



Tiered Lakehouses for Real-Time MLOps

Hiroshi Tanaka

Asst Professor

Abstract

As large-scale machine learning in production becomes ubiquitous across industries, both engineering and research teams are pushing for faster, continuously retrained models. Yet building and running a model — even with a well-established MLOps framework — is often only a small part of the overall task. Feature engineering and feature selection typically consume the majority of a project's data and ML effort. A Medallion Architecture purpose-built for real-time feature engineering, paired with an enterprise Feature Store, can accelerate this process by integrating the latest data insights into models and bringing both new features and models under continuous integration pipelines.

This architecture incorporates key DataOps and MLOps best practices for target labeling, feature quality, and CI/CD pipelines, while allowing features to be built in parallel — either as a dedicated MLOps modeling step or as part of standard product development. We present ML pipelines and Feature Stores for scalable, real-time feature engineering, covering the Medallion Architecture and a large-scale implementation of it, along with best practices for real-time streaming data ingestion and processing and continuous integration/continuous delivery (CI/CD) for features. The resulting architecture is designed for data scale, balancing the engineering demands of feature development with horizontal scaling through sharding strategies aligned to both the features being built and cost objectives for exploratory data analysis and feature selection.

Keywords :Feature stores, DataOps, Continuous Integration, Continuous Delivery, Cloud-Native platforms, Medallion Architecture, Climate Simulation, City of New York, NYC Citywide Climate Networks, NYC City Sensor Data, Medium Data, Continuous Data Engineering .

1. Introduction


The evolving landscape of artificial intelligence and big data technologies poses fresh challenges for the deployment of Machine Learning models and the Productionisation of AI-stacks. One such challenge is to efficiently enable real-time prediction serving systems as part of the MLOps (Machine Learning Operations) for cloud-native platforms. Mainstream DataOps practices and Cloud Data Platforms have created common yet versatile approaches to ensuring timely availability of structured datasets. Such DataOps solutions focus on recognising the patterns of preparing structured data elements for analytical consumption. Most Cloud Data Platforms support horizontal scaling techniques that help tackle the extremes of data engineering and DataOps workloads without compromising performance. Nevertheless,

addressing streaming operations when processing and delivering ML-ready features for prediction serving still often necessitate workarounds that do not seem to fit nicely within the main DataOps workspace. These tips and patterns have primarily been embraced to integrate DataOps and real-time prediction serving.

The Medallion architecture consists of three main processing and access patterns – Bronze, Silver and Gold. The Bronze layer contains the Ingest pattern to process diverse raw data on object storage. The available raw data can be consumed or processed into more structured datasets within the Silver layer, which is realised using the Recon pattern.

Governance Component	Primary Function	Key Technologies	Expected Outcome
Unified Control Plane	Centralized orchestration of governance	Policy-as-Code, APIs	Automated governance management
Risk Analytics Engine	Detect and analyze enterprise risks	AI/ML Analytics	Real-time risk visibility
Compliance Automation Layer	Regulatory compliance enforcement	Workflow Automation	Reduced audit overhead
Monitoring & Telemetry	Continuous infrastructure observation	SIEM, Log Analytics	Faster incident detection
Identity & Access Governance	Secure access management	IAM, MFA, RBAC	Reduced unauthorized access
Remediation Engine	Automated corrective actions	SOAR, AI Agents	Faster risk mitigation

Table 1: Cloud Governance Components and Functions



Intelligent Risk & Compliance Orchestration – Double Layer Flowchart

2. Background and Motivation

Data science and machine learning models have gained significant attention over the last years, and companies are now actively hiring data engineers to build production-ready deployment pipelines. An ML-based operational management supports data-driven decision making; however, implementations remain in the early stages. Nevertheless, many organizations begin considering investing in a Feature Store solutions, although these services don't seem to have become the norm. A medallion architecture incorporates these characteristics but lacks a formal definition. This paper aims to have a textual definition for it, along with guidelines and considerations during development to help organizations fulfill requirement changes within the data and data science areas.

A Feature Store covers aspect of a traditional Data Warehouse (DW) and serves as a persistent storage for features, data



slices, feature-versioning control, and feature validation. Data science is a continuous process: changes on the data, models performances, and business requirements lead to the need for new features. The “one-off” implementation strategy of ML models in production is no longer sufficient and new solutions must incorporate CI/CD pipelines. Organizations that have the skill to support these practices can implement a Feature Store solution and integrate it within a larger Medallion architecture.

Risk Scenario	Detection Technique	AI/ML Role	Mitigation Strategy
Vulnerability Exposure	Behavioral Analysis	Pattern Recognition	Automated Patch Deployment
Data Leakage	Anomaly Detection	Data Classification Models	Encryption Enforcement
Unauthorized Access	User Behavior Analytics	Access Pattern Learning	Dynamic Access Revocation
DDoS Attacks	Traffic Modeling	Predictive Detection	Auto-Scaling & Filtering
Misconfigured Cloud Assets	Configuration Mining	Compliance Drift Detection	Policy Reconfiguration
Credential Theft	Threat Correlation	Risk Scoring Models	Secret Rotation

Table 2: AI-Driven Risk Detection Scenarios


2.1. Feature Store Integration and Consistency

Cloud-native meta-architecture enables horizontal scalability, cost-effectiveness, and data security in integrated data and Machine Learning (ML) pipelines. For greater consistency and faster responses, aligning low-latency, real-time, and

business-critical applications with feature engineering via a medallion architecture is advantageous. Real-time streaming, batch ingestion, storage, pre-calculation, union, and broadcasting of features within the design's bronze, silver, and gold layers foster MLOps practices supporting quality assurance and delivery pipelines.

Medallion Architecture enables Feature Store integration via feature tables, modeling datasets for offline jobs, deployment validation, shared-data components, and scikit-learn Pipeline workflows. These mutually beneficial connections cut costs and delay by maintaining functions and Automation with Continuous Delivery. Metalogic Flow upgrades and downgrades Wilson Data Platform Layer Data by embedding metadata, connection logic, and small Data dependencies, slashing manual tasks and expediting information transfer across three features. Transforming structurally heterogeneous Data into flat Data formats enhances processing efficiency and outcome reliability by curbing missing value management.

MLOps becomes unattainable when dashboards rely on Data Warehouse Data, Data Lakes saturate, Data Lakes fail to refresh automatically, and re-integration into Cohere documentation becomes mandatory. Medallions must therefore integrate managed Data Lakes and Data Warehouses, and cope with ML and Dashboards. Maintenance and pre-refresh fixes introduce high maintenance and operational risk. Consistent dumping and consistency of the three Metalogic transports greatly simplify and align online MLOps with global sources.



Derived from section word counts in the uploaded article.



AI-Driven Cloud Governance Architecture – Double Layer Flowchart

3. Medallion Architecture: Core Concepts

The Progressive Medallion Architecture for Cloud-Native Data Platforms integrates established data engineering practices—feature engineering and MLOps—into a single DataOps framework. These elements have evolved independently during the rapid maturation of Cloud-Native Data Platforms. Precedence for a holistic architecture, however, has been set by other initiatives that address Integration and Continuous Delivery within a single coherent system. The Medallion Model, which emphasizes Bronze, Silver and Gold layers, provides a logical foundation for such an approach. Its proven efficacy in enabling DataOps for batch use cases can be extended to support both Real-Time Feature Engineering and MLOps Practices. Careful integration of DataStream Processing techniques into the architecture now enables programmers to harness the full sub-second precision of a Streaming Platform.

Cloud-Native Data-Processing Platforms excel in ingesting vast quantities of data from diverse sources—regularly, if not constantly—acting as the central repository, or Data Lake, for the entire organization. Minimal preparation and no semantic enrichment are required in the Bronze Layer of a Medallion Architecture. Tableau’s Medallion Architecture Moves to the Cloud explains how data structures, formats and storage paths are designed to maximize ingestion throughput. The DataSource Template Library contains ready-to-deploy templates for the streaming ingestion and continuous preparation of CyberSecurity, Click-Stream, Industrial IoT and Financial Services Data.

Challenge Area	Description	Impact	Governance Response
Shared Responsibility Complexity	Split responsibilities across providers	Compliance ambiguity	Shared governance framework
Configuration Drift	Continuous infrastructure changes	Security gaps	Continuous monitoring
Third-Party Access	External vendor integrations	Increased attack surface	Access governance controls
Hybrid Cloud Visibility	Limited centralized oversight	Operational blind spots	Unified telemetry systems
Regulatory Diversity	Different regional compliance laws	Audit complexity	Automated compliance mapping
Resource Elasticity	Dynamic scaling environments	Inconsistent policies	Policy-as-Code orchestration


Table 3: Multi-Cloud Governance Challenges

3.1. Bronze Layer: Ingest and Raw Data Management


General-purpose data ingest, be it batch or streaming, constitutes a fundamental requirement of any Medallion or Delta architecture. Volume and velocity of ingest must therefore be processed and managed accordingly to ensure consistency with Delta Lake ACID guarantees.

For streaming ingest, the Delta framework adds considerable benefit when using the streaming query mechanism to decouple the ingest from processing logic. In addition to the charted batch ingestion responsibilities, the processing must

also manage the overall lifecycle of the ongoing incoming changes. Monitoring of data coalescing is achieved by checking incoming volumes against predefined thresholds.



Derived by counting representative keywords for each theme.



Multi-Cloud Compliance Monitoring – Double Layer Flowchart

4. Real-Time Feature Engineering in Cloud-Native Environments

Real-Time Feature Engineering in Cloud-Native Environments

The wide adoption of cloud technologies combined with the democratization of data accessibility has made it possible to run data pipelines in cloud-native environments. Serverless solutions allow processing of batch and streaming data with near-zero maintenance effort at almost no cost on a responsibility level based on demand. The operational support reduces maintenance efforts and allows the data engineer to focus on more valuable tasks such as data quality, lineage, architecture, and process optimization.

Recently emerged and actively developed serverless prototypes make it possible to natively design cloud data platforms with real-time DataOps and MLOps. Continuous Integration-Continuous Delivery (CI-CD) pipelines built into the proprietary Digital Warehouse can easily integrate data with a Data Lake or Feature Store and then expose it as accessible third-party FaaS or API interfaces. In a pure cloud environment, DataOps and MLOps can co-exist within a single data pipeline, where DataOps delivers real-time data directly to the Feature Store and where CI-CD processes for Features are initiated automatically based on Data Quality checks.

KRI Metric	Measurement Objective	Threshold Type	Governance Significance
Encrypted Data Coverage	% of encrypted data assets	Minimum Coverage	Data protection compliance
Secrets Rotation Frequency	Rotation interval of credentials	Time-Based	Reduced credential compromise
Vulnerability Severity Count	Number of high-risk vulnerabilities	Severity Threshold	Risk prioritization
Misconfigured Resources	Count of non-compliant assets	Resource Threshold	Operational governance quality

KRI Metric	Measurement Objective	Threshold Type	Governance Significance
Unauthorized Access Attempts	Suspicious login activities	Event Threshold	Security posture monitoring
Compliance Drift Rate	Policy deviations over time	Trend Threshold	Governance effectiveness

Table 4: Key Risk Indicators (KRIs) for Cloud Governance

Mathematical Formulas:

1. Feature Engineering Transformation

$$F_t = T(D_r)$$

Where:

- F_t = transformed feature set
- T = transformation function
- D_r = raw data input

2. Streaming Data Processing Rate

$$R_p = \frac{N_e}{T}$$

Where:

- R_p = processing rate
- N_e = number of events

- T = processing time

3. Feature Store Update Function

$$FS_{t+1} = FS_t + \Delta F$$

Where:

- FS_t = current feature store state
- ΔF = incoming feature updates

4. Model Prediction Function

$$\hat{Y} = f(X, F)$$

Where:

- $\hat{Y}$ = predicted output
- X = input data
- F = engineered features

5. CI/CD Deployment Efficiency

$$E_d = \frac{S_v}{T_d}$$

Where:

- E_d = deployment efficiency
- S_v = successful validations
- T_d = deployment duration

6. Horizontal Scaling Capacity



$$C_s = N_n \times P_n$$

$$TP = \frac{F_n}{T_p}$$

Where:

- C_s = total scaling capacity
- N_n = number of nodes
- P_n = processing power per node

7. Streaming Latency Equation

$$L_s = T_r - T_i$$

Where:

- L_s = streaming latency
- T_r = response time
- T_i = ingestion time

8. Data Quality Metric

$$Q_d = \frac{D_v}{D_t}$$

Where:

- Q_d = data quality score
- D_v = valid records
- D_t = total records

9. Feature Pipeline Throughput

Where:

- TP = throughput
- F_n = number of generated features
- T_p = processing time

10. Resource Utilization in Cloud Pipeline

$$U_r = \frac{R_u}{R_a} \times 100$$

Where:

- U_r = utilization percentage
- R_u = used resources
- R_a = available resources


4.1. Streaming Data Ingestion and Processing

The real-time capability provided by streaming data services, such as Kafka or Azure Stream Analytics, allows a clear repartition of the workload: (i) continuous processing of incoming data into the feature store and (ii) handling of incoming user requests against the feature store. A typical streaming data ingestion pattern can therefore be separated into an ETL (Extract, Transform and Load) pattern, where incoming data is processed through the boundary of the feature store, and an EAS (Easy Access Service) pattern, where stored data is queried and delivered for instantaneous consumption by an API (Application Programming Interface). Although these patterns can be treated separately for conceptual clarity, they will normally be deployed



together in a cloud-native environment. The workload of streaming data ingestion is hosted in a serverless-processing paradigm, and it simultaneously enables the delivery of processed data into a feature store and supports the prediction request feature in a MLOps setting.

Using the traditional microservices model and considering the inevitable, concurrent spikes on the request path, a cost-effective solution is to leverage the reactive stream enabling technology offered by a cloud-provider data-messaging service. The enabling logic is then added to the microservice that provides the EAS pattern, separating the enabling code from the pure business logic. The reactive stream enables access to the incoming data flow for prediction requests against a model, and this enables low-latency prediction responses. The enabling logic is responsible for using the incoming prediction request to access the corresponding data stored in the feature store.



Shows detailed coverage across all major and subsection headings.


5. MLOps Practices within the Medallion Framework

MLOps enables the integration of Data Science and Machine Learning into the DevOps pipeline. MLOps encompasses a set of practices that deploys and maintains machine learning models in production reliably and efficiently. Monitoring of machine learning models, code, and processes needs to be measured accurately. Test-driven development (TDD) delivers high-quality code.

Unfortunately, all these MLOps best practices are not considered while designing and implementing the medallion architecture of large-scale cloud-native data platforms. The time scaling of training and retesting machine learning models is often not measured accurately and filling various data from the different predefined sources is still not tapped. Concepts of CI/CD for features are also not exercised in these data platforms.

Framework Layer	Responsibilities	Technologies Used	Automation Benefit
Policy Management	Define compliance rules	Policy-as-Code	Standardized governance
Compliance Monitoring	Detect policy violations	AI Monitoring Systems	Continuous compliance
Audit Orchestration	Prepare audit evidence	Workflow Engines	Faster certification
Incident Response	Coordinate remediation	SOAR Platforms	Reduced response time
Reporting & Analytics	Compliance dashboards	BI & Analytics Tools	Executive visibility
Continuous Improvement	Optimize governance maturity	ML Optimization	Adaptive governance

Table 5: Enterprise Compliance Automation Framework




Derived from years in the reference list.

5.1. Continuous Integration and Delivery for Features

Practising MLOps principles within the medallion framework involves engineering and serving features through the same CI/CD path as models. While a feature store provides compositionality, sharing and governing knowledge across groups, it does not deliver features in an on-demand way corresponding to the ingestion of labels. Thus, feature engineering must also be integrated into a CI/CD pipeline, with validation and testing practices ensuring timely delivery of supervised training data free from feature leakage and informative for the prediction task. Each ready-to-predict model requires one or more feature specifications, detailing the features and the data from which they are derived, as well as the expected target record time range. Feature engineering is then triggered and logged as part of the model's CI/CD job, providing the feedback loop necessary for data-centered MLOps.

Direct and indirect dependencies between models and features can be tracked in managed graphs. Enriching labels or prediction data with features from dedicated tables is a separate step, with features managed and governed in their own CI/CD pipelines through the feature store API. Extensive testing allows the creative exploration and engineering of supervised learning features, with the managed graphs focused on the continual delivery of features that are up to

date for model retraining. The triggers respond to the constraint that timely delivery is essential for the labels, and hence should be controlled and maintained separately from the feature engineering logic. When the features are directly mapped to the underlying data, they can even be produced and validated via real-time feature stores.



Data Encryption & Secrets Governance – Double Layer Flowchart

6. Scalability and Performance Considerations

A medallion architecture enhances the production readiness and proper integration of streaming features with an evolving cloud-based feature store, though challenges persist in ensuring performance, production readiness, and automated continuous integration and deployment of feature engineering. Well-designed pipelines automatically transform



and enrich data crucial for ML and AI. Such pipelines are typically horizontal scalable and capable of handling very large amounts of data with acceptable performance. A horizontal scalability and sharding strategy for the feature pipeline is required to make it operationally efficient and cost-effective.

Various feature extractors that represent a reference architecture serve specific groups of machine learning (ML) and artificial intelligence (AI) model input types. Each Feature Service acts as a single source of truth for a group of ML and AI models and is responsible for ensuring that these features are consistently updated and refreshed in the production environment. By implementing a horizontal scalability mechanism with clear boundaries, the Medallion architecture supports an evolving architecture alongside feature engineering pipelines.


Maturity Level	Governance Characteristics	Automation Capability	Risk Posture
Level 1 – Initial	Manual governance activities	Minimal	High Risk
Level 2 – Managed	Basic policy enforcement	Partial Automation	Moderate Risk
Level 3 – Defined	Centralized governance controls	Standardized Automation	Controlled Risk
Level 4 – Intelligent	AI-assisted orchestration	Advanced Automation	Predictive Governance
Level 5 – Autonomous	Self-healing governance ecosystem	Fully Autonomous	Optimized Risk Management

Table 6: Cloud Security and Compliance Maturity Levels

6.1. Horizontal Scalability and Sharding Strategies

The throughput and capacity of the complete data architecture remain limited by physical resources. While individual resources can be provisioned to fit specific capacity and latency requirements, designs should support horizontal scaling and elastic provisioning for different layers to cope with different workloads. Streaming and batch workload patterns that can scale independently allow allocation of dedicated resources for realtime or near-realtime workloads. For the streaming pipeline, sharding typically follows the dimension used for partitioning the stream, either from the source or by pre-sharding before publishing to a Message Queue service. Request-Rate-Limiting strategies should then be implemented in the feature serving layer to ensure available processing capacity to consume requests.

Batch workloads for feature generation typically process multiple features across different feature sets, and auto-scaling can be applied readily. Parallelization of the feature batch processing can be achieved at the feature or feature-set granularity level, by logically partitioning request-target feature sets, or at both levels. Parallel processing should align the workload shape to make effective use of the provisioning ratio by pre-sharding the request. The design also separates feature discovery from feature generation and serving, allowing each to be scheduled dynamically based on independently changing workloads.



Complements the theme pie chart with absolute mention counts.

7. Conclusion

The proposed simple, yet scalable, Medallion Architecture allows feature engineering to be conveniently formalized as a series of DAGs, browsing through the advantages of automated feature discovery and cataloging while remaining compatible with common MLOps practices. Special care is taken to ensure that features engineered at different latencies remain consistent, enabling their use in real-time scenarios without the need for periodic backfilling or training set preparation. The Medallion specifications allow the design of dedicated data-pipelines for real-time feature engineering, and their implementation in native cloud platforms can proceed transparently in the context of the cloud-native paradigm.

Such Medallion principles seem especially well-suited for cloud platforms, where services such as Cloud Pub/Sub, Cloud Dataflow or AWS Kinesis can be leveraged for ingesting and processing changing streams of data, other services such as Cloud Bigtable can be used to store such streams in a low-latency fashion, and Autoscale can seamlessly match change workloads and resource scaling.



Continuous Risk Improvement Lifecycle – Double Layer Flowchart

References

1. Armbrust, M., Das, T., Sun, L., Yavuz, B., Zhu, S., Murthy, M., Torres, J., van Hovell, H., Ionescu, A., Łuszczak, A., Świtakowski, M., Szafranski, M., Li, X., Ueshin, T., Mokhtar, M., Boncz, P., Ghodsi, A., Paranjpye, S., Senster, P., Xin, R., & Zaharia, M. (2020). Delta Lake: High-performance ACID table storage over cloud object stores. *Proceedings of the VLDB Endowment*, 13(12), 3411–3424.
2. Armbrust, M., Ghodsi, A., Xin, R., & Zaharia, M. (2021). Lakehouse: A new generation of open platforms that unify data warehousing and advanced analytics. In *Proceedings of the 7th Biennial Conference on Innovative Data Systems Research (CIDR 2021)*.
3. MANGALAMPALLI, B. M., KOLLA, S. H., APPA RAO NAGUBANDI, D. R., & SEGIREDDY, A. R. (2025). AN INTELLIGENT, REAL-TIME DIGITAL FABRIC FOR HEALTHCARE AND FINANCIAL ECOSYSTEMS USING AUTONOMOUS LEARNING AND GENERATIVE SYSTEMS. *TPM–Testing, Psychometrics, Methodology in Applied Psychology*, 32(S9 (2025): Posted 15 December), 3070-3086.
4. Munappy, A. R., Mattos, D. I., Bosch, J., Olsson, H. H., & Dakkak, A. (2020). From ad-hoc data analytics to DataOps. In *Proceedings of the International Conference on Software and Systems Processes (ICSSP 2020)*, 165–174.
5. Tamburri, D. A. (2020). Sustainable MLOps—Trends and challenges. In *Proceedings of the 22nd International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2020)*, 17–23.
6. Karlas, B., Interlandi, M., Renggli, C., Wu, W., Zhang, C., Babu, D. M. I., Edwards, J., Lauren, C., Xu, A., & Weimer, M. (2020).



- Building continuous integration services for machine learning. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 1163–1173.
7. Mangalampalli, B. M., Bandi, V. D. V. K., Kolla, S. K., & Kumar, M. V. K. (2025). Towards Self-Evolving Healthcare Intelligence: Integrating Advanced Learning Systems with Real-Time Clinical Data Pipelines. *Cultura: International Journal of Philosophy of Culture and Axiology*, 22(12s), 464-486.
 8. Lwakatare, L. E., Crnković, I., Rånge, E., & Bosch, J. (2020). From a data science driven process to a continuous delivery process for machine learning systems. In I. Crnkovic, M. Chechik, & P. Grünbacher (Eds.), *Product-focused software process improvement* (pp. 165–181). Springer.
 9. Zhang, J. M., Harman, M., Ma, L., & Liu, Y. (2020). Machine learning testing: Survey, landscapes and horizons. *IEEE Transactions on Software Engineering*, 48(1), 1–36.
 10. Renggli, C., Rimanic, L., Gürel, N. M., Karlaš, B., Wu, W., & Zhang, C. (2021). A data quality-driven view of MLOps. *IEEE Data Engineering Bulletin*, 44(1), 11–23.
 11. Mäkinen, S. S., Skogström, H., Laaksonen, E., & Mikkonen, T. (2021). Who needs MLOps: What data scientists seek to accomplish and how can MLOps help? In 2021 IEEE/ACM 1st Workshop on AI Engineering—Software Engineering for AI (WAIN), 110–112.
 12. Hai, R., Koutras, C., Quix, C., & Jarke, M. (2023). Data lakes: A survey of functions and systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(12), 12571–12590.
 13. Kreuzberger, D., Kühl, N., & Hirschl, S. (2023). Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11, 31866–31879.
 14. Hewage, N., & Meedeniya, D. (2022). Machine learning operations: A survey on MLOps tool support. In Proceedings of the 2022 IEEE International Conference on Industrial and Information Systems.
 15. Paleyes, A., Urma, R.-G., & Lawrence, N. D. (2022). Challenges in deploying machine learning: A survey of case studies. *ACM Computing Surveys*, 55(6), Article 114, 1–29.
 16. Behm, A., Palkar, S., Agarwal, U., Armstrong, T., Cashman, D., Dave, A., Greenstein, T., Hovsepian, S., Johnson, R., Krishnan, A. S., Leventis, P., Luszczak, A., Menon, P., Mokhtar, M., Pang, G., Paranjpye, S., Rahn, G., Samwel, B., van Bussel, T., van Hovell, H., Xue, M., Xin, R., & Zaharia, M. (2022). Photon: A fast query engine for lakehouse systems. In Proceedings of the 2022 International Conference on Management of Data, 2326–2339.
 17. Mangalampalli, B. M., Kolla, S. K., Bandi, V. D. V. K., Yandamuri, U. S., & Rani, P. S. (2025). Designing Intelligent Healthcare Ecosystems through Adaptive Data Integration and Autonomous Learning Systems. *Vascular and Endovascular Review*, 8(20s), 330-347.



18. Schneider, J., Gröger, C., Lutsch, A., Schwarz, H., & Mitschang, B. (2024). The lakehouse: State of the art on concepts and technologies. *SN Computer Science*, 5, Article 449.
19. AbouZaid, A., Barclay, P. J., Chrysoulas, C., & Pitropakis, N. (2025). Building a modern data platform based on the data lakehouse architecture and cloud-native ecosystem. *Discover Applied Sciences*, 7, Article 166.
20. Tagliabue, J., & Greco, C. (2024). Reproducible data science over data lakes: Replayable data pipelines with Bauplan and Nessie. In *Proceedings of the 2024 International Workshop on Data Management for End-to-End Machine Learning*.
21. Bollineni, S. (2022). Enhancing data lakehouse architecture with DevOps and MLOps practices. *Journal of Mathematical & Computer Applications*, 1(2), 1–5.
22. Bollineni, S. (2023). DevOps approaches to managing and deploying machine learning algorithms in lakehouse environments. *International Journal of Business, Management and Economics Research*, 7(8).
23. Paleyes, A., Urma, R.-G., & Lawrence, N. D. (2022). Machine learning deployment: Challenges and opportunities in production environments. *ACM Computing Surveys*, 55(6), 1–29.
24. Renggli, C., Rimanic, L., Gürel, N. M., Karlaš, B., Wu, W., & Zhang, C. (2021). Data quality in machine learning systems: A systematic perspective on MLOps. *IEEE Data Engineering Bulletin*, 44(1), 11–23.
25. Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2021). Understanding deep learning requires rethinking generalization. *Communications of the ACM*, 64(3), 107–115.
26. Xu, Y., Liu, X., Cao, X., Huang, C., Liu, E., Qian, S., & Zhang, Y. (2021). Artificial intelligence: A powerful paradigm for scientific research. *The Innovation*, 2(4), 100179.
27. Sculley, D., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2020). Machine learning: The high-interest credit card of technical debt. In *Proceedings of the 2020 International Conference on Machine Learning and Applications*.
28. Amershi, S., Weld, D., Vorvoreanu, M., Fournay, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2020). Software engineering for machine learning: A case study. In *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice*, 291–300.
29. Tamburri, D. A., Kazman, R., & Fahimi, H. (2020). Cloud-native software architecture for machine learning operations. In *Proceedings of the International Conference on Cloud Computing and Services Science*.
30. Sato, D., Wider, A., & Windheuser, C. (2020). Continuous delivery for machine learning: MLOps and machine learning pipelines. *IEEE Software*, 37(5), 31–37.



31. Karlaš, B., Wu, W., Renggli, C., Zhang, C., & Weimer, M. (2021). Towards data-centric machine learning operations: Managing data quality and lifecycle. *IEEE Data Engineering Bulletin*, 44(1), 24–35.
32. Mangala, N., & Kolla, S. H. (2025). Real-Time Feature Engineering for Streaming AI Workloads Using PySpark and Azure Event Hubs. *International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management*, 1(6), 93-105.
33. Hai, R., Quix, C., & Jarke, M. (2021). Data lakes: A survey of architectures, management, and applications. *Information Systems*, 100, 101799.
34. AbouZaid, A., Barclay, P. J., Chrysoulas, C., & Pitropakis, N. (2025). DataOps and cloud-native approaches for resilient lakehouse platforms. *Discover Applied Sciences*, 7, Article 166.