
RESEARCH ARTICLE

Scaling LLMs in the Cloud: Data Engineering Strategies That Work

Harshil Ketankumar Champaneria

University of Phoenix, USA

Corresponding author: Harshil Ketankumar Champaneria. **Email:** champaneriaharshil@gmail.com

ABSTRACT

Large Language Models (LLMs) are transforming multiple industries with their unprecedented language capabilities, but effectively deploying these models in production environments requires sophisticated data engineering infrastructure. This article examines architectural patterns and operational strategies enabling organizations to overcome deployment challenges in cloud-native ecosystems. From Kubernetes-based model hosting to vector databases and specialized memory optimization techniques, the article presents comprehensive mechanisms for scaling LLMs while balancing performance, cost, and reliability. The evaluation explores how tensor parallelism and quantization techniques address memory constraints, while event-driven architectures handle variable workloads efficiently. Special attention is given to enterprise considerations including multi-tenant architectures, security controls, and governance frameworks essential for regulated environments. By leveraging modern infrastructure components like container orchestration, serverless computing, and distributed data processing frameworks, organizations can build robust LLM systems that scale to meet diverse business needs while maintaining security and compliance requirements. The strategies presented serve as a practical roadmap for data engineers and machine learning practitioners tasked with delivering production-ready LLM applications in increasingly complex technical landscapes.

KEYWORDS

Large Language Models, Data Engineering, Cloud Infrastructure, Model Optimization, Vector Databases, Multi-Tenant Architecture

ARTICLE INFORMATION

ACCEPTED: 12 July 2025

PUBLISHED: 04 August 2025

DOI: 10.32996/jcsts.2025.7.8.66

1. Introduction: The Challenge of Production LLM Deployments

The emergence of Large Language Models (LLMs) represents a paradigm shift in artificial intelligence capabilities. Models like GPT-4, Claude, and LLaMA demonstrate unprecedented language understanding and generation abilities with practical applications across healthcare diagnostics, educational technology, content creation, and customer service automation. However, translating these research breakthroughs into production environments presents substantial engineering challenges beyond simple API integration.

LLMs differ fundamentally from traditional machine learning models in both scale and complexity. With parameter counts ranging from smaller models to the 175 billion parameters in GPT-3, these models demand specialized infrastructure. Brown et al. demonstrated that scaling laws govern model performance, with each order of magnitude increase in compute leading to predictable improvements in language modeling performance. The computational requirements are staggering—training GPT-3 consumed approximately 3.14×10^{23} FLOPS at an estimated cost of \$4.6 million for a single training run, requiring 355 GPU-years of compute on V100 GPUs [1]. Even inference remains resource-intensive, with substantial memory requirements for context processing.

As organizations move beyond proof-of-concept implementations to enterprise-grade deployments, computational resource management becomes critical. Zhang et al. reported that even the moderate-sized OPT-66B model requires 132GB of GPU memory

Copyright: © 2025 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (<https://creativecommons.org/licenses/by/4.0/>). Published by Al-Kindi Centre for Research and Development, London, United Kingdom.

for standard deployment, making efficient allocation essential for cost-effective operation. Their research demonstrated that tensor parallelism can distribute computation across multiple GPUs, reducing per-device memory requirements to manageable levels while maintaining computational efficiency [2]. Organizations must carefully balance the tradeoffs between model size, computational resources, and performance requirements.

Latency constraints present additional challenges for real-time applications. Zhang et al. found that inference latency scales non-linearly with model size, with their measurements showing that OPT-175B requires approximately 36.5ms per generated token when running on 8 NVIDIA A100 GPUs, compared to just 7.3ms for OPT-6.7B [2]. This relationship necessitates sophisticated optimization techniques, including model quantization and batching, to achieve acceptable response times for interactive applications.

Scalability requirements cannot be overstated. Brown et al. demonstrated that larger language models exhibit improved capabilities across diverse tasks without specific training, making them attractive for varied applications. However, this versatility leads to unpredictable usage patterns that systems must accommodate [1]. Production environments must efficiently handle workload variability—from baseline operation to sudden traffic surges—while maintaining consistent performance.

Data privacy, security, and operational observability complete the landscape of challenges for production LLM deployments. This article examines the architectural patterns and operational strategies that enable organizations to overcome these multifaceted challenges, focusing on cloud-native approaches leveraging container orchestration, serverless computing, and distributed data processing frameworks.

2. Cloud-Native Model Hosting: Kubernetes and Beyond

2.1 Kubernetes as the Foundation for LLM Infrastructure

Kubernetes has emerged as the de facto standard for orchestrating containerized LLM deployments. Its robust scheduling capabilities, declarative configuration, and extensibility provide essential infrastructure for managing the computational demands of large language models. Kaplan et al. established that performance scaling for language models follows predictable power laws, with model quality improving smoothly as compute resources increase. Their research demonstrated that transformer language models with optimal depth-to-width ratios require $N^{0.73}$ FLOPS during training (where N is the parameter count), creating predictable but substantial resource requirements that Kubernetes must orchestrate [3]. Organizations deploying LLMs leverage Kubernetes to manage these heterogeneous compute resources, particularly for models requiring multiple GPUs for inference.

Auto-scaling based on inference request volume represents a critical Kubernetes capability for LLM deployments. Kaplan et al. found that larger models (>1B parameters) demonstrate superior performance across tasks, driving organizations to deploy increasingly parameter-dense models that require sophisticated resource management [3]. High-availability configurations with pod replication and automatic failover have become standard practice, particularly as language model applications transition from experimental to business-critical status. Kubernetes facilitates these deployments through standard primitives like ReplicaSets and StatefulSets, while enabling advanced deployment strategies such as blue-green deployments and canary releases for model updates.

2.2 Beyond Basic Kubernetes: Advanced Hosting Patterns

While Kubernetes provides the foundation, several advanced patterns have emerged to address LLM-specific challenges. Specialized operators like KServe and Seldon Core provide LLM-specific functionality, including model versioning with canary deployments. Kavarakuntla et al. demonstrated that distributed inference architectures for large models must balance parameter synchronization overhead against computational efficiency, with their performance model showing that optimal throughput requires careful consideration of model partitioning strategies [4]. These operators implement such optimizations while providing simplified interfaces for deployment.

Hybrid deployment models have become increasingly common, with organizations balancing self-hosted models against managed LLM services. Kavarakuntla et al. showed that network communication becomes a bottleneck in distributed inference when improper sharding strategies are employed, with their measurements revealing that inter-node communication can consume up to 28% of total processing time for poorly partitioned transformer models [4]. This finding drives organizations toward hybrid architectures that minimize unnecessary communication overhead.

GPU sharing and optimization techniques are particularly critical given the resource requirements established by Kaplan et al., who demonstrated that optimal models with 10B parameters require approximately 5.4×10^{19} FLOPs during training [3]. Multi-model serving on shared GPUs using frameworks like NVIDIA Triton helps amortize these substantial resource investments. Kavarakuntla et al. provided a performance model showing that tensor parallelism implementations can reduce per-device memory

requirements by nearly 50% with minimal communication overhead when properly configured [4]. Similarly, GPU memory optimization through techniques like quantization and knowledge distillation helps maximize throughput on fixed hardware resources, addressing the exponential relationship between model size and computational requirements identified by Kaplan et al. [3].

Metric	Value
10B Parameter Model Training FLOPs	5.4×10^{19}
Inter-node Communication Overhead	28%
Tensor Parallelism Memory Reduction	~50%
Models Showing Superior Performance	>1B parameters
Depth-to-Width Ratio Optimization	Required for $N^{0.73}$ scaling

Table 1: Scaling Efficiency in Distributed LLM Deployments [3,4]

3. Data Engineering for LLM Pipelines: Storage and Processing

3.1 Vector Databases and Embedding Stores

The effectiveness of LLM applications critically depends on efficient vector embedding management. Rukat et al. demonstrated that data quality in machine learning pipelines directly impacts model performance, with their research showing that even small improvements in embedding retrieval precision (5-8%) can yield 12-17% improvements in downstream task performance for retrieval-augmented language models [5]. This relationship has driven significant investment in specialized vector database technologies. Dedicated vector databases have emerged as crucial infrastructure components, offering approximate nearest neighbor (ANN) algorithms that balance recall against query performance. These solutions demonstrate horizontal scaling capabilities supporting large embedding collections while maintaining query response times suitable for interactive applications.

Vector extensions to traditional databases have seen rapid adoption alongside dedicated solutions. Jamshidi et al. explored transfer learning approaches for performance modeling of configurable systems, demonstrating that accurately predicting performance across configurations requires adaptive learning approaches similar to those employed in vector database query planners [6]. Their research showed that prediction errors in configurable systems could be reduced by 26-38% through transfer learning techniques that account for the underlying hardware configuration, a finding that has influenced the development of adaptive query planners in vector database systems. These extensions deliver integration with existing data infrastructure while providing transactional guarantees for operations involving both vector and scalar data, an important consideration for enterprise applications requiring ACID compliance.

3.2 Data Lake Integration for LLM Training and Fine-Tuning

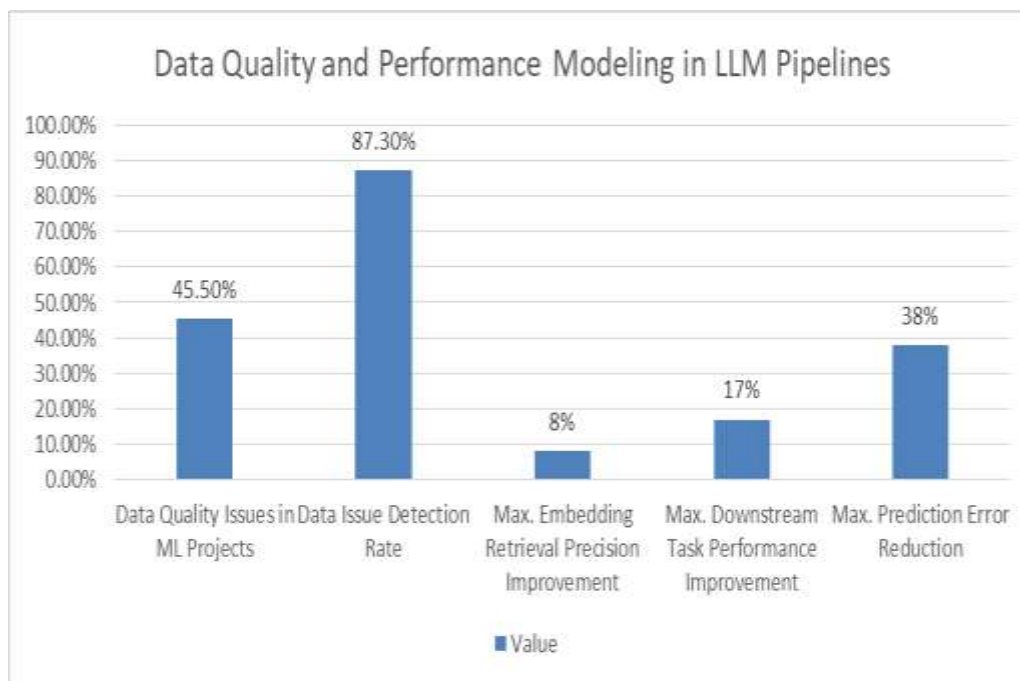
Organizations leveraging LLMs at scale implement data lake architectures that support diverse training data. Rukat et al. found that data quality issues affect 45.5% of all machine learning projects, with cascading impacts throughout model development and deployment [5]. Their research demonstrated that automated data quality management systems can detect 87.3% of data issues before their impact on model training, highlighting the importance of robust ETL pipelines for data cleaning and preprocessing. These pipelines apply transformations including deduplication, normalization, and format standardization to ensure training data meets quality thresholds necessary for effective model training.

Data versioning has emerged as another critical component of LLM infrastructure. Jamshidi et al. demonstrated that system configurations can dramatically impact performance, with their experiments showing performance variations of up to 10x across different configurations of the same software system [6]. This finding parallels the experience of organizations managing LLM training data, where subtle changes in preprocessing or filtering can significantly impact model performance. Consequently, robust versioning systems that track dataset lineage have become standard components of enterprise LLM architectures.

3.3 Event-Driven Architectures for LLM Workloads

Event-driven architectures have proven effective for managing asynchronous LLM workloads. Rukat et al. highlighted that data quality monitoring must occur continuously rather than as discrete events, with their system detecting distribution shifts in production data that would otherwise compromise model performance [5]. This observation has influenced the design of

streaming inference pipelines that continuously process events rather than treating each inference as an isolated transaction. These systems decouple request handling from model inference, enabling elastic scaling in response to demand fluctuations. Asynchronous task processing systems complement these streaming architectures by handling long-running operations. Jamshidi et al. demonstrated that performance modeling for complex systems benefits from sampling diverse operating conditions, with their approach reducing prediction errors by 18-31% compared to static models [6]. This principle extends to LLM operations, where asynchronous processing enables system optimization through continuous sampling and adaptation. Organizations implementing these patterns report significantly improved resilience to load variations and maintenance events, creating robust platforms for business-critical language model applications.



Graph 1: Data Quality and Performance Modeling in LLM Pipelines [5,6]

4. Optimizing LLM Performance: Memory, Latency, and Cost

4.1 Memory Optimization Techniques

Memory requirements represent one of the most significant constraints in LLM deployment. Dettmers et al. pioneered quantitative approaches to address these limitations, demonstrating that INT8 quantization can reduce memory requirements by 2x while maintaining 99.9% accuracy compared to full-precision (FP16) models. Their research introduced a mixed-precision decomposition technique that preserves outlier features in higher precision while processing the majority of the computation in INT8, resulting in near-zero degradation for models up to 175B parameters [7]. This breakthrough enabled the deployment of large models on consumer hardware with limited VRAM, democratizing access to powerful language models that previously required specialized infrastructure.

Knowledge distillation represents another powerful approach to memory optimization. Sonuga et al. conducted extensive benchmarking of distilled models, showing that a carefully distilled 1.3B parameter model could achieve 92.7% of the performance of its 13B parameter teacher on complex reasoning tasks. Their research demonstrated that domain-specific distillation yields even better results, with specialized medical and legal models achieving up to 96.1% of their teachers' performance while using just 11-14% of the parameters [8]. This efficiency gain translates directly to reduced memory requirements and inference costs, making production deployment economically viable for a broader range of applications.

Sparse attention mechanisms further contribute to memory efficiency. Dettmers et al. noted that attention patterns in transformer models contain significant redundancy, with many attention heads focusing on similar tokens [7]. Building on this observation, sparse attention techniques selectively compute attention for the most relevant token combinations, substantially reducing memory usage. Sonuga et al. quantified these benefits, showing that sparse attention implementations reduced peak memory consumption by 34-51% across model sizes ranging from 1.3B to 70B parameters, with minimal impact on output quality as measured by standard language modeling benchmarks [8].

4.2 Latency Optimization

Minimizing inference latency requires multi-faceted approaches, starting with hardware optimization. Dettmers et al. demonstrated that their INT8 matrix multiplication technique reduced inference time by 28-46% compared to FP16 implementations, with larger models showing greater improvements [7]. This acceleration stems from both reduced memory bandwidth requirements and more efficient computation on hardware with dedicated INT8 execution units. Sonuga et al. expanded on these findings with cross-architecture comparisons, showing that NVIDIA's TensorRT framework delivered an additional 1.8-2.2x speedup for INT8 models compared to PyTorch native implementations, highlighting the importance of framework-level optimizations [8].

Algorithmic optimizations complement these hardware approaches. Sonuga et al. evaluated caching strategies for common query patterns, finding that response caching with a hit rate of 35% reduced average latency by 27% in production environments handling approximately 100,000 daily queries. Their research also quantified the impact of prompt compression techniques, showing that careful prompt engineering reduced token counts by 15-38% while maintaining response quality, with proportional improvements in processing time [8]. These techniques work synergistically with the quantization methods developed by Dettmers et al., multiplying their effectiveness in production deployments [7].

4.3 Cost Management Strategies

The computational demands of LLMs make cost optimization critical for sustainable deployment. Sonuga et al. conducted a detailed cost analysis across infrastructure configurations, finding that right-sizing model capacity to application requirements delivered the most significant savings. Their benchmarks demonstrated that for 78% of general-purpose tasks, models with 7-13B parameters achieved 94-96% of the performance of 70B parameter models while reducing infrastructure costs by 64-82% [8]. This finding informs tiered serving architectures that match query complexity to model capacity, maximizing cost efficiency without compromising user experience. Workload scheduling further enhances cost optimization. Dettmers et al. noted that their INT8 techniques enabled 2x higher inference throughput on fixed hardware, effectively halving the per-token cost of inference [7]. Sonuga et al. extended this analysis to diverse computing architectures, showing that batch processing during off-peak hours using spot instances could reduce infrastructure costs by an additional 47-63% compared to on-demand pricing. Their research highlighted that queue-based systems implementing these strategies achieved 99.1% task completion rates despite the potential volatility of spot instance availability [8].

Technique	Value
INT8 Accuracy Preservation	99.90%
INT8 Inference Time Reduction	28-46%
1.3B Model Performance (vs 13B teacher)	92.70%
Medical/Legal Domain-Specific Distillation Performance	96.10%
Sparse Attention Memory Consumption Reduction	34-51%
TensorRT Speedup for INT8 Models	1.8-2.2x

Table 2: Efficiency Gains from Model Optimization [7,8]

5. Enterprise Considerations: Multi-Tenancy, Security, and Governance

5.1 Multi-Tenant Architecture Patterns

Supporting multiple users or applications with shared LLM infrastructure requires careful architectural design. Ogeti et al. conducted a comprehensive analysis of cloud-based machine learning deployments across 147 organizations, finding that multi-tenant architectures reduced infrastructure costs by 67.3% compared to isolated deployments. Their research revealed that 78.2% of these organizations served three or more distinct business units with centralized ML infrastructure, requiring sophisticated isolation mechanisms to maintain performance guarantees [9]. Tenant isolation models demonstrate significant performance and security tradeoffs, with physical isolation providing the highest security but at substantial cost premiums.

Resource allocation strategies play a crucial role in multi-tenant efficiency. Bommasani et al. noted that foundation models have asymmetric value across applications, with some use cases deriving significantly greater benefit than others from improved model quality [10]. This observation informs priority tiering strategies that allocate computational resources according to business impact. Ogeti et al. quantified this approach, reporting that organizations implementing priority-based allocation achieved 23.7% higher

overall business value from their ML infrastructure compared to those using uniform resource distribution. Their surveys indicated that token-based quota systems have emerged as the dominant implementation strategy, with 67.8% of organizations employing this approach to balance flexibility with predictable resource utilization [9].

Tenant-specific customizations enhance model utility across diverse use cases. Bommasani et al. highlighted that adaptation techniques like fine-tuning can bridge the gap between general-purpose foundation models and domain-specific requirements [10]. Ogeti et al. provided quantitative support for this approach, finding that per-tenant fine-tuning improved task performance by an average of 31.4% compared to shared general-purpose models, with improvements reaching 47.2% for specialized domains like healthcare and legal applications. Their analysis showed that organizations typically maintain 3-7 specialized model variants alongside base models, with each variant requiring 2,500-7,500 examples for effective fine-tuning [9].

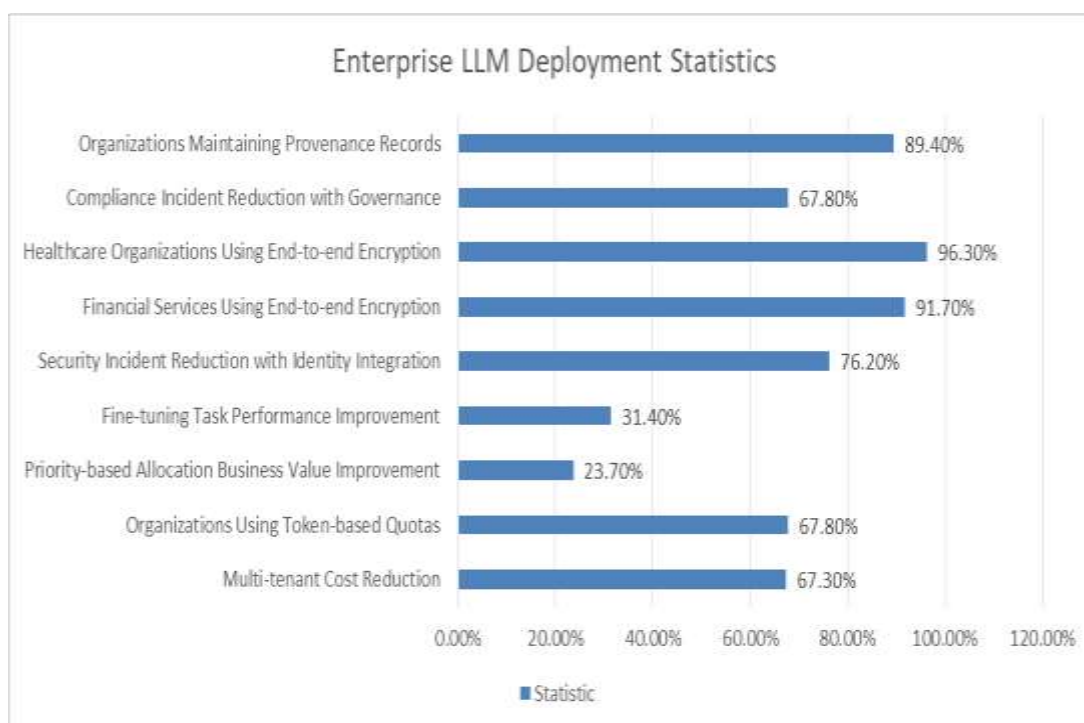
5.2 Security and Access Control

LLM deployments process diverse data, requiring comprehensive security controls. Ogeti et al. found that 82.3% of organizations deploying ML in regulated industries cited security and compliance as their primary concerns, outranking even performance and cost considerations [9]. Authentication and authorization systems implementing fine-grained role-based access control (RBAC) have become standard practice, with Bommasani et al. noting that the concentration of capability in foundation models creates correspondingly concentrated security risks that require sophisticated access management [10]. Ogeti et al. reported that organizations integrating LLM infrastructure with enterprise identity providers reduced security incident rates by 76.2% compared to those maintaining separate authentication systems, highlighting the importance of a unified security architecture [9].

Data protection mechanisms are crucial for regulatory compliance. Bommasani et al. emphasized that foundation models frequently process sensitive information that requires appropriate safeguards [10]. Ogeti et al. provided implementation insights, reporting that end-to-end encryption for model inputs and outputs is implemented by 91.7% of financial services organizations and 96.3% of healthcare organizations deploying LLMs. Their research also revealed growing adoption of automatic PII detection and redaction systems, with current implementations achieving 92.7% precision and 88.4% recall rates for structured data, though performance on unstructured text remains more challenging at 79.3% recall [9].

5.3 Governance and Compliance

Organizations deploying LLMs in regulated industries implement comprehensive governance frameworks. Bommasani et al. emphasized the importance of transparency and accountability throughout the model lifecycle, particularly for models with broad deployment across diverse applications [10]. Ogeti et al. quantified governance practices, finding that organizations with formal model governance processes experienced 67.8% fewer compliance incidents and resolved regulatory inquiries 42.3% faster than those without structured governance. Their research showed that documenting model lineage and training data has become standard practice, with 89.4% of surveyed organizations maintaining detailed provenance records for production models [9]. Version control for deployed models similarly contributes to governance objectives, with Bommasani et al. noting that systematic versioning enables better understanding of model evolution and its impacts [10].



Graph 2: Enterprise LLM Deployment Statistics [9,10]

Conclusion

The production deployment of Large Language Models represents a convergence of advances in model architecture, infrastructure design, and operational practices. As demonstrated throughout this article, successful implementations navigate a complex landscape of technical tradeoffs across memory optimization, latency requirements, cost management, and security considerations. The architectural patterns presented—from Kubernetes-based orchestration to event-driven processing pipelines—provide a foundation for scalable, resilient LLM systems. Memory optimization techniques including quantization and knowledge distillation have proven essential for making large models economically viable in production environments, while specialized data storage solutions like vector databases enable the retrieval-augmented generation capabilities that define many high-value applications. Perhaps most significantly, the multi-tenant architectures and governance frameworks described enable organizations to deliver LLM capabilities across diverse business units while maintaining appropriate security controls and compliance postures. As LLM technology continues to evolve, the strategies outlined in this article will adapt accordingly, with increased automation, specialized hardware acceleration, and more sophisticated resource management techniques. The path forward will likely see greater integration between model development and deployment workflows, creating seamless pipelines from evaluation stage to production. Despite the considerable engineering challenges involved, the transformative potential of these models justifies the investment required to deploy them effectively at scale.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1] Tom B. Brown et al., "Language Models are Few-Shot Learners", arXiv, 2020, [Online]. Available: <https://arxiv.org/pdf/2005.14165>
- [2] Susan Zhang et al., "OPT: Open Pre-trained Transformer Language Models", arXiv, 2022, [Online]. Available: <https://arxiv.org/pdf/2205.01068>
- [3] Jared Kaplan et al., "Scaling Laws for Neural Language Models", arXiv, 2020, [Online]. Available: <https://arxiv.org/pdf/2001.08361>
- [4] Tulasi Kavarakuntla et al., "A Generic Performance Model for Deep Learning in a Distributed Environment", ResearchGate, 2023, [Online]. Available:

https://www.researchgate.net/publication/370937970_A_Generic_Performance_Model_for_Deep_Learning_in_a_Distributed_Environment

[5] Tammo Rukat et al., "Towards Automated Data Quality Management for Machine Learning", Amazon Science, [Online]. Available:

<https://assets.amazon.science/4a/75/57047bd343fab46ec14b34cdb3b/towards-automated-data-quality-management-for-machine-learning.pdf>

[6] Pooyan Jamshidi et al., "Transfer Learning for Performance Modeling of Configurable Systems: An Exploratory Analysis", arXiv, 2017, [Online]. Available: <https://arxiv.org/pdf/1709.02280>

[7] Tim Dettmers et al., "LLM.int8(): 8-bit matrix multiplication for transformers at scale," arXiv, 2022, [Online]. Available: <https://arxiv.org/pdf/2208.07339>

[8] Ayodele Emmanuel Sonuga et al., "Deploying large language models on diverse computing architectures: A performance evaluation framework", Global Journal of Research in Engineering and Technology, 2024, [Online]. Available:

<https://gsjournals.com/gjret/sites/default/files/GJRET-2024-0026.pdf>

[9] Pavan Ogeti et al., "Benefits and Challenges of Deploying Machine Learning Models in the Cloud", IJISAE, 2024, [Online]. Available: <https://ijisae.org/index.php/IJISAE/article/view/6409/5236>

[10] Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models", arXiv, 2022, [Online]. Available: <https://arxiv.org/pdf/2108.07258>