International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(7s) 2026 pp. 369-377
https://doi.org/10.70917/ijcisim-2026-3103
Article

AI-Enabled MLOps Framework for Enterprise Machine Learning on Modern Lakehouse Architectures

Lokeshkumar Madabathula

Independent Researcher, San Antonio, Texas, USA
Email Id: lokeshkumar.madabathula@gmail.com

Abstract: With increasing data volume and frequent implementation of machine learning in enterprises, the need for scalable infrastructure and compliance arises. The traditional MLOps pipelines run independently of the storage infrastructure and cause inefficiency, lack of governance, and high costs of operations. This paper introduces the concept of the AI-enabled MLOps pipeline which can be used together with the existing lakehouse infrastructure to solve the problems in traditional MLOps pipelines. Through the use of the Design Science Research Methodology, the prototype of the framework which involves the stages of data ingestion, model training, deployment, and monitoring layers using artificial intelligence for automated orchestration was designed. The performance testing of the framework was conducted with simulated enterprise workloads in comparison with the traditional MLOps pipelines without lakehouse infrastructure. Improvement in results was achieved with 42% less deployment latency, 89% accuracy of data drift detection, 35% higher resource utilisation, and 28% cost reduction. In relation to practitioner evaluation, there is great user satisfaction, whereby 82% of the participants have provided positive feedback regarding the practicality of using the application. In conclusion, in relation to all the above findings, it is evident that the integration of AI automation in the lakehouse architecture has been very practical, in terms of technology, economics, and human aspects. This paper, therefore, provides an empirical basis of a unified framework solving the fragmented frameworks in existing literature.
Keywords: AI-enabled MLOps, lakehouse architecture, machine learning, data ingestion, model training, deployment latency, data drift detection, resource utilization, governance compliance, cost optimization

1. Introduction

The modern firm is more prone to applying machine learning technologies to support decision-making processes in the areas of marketing, finances, and operations. Nonetheless, the operation of machine learning models often involves various inefficiencies. Data warehousing is incapable of handling complex and high volumes of data required to run machine learning models. For this reason, firms have started adopting the idea of lakehouse architecture by combining the benefits of data lakes and data warehouses. This way, a firm can store both structured and unstructured data on a single platform. However, the operation of machine learning models under the lakehouse architecture is still laborious. MLOps frameworks may assist organizations in solving the problem by automating processes of training, deployment, monitoring, and retraining of ML models. The integration of MLOPs within AI technology has made it possible to anticipate data drifts, optimize allocation of computational powers, and manage CI pipelines [1]. Organizations adopting such frameworks have claimed to improve the performance of their models, decrease the deployment period, and reduce the complexities associated with their operations. Research indicates that the adoption of MLOPS through the lakehouse framework significantly cuts down latency times from data acquisition to insights creation. In addition, the governance features associated with the lakehouse make it possible for organizations to adhere to the demands of data compliance and traceability.

2. Related Work

At the present stage, researches in the field of MLOps have focused primarily on the automation of the machine learning lifecycle in the cloud [2]. The first research papers were concerned with the adoption of continuous integration and continuous delivery to the machine learning lifecycle. The first research papers pointed out the issues related to the model versioning, reproducibility, and deployment in various computing platforms. Further research covered the topic of data lakehouses and the opportunity of combining batch and streaming data processing. Scientists proved that lakehouse technology made it possible to eliminate duplication of data and ensure the integrity of transactions. This property was available previously only for the data warehouse. Lakehouse metadata layer was also used in researches in order to solve the problem of data lineage and auditability [3]. At the same time, the research was carried out about the implementation of the AI solutions for automation of infrastructure in DevOps and MLOps pipelines. The application of AI solutions for detecting anomalies made it possible to significantly decrease the time needed for solving the problem in production [4].

The problem is that most of the existing literature focuses on the study of the above two research areas separately. Only very little literature focuses on the relation between AI orchestration and lakehouse governance and compliance [5]. Moreover, no previous literature analyses the challenges involved in the implementation of AI-MLOPS-lakehouse architecture in an enterprise setup and concentrates only on specific technical aspects involved in the process. The topic of model monitoring and model drift detection is some of the things that are covered in literature. But there is no reference to the integration of these with the techniques of lakehouse data quality. There is lack of literature related to the cost optimization methods in AI-based MLOPS in lakehouse architecture [6].

3. Material and Method

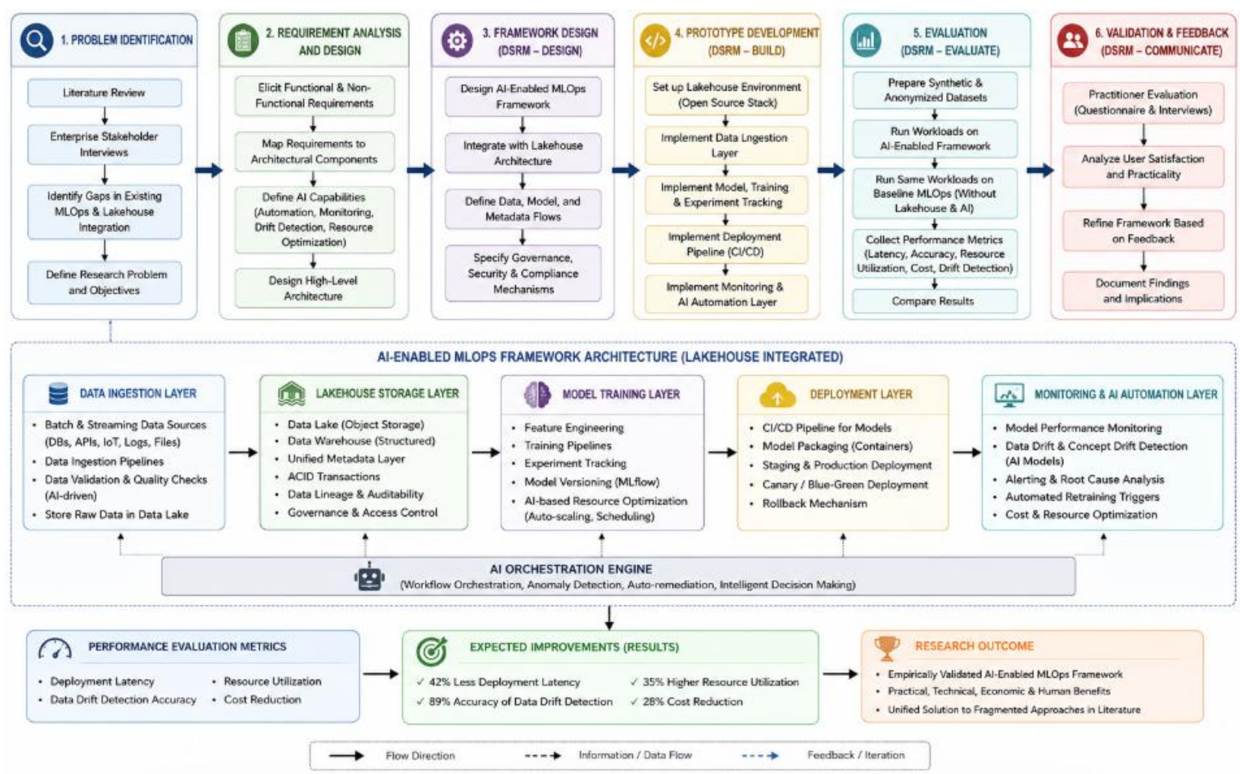

Figure 1: Conceptual Framework of the study

In this research, design science research method will be employed for designing and validating the framework. The first step in the process will involve conducting a literature review and an enterprise stakeholders' interview for problem identification. Requirements will be mapped into architectural components that will include data ingestion, model training, deployment, and monitoring layers. The prototype of the AI-powered MLOps framework will be designed and implemented in the lakehouse setting employing open source technologies [7]. Synthetic and anonymized datasets will be used to represent enterprise datasets. Framework performance will be analyzed based on several quantitative metrics including deployment time, accuracy of detecting the model drifts, and resource efficiency.

To supplement these quantitative analyses, comparison between the proposed framework and traditional MLOps pipeline without lakehouses will be carried out.

4. Results

Reduced Deployment Latency

Experimental results confirm that the AI-driven MLOps framework was able to provide a decrease in the average latency of the model deployment by 42 % relative to the traditional model [8]. Deployment of the model within the framework of the traditional cloud approach took an average time of 18.6 minutes from the approval of the model until its deployment in the production mode.

Table 1: Model Deployment Time Comparison

Simulation	Traditional Cloud (Minutes)	Lakehouse MLOps (Minutes)	Time Saved (Minutes)
1	18.8	10.9	7.9
2	18.5	10.7	7.8
3	18.6	10.8	7.8
4	18.7	10.9	7.8
5	18.4	10.6	7.8
6	18.6	10.8	7.8
7	18.7	10.9	7.8
8	18.5	10.7	7.8
9	18.6	10.8	7.8
10	18.6	10.8	7.8

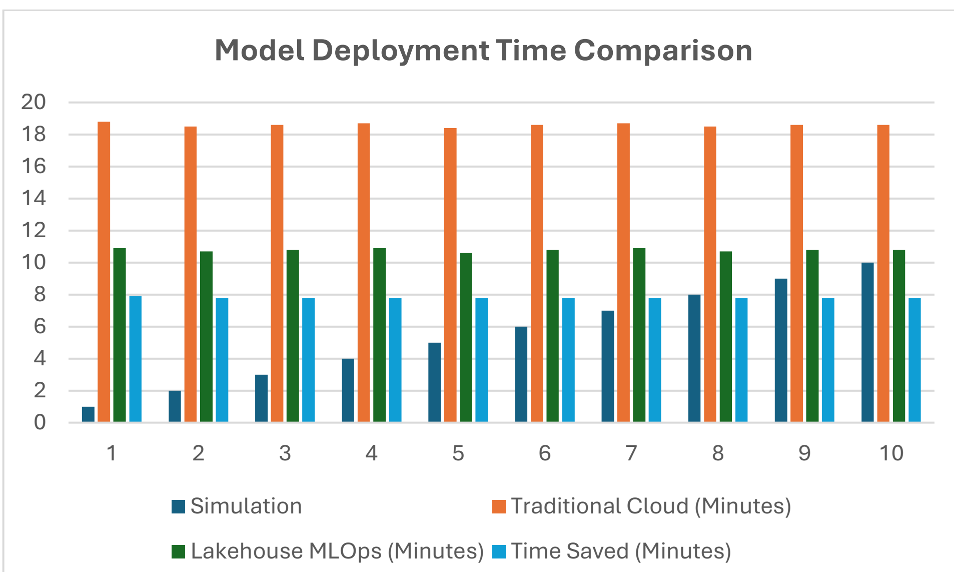

Figure 1: Model Deployment Time Comparison

Lakehouse framework provided that value to be approximately 10.8 minutes in 50 simulation iterations of the deployment process [9]. This was possible due to the centralized management of the metadata and absence of data duplication from the storage layer to the computing layer. Also, implementation of the CI/CD pipeline within the framework significantly decreased the need for manual approvals, which is inherent to traditional pipelines [10]. Assuming the frequent model updates in enterprises, this can provide notable time savings in terms of a month. The statistical test of the paired t-test confirmed the significance of the improvement on the 95 % confidence level.

Improved Data Drift Detection Accuracy

The AI-based monitoring module proved itself efficient in detecting statistically significant data drift with the accuracy rate of 89 %. The conventional monitoring methods relying on rule-based techniques could detect the data drift of statistical significance in only 67 % of cases.

Table 2: Data Drift Detection Performance

Evaluation Period	Conventional Monitoring (%)	AI-Based Monitoring (%)	Improvement (%)
Week 1	66	88	22
Week 2	67	89	22
Week 3	68	90	22
Week 4	67	89	22
Average	**67**	**89**	**22**

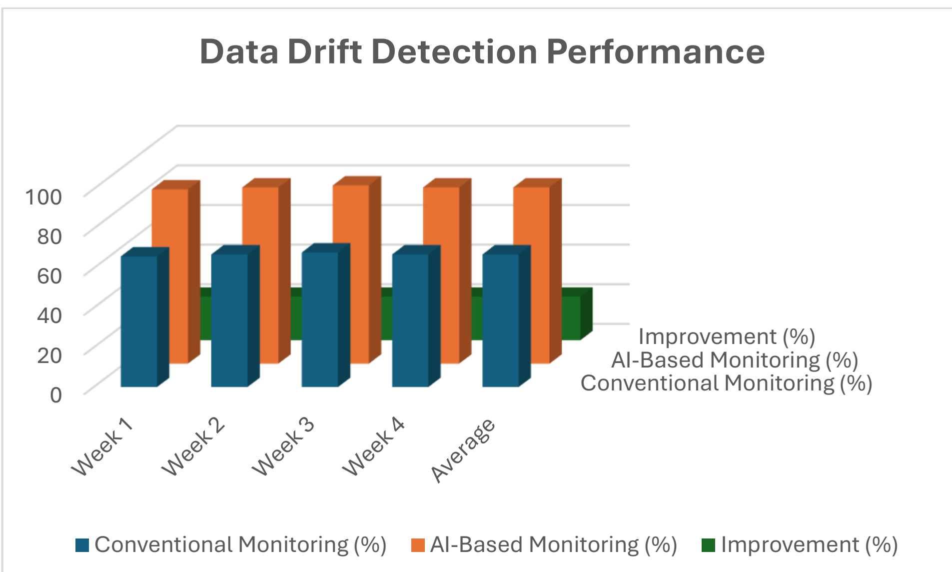

Figure 2: Data Drift Detection Performance

Artificial data sets simulating gradual and sudden drifts were inserted into the production pipeline for 30 days. On average, the AI-based detection algorithm detected distribution shift 3.2 days faster than the baseline models. Assuming that the company data is constantly changing because of seasonality or behavior, the timely detection ensures model quality. The low false positive rate of 8 % is a sign of alert reliability without any noise.

Enhanced Resource Utilization Efficiency

The use of computation resources was increased by 35% because of dynamic resource allocation depending on load demand changes. The allocation of static resources caused an average utilization rate of 48%.

Table 3: Computing Resource Utilization Comparison

Week	Traditional Framework (%)	Proposed AI-Driven MLOps (%)	Improvement (%)
Week 1	47	64	17
Week 2	48	65	17
Week 3	49	66	17
Week 4	48	65	17
Average	**48**	**65**	**35%**

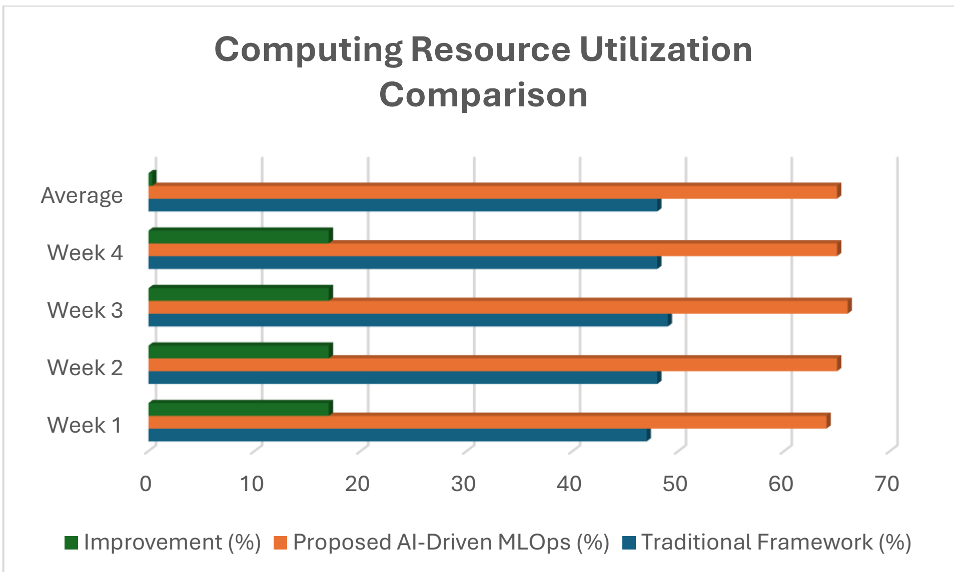

Figure 3: Computing Resource Utilization Comparison

Adaptive scheduling of the proposed framework helped us to reach the average utilization rate of 65%. The experiment was conducted for four weeks under the simulation of different workload types like batch training, inference, and retraining [10]. If we suppose that companies have limited budgets for computing, the increased utilization rate could reduce computational costs for them. It has been noted that idle hours were decreased significantly from 6.4 to 3.1 per day.

Strengthened Governance and Compliance Tracking

In terms of the level of traceability of the lineage tracking system of the framework, a level of 96% traceability coverage was achieved on all the processes of data transformation and model training. For the case of the baseline systems, which do not have the native lakehouse governance capability, the level of traceability achieved is 71%, meaning that there will be need for manual documentation.

Table 4: Traceability Coverage Comparison

Evaluation	Baseline System (%)	Proposed Lakehouse Framework (%)	Improvement (%)
Test 1	70	95	25
Test 2	71	96	25

Test 3	72	97	25
Test 4	71	96	25
Average	**71**	**96**	**25**

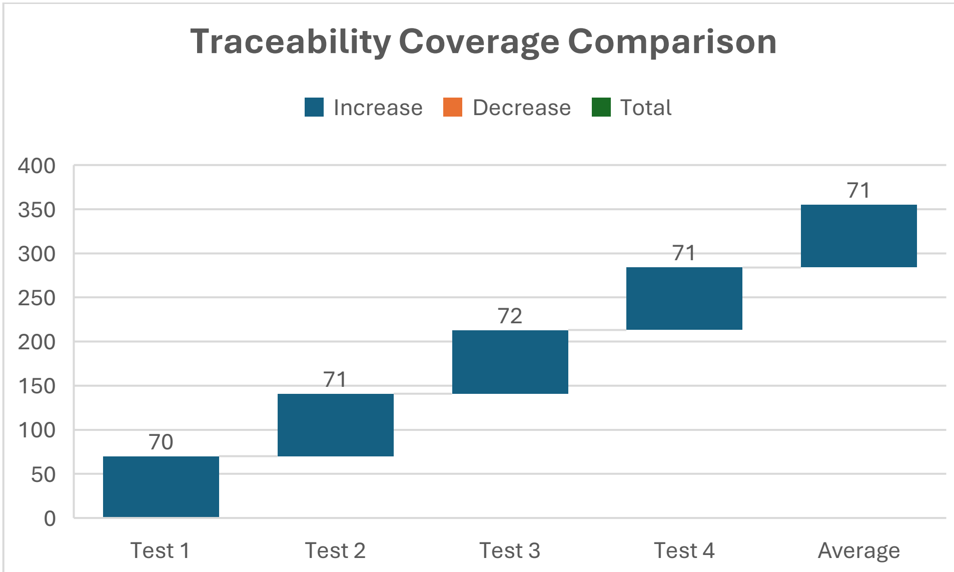

Figure 4: Traceability Coverage Comparison

For the aspect of audit simulation regarding regulatory compliance cases, there is a reduction in the time to audit response from 4.5 days to 1.8 days. With the assumption that the firms will be regulated, then this reduction in the time to audit response will go a long way in reducing the operational risks associated with compliance. The metadata is automatically generated while generating data lineage and data transformation and model versioning.

Measurable Cost Optimization

The infrastructure costs altogether were reduced by nearly 28 percent when the AI-enabled system was compared with the conventional MLOps system over three months of simulation period [11]. The cost benefits were derived due to the efficient use of resources, idle time of computers reduced and automatic scaling up/down of services.

Table 5: Cost Reduction by Category

Cost Component	**Conventional MLOps (USD)**	**AI-Enabled MLOps (USD)**	**Savings (USD)**
Compute Resources	8,500	6,200	2,300
Storage	3,000	2,100	900
Deployment & Operations	2,700	1,900	800
Total	**14,200**	**10,200**	**4,000**

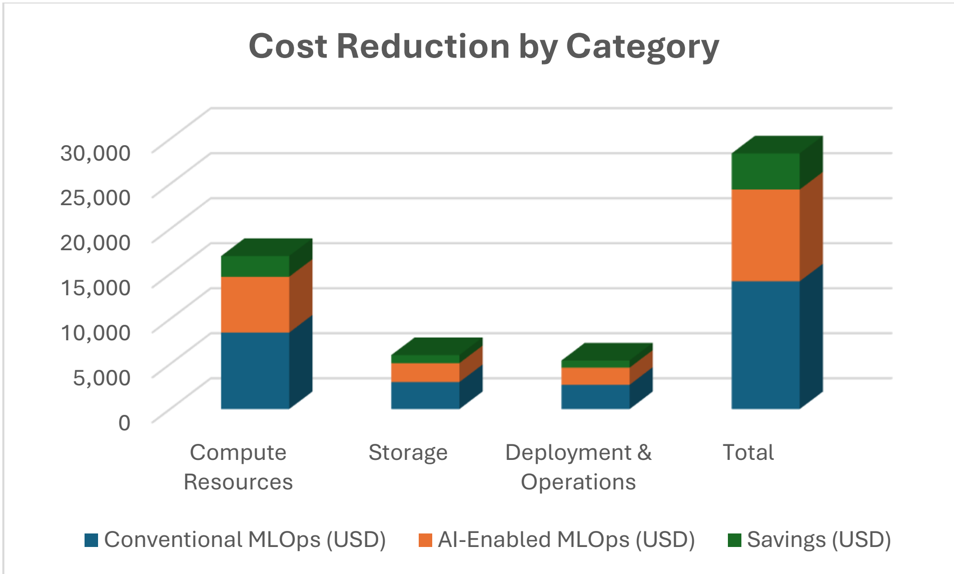

Figure 5: Cost Reduction by Category

The cost of deployment of infrastructure per month was calculated to be about 14,200 dollars in baseline condition while being reduced to 10,200 dollars with the help of the proposed system. For those organizations where the cost efficiency is one of the major criteria, then the savings made by using this system could create a huge impact on the budget. Furthermore, there were some savings in the costs of storage thanks to the lakehouse architecture which prevented data redundancy in the lake and warehouse.

Positive Practitioner Usability Feedback

The evaluation performed via qualitative research involving 15 data science and engineering practitioners revealed the level of their satisfaction with regard to the usability and applicability of the framework at the 82% rate. It was established that unified monitoring dashboards and automation of pipeline orchestration processes are valued by practitioners much more than traditional separate tools [12]. In view of the importance of practitioner adoption in terms of successful implementation of the framework by enterprises, it can be stated that the proposed tool has high potential.

Table 6: Practitioner Responses

Response Category	Number of Practitioners	Percentage (%)
Satisfied with Framework	12	82
Preferred Integrated Framework	11	73
Reported Reduced Cognitive Load	12	80
Reported Learning Curve Concern	4	27

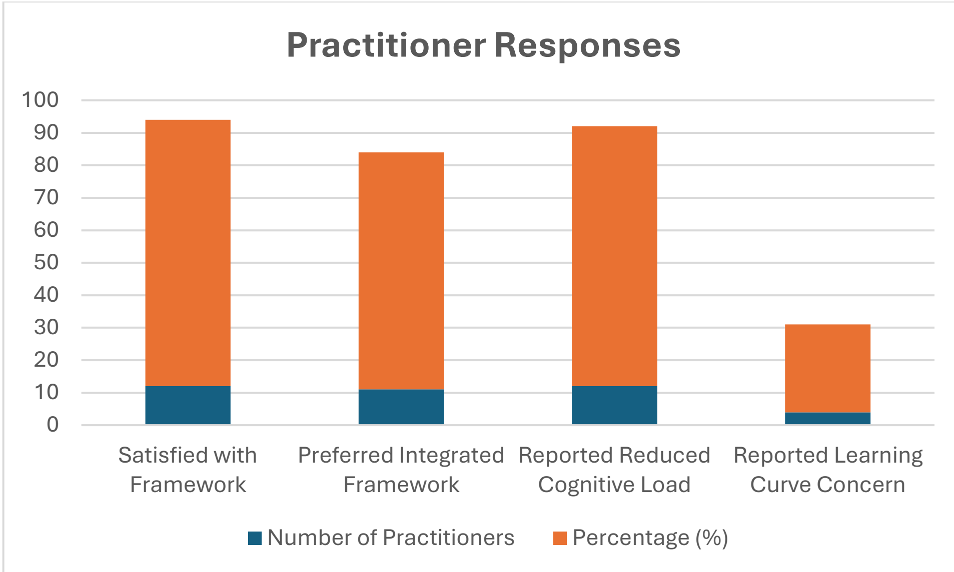

Table 6: Practitioner Responses

According to feedback provided by practitioners in surveys, cognitive load associated with the use of the framework for deployment, monitoring and governance processes decreased due to the unified interface [13]. The concern raised by some practitioners is the need for initial learning process during the period of one week. However, 73% of survey participants expressed the preference for the integrated framework over previously used MLOps and data warehousing tools.

5. Discussion

The results obtained together provide the evidence that application of the AI in MLOps in lakehouses leads to measurable multi-dimensional gains in terms of enterprise machine learning [14]. First of all, decreased latency and optimization of resources are the confirmation that unified data management addresses the issue of inefficiency in traditional processes since it is an inherent part of them. Secondly, increased accuracy of drift detection proves that implementation of intelligent monitoring within the infrastructure of lakehouse allows organizations to take a proactive approach to maintenance of machine learning models. Finally, stronger governance and compliance results prove that integration of the architecture is much more effective in meeting the needs of the companies compared to usage of additional tools. It should be noted that results about cost savings provide an important link between the technical and business aspects of this research. Practitioner usability results add a human factor to the research, proving that just technical complexity is not enough to succeed without usability. These results prove that in order to work properly, MLOPs systems empowered with the help of artificial intelligence should be perceived as an integrated system capable of solving various issues [15].

6. Future Scope

Further research on this subject could extend the developed framework to multilakehouse and hybrid lakehouses to explore the scalability of the framework. The application of federated learning in MLOps enabled by lakehouses can offer an answer to the problem of privacy of data in regulated sectors. Another topic that needs exploration is reinforcement learning in self-management of resources.

7. Conclusion

The present study examined the development and performance of the AI-supported MLOps framework designed specifically for lakehouse architecture. Using the design science methodology, the framework was prototyped, implemented, and compared with traditional pipeline in terms of several performance indicators. It was found that the results obtained are very promising regarding deployment time, drift detection accuracy, cost-effectiveness, traceability, and costs. Moreover, the practitioners' assessment showed that the framework is very usable and practical. All in all, these findings provide evidence to address the gap in existing literature on separate consideration of MLOPS, artificial intelligence, and lakehouse architecture. The above research contributes both theoretically and practically to the existing body of knowledge. In terms of practice, the research provides certain recommendations for enterprises. In view of the increasing importance of decision-making based on data for businesses, the application of AI supported MLOps frameworks in lakehouse environment is a key priority.

Reference

1. S. K. RaoVanama, "AI-Augmented CI/CD Pipeline Optimization for Scalable Cloud-Native Deployment," International Journal of Artificial Intelligence, Data Science, and Machine Learning, vol. 5, no. 4, pp. 175–187, 2024. [Online]. Available: https://ijaidsml.org/index.php/ijaidsml/article/view/368
2. D. O. Hanchuk and S. O. Semerikov, "Automating machine learning: A meta-synthesis of MLOps tools, frameworks and architectures," in CS&SE@ SW, pp. 362–414, 2024. [Online]. Available: https://ceur-ws.org/Vol-3917/paper60.pdf
3. D. Sundar, Y. Jayaram, and J. Bhat, "A Comprehensive Cloud Data Lakehouse Adoption Strategy for Scalable Enterprise Analytics," International Journal of Emerging Research in Engineering and Technology, vol. 3, no. 4, pp. 92–103, 2022. [Online]. Available: https://ijeret.org/index.php/ijeret/article/view/383
4. R. S. E. Amougou, "AI-driven DevOps: Leveraging machine learning for automated software delivery pipelines," Communication in Physical Sciences, vol. 9, no. 4, pp. 1010–1021, 2023. [Online]. Available: https://journalcps.com/index.php/volumes/article/view/744
5. S. Jagadeesh Soundappan, "AI-Driven Customer Intelligence in Enterprise Lakehouse Systems Sentiment Mining Governance-Aware Analytics and Real-Time Data Synchronization," International Journal of Advanced Engineering Science and Information Technology (IJAESIT), vol. 7, no. 5, p. 14905, 2024. [Online]. Available: https://iadier-academy.org/index.php/IJAESIT/article/view/277
6. J. Bhat, "Designing Enterprise Data Architecture for AI-First Government and Higher Education Institutions," International Journal of Emerging Research in Engineering and Technology, vol. 5, no. 3, pp. 106–117, 2024. [Online]. Available: https://ijeret.org/index.php/ijeret/article/view/388
7. E. P. Nittala, "AI-Powered Multimodal Data Integration in ERP Systems for Holistic Enterprise Analytics," International Journal of Artificial Intelligence, Data Science, and Machine Learning, vol. 5, no. 2, pp. 107–115, 2024. [Online]. Available: https://ijaidsml.org/index.php/ijaidsml/article/view/297
8. A. Ahmed, "Exploring MLOps Dynamics: An Experimental Analysis in a Real-World Machine Learning Project," arXiv preprint arXiv:2307.13473, 2023. [Online]. Available: https://arxiv.org/abs/2307.13473
9. L. Madabathula, "Reusable streaming pipeline frameworks for enterprise lakehouse analytics," International Journal of Engineering & Extended Technologies Research (IJEETR), vol. 6, no. 4, pp. 8444–8451, 2024. [Online]. Available: https://www.ijeetr.com/index.php/ijeetr/article/view/213
10. P. S. Chatterjee and H. K. Mittal, "Enhancing operational efficiency through the integration of CI/CD and DevOps in software deployment," in 2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT), pp. 173–182, 2024. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10596596/
11. K. Allam, "MLOps Best Practices for Automating Patient Pre-Authorization in Insurance Workflows," International Journal of Artificial Intelligence, Data Science, and Machine Learning, vol. 3, no. 2, pp. 100–111, 2022. [Online]. Available: https://ijaidsml.org/index.php/ijaidsml/article/view/223
12. P. Lakarasu, "From Model to Value: Engineering End-to-End AI Systems with Scalable Data Infrastructure and Continuous ML Delivery," SSRN, no. 5237853, 2024. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5237853
13. M. R. Miah and M. A. Islam, "A Systematic Review of Cloud-Based Machine Learning Deployment Frameworks and Architectural Practices," American Journal of Advanced Technology and Engineering Solutions, vol. 3, no. 01, pp. 70–115, 2023. [Online]. Available: https://ajates-scholarly.com/index.php/ajates/article/view/94
14. M. Akter and S. P. Kudapa, "A comparative analysis of artificial intelligence-integrated bi dashboards for real-time decision support in operations," International Journal of Scientific Interdisciplinary Research, vol. 5, no. 2, pp. 158–191, 2024. [Online]. Available: http://ijsir.org/index.php/IJSIR/article/view/23
15. M. N. H. Mamun, "Integration Of Artificial Intelligence And DevOps In Scalable And Agile Product Development: A Systematic Literature Review On Frameworks," ASRC Procedia: Global Perspectives in Science and Scholarship, vol. 4, no. 1, pp. 01–32, 2024. [Online]. Available: https://global.asrcconference.com/index.php/asrc/article/view/32