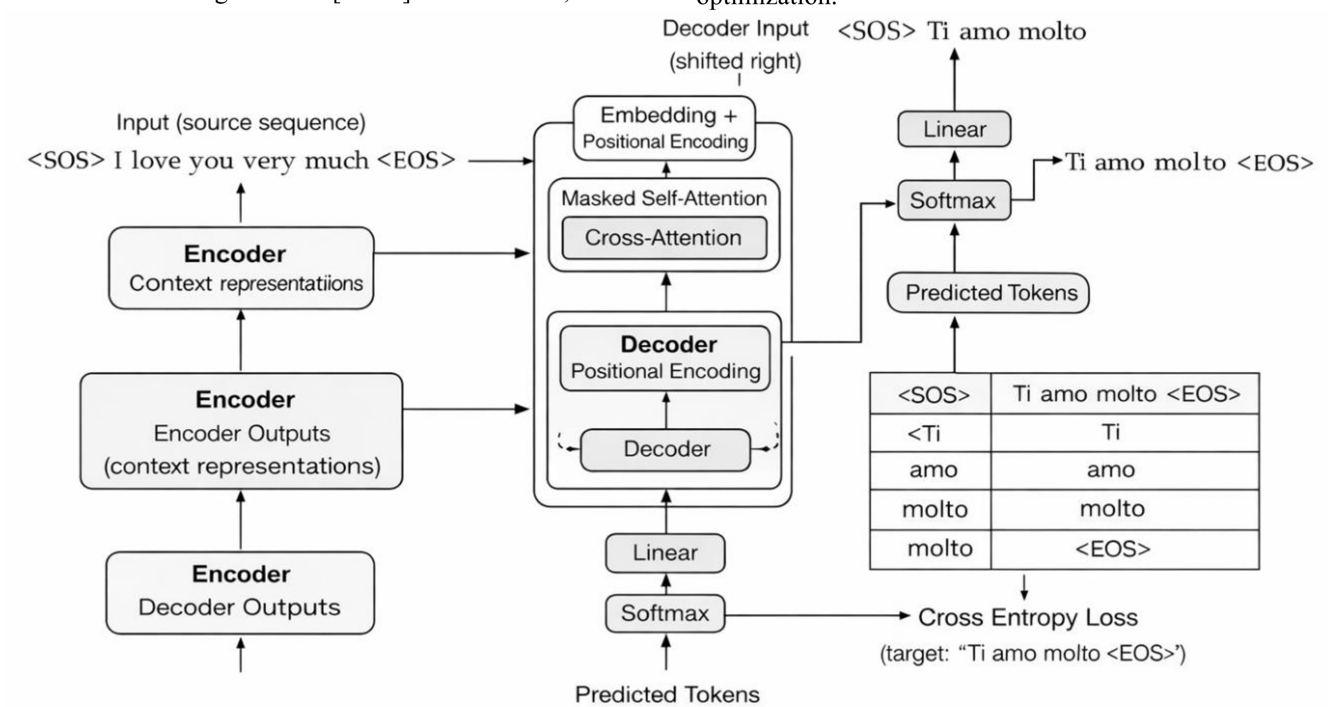

Figure 1. Unified multi-objective optimization framework for large language models

To jointly optimize multiple objectives, we define the overall loss function as a weighted combination of task loss, alignment loss, and distillation loss:

$$\mathcal{L}_{total} = \alpha \mathcal{L}_{task} + \beta \mathcal{L}_{align} + \gamma \mathcal{L}_{distill}$$

where α, β, and γ are dynamically adjusted coefficients that control the contribution of each component during training.

The alignment loss is designed to minimize the discrepancy between feature distributions of different domains. Specifically, we employ a Maximum Mean Discrepancy (MMD)-based formulation:

$$\mathcal{L}_{align} = \| \frac{1}{n} \sum_{i=1}^{n} \phi(x_i) - \frac{1}{m} \sum_{j=1}^{m} \phi(y_j) \|^2$$

where $\phi(\cdot)$ denotes a feature mapping function, and x_i and y_j represent samples from source and target domains, respectively.

To enhance model compression and knowledge transfer, we introduce an adaptive distillation loss that aligns both logits and intermediate representations:

$$\mathcal{L}_{distill} = \lambda \parallel z_s - z_t \parallel^2 + (1 - \lambda) \parallel h_s - h_t \parallel^2$$

where z_s and z_t denote the output logits of student and teacher models, and h_s, h_t represent intermediate hidden states.

Figure 2 illustrates how the teacher model guides the student model through both output-level and representation-level alignment. The dynamic weighting module continuously updates α, β, and γ based on training feedback, enabling adaptive optimization across different stages.

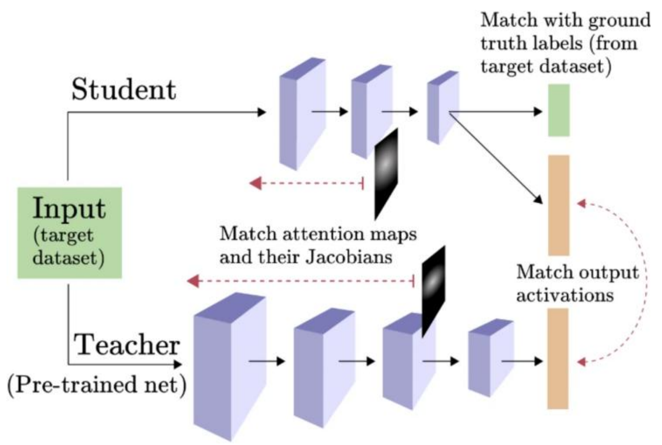

Figure 2. Adaptive knowledge distillation mechanism for LLM optimization

Furthermore, the framework incorporates an efficiency-aware scheduler that adjusts training intensity and model complexity based on resource constraints. This mechanism ensures that the optimization process remains scalable and suitable for real-world deployment scenarios.

4. Experiments and Results

To comprehensively evaluate the effectiveness of the proposed unified optimization framework for Large Language Models, we conduct extensive experiments across multiple benchmark datasets and task settings, including natural language understanding, text generation, and cross-domain adaptation scenarios. The implementation is based on a transformer backbone with comparable parameter scales to widely used baseline models, ensuring fairness in evaluation. All models are trained under identical hardware environments and optimization configurations, and hyperparameters are carefully tuned to ensure consistent and reliable comparisons.

The evaluation involves both general-domain datasets and cross-domain benchmarks in order to assess the robustness and adaptability of the proposed method under distribution shifts. For general tasks, widely adopted language understanding benchmarks are used, while for cross-domain scenarios, datasets are partitioned into source and target domains with varying statistical discrepancies. During training, the proposed framework jointly optimizes task loss, alignment loss, and distillation loss, while the adaptive weighting module dynamically adjusts their contributions to balance representation learning and task-specific optimization.

The quantitative comparison is summarized in Table 1, where the proposed method demonstrates clear advantages over baseline approaches. Specifically, the accuracy improves to 91.2%, outperforming both the baseline LLM and standard fine-tuning methods. At the same time, the training time is significantly reduced, indicating that the integration of knowledge distillation and adaptive optimization leads to more efficient learning. The reduction in MMD further confirms that the alignment mechanism effectively minimizes distribution discrepancies between domains, thereby improving generalization performance. Additionally, the H-Score shows a consistent improvement, reflecting a better balance between accuracy and robustness.

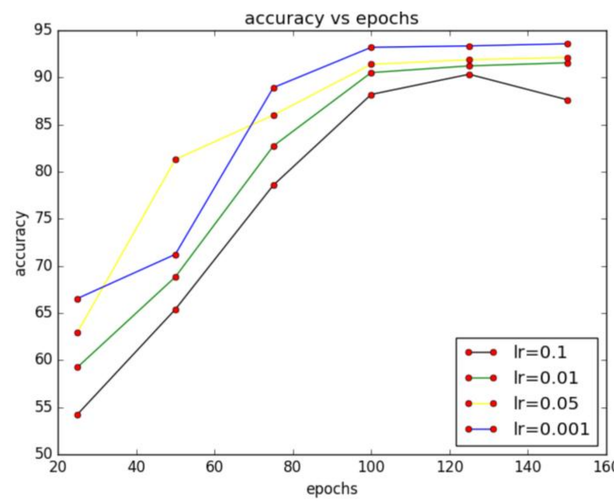

Figure 3. Training convergence and performance comparison between baseline and proposed methods

To further analyze the optimization process, the convergence behavior is illustrated in Figure 3. It can be observed that the proposed framework converges faster than baseline methods, particularly in the early training stages. This improvement is mainly attributed to the distillation component, which provides additional supervisory signals and guides the model toward more informative representations. Compared to conventional single-objective training, the proposed approach exhibits smoother convergence curves and reduced oscillation, indicating enhanced training stability.

Table 1: Performance comparison across different models

Model	Accuracy (%)	Training Time (h)	MMD ↓	H-Score ↑
Baseline LLM	88.2	12.5	0.132	0.781
Fine-tuned LLM	89.6	11.2	0.118	0.802

Distillation Model	90.3	9.6	0.11	0.827
Proposed Framework	91.2	8.2	0.089	0.854

Another important observation is the effectiveness of the dynamic weighting mechanism. In the early phase of training, the model prioritizes alignment and distillation losses, which helps reduce domain discrepancies and accelerate representation learning. As training progresses, the weight gradually shifts toward the task loss, allowing the model to refine task-specific performance. This adaptive adjustment mitigates the issue of conflicting optimization objectives and leads to more efficient convergence compared to fixed-weight strategies.

To further validate the contributions of different components, ablation experiments are conducted and the results are consistent with the trends shown in Table 1 and Figure 3. When the alignment loss is removed, the model exhibits higher domain discrepancy and reduced cross-domain accuracy, indicating that distribution alignment plays a crucial role in generalization. Similarly, removing the distillation component results in slower convergence and lower final performance, demonstrating its importance in knowledge transfer and training acceleration. Furthermore, replacing the adaptive weighting mechanism with static coefficients leads to unstable training behavior and suboptimal results.

In cross-domain evaluation, the proposed framework maintains strong performance even under significant distribution shifts, whereas conventional fine-tuning methods suffer from noticeable degradation. This robustness can be attributed to the joint optimization of alignment and distillation objectives, which enables the model to learn domain-invariant and transferable representations. Such capability is essential for real-world applications where data distributions are often non-stationary.

Finally, the proposed framework achieves a favorable trade-off between performance and efficiency, as reflected in both Table 1 and Figure 3. By integrating knowledge distillation and adaptive scheduling, the model not only improves accuracy but also reduces computational overhead, making it suitable for deployment in resource-constrained environments. The overall results demonstrate that the proposed method effectively addresses key challenges in LLM optimization, including multi-objective learning, domain adaptation, and training efficiency, while maintaining strong empirical performance.

5. Conclusion

This paper presents a unified optimization framework for Large Language Models that integrates multi-objective learning, adaptive knowledge distillation, and distribution alignment. By jointly optimizing task performance, feature consistency, and efficiency, the proposed method addresses key limitations of existing approaches. Experimental results demonstrate significant improvements in accuracy, convergence speed, and robustness across diverse benchmarks. Future work will explore the extension of this framework to multimodal models and real-time deployment scenarios.

References

[1] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, “Scaling laws for neural language models,” arXiv preprint arXiv:2001.08361, 2020.

[2] J. Howard and S. Ruder, “Universal language model fine-tuning for text classification,” in Proc. 56th Annu. Meeting of the Association for Computational Linguistics (ACL), 2018, pp. 328-339.

[3] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” arXiv preprint arXiv:1503.02531, 2015.

[4] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in Proc. Int. Conf. Machine Learning (ICML), 2020, pp. 1597-1607.

[5] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, “Adversarial discriminative domain adaptation,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 7167-7176.

[6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, and I. Polosukhin, "Attention is all you need," Advances in Neural Information Processing Systems, vol. 30, 2017.

[7] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," Proceedings of NAACL-HLT, pp. 4171 – 4186, 2019.

[8] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, and V. Stoyanov, "RoBERTa: A robustly optimized BERT pretraining approach," arXiv preprint arXiv:1907.11692, 2019.

[9] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," Journal of Machine Learning Research, vol. 21, no. 140, pp. 1 – 67, 2020.

[10] R. Pope, S. Douglas, A. Chowdhery, J. Devlin, J. Bradbury, J. Heek, and J. Dean, "Efficiently scaling transformer inference," Proceedings of Machine Learning and Systems, vol. 5, pp. 606 – 624, 2023.

[11] D. Lepikhin, H. Lee, Y. Xu, D. Chen, O. Firat, Y. Huang, and Z. Chen, "GShard: Scaling giant models with conditional computation and automatic sharding," arXiv preprint arXiv:2006.16668, 2020.

[12] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, and A. Ray, "Training language models to follow instructions with human feedback," Advances in Neural Information Processing Systems, 2022.

[13] Y. Bai, S. Kadavath, S. Kundu, A. Askell, J. Kernion, A. Jones, and J. Kaplan, "Constitutional AI: Harmlessness from AI feedback," arXiv preprint arXiv:2212.08073, 2022.

[14] J. Yang, S. Sun, Y. Wang, Y. Wang, X. Yang, and C. Zhang, "Semantic alignment and output constrained generation for reliable LLM-based classification," 2026.

[15] C. Wen, A. Zhu, R. Long, H. Huang, J. Jiang, and C. S. Lee, "CalibJudge: Calibrated LLM-as-a-judge for multilingual RAG with uncertainty-aware scoring," 2026.

[16] C. Shao, Y. Zi, Y. Deng, H. Liu, C. Zhang, and Y. Ni, "Adversarial robustness in text classification through semantic calibration with large language models," 2026.

[17] Y. Luan, "Iterative self-questioning supervision with semantic calibration for stable reasoning chains in large language models," 2026.

[18] Y. Li, Y. Tang, K. Wu, Y. Yang, Y. Li, and Y. Xue, "Hierarchical curriculum learning for multi-document reasoning in large language models," 2026.

[19] K. Gao, H. Zhu, R. Liu, J. Li, X. Yan, and Y. Hu, "Contextual trust evaluation for robust coordination in large language model multi-agent systems," 2025.

[20] J. Chen, F. Wang, T. Guan, Y. Ma, L. Yang, and Y. Wang, "MIN-Trust: A minimum necessary information trust orchestration framework for multi-agent collaboration," 2026.

[21] F. Wang, Y. Ma, T. Guan, Y. Wang, and J. Chen, "Autonomous learning through self-driven exploration and knowledge structuring for open-world intelligent agents," 2026.

[22] L. Yang, T. Guan, Y. Ma, Z. Li, Z. Fang, and F. Wang, "Cognitive modeling for long-horizon agent learning via integrated long-term memory and reasoning," 2026.

[23] R. Meng, S. Y. Huang, X. Zhang, Z. Huang, K. Zeng, and Y. Yang, "Constraint-consistent skill composition for reliable zero-shot task generalization in LLM agents."

[24] O. Sener and V. Koltun, "Multi-task learning as multi-objective optimization," Advances in Neural Information Processing Systems, vol. 31, 2018.

[25] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses," Proceedings of CVPR, pp. 7482 – 7491, 2018.

[26] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, and V. Lempitsky, "Domain-adversarial training of neural networks," Journal of Machine Learning Research, vol. 17, no. 59, pp. 1 – 35, 2016.

[27] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," Proceedings of ICML, pp. 97 – 105, 2015.

[28] C. Hu, Z. Cheng, D. Wu, Y. Wang, F. Liu, and Z. Qiu, "Structural generalization for microservice routing using graph neural networks," Proceedings of AIAC, pp. 278 – 282, 2025.

[29] J. Jiang, C. Shao, C. Zhang, N. Lyu, and Y. Ni, "Adaptive AI spatiotemporal modeling with dependency drift awareness for anomaly detection in large-scale clusters," 2025.

[30] C. Chiang, D. Li, R. Ying, Y. Wang, Q. Gan, and J. Li, "Deep learning-based dynamic graph framework for robust corporate financial health risk prediction," Proceedings of ICMLML, pp. 98 – 105, 2025.

[31] X. Liang, Y. Zhao, M. Chang, R. Zhou, K. Cao, and Y. Zheng, "Spatiotemporal risk representation learning using transformers and graph structure," 2026.

[32] Q. Zhang, "Adaptive resource scheduling in distributed computing via multi-agent reinforcement learning and graph convolutional modeling," Transactions on Computational and Scientific Methods, vol. 4, no. 11, 2024.

[33] N. Lyu, Y. Wang, Z. Cheng, Q. Zhang, and F. Chen, "Multi-objective adaptive rate limiting in microservices using deep reinforcement learning," Proceedings of AIIP, pp. 862 – 869, 2025.

[34] C. Wang, C. S. Lee, X. Yang, Z. Qiu, and Y. Tang, "Deep reinforcement learning guided by game-theoretic structure for multi-agent resource allocation and scheduling."

[35] H. Feng, Y. Wang, R. Fang, A. Xie, and Y. Wang, "Federated risk discrimination with siamese networks for financial transaction anomaly detection," Proceedings of DECS, pp. 231 – 236, 2025.

[36] C. Chen, R. Fang, and J. Lai, "Causal representation learning for robust and interpretable audit risk identification in financial systems," Proceedings of ICEMME, 2026.

[37] N. Chen, S. Sun, Y. Wang, Z. Li, A. Zhu, and Y. Lu, "Few-shot financial fraud detection using meta-learning and large language models," Proceedings of CSMT, pp. 822 – 826, 2025.

[38] X. Yang, S. Sun, Y. Li, Y. Xing, M. Wang, and Y. Wang, "CaliCausalRank: Calibrated multi-objective ad ranking with robust counterfactual utility optimization," arXiv preprint arXiv:2602.18786, 2026.

[39] Q. Liu, Y. Zhang, S. Chen, Z. Liu, Y. Xu, and H. Cui, "Uncertainty-aware marketing attribution inference and budget decision-making with intelligent agents," 2026.

[40] X. Song, Y. Huang, J. Guo, Y. Liu, and Y. Luan, "Multi-scale feature fusion and graph neural network integration for text classification with large language models," arXiv preprint arXiv:2511.05752, 2025.

[41] X. Song, Y. Liu, Y. Luan, J. Guo, and X. Guo, "Controllable abstraction in summary generation for large language models via prompt engineering," arXiv preprint arXiv:2510.15436, 2025.

[42] R. Ying, Q. Liu, Y. Wang, and Y. Xiao, "AI-based causal reasoning over knowledge graphs for data-driven enterprise performance analysis," Proceedings of CSMT, pp. 770 – 775, 2025.

[43] H. Chen, Y. Lu, Y. Wei, J. Lyu, R. Wu, and C. Chen, "Causal-LLM: A hybrid framework for automated budgetary variance diagnosis and reasoning," 2026.

[44] R. Yan, Y. Ou, S. Sun, N. Chen, K. Zhou, and Y. Shu, "DualShiftNet: Joint class-imbalance and distribution-shift aware learning for business risk prediction," 2026.

[45] L. Yan, Q. Wang, and J. Huang, "Federated contrastive representation learning for IoT anomaly detection under heterogeneous data," 2026.

[46] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," arXiv preprint arXiv:1510.00149, 2015.

[47] J. Frankle and M. Carbin, "The lottery ticket hypothesis," Proceedings of ICLR, 2019.

[48] T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré, "FlashAttention: Fast and memory-efficient exact attention," Advances in Neural Information Processing Systems, vol. 35, 2022.

[49] X. Chen, S. U. Gadgil, and J. Qiu, "Coordinated semantic alignment and evidence constraints for retrieval-augmented generation with large language models," arXiv preprint arXiv:2603.04647, 2026.