ICRAMCS 2026

THE EIGHTH EDITION OF THE INTERNATIONAL CONFERENCE ON RESEARCH IN APPLIED MATHEMATICS AND COMPUTER SCIENCE

April 23-24-25, 2026 | Marrakech, Morocco

From MLOps to LLM-Augmented Operations: A Survey of Foundations, Challenges, and Opportunities for Generative AI in ML Pipelines

Communication Info

Authors:
Ichraq ESSADEQ[1]
Said NOUH [1]
Khalid KANDALI [1]

[1] *Laboratory of Information Technology and Modeling (LTIM), TCA Team, Faculty of Sciences Ben M'Sick, Hassan II University of Casablanca, Casablanca 20360, Morocco.*

Keywords:
MLOps
LLM-Augmented Operations
Generative AI
Large Language Models
DevOps
Machine Learning Lifecycle
CI/CD Automation
Concept Drift Detection
Feature Engineering
AI-Driven Optimization
ML Systems Governance
Foundation Models

Abstract

The large-scale deployment of machine learning (ML) systems has exposed persistent gaps between experimentation and reliable production. Machine Learning Operations (MLOps) extends DevOps to data-centric pipelines, promoting automation, reproducibility, and monitoring [1], [2]. However, key challenges remain, including technical debt, limited reproducibility, and concept drift [2], [3].

This paper presents a structured survey of DevOps and MLOps foundations based on literature published between 2015 and 2024, highlighting architectural and operational distinctions between software-centric and data-centric systems [1], [4].

We then propose a conceptual framework mapping these challenges to emerging generative AI capabilities. We argue that large language models (LLMs) can act as optimization tools for automation, semantic-aware drift analysis, and collaborative support [3], [5].

This work establishes a foundation for LLM-augmented MLOps and outlines research directions for intelligent ML operations.

© ICRAMCS 2026 Proceedings ISSN: 2605-7700

References

[1] D. Kreuzberger, N. Kühl, and S. Hirschl, "Machine Learning Operations (MLOps): Overview, Definition, and Architecture," *IEEE Access*, vol. 10, pp. 31866–31879, 2022, doi: 10.1109/ACCESS.2022.3156391.

[2] A. Paleyes, R.-G. Urma, and N. D. Lawrence, "Challenges in Deploying Machine Learning: A Survey of Case Studies," *ACM Comput. Surv.*, vol. 55, no. 6, pp. 1–29, Jul. 2023, doi: 10.1145/3533378.

[3] D. Sculley *et al.*, "Hidden Technical Debt in Machine Learning Systems," in *Advances in Neural Information Processing Systems (NeurIPS)*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds., Montreal, Canada: Curran Associates, Inc., 2015. Accessed: Feb. 26, 2026. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2015/file/86df7dcfd896fcaf2674f757a2463eba-Paper.pdf

[4] S. Amershi *et al.*, "Software Engineering for Machine Learning: A Case Study," in *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, IEEE, May 2019, pp. 291–300. doi: 10.1109/ICSE-SEIP.2019.00042.

[5] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Comput. Surv.*, vol. 46, no. 4, pp. 1–37, Apr. 2014, doi: 10.1145/2523813.