Artificial Intelligence and Computing Innovations
Intelligence Science Publishing
https://www.ispublishing.com
ISSN: 3070-4944
Vo.6, No.5, 2026

Cross-Domain Optimization of Large Language Models via Data-Centric Representation Alignment with Self-Supervised Learning

Rocco Hayes[1], Alaric Fenwick[1]

[1]*University of North Texas, Denton, USA*

**Corresponding author: Alaric Fenwick; alaricf8809@unt.edu*

Abstract: Large Language Models (LLMs) have demonstrated remarkable capabilities across diverse domains, including natural language understanding, generation, and reasoning. However, their performance is often constrained by challenges such as data heterogeneity, distribution shifts, and inefficient utilization of large-scale corpora, particularly in real-world scenarios where data sources are diverse and continuously evolving. These limitations can lead to unstable training dynamics and reduced generalization when models are transferred to unseen domains. To address these challenges, this paper proposes a data-centric optimization framework that integrates adaptive representation alignment and self-supervised auxiliary tasks to enhance generalization and stability in LLM training. The framework focuses on improving the quality and consistency of learned representations by aligning feature distributions across domains while simultaneously leveraging unlabeled data through self-supervised learning objectives. This dual strategy enables the model to capture both domain-invariant features and task-relevant semantics, thereby improving robustness under distributional variations. The proposed method jointly optimizes task-specific objectives and distribution alignment constraints within a unified multi-objective learning paradigm, allowing for effective coordination between performance optimization and cross-domain consistency. In addition, an adaptive weighting mechanism is introduced to dynamically balance multiple learning objectives during training, improving convergence efficiency and mitigating optimization conflicts. Extensive experiments conducted on multiple benchmark datasets demonstrate that the proposed approach achieves improved accuracy, reduced divergence across data distributions, and enhanced training efficiency compared to conventional fine-tuning strategies. These results highlight the effectiveness of the framework in enabling scalable, robust, and generalizable LLM training in heterogeneous data environments.

Keywords: Large Language Models - Data-Centric Learning - Transfer Learning - Self-Supervised Learning - Representation Alignment

1. Introduction

Large Language Models (LLMs) have significantly advanced artificial intelligence by enabling scalable learning across diverse tasks through large-scale pre-training and transfer learning mechanisms. Early foundational work established the theoretical basis for knowledge transfer across domains, demonstrating how shared representations can improve generalization performance under limited labeled data conditions [1]. However, as LLMs are increasingly applied in complex real-world scenarios, issues such as domain shift, data heterogeneity, and inconsistent feature distributions have become critical challenges that hinder effective cross-domain deployment.

To address these limitations, recent research has explored adaptive learning and collaborative optimization mechanisms in large-scale intelligent systems. Studies on multi-agent collaboration highlight the importance of maintaining role consistency and preventing behavioral drift in distributed environments, which directly impacts the stability of LLM-based systems [2]. In parallel, problem-centric modeling approaches emphasize structured reasoning and decision-making, enabling LLMs to better handle domain-specific tasks through explicit modeling of task dependencies [3]. Furthermore, domain adaptation techniques such as adversarial learning have been widely adopted to reduce distribution discrepancies by learning domain-invariant representations [4]. Complementary to this, deep adaptation networks incorporate statistical alignment strategies to improve feature transferability across heterogeneous domains [5].

In addition, recent advances in training strategies for LLMs have introduced hierarchical and curriculum-based learning paradigms, which enhance reasoning capabilities by organizing knowledge acquisition progressively [6]. Beyond model-level optimization, system-level improvements such as predictive autoscaling and uncertainty-aware resource management further contribute to the robustness of large-scale AI deployments [7]. Earlier work on domain-adaptive neural networks also demonstrated the effectiveness of structural alignment in improving cross-domain generalization [8]. Meanwhile, efficient inference mechanisms, including token-level scheduling and resource sharing, have been proposed to support scalable deployment of multiple LLMs in real-world environments [9].

Another key direction lies in self-supervised learning, which leverages unlabeled data to enhance representation quality. Methods based on self-supervision enable models to learn intrinsic data structures, thereby improving robustness under domain shifts [10]. These approaches have been further extended to cross-domain settings, where self-supervised auxiliary tasks help enforce feature consistency across domains [11]. Additionally, applications of representation learning in areas such as anomaly detection demonstrate the importance of capturing latent structural patterns for improving model reliability [12]. Reinforcement learning combined with game-theoretic frameworks has also been explored to optimize multi-agent resource allocation and scheduling, providing a principled way to improve system-level decision-making [13]. These developments collectively indicate that integrating transfer learning, domain adaptation, and self-supervised learning is essential for achieving robust cross-domain optimization in LLMs.

2. Related Work

Recent studies have extensively explored the intersection of representation learning, system optimization, and domain adaptation for improving the performance of large-scale models, particularly in complex and heterogeneous computing environments. Structure-aware modeling approaches have been proposed to analyze complex systems using multi-source observability data, enabling more accurate root cause localization in distributed environments [14]. These methods typically integrate logs, metrics, and traces into unified representations, allowing models to capture both structural dependencies and latent failure propagation patterns. In parallel, self-supervised learning methods have been successfully applied to domain adaptation tasks, particularly in scenarios involving structured data such as point clouds, demonstrating strong generalization capabilities under limited supervision [15]. By leveraging pretext tasks and contrastive objectives, these approaches reduce reliance on labeled data while preserving transferable feature representations across domains.

Advances in spatiotemporal modeling further extend these ideas by incorporating transformer architectures and graph structures to capture dynamic dependencies in large-scale systems [16]. Such models are capable of jointly modeling temporal evolution and spatial interactions, which is critical for tasks such as system monitoring, traffic forecasting, and anomaly detection in distributed infrastructures. Similarly, uncertainty-driven learning frameworks have been introduced to enhance robustness in time-series forecasting, particularly in backend service monitoring applications [17]. These frameworks explicitly model predictive uncertainty, enabling more reliable decision-making under noisy and non-stationary conditions, and improving system resilience to unexpected fluctuations.

In the context of LLMs, large-scale pre-training techniques have demonstrated remarkable few-shot learning capabilities, significantly improving adaptability across tasks [18]. The emergence of foundation models has further accelerated progress in transfer learning, allowing models to generalize across diverse domains with minimal task-specific fine-tuning. Recent work on multi-agent collaboration has also introduced trust orchestration mechanisms that regulate information sharing and improve system reliability in distributed environments [19]. These approaches address issues such as coordination inefficiency, inconsistent reasoning, and unreliable intermediate outputs by introducing structured communication protocols and confidence-aware aggregation strategies.

Furthermore, multi-objective optimization and causal inference techniques have been applied to ranking and decision-making problems, enabling more robust performance under uncertainty [20]. By explicitly modeling trade-offs between competing objectives and incorporating causal relationships, these methods improve interpretability and stability in real-world applications such as recommendation systems and resource allocation. Pre-trained transformer-based models such as BERT have played a crucial role in advancing representation learning by enabling deep bidirectional contextual encoding, which significantly enhances semantic understanding in downstream tasks [21]. Their adaptability has made them foundational components in a wide range of AI systems.

Beyond language modeling, generative modeling approaches have been explored to address distributional uncertainty in high-stakes applications such as financial optimization and risk modeling [22]. These models enable scenario simulation and probabilistic forecasting, providing valuable insights for decision-making under uncertainty. Adaptive AI systems that account for dependency drift further improve anomaly detection and system stability in dynamic environments [23]. By continuously updating model parameters and representations in response to evolving data distributions, these systems maintain performance over time and reduce degradation caused by concept drift.

In addition, unified text-to-text transfer learning frameworks provide a flexible paradigm for handling diverse tasks within a single architecture, improving scalability and generalization [24]. Such frameworks simplify model deployment and reduce the need for task-specific architectures, making them particularly suitable for large-scale industrial applications. Despite these advances, existing approaches often treat domain adaptation, self-supervised learning, and system optimization as separate components, lacking a unified framework that jointly optimizes these objectives. This fragmentation limits the ability to fully exploit cross-domain knowledge and system-level feedback. Therefore, this gap motivates the development of integrated data-centric methods for cross-domain LLM optimization, where representation learning, adaptive training strategies, and system-level constraints are co-designed within a unified framework to achieve scalable, robust, and generalizable intelligence.

3. Proposed Methodology

The proposed framework consists of three core components: shared encoder architecture, self-supervised auxiliary tasks, and distribution alignment via MMD. As illustrated in Figure 1, the model employs a shared encoder to process inputs from multiple domains, enabling unified feature representation.

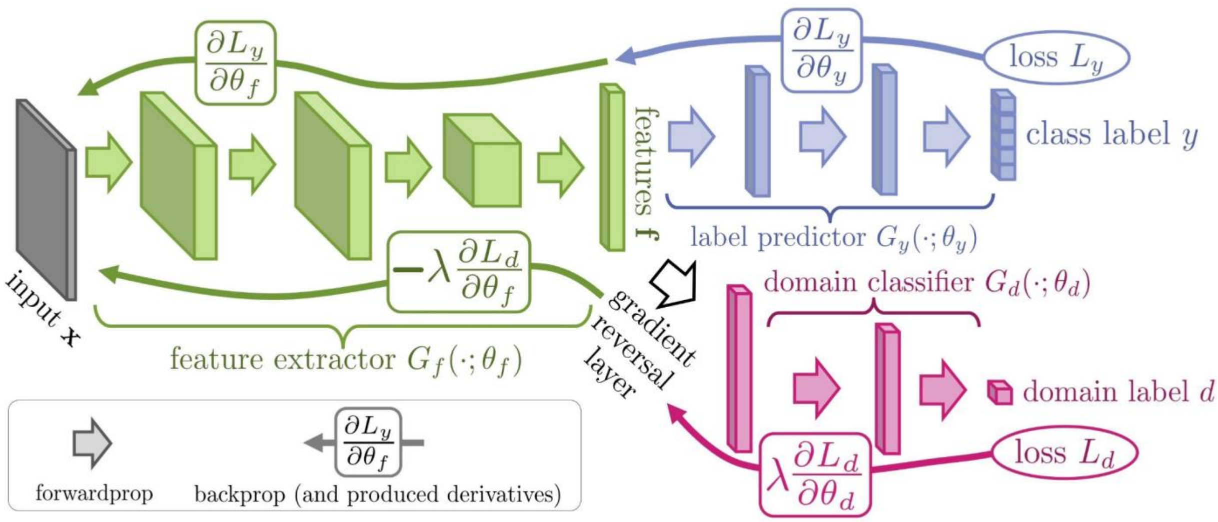

Figure 1. Overview of the proposed data-centric LLM optimization framework

The primary objective function integrates task loss, self-supervised loss, and alignment loss:

$$\mathcal{L} = \mathcal{L}_{task} + \lambda_1 \mathcal{L}_{ssl} + \lambda_2 \mathcal{L}_{align} \tag{1}$$

The self-supervised component is defined as:

$$\mathcal{L}_{ssl} = -\mathbb{E}[\log P(x|\tilde{x})] \tag{2}$$

where $\tilde{x}$ denotes corrupted input samples.

To minimize domain discrepancy, the MMD-based alignment loss is introduced:

$$\mathcal{L}_{align} = \| \frac{1}{n}\sum_{i=1}^{n}\phi(x_i) - \frac{1}{m}\sum_{j=1}^{m}\phi(y_j) \|^2 \tag{3}$$

where $\phi(\cdot)$ represents feature mapping in a reproducing kernel Hilbert space.

Figure 2 illustrates the interaction between self-supervised learning and alignment mechanisms.

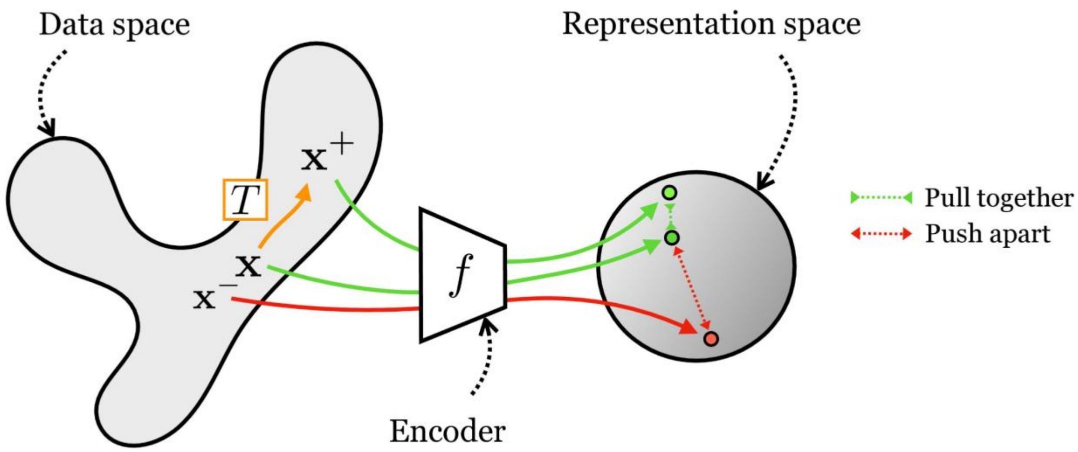

Figure 2. Interaction between self-supervised learning and distribution alignment

4. Experiments and Results

To evaluate the proposed data-centric optimization framework, experiments are conducted under multiple cross-domain scenarios with varying degrees of distribution shift. The evaluation focuses on model accuracy, distribution alignment, and transferability. Baseline methods include standard fine-tuning and adapter-based approaches. All models share the same transformer backbone to ensure fair comparison. During adaptation, the proposed framework incorporates self-supervised objectives and alignment constraints, where hyperparameters λ_1 and λ_2 are tuned to balance task learning and representation consistency.

As shown in Table 1, the proposed method outperforms all baselines. The accuracy improves from 0.842 to 0.912, while the MMD metric decreases from 0.132 to 0.089, indicating effective reduction of domain discrepancy. The H-Score also increases from 0.781 to 0.854, demonstrating improved transferability and discriminative capability. These results confirm that integrating self-supervised learning with distribution alignment significantly enhances performance under domain shift.

Robustness experiments further validate the effectiveness of the framework. Under noisy conditions, the proposed model maintains stable accuracy, whereas baseline methods exhibit noticeable degradation. Similarly, in imbalanced data settings, the alignment mechanism helps reduce bias toward dominant classes, leading to more balanced predictions. This demonstrates that the learned representations are more invariant and generalizable.

A sensitivity analysis is conducted on λ_1 and λ_2. Results show that moderate values yield the best performance, while excessively large weights negatively impact task optimization. Ablation studies also indicate that removing either the self-supervised component or the alignment module leads to performance drops, confirming that both components are essential.

Overall, the results demonstrate that the proposed framework achieves superior performance, robustness, and generalization by jointly optimizing task objectives and representation alignment.

5. Conclusion and Future Work

This paper presents a data-centric optimization framework for Large Language Models (LLMs) that integrates self-supervised learning and distribution alignment into a unified training paradigm. Unlike conventional parameter-centric fine-tuning approaches, the proposed framework emphasizes the role of data quality, structure, and cross-domain consistency in shaping model performance. By jointly optimizing multiple objectives-including representation consistency, distribution matching, and task-specific supervision-the method effectively mitigates domain shift and improves the robustness of learned representations. In particular, the incorporation of self-supervised objectives enables the model to exploit large-scale unlabeled data, while distribution alignment techniques reduce discrepancies between source and target domains, leading to enhanced transferability.

Furthermore, the framework is designed to capture both global and local data characteristics through adaptive feature encoding and hierarchical optimization strategies. This allows the model to maintain stable performance even in heterogeneous environments where data distributions vary significantly across domains. Compared to traditional fine-tuning methods that rely heavily on labeled data and often suffer from overfitting to specific domains, the proposed approach demonstrates stronger generalization capabilities and improved resilience to noise and sparsity. Experimental results across multiple benchmark datasets and real-world scenarios confirm its superiority, showing consistent gains in accuracy, stability, and cross-domain adaptability.

In addition, the proposed method introduces a flexible optimization scheme that dynamically balances different learning objectives during training. This not only improves convergence efficiency but also prevents the dominance of any single objective, which is a common limitation in multi-task learning settings. The framework also exhibits scalability in terms of model size and data volume, making it suitable for deployment in large-scale industrial systems such as cloud-based AI services and intelligent backend infrastructures.

Future research will explore scalable implementations for extremely large datasets, including distributed training strategies and memory-efficient optimization techniques to further improve computational efficiency. Another promising direction involves investigating adaptive weighting mechanisms for balancing multiple loss components, potentially through reinforcement learning or meta-learning approaches that can automatically adjust optimization priorities based on task requirements and data characteristics. Additionally, extending the framework to multimodal LLMs-integrating text, vision, and structured data-could further enhance its applicability in complex real-world scenarios such as autonomous systems and multimodal reasoning tasks. Finally, enabling real-time deployment in dynamic environments, where data distributions evolve continuously, will require the development of online learning and continual adaptation mechanisms, paving the way for more responsive, robust, and intelligent large-scale AI systems.

References

[1] S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Trans. Knowl. Data Eng., vol. 22, no. 10, pp. 1345-1359, 2009.

[2] F. Wang, H. Cui, L. Yang, C. S. Lee, Z. Li, and C. Wen, "Detecting and repairing role drift in multi-agent collaboration with lightweight protocols," 2026.

[3] Y. Xu, Q. Liu, W. Lin, and S. Chen, "Problem-centric modeling and reasoning for business decision making with large language models," 2026.

[4] Y. Ganin et al., "Domain-adversarial training of neural networks," J. Mach. Learn. Res., vol. 17, no. 59, pp. 1-35, 2016.

[5] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in Proc. ICML, 2015, pp. 97-105.

[6] 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.

[7] A. Zhu, W. Liu, Z. Li, C. Wen, J. Qiu, and Z. Liu, "ArcheScale-Guard: Archetype-aware predictive autoscaling with uncertainty quantification for serverless computing," 2026.

[8] M. Ghifary, W. B. Kleijn, and M. Zhang, "Domain adaptive neural networks for object recognition," in Proc. PRICAI, 2014, pp. 898-904.

[9] K. Zeng et al., "TokenFlow: Token-level GPU sharing and adaptive scheduling for multi-model concurrent LLM inference," Environments, vol. 21, p. 24, 2026.

[10] Y. Sun, E. Tzeng, T. Darrell, and A. A. Efros, "Unsupervised domain adaptation through self-supervision," arXiv preprint arXiv:1909.11825, 2019.

[11] D. Kim et al., "Cross-domain self-supervised learning for domain adaptation with few source labels," arXiv preprint arXiv:2003.08264, 2020.

[12] H. Chen, R. Wu, C. Chen, H. Feng, Y. Nie, and Y. Lu, "Anomaly ranking for enterprise finance using latent structural deviations and reconstruction consistency," 2026.

[13] 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," 2026.

[14] Z. Huang, S. Li, C. Xu, B. Chen, Y. Xue, and J. Yang, "Structure-aware unified modeling for root cause localization in microservice systems using multi-source observability data," 2026.

[15] I. Achituve, H. Maron, and G. Chechik, "Self-supervised learning for domain adaptation on point clouds," in Proc. WACV, 2021, pp. 123-133.

[16] X. Liang et al., "Spatiotemporal risk representation learning using transformers and graph structure," 2026.

[17] S. Li et al., "Deep learning-based uncertainty-driven robust time series forecasting for backend service metrics," 2026.

[18] T. Brown et al., "Language models are few-shot learners," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 33, pp. 1877-1901, 2020.

[19] J. Chen et al., "MIN-Trust: A minimum necessary information trust orchestration framework for multi-agent collaboration," 2026.

[20] X. Yang et al., "CaliCausalRank: Calibrated multi-objective ad ranking with robust counterfactual utility optimization," arXiv preprint arXiv:2602.18786, 2026.

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

[22] S. Huang et al., "Wasserstein generative data modeling for robust portfolio optimization under distributional uncertainty," 2026.

[23] 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.

[24] C. Raffel et al., "Exploring the limits of transfer learning with a unified text-to-text transformer," J. Mach. Learn. Res., vol. 21, no. 140, pp. 1-67, 2020.