



An analysis of the challenges in the adoption of MLOps

Chintan Amrit^{*}, Ashwini Kolar Narayanappa

Amsterdam Business School, University of Amsterdam, P.O. Box 15953 1001 NL Amsterdam, the Netherlands

ARTICLE INFO

JEL classification:

C88
D83
L86
L17
O32

Keywords:

Machine learning operations (MLOps)
Grounded theory
Data science

ABSTRACT

The field of MLOps (Machine Learning Operations), which focuses on effectively managing and operationalizing ML workflows, has grown because of the advancements in machine learning (ML). The goal of this study is to examine and contrast the difficulties encountered in the implementation of MLOps in enterprises with those encountered in DevOps. An SLR (Systematic Literature Review) is the first step in the research process to find the issues raised in the literature. The results of this study are based on qualitative content analysis using grounded theory and semi-structured interviews with 12 ML practitioners from different sectors. Organisational, technical, operational, and business problems are the four distinct aspects of challenges for MLOps that our study highlights. These challenges are further defined by eleven different themes. Our research indicates that while some issues, such as data and model complexity, are unique to MLOps, others are shared by DevOps and MLOps as well. The report offers suggestions for further research and summarises the difficulties.

Introduction

The increasing use of machine learning (ML)-based techniques has made it more difficult to integrate them into production systems while maintaining the efficiency and dependability of constantly changing ML projects (Rzig et al., 2022). According to VentureBeat research (VB Staff, 2019), just 13% of machine learning projects in the business can reach production. Prototyping and production deployment can take up to 90% of the project's time, despite it seeming like the final 10% (Flaounas, 2017). In response to these challenges, the notion of MLOps (Machine Learning Operations) has emerged as a comprehensive collection of procedures intended to guarantee the dependable and effective implementation and upkeep of machine learning (ML) models in operational settings (Alla & Adari, 2020). MLOps is an adaptation of the DevOps discipline, which was created to address comparable continuous deployment problems for standard software and has been in existence for more than ten years. MLOps, as opposed to DevOps, tries to address problems specific to machine learning, like continuous training, model monitoring, testing, and versioning of data and models.

Although MLOps are becoming more and more popular, there isn't much research on how automation is adopted and how that affects changes in ML-enabled systems (Calefato et al., 2022). Nowadays, all we know about MLOps comes from a fragmented field of white papers, stories, and opinion articles (Shankar et al., 2022). According to John et al. (2021), a significant amount of the literature on software

engineering (SE) best practices for machine learning applications is non-peer-reviewed, or grey literature. This category includes presentation slides, blog posts, and white papers (Serban et al., 2020). Most writers agree that MLOps is challenging. 90% of ML models are never put into production, and 85% of ML initiatives are ineffective (Shankar et al., 2022).

Very few production-grade machine learning projects were discovered by Calefato et al. (2022) during their examination of MLOps projects on GitHub. They also drew attention to the dearth of MLOps tool-using open-source machine learning systems (Calefato et al., 2022). There are very few articles on case studies from organisations. Some papers examine the technical obstacles of applying MLOps and offer solutions (Cardoso Silva et al., 2020; Symeonidis et al., 2022). We see the need for a comprehensive framework that discusses the challenges organizations face in implementing MLOps. While existing research has explored various challenges in the implementation of MLOps, there is a notable gap in providing a comprehensive framework of the different challenges. Previous studies have often focused on specific challenges in deploying and operating machine learning in practice or surveyed individual case studies of ML deployment (for e.g.: Baier et al. (2019), Diaz-de-Arcaya et al. (2023), Paleyes et al. (2022)) but a unified framework that integrates all essential elements is still missing. This lack of a comprehensive overview has led to inconsistent understandings and implementations of MLOps across different organizations and projects. Additionally, the rapid evolution of ML technologies

^{*} Corresponding author.

E-mail addresses: c.amrit@uva.nl (C. Amrit), ashwinikn30@gmail.com (A.K. Narayanappa).

and practices has outpaced the academic literature, resulting in a disconnect between theoretical concepts and practical applications. This research aims to bridge this gap by providing a thorough, up-to-date, and practice-oriented conceptualization of the challenges in MLOps implementation, synthesizing insights from both academic literature and industry expertise.

Since there isn't much research that contextualises the literature on MLOps, exposes the challenges related to MLOps, and offers a thorough overview of recent material, we focused on evaluating the challenges that companies face while implementing MLOps in this study. We therefore ask the following research question:

What are the key challenges organisations face while implementing MLOps?

The results from our structured literature review (SLR) show many implementation problems in addition to a general lack of empirical data on MLOps. The results of our interviews show that there are four major types of obstacles: technical, operational, business, and organisational. These challenges are further divided into eleven topics for explanation. Also, one of the primary contributions of our work is arriving at a Typology of the challenges that represent Type 1 theory (Gregor, 2006), which is the central theoretical contribution of this paper.

The remainder of this paper is structured as follows. Section 2 gives a background of the existing literature on MLOps, section 3 describes our research design, and section 4 provides results from the Systematic Literature Review. Section 5 presents and analyses the results of the Interviews. We discuss these results in Section 6 and we conclude the paper in Section 7.

Literature background

We introduce and review existing literature on MLOps, which is a modification of DevOps. The following subsection highlights their differences.

DevOps vs MLOps

The definition of DevOps is the subject of debate in research (Krey et al., 2022). Different views and stances exist in the scientific community. Macarthy and Bass (2020) highlight two conflicting perspectives: one views DevOps as a cultural movement for rapid software development, while the other views it as a job title requiring both development and IT operation expertise. In this paper, we will consider the following definition: DevOps is a paradigm that aims to bring innovative products and features to the market faster (Ebert et al. 2016) by integrating the development, testing, and operational software development teams through automation, tools, (Macarthy & Bass, 2020) and cultural philosophy that emphasizes team empowerment, cross-team communication, collaboration, and technology automation.

Academic studies focusing on the principles, techniques, and advantages of DevOps are well-documented. According to research studies, the fundamental components of DevOps include continuous delivery, software development process automation, and communication and collaboration between different actors delivering the software (Subramanya et al. 2022; Davies & Daniels, 2016). MLOps is not the same as DevOps as it has more components involved than just the software application. In DevOps, there needs to be a close collaboration with the development and operations teams, whereas in MLOps, it also extends to the data science team. Makinen et al. (2021) also confirmed in their survey research that, like DevOps for traditional software, MLOps, or continuous delivery of machine learning software, is becoming a prerequisite for businesses using ML in production, especially for more ML-mature organizations. Since cloud-based DevOps and containerised microservices have shown success in production deployments (Hui Kang et al., 2016), organisations are adopting

continuous practices in ML system development through the adoption of DevOps concepts for end-to-end automation (John et al., Sep 2021). MLOps is presently in the "Peak of Inflated Expectations" phase, per Gartner's (2022) hype cycle for data science (DS) and machine learning (ML) (Choudhary & Krensky, 2022).

MLOps, like DevOps, has no formal definition but can be seen as the meeting point of ML and DevOps practices (Matsui & Goya, 2022). An MLOps project encompasses infrastructure management, integration, testing, release, and deployment, as well as automation, integration, and monitoring throughout the entire process of constructing an ML system (Testi et al. 2022a). The goal of MLOps is to create a set of procedures for quickly and efficiently creating machine learning models using tools, deployment flows, and work processes.

While DevOps primarily deals with software actions, MLOps involves constructing, training, and tuning machine learning models using hyperparameters and datasets (Liu et al. 2020). Another significant distinction between MLOps and DevOps is that while the former emphasizes optimizing delivery procedures and standardizing development environments, the latter places a strong emphasis on data from which applications are derived. As a result, data scientists and software engineers collaborate closely in MLOps (Mboweni et al., 2022).

MLOps brings in newer challenges than traditional software with its complexity. Understanding the challenges of MLOps involves grasping the steps required for training and deploying ML modules. This includes data preparation, dividing data into training, testing, and cross-validation sets, selecting an ML model and hyperparameters, training the model with iterative adjustments, validating the model, and deploying it. Once deployed, ML features require monitoring, considering ML-specific factors like biases and drift. Additionally, techniques for improving the model in real-time while in use must be incorporated into the monitoring system. In summary, the pre-model completion phases in ML resemble a waterfall approach while operationalizing the model aligns with conventional software practices (Makinen et al., 2021; Granlund et al., 2021).

MLOps, like DevOps, aims to automate software delivery, ensuring continuous delivery and feedback loops. However, integrating ML applications into a DevOps-based CI/CD environment requires additional steps due to variations in ML methodologies. The adaptation of MLOps is still in its early stages, with limited research compared to DevOps. Data scientists and operations teams are working on automating the end-to-end ML lifecycle using DevOps concepts, facing challenges in generalizing MLOps components. Different designs and processes, such as Iterative/Incremental and Continuous delivery for ML(CD4ML), have been proposed (Subramanya et al., 2022). It is crucial to include data, infrastructure, and core ML code in the MLOps lifecycle, as highlighted by Sculley et al. (2015).

Challenges in ML systems, like model complexity and reproducibility, are different from DevOps components. However, MLOps shares typical DevOps limitations, including the oversight of human factors in technology development and adoption (Mucha et al., 2022). Addressing versioned data, ML models, and dependencies is a challenge, as existing tools and processes in DevOps are inadequate. Expertise in data science, AI, and software engineering is required for successful MLOps implementation (Lwakatare et al., 2020). The literature background confirms that there is a growing need for MLOps in businesses that use ML in production.

MLOps tools

Many organizations have the option to either build their tools and platforms or choose from the available tools in the market. MLOps process and its components can be built using a combination of open-source tools and enterprise solutions, which allows for flexibility and customization. It is possible to leverage both enterprise and open-source tools together to achieve MLOps objectives effectively (Kreuzberger et al., 2023).

Several MLOps platforms have emerged in recent years to offer guidelines for building enterprise-level AI/ML applications (Garg et al., 2022). Approximately 50% of IT businesses are using these tools, which is evident from the staggering numbers of projects, developers, and companies actively engaged in platforms like GitHub. As of January 26, 2022, 200 million projects are being worked on by 65 million developers and 3 million companies (Symeonidis et al., 2022).

Kubeflow, MLFlow, and Kedro are some of the open-source MLOps technologies. These platforms are open-source projects that contain a curated selection of compatible machine-learning tools and frameworks to simplify the automated development process (Garg et al., 2022). The most well-known MLaaS include Azure's Machine Learning, Google's AI Platform, and Amazon's SageMaker (Maya & Felipe, 2021).

In the next section we discuss our research design for understanding the key challenges organizations face in their implementation of MLOps.

Research design

When there is minimal previous research or when the empirical environment is new or understudied, inductive reasoning based on qualitative data is particularly applicable (Bansal et al., 2018). Therefore, to obtain a comprehensive overview of pertinent research, we first performed a systematic literature review (SLR). Concurrently, we conducted semi-structured interviews with ML specialists from a range of businesses. To reach pertinent conclusions, the results of the SLR were then contrasted with findings from expert interviews (in the discussion portion).

Individual interviews with participants who are knowledgeable about MLOps and can offer an opinion on it were used to gather data for the study. This is a typical approach to gathering data (Beitin, 2012).

To discover individuals with experience working with machine learning and/or its operations, we used a purposeful sampling technique to find participants with experience working with machine learning and/or its operations with job titles such as Data scientist, ML Manager, MLOps engineer, or comparable. The respondents were found through multiple channels like in the personal network of the first author, LinkedIn, and online ML communities. We got 12 respondents who agreed to the interview. We used an in-depth, semi-structured interview method which allowed us to change the ordering of questions or ask about the relevant finding of one interview to another (Qu and Dumay, 2011). All the interviews were performed online using Microsoft Teams, and they were all recorded after getting the participant's consent. Interviews typically lasted between 25 and 60 minutes. To ensure authentic responses, interview questions were not revealed before the interviews, and the interview started off with a short introduction to the purpose of this research. The first draft was created using the transcribing feature of Microsoft Teams, and all adjustments were carefully reviewed and modified. The first author also made notes during the interview, which made the participants think and add to the answer during the silence (Qu & Dumay, 2011).

Fig. 1 shows our overall research design. As a result, we conducted the semi-structured interviews and the structured literature review simultaneously. The insights from the SLR and the interviews were combined to create the typology of challenges that we arrive at in Fig. 2.

We start with the details of the SLR in the next section.

Structured literature review

To better understand the obstacles that organisations encounter while deploying MLOps, we started by performing a thorough analysis of existing literature. Setting a research topic, choosing pertinent studies, extracting data, and synthesising the data were all steps in the methodology we followed to perform the SLR (Kitchenham, 2004). We employed a manual search strategy using electronic research resources for this SLR. To find pertinent material, we searched both Scopus and ScienceDirect, two important databases. We then complemented this search with a search using Google Scholar and the snowballing technique (Wohlin, 2014).

The release of "Hidden Technical Debt in Machine Learning Systems" in 2015 (Sculley et al., 2015) helped to popularise the relatively young area of MLOps. We created inclusion and exclusion criteria, which are outlined in Table 3 and Table 4 (in Appendix A), and limited our search to publications published between 2015 and 2023 to make sure that we concentrate on pertinent information. When it came to the challenges associated with introducing MLOps, we prioritised organisational research. The initial results were limited to papers published between 2018 and 2023 and were produced using the search phrases "MLOps OR machine learning in production OR machine learning operations". Out of 175 items, we only included research articles and conference papers with library access. The first author manually screened the abstracts of these papers and selected ten articles for data extraction and synthesis. The process is depicted in Fig. 3 (in Appendix B).

Although not all the publications we found addressed MLOps specifically, nine of them addressed MLOps' challenges. We only included sections that were directly related to the difficulties faced by MLOps to ensure their relevancy. Table 1 displays the concept matrix for various tasks.

When deploying ML models in production, organisations have distinct problems related to data quality and quantity, as noted by Baier et al. (2019) in their analysis of MLOps in practice. Long-term data validity is especially challenging to address. The data utilised for training must be of a high calibre and adequate throughout time to guarantee correct predictions when the ML model is presented with fresh data for testing in production. The authors also discovered that managing concept drift is an additional special difficulty related to MLOps. Concept drift is the process via which the model becomes obsolete as underlying data, user behaviour, and/or statistical characteristics of the data change over time. Baier et al. (2019) identified non-technical challenges as well as technical ones. These challenges included adhering to organisational and national standards of ethical and legal compliance, integrating and streamlining the ML model into existing business operations, and a lack of collaboration and communication among technical and non-technical teams.

According to Tamburri (2020), there are several internal and external hurdles that organisations have when implementing MLOps. These include integrating MLOps with the current IT infrastructure, which may lead to problems with scalability and sustainability, and making sure MLOps are in line with the organization's objectives and core values. To ensure adherence to pertinent laws and standards, the MLOps pipeline's safety and security are also major concerns. Finding qualified candidates with the MLOps skill set who are motivated,




Fig. 1. The research design adopted in this paper.

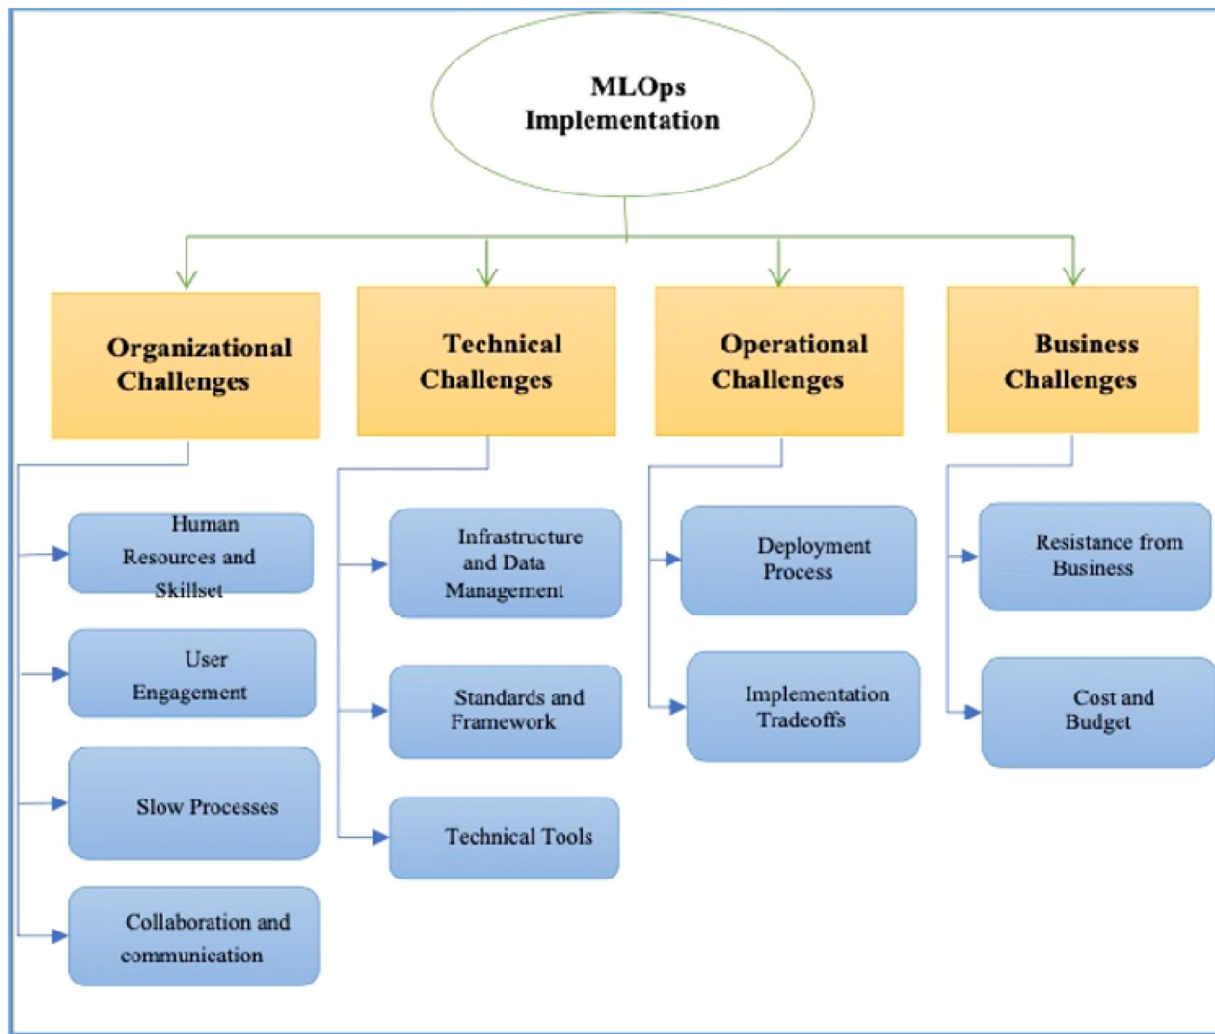


Fig. 2. The typology of challenges in MLOps implementation.

experienced, and relevant is another difficulty.

Different organisations have different obstacles when putting MLOps into practice. [Granlund et al. \(2021\)](#) pointed out that the generation, integration, and discovery of new features that are contributed to the feature set fed to the ML model all result in the constant updating of ML models. Consequently, MLOps have an urgent issue in developing the capacity to deploy an updated ML model that enables efficient and instantaneous deployment. The Integration Challenge and the Scaling Challenge are two more significant MLOps difficulties that are covered in this study. The Integration Challenge arises when there are disparate contractual duties, data formats, and machine learning features among several organisations, resulting in incompatible APIs.

In organisational frameworks, [Lima et al. \(2022\)](#) examine the different MLOps methods and difficulties. One of these difficulties is ML model versioning, which has to do with selecting the right model among a variety of versions of the ML models created along the ML pipeline. The authors also note that a major source of resistance to MLOps, especially among internal teams within businesses, is the absence of standardisation. Given that a multitude of tools across various platforms can be used to create ML models, datasets, and feature sets ([Lima et al., 2022](#)).

Challenges related to the explainability, and transparency of machine learning models and decisions were discussed by [Testi et al. \(2022\)](#). In the context of MLOps for organisations, a correct knowledge of the ML model through transparency and explainability is essential since the elements influencing the ML model's various decisions must be

precisely recognised. In addition to a dearth of standardisation and best practices, businesses implementing MLOps have a serious issue with limited and isolated research. Furthermore, [Testi et al. \(2022\)](#) discovered that companies were struggling to deal with the continuously changing data, which occasionally necessitates retraining or even a total rebuild of the model in the machine learning pipeline.

[Painoli and Datrike \(2021\)](#) concentrate on general AI problems in businesses. The authors draw attention to businesses that are having trouble selecting which qualities to include in their model, how to clean their data, and what algorithm is appropriate. The authors also point out that successful deployment requires overcoming challenges with infrastructure, data privacy, and security. Lack of qualified personnel also makes it difficult for organisations to understand and share the model's results. The authors conclude that making sure the model is fair, accurate, and transparent when making judgements is one of the largest problems in applying AI/ML ([Painoli & Datrike, 2021](#)).

[Paleyes et al. \(2022\)](#) survey case studies to identify challenges associated with implementing machine learning models. They note that feature engineering, model selection, and data management are issues that enterprises frequently deal with. The authors also emphasise the importance of the models' interpretability and explainability as well as the need for rigorous testing and validation procedures. They also add that a lack of infrastructure and skilled personnel may make it difficult to use machine learning models. When applying machine learning models, the authors point out the growing significance of data privacy and ethical considerations including justice and transparency ([Paleyes](#)

Table 1

Concept matrix of challenges faced in MLOps.

Articles	MLOps Challenges – Concept matrix									
	Model Issues (Scalability, accuracy, Versioning, Monitoring)	Data Issue (Availability, quality, privacy)	Integration and Infrastructure	Regulatory Compliance	Standardization	Tools Support for MLOps	Testing	Lack of Talent	Collaboration & communication	Lack of Knowledge
Paleyes et al. (2022)	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
Granlund et al., (2021)	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓
Lima et al., (2022)	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗
Baier et al., (2019)	✓	✓	✓	✓	✓	✓	✗	✗	✓	✓
Tamburri (2020)	✗	✓	✗	✓	✗	✗	✗	✓	✗	✓
Painoli & Datrke (2021)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
Schröder and Schulz (2022)	✗	✓	✓	✗	✗	✗	✓	✗	✗	✗
Testi et al., (2022)	✓	✓	✗	✗	✓	✗	✗	✗	✗	✓
Zhang et al. (2020)	✓	✓	✗	✗	✗	✗	✗	✓	✗	✓
Kreuzberger et al. (2023)	✓	✗	✓	✗	✗	✗	✗	✓	✓	✓
Diaz-De-Zrcaya et al. (2023)	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓
Total	9	10	8	3	3	1	1	4	4	7

et al., 2022).

Organisational challenges in monitoring machine learning models are explained by Schröder and Schulz (2022) while emphasising operations. Issues pertaining to data drift, mislabelling, and quality are emphasised. The authors also emphasise how important it is to keep an eye on the model's performance and address issues with fairness, bias, and interpretability. Additionally, they note that companies may find it challenging to implement and integrate monitoring systems into already-existing infrastructures (Schröder & Schulz, 2022).

Conversely, Zhang et al. (2020) contend that the development of machine learning artificial intelligence systems is particularly challenging in knowledge-intensive fields. Ineffective adoption and resistance to change may arise from a lack of understanding regarding the integration of AI with existing organisational systems and protocols. A further challenge they identify is choosing which models and algorithms to utilise from the many accessible models and algorithms while maintaining a balance between accuracy, accessibility, and complexity. Data quality issues could negatively impact AI system performance because they require large amounts of data to learn. This presents another challenge. Furthermore, while multidisciplinary teams with a range of skill sets are necessary for developing AI systems, they can be challenging to locate and manage (Zhang et al., 2020).

Kreuzberger et al. (2023) identify several key challenges for implementing MLOps, categorized into organizational, ML system, and operational challenges. Organizational challenges include the need for a culture shift towards product-oriented machine learning, a lack of skilled experts in roles like ML engineers and DevOps engineers, and the necessity for multidisciplinary teamwork. ML system challenges involve designing for fluctuating demand, especially in ML training processes, and managing potentially voluminous and varying data. Operational challenges include the difficulty of manually operating ML due to complex software and hardware stacks, the need for robust automation to handle constant data streams and retraining, governance of numerous artifacts, versioning of data, models, and code, and the complexity of resolving support requests due to the involvement of multiple components and parties.

The systematic review by Diaz-de-Arcaya et al. (2023) examines challenges, opportunities, and trends in MLOps and AIOps. Where they define AIOps as a solution for handling growing data and IT infrastructures. The key challenges they describe include cross-domain expertise requirements, data management complexities, and organizational culture barriers. They highlight opportunities, such as, continuous delivery in MLOps, AI applications across industries, and edge computing utilization. They discuss frameworks facilitating MLOps and AIOps adoption, emphasizing lifecycle management tools, and state that while MLOps is more prevalent in traditional industries, AIOps is gaining traction in emerging technologies like 5G. The authors conclude that both methodologies are crucial for successful AI implementation in production environments, requiring collaborative culture and cross-functional skills to overcome identified challenges.

The literature review, as summarized in Table 1, reveals a range of challenges organizations face when implementing MLOps. The most prevalent issues include model-related challenges (such as scalability, accuracy, versioning, and monitoring), data issues (availability, quality, privacy), and integration and infrastructure concerns (the first three columns of Table 1). Regulatory compliance, standardization, and lack of tool support for MLOps were also identified as significant hurdles. Furthermore, the review highlights human-centric challenges like lack of talent, collaboration difficulties, and knowledge gaps. However, the impact of tool support and testing has been given little attention in the literature. We found they were mentioned only once in the literature we reviewed (Table 1). In general, while studies provide valuable insights into specific aspects of MLOps implementation, they often focus on specific components or practices. This fragmented approach underscores the need for a more comprehensive framework that integrates all essential elements of MLOps. In the next section, we address this gap by using semi structured interviews and inductively creating a typology (a Type 1 theory) of the challenges in implementing MLOps.

In the next section, we discuss the results from the semi-structured interviews that we compare with the SLR findings.

Results from the semi structured interviews

We used an in-depth, semi-structured interview method that allowed us to change the order of questions or ask about the relevant findings of one interview to another. To organize, analyze, and visualize semi-structured data, we used ATLAS.ti, a computer-aided qualitative data analysis software. Table 5 in Appendix C provides an overview of the interview participants. The data were analyzed inductively using a grounded theory approach while reflecting on the research question. After transcribing the first two interviews, we started coding the quotations, which is also known as open coding. Coding involves giving short labels to groups of data. We added codes and compared them to the prior data after finishing the transcription of each interview, ensuring that the data coding and the analytic process were consistent. We initially had 102 open codes, but after multiple iterations and focusing on the research question, we reached saturation and ended up with 51 first-order codes. The 11th and 12th interview transcriptions did not yield any new codes, but we utilized the quotes from them to broaden the study's applicability.

In Fig. 4 (Appendix D), we see a Sankey diagram representing the codes generated from the different interview participants. Participants 1 and 7 had more codes, and the commonality between them was that they had implemented MLOps in multiple organizations. Moreover, it clearly shows that the last two interviews did not generate any new codes. In the next step, we combined related first-order concepts into eleven focus codes or second-order themes. The second-order themes were then categorized into four aggregated dimensions (Fig. 5, Appendix D). Fig. 6 (Appendix D) provides an impression of our codebook. The study's findings will comprise a thorough narrative supported by data that uses 2nd-order themes and aggregated dimensions, frequently quoting the first-order statements of the informants.

Following analysis, we created a typology in Fig. 2 that reflected the pertinent themes (blue) beneath the aggregate dimensions (yellow). Although they are not displayed here, the coding system (Fig. 5 in Appendix D) contains the first-order notions.

Organizational challenges

As shown in Table 2, the data was distilled into high-density codes for organisational problems, which comprise four topics that are elaborated upon below.

Human Resources and Skillset: The area of machine learning is relatively new to the corporate world, posing unique resource challenges. The fact that resources employed for machine learning engineering frequently lack relevant experience is one of the main problems. It's possible that what they learned in school or college won't apply in practical situations. This mismatch between academic training and industrial demands is a problem since it takes a lot of time for these

resources to become up to speed and acquire the knowledge and expertise needed to accept responsibility for data.

"We have been facing the problem for the last 5–6 years regarding the resources not being ready, and sometimes what they have learned in college or school is not practical enough. So, the onboarding time for machine learning engineers is 6 to 7 months. Only after six to seven months can you trust that they can take data ownership," - Participant 1

The lack of relevant skills needed to build MLOps pipelines in one team or sometimes even within the organization can lead to knowledge gaps.

User Engagement and Resistance: Even if some current team members may resist learning new skills because they lack competence, it is nevertheless crucial to train them.

"Because along with our pipeline came all the new learning to do, which was challenging for some team members to adapt to these new tools." - Participant 11. Teams working on machine learning applications might not see the benefits of unused MLOps pipelines. As Participant 7 said, *"I would say the biggest challenge was keeping the teams that we worked with, involved and interested. We spend a lot of time involving different teams and making them feel engaged in the process because our biggest challenge was not making the model or making the pipeline, but getting people to just think about using it and making people see the benefits."*

Resistance to adoption can be reduced through the Organizational culture, as Participant 1 explains, *"Culture is good here because it is a data-based company, and that is why the situation is better,"* demonstrating that teams did not oppose their use of the newest MLOps tools.

Slow Processes: One recurring topic in talks about MLOps is the importance of organisational and team structure. The successful implementation of MLOps typically requires the participation of several teams. These teams might, however, have different priorities, which could lead to dependencies that continue longer than expected. Participant 1 shared, *"Very often, it's not one team that manages this whole thing, and that means that if you, for example, have a data engineer who works with you from a different team, they're also working on other projects. They need to manage all of those, and that means that they don't always have time to help you out with whatever you need at that time, causing delays."*

Collaboration & Communication: Our results show that poor understanding of machine learning models and data within an organisation might cause problems with collaboration. Effective coordination is difficult since MLOps involves varied teams including data scientists, developers, and ML engineers who have different priorities and objectives. Problems with teamwork and miscommunication can occur when a team lacks the essential competencies; Participant 7 shared, *"Very often it's not one team that manages everything which means that they don't always have time to help you out with whatever you need making collaboration difficult."*

It can be hard to communicate changes to a model to non-ML clients and to numerous teams at once.

"You get a lot of miscommunications between those teams and a lack of accountability or responsibility. Because it's a lot easier to say. Yeah, there's a problem. That's not my problem. That's their problem. And you just start going in circles." -Participant 7

Technical challenges

Organisations looking to put MLOps into practice usually face a variety of technological difficulties. These difficulties can be caused by a variety of factors, including the intricacy of machine learning models, a lack of data, or the fusion of various tools and technologies. After discussing the organisational problems, this has been explained under the following three themes. More codes are produced in the data as a result, creating more technological difficulties.

Infrastructure and Data Management: The underlying infrastructure is essential to the complicated task of building pipelines for machine

Table 2
Code density.

Category	Themes	Codes	Total density
Organizational Challenges	Human Resources and Skillset	7	23
	User Engagement and Resistance	4	
	Slow Processes	8	
	Collaboration & Communication	4	
Technical Challenges	Infrastructure and Data Management	7	14
	Standards and Framework	4	
	Technical tools	3	
Operational Challenges	Difficulties in Pipeline	4	7
	Implementation Trade-offs	3	
Business Challenges	Business Value	2	7
	Cost and Budget Constraints	5	

learning models. Dependable infrastructure is required for the production implementation of machine learning models. To run models at scale, organisations need to make sure that the deployment environment can support the extra processing and storage requirements. It might be difficult to manage dispersed systems, maximise hardware efficiency, and scale up resources. Integration problems arise if the infrastructure configuration does not match the technical stack needed for machine learning applications. In MLOps, efficient data governance and management are essential. The confidentiality, quality, and accessibility of the data used for inference and training must be guaranteed by organisations. As Participant 6 says, *"I think the big challenge we have is the enormous variation in data. As I mentioned, we have half a million models, so it's very hard. If the data sets are very different, then the data quality is very different. Customers are not interested in low-quality results"*. The quality of MLOps depends on the data supplied by the pipeline; Participant 3 elaborates, *"The entire MLOps output depends on how good your data quality is and rather not on MLOps."*

Standards and Framework: There seems to be a *"gap in the best practices provided by industry leaders and the development of the technology"*, says Participant 10. Standards and guidelines should be established by leaders in the industry and solution suppliers, but sadly, this isn't the case everywhere at the moment. Teams find it challenging to comprehend and use open-source tools efficiently when there is a dearth of clear and easily accessible documentation, which can result in inconsistent results and problems during implementation.

Writing unit tests is standard procedure in traditional software development to guarantee code quality. However, because machine learning models are non-deterministic, testing becomes more difficult in this environment. The standardisation of testing methodologies is made more difficult by the need for unique methods and techniques for testing machine learning systems, which go beyond simple software testing. *"For regular software, I'm used to writing unit tests. For machine learning, these tests are a bit more complex because it's not deterministic."* – Participant 3

Technical Tools: Having the appropriate technological tools is one of the most important needs for implementing MLOps, according to several interviewees. Two interviewees claim that although there are many tools accessible, it can be difficult to select the best one depending on business requirements. Furthermore, the cost of locating a quality MLOps platform can be high.

Technical tool challenges draw attention to the challenges that organisations encounter. The efficacy and efficiency of MLOps procedures might be hampered by problems including instability, poor integration, restricted capabilities, and the requirement for customisation. This demonstrates how important it is to carefully assess and test tools, keep abreast of new developments, and take the organization's MLOps-specific objectives and limits into account.

Operational challenges

Two themes of operational challenges that are more focused on day-to-day operations and closely relate to technological challenges surfaced during the interview.

Deployment Process: As was already established, there are certain technical difficulties with data, tools, and infrastructure in the deployment process. But interviewees have also pointed up certain operational difficulties. One such operational difficulty is the increased complexity of continuous integration in MLOps as opposed to DevOps. The interviewees stated that MLOps necessitates a more intricate and thorough method of continuous integration.

Automating and generalising the deployment process presents another challenge. It is difficult to create a deployment strategy that functions well for several datasets with different statistical properties. Careful planning and research are needed to automate the training, algorithm selection, fine-tuning, and other processes to achieve acceptable results across a variety of datasets.

In MLOps, relying on other teams can provide difficulties. One interviewee stated that their team depends on other teams' deployment pipelines and API hosting providers. Any problems these external dependencies may experience could have a big effect on how the MLOps team operates.

Implementation trade-offs: One of the most important things that the interviewees mentioned is that ML developers often have to make sacrifices while implementing MLOps. When they have to choose between conflicting priorities, like cost and forecast speed, they encounter difficulties.

Businesses need to weigh the expenses of off-the-shelf products against their dependability and benefits for integration. Over-engineering solutions is a common trend in MLOps, and it can lead to wasteful use of time and resources. As a result, it is imperative to continuously optimise and enhance; nevertheless, there should be reasonable limitations to prevent over-engineering, which can lower efficiency.

Business challenges

According to some respondents, a common obstacle to adopting MLOps projects is resistance from enterprises. Getting support and buy-in at the highest level is crucial, particularly because the processes involved are getting more complicated. Explaining the return on investment to businesses can be difficult, though.

However, some companies tend to think of AI as a panacea. This viewpoint is frequently stated by non-technical stakeholders who don't get the context and just use buzzwords. This kind of thinking can make it harder to control expectations and inform the company about the constraints and practical results of MLOps efforts. Regarding a "simple descriptive analytics project," one of the interviewees advises against the organisation searching for AI/ML solutions. Creating a dashboard should be sufficient instead.

Consumers typically have high expectations for the system, counting on complete problem-solving and 100% accuracy. It can be difficult to control these expectations, therefore it's important to inform clients about the dangers and restrictions associated with machine learning models.

Cost and Budget Constraints: Companies frequently undervalue MLOps, which leads to inadequate budget allocation for this crucial role. An respondent said that occasionally, money is spent more on product marketing than on infrastructure improvements. This makes it difficult to update tools or underlying legacy systems since the company can adopt a sunk cost fallacy.

Participant 7 explains that Management and Business may support MLOps until a budget is requested., *"When we started, it was open and free for us to experiment. I didn't feel any restrictions. Later, when we posed the question: could we have a managed Kubeflow environment? And then we did feel the restrictions in terms of budget."*

MLOps vs DevOps

MLOps is a DevOps extension that concentrates on machine learning model deployment. The insertion of components specific to the model deployment is the primary distinction between the two. The difficulty is in using current DevOps procedures while comprehending and adjusting to these novel components. Similar to DevOps, MLOps has challenges in building a strong atmosphere and promoting teamwork. It's challenging to make sure that different parties coordinate their actions, integrate seamlessly, and communicate effectively. However, because machine learning models introduce unfamiliar and sophisticated methodologies, MLOps presents special challenges. Planning and careful consideration are necessary for handling and accounting for these risks.

“MLOps is DevOps that is applied to machine learning. But the challenge is there are a lot of unknown processes that make it a bit more difficult.” – Participant 9

In the next section we discuss the above findings in the context of published literature in more detail.

Discussion

As we start our discussion, we consider the primary characteristics and themes indicated in Fig. 2 as we analyse our findings. We then go over the implications, restrictions, and suggestions for additional study. We also go over how the findings connect to the review of the literature.

Organisational challenges

Organisational context has been widely employed in IS/IT research papers (Gaskin et al., 2018) and technology adoption models (Gangwar et al., 2015). Research shows that many IS/IT implementations fail due to user resistance (Kim & Kankanhalli, 2009), tools, skills or organisational culture (Bunker et al., 2008). The organisational context focuses on descriptive measures, which include, among other things, resource availability and skill in use, firm scope, firm size, slack resources, social influences, culture, structural configurations, and managerial beliefs (Awa et al., 2017).

In this study, organisational challenges consist of four themes, as shown in Fig. 2 of the previous section. The issue of having fewer human resources who know ML and data engineering is a recognised challenge both in the expert interview as well as research papers. To fully harness the advantages of MLOps, businesses face challenges in accessing personnel experienced in artificial intelligence and machine learning (Zhang et al., 2020) and stopping the turnover of IT professionals who have other perceived job opportunities (Joseph et al., 2007).

This lack of expertise makes it difficult for companies to align their MLOps strategy with overall goals (Painoli & Datrika, 2021), and the shortage of fully specialised data engineering talent within the human resources department further exacerbates the situation. The limited educational output and the insufficient expertise level of graduates fail to meet industry expectations, particularly in terms of quality of skills, with a strong emphasis on the engineering aspect (Tamburri, 2020). This confirms the lack of data and ML engineering talent and the long onboarding time interviewees mentioned.

Participants in our interviews explained the time taken to keep the users engaged, as some of them did not perceive the usefulness of MLOps. Employee resistance to technology implementation is recognised in the literature (Lapointe & Rivard, 2005) which may be due to individual issues, organisational issues, system issues, process-related factors (Klaus & Blanton, 2010) or the complex nature of work (Aubert et al., 2008). The application of ML models for people with little knowledge of data science is quite challenging, and hence employees may be resistant to change when it comes to established corporate procedures (Baier et al., 2019; Painoli & Datrika, 2021).

It might be difficult to create machine learning AI systems for knowledge-intensive workplaces due to the need to change how an organisation values innovation and welcomes new ways of doing things. This shift in mindset is crucial for MLOps implementation because it helps businesses overcome reluctance to change and successfully incorporate AI into their operations (Zhang et al., 2020).

Interview participants explained the delays in implementing MLOps due to teams having different priorities, organisations with longer approval chains, and people's conservative mindsets. Baier et al. (2019) talk about how digitalisation, in general, is slower in the healthcare industry as some data is not even available in digital format. Hence the process of data collection to build MLOps becomes very slow. Usually, data scientists and engineers are not part of the same team, which brings challenges of dependencies and wait times. Paleyes et al. (2022) suggest

including both roles in the same team to avoid such a dependency.

Communication and collaboration issues were recognised by multiple interview participants, which are reflected in the generic data science challenge (Cao, 2017) as well as in DevOps adoption challenges (Lassenius et al., 2015). Granlund et al. (2021) briefly discuss the need for seamless communication between teams and collaborations across the organisation, which is rather a difficult challenge like that of DevOps. Interview participants relate this issue to having less knowledge about models or machine learning among team members they must collaborate with. Baier et al. (2019) recognise the challenges in customer communication along with expectation management as customers want transparency in the models, which are complex to explain.

The root cause of the four Organizational challenges (in Fig. 2) can be traced back to a lack of organizational readiness and strategic alignment for MLOps implementation. To address these challenges, organizations should adopt a comprehensive approach that includes investing in talent development through training programs and collaborations with educational institutions (Tamburri, 2020), implementing change management strategies with clear communication of MLOps benefits (Kim & Kankanhalli, 2009), creating cross-functional teams that integrate data scientists and engineers (Paleyes et al., 2022), and establishing platforms for knowledge sharing across the organization (Granlund et al., 2021). Additionally, promoting a culture of innovation and continuous learning is crucial for overcoming resistance to change and fostering better collaboration (Zhang et al., 2020). By focusing on these solutions, organizations can create an environment conducive to MLOps adoption, leading to more successful implementations and better utilization of machine learning technologies (Baier et al., 2019).

Technical challenges

There is a certain degree of complexity in every technology, where *complexity* is defined as the perceived difficulty in learning and implementing a system (Sonnenwald et al., 2001). Implementing MLOps is difficult because of data quality and availability issues, as stated by Paleyes et al. (2022). Participants confirmed that MLOps implementation should begin only after making sure that there is enough data available to create models.

Even the most advanced ML technology cannot be effectively leveraged without an established data infrastructure and the high-quality data it delivers (Shollo et al., 2022). According to the case studies examined by Paleyes et al. (2022), data-related worries are a major impediment to implementing machine learning models. Incomplete, skewed, or false information may lower the quality of machine learning models, and this is a common problem for Machine learning and not specific to MLOps implementation. One of the key issues with ML model deployment is infrastructure. Shollo et al. (2022) suggest that based on the availability of the technical infrastructure and the process maturity, businesses should adapt their machine learning approach. Setting up pertinent data infrastructures is a challenge in addition to deploying infrastructures for running ML models (Baier et al., 2019). Furthermore, as Painoli et al. (2021) and Zhang et al. (2020) point out, tackling these data and model difficulties often involves considerable technical skills, which may be difficult for enterprises to acquire.

The existing literature focuses more on technical challenges such as data availability, data drift, model versioning, scalability, and model monitoring (Baier et al., 2019; Lima et al., 2022; Testi et al., 2022). However, in interviews, it was found to be more about the integration of tools into existing infrastructure, managing data privacy, and not having enough standardisation of MLOps tools. Data science, according to Cao (2017), could be enhanced to incorporate social issues, including privacy, security, and trust. There needs to be more standardisation for the application of MLOps, its tools, and documentation within organisations and in the industry. Since the research is scattered and isolated and ML models, datasets, and feature sets can be produced using numerous tools on a wide range of platforms, lack of standardisation is one of the most

important MLOps challenges within the organisation (Lima et al., 2022; Testi et al., 2022).

The root cause of the technical challenges in MLOps implementation can be attributed to the complexity of integrating advanced machine learning systems with existing data infrastructure and processes. This complexity manifests in issues such as data quality and availability, infrastructure setup, technical skill requirements, and lack of standardization. To address these challenges, organizations should focus on establishing robust data infrastructure before implementing MLOps, investing in technical skill development, and adopting standardized MLOps tools and practices. Additionally, organizations should prioritize data privacy and security measures, implement effective model versioning and monitoring systems, and develop strategies to handle data drift and scalability issues (Paleyes et al. (2022) and Shollo et al. (2022)

Operational challenges

Manually operating ML in a production setting is difficult owing to the various stacks of software and hardware components and their interplay (Ruf et al., 2021). Hence automation through the MLOps pipeline is a must. Practitioners mentioned having difficulty regarding making these pipelines generic and maintaining them while data and models are changing dynamically. An organisation's structure and communication patterns have a direct impact on the design and structure of the systems or software it develops (Conway, 1968). When there are multiple teams involved in MLOps work, this creates borders and dependencies, as mentioned by participants. Also, there is always a trade-off that needs to be made while implementing this pipeline and choosing the right one is a difficult choice.

The root cause of the operational challenges in MLOps implementation can be attributed to the complexity of managing diverse software and hardware components in production environments, the difficulty in creating and maintaining generic pipelines, and the impact of organizational structure on system design. To address these challenges, organizations should focus on automating MLOps pipelines, designing flexible and adaptable systems, improving cross-team communication and collaboration. Additionally, organizations should consider restructuring teams to align with Conway's Law, ensuring that the system architecture reflects the desired communication patterns and reduces dependencies between teams (Conway, 1968).

Business challenges

Creating real-world business value of ML solutions is an important challenge mentioned in the literature. Even though it is not directly related to MLOps, it is a challenge for ML applications. While data science has made it possible to solve many complicated scientific problems, not all scientific answers are equally useful to businesses (Cao, 2017). As few participants mentioned, when businesses try to apply AI/ML solutions for a problem which could be done in traditional software, it does not add any value. It is also noted that developing useful digital services or products based on ML model findings remains difficult (Baier et al., 2019). With MLOps, the process and pipelines get automated, but the value proposition that is delivered to the customer remains unchanged (Shollo et al., 2022). And participants confirmed that it is hard to convince management to invest in MLOps as there is no functionality changes is made.

The root cause of the business challenges in MLOps implementation can be attributed to the difficulty in creating tangible business value from ML solutions and the struggle to justify investments in MLOps to management. To address these challenges, organizations should focus on carefully selecting ML projects that provide clear business value beyond traditional software solutions, developing strategies to translate ML model findings into useful digital services or products, and effectively communicating the long-term benefits of MLOps to management. Additionally, organizations should prioritize projects that demonstrate

measurable improvements in efficiency, accuracy, or customer experience, even if they don't result in immediate functionality changes (Cao, 2017; Shollo et al. 2022).

DevOps vs MLOps

Participants identify the overlap of MLOps with DevOps, which is also mentioned in several other studies (Cardoso Silva et al., 2020; Flaounas, 2017; Jordan & Mitchell, 2015; Kreuzberger et al., 2023; Testi et al., 2022). MLOps, like DevOps, lacks skill and knowledge (Khan et al., 2022) and has challenges in communication and collaboration (Lassenius et al., 2015). For both MLOps and DevOps, organisational culture is a crucial element (Abrahamsson et al., 2016). When it comes to the effectiveness of DevOps, there is no standard quantitative standard (Erich et al., 2017), there are some researchers who have proposed maturity models, but most firms hardly ever measure or have access to information on their software delivery practices (Leite et al., 2020). DevOps implementation in organisations has been happening for over a decade (Lwakatare et al., 2019) yet, challenges in adopting it prevail (Tanzil et al., 2023).

The challenges which come with model scalability, large data, generalising the pipeline, Testing, and hyper-parameter setting are unique to MLOps which is explained in previous sections.

Implications of the research

Our research highlights a gap in the existing literature regarding a unified framework that addresses all aspects of MLOps challenges. This exploratory study was designed to add knowledge about factors that hinder the implementation of MLOps. Since the topic is new, this inductive study was a good choice. Relating the themes to common challenges recognised in the other technological implementations adds depth to the literature review. Literature focuses more on technical challenges such as data availability, data drift, model versioning, scalability, and model monitoring (Baier et al., 2019; Lima et al., 2022; Testi et al., 2022). However, this paper provides insights into the issues related to integration of tools into existing infrastructure, managing data privacy, and not having enough standardization of MLOps tools. It also provides insights into the challenges which are non-technical as well. Lack of knowledge, human resources, budget constraints, and unclear business value are some of those. The typology of challenges we arrive at in Fig. 2 is a Type 1 theory according to Gregor (2006). This is the primary contribution of our paper.

There is a need for more empirical studies and case analyses to validate the theoretical findings and frameworks proposed in the literature. Researchers should focus on gathering data from diverse industries to understand how different sectors face and overcome MLOps challenges.

Our study identifies a lack of focus on tool support and testing within MLOps literature. Future research should explore the development and evaluation of tools that facilitate MLOps processes, emphasizing automation, scalability, and integration capabilities. Given the multidisciplinary nature of MLOps, involving data science, software engineering, and IT operations, research should adopt interdisciplinary approaches to address the complex challenges in MLOps environments effectively.

Implications for practice

The implications for practice derived from this research on MLOps challenges are multifaceted, offering valuable insights for organizations seeking to implement or improve their MLOps processes.

One of the most significant implications is the need for comprehensive organizational change management. As MLOps represents a paradigm shift in how machine learning models are developed, deployed, and maintained, organizations must be prepared to undergo substantial cultural and structural changes. This involves fostering a collaborative

environment where data scientists, software engineers, and IT operations professionals can work seamlessly together. Organizations should focus on breaking down silos between these traditionally separate departments and encourage cross-functional teams that can address the complex, interdisciplinary challenges of MLOps.

Another critical implication is the urgent need for skill development and training programs. This research highlights a significant skills gap in the industry, with many organizations struggling to find professionals who possess the unique blend of skills required for effective MLOps implementation. To address this, companies should invest heavily in upskilling their existing workforce and developing comprehensive training programs that cover not only technical aspects of MLOps but also the necessary soft skills for cross-functional collaboration. Additionally, organizations may need to reevaluate their hiring strategies to attract talent with MLOps expertise or the potential to develop such skills.

Process standardization emerges as another crucial implication for practice. Our research indicates that many organizations face challenges due to inconsistent MLOps implementations across different teams or projects. To overcome this, companies should work towards establishing standardized processes and best practices for model development, deployment, monitoring, and governance. This standardization should be flexible enough to accommodate the specific needs of different projects while ensuring a consistent approach to MLOps across the organization.

Infrastructure investment is another key implication highlighted by our research. Successful MLOps implementation requires robust, scalable infrastructure that can support the unique demands of machine learning workflows. Organizations should carefully evaluate their current IT infrastructure and be prepared to make significant investments in tools and technologies that facilitate automation, scalability, and integration of ML models into production environments. This may involve adopting cloud-based solutions, implementing containerization technologies, or developing custom infrastructure tailored to the organization's specific MLOps needs.

The research also underscores the importance of data management and governance in MLOps practice. Organizations must develop comprehensive strategies for data quality assurance, versioning, and privacy protection. This involves implementing robust data pipelines, establishing clear data governance policies, and ensuring compliance with relevant regulations. Companies should also focus on developing strategies to handle data drift and ensure the long-term validity of their ML models.

By comparing MLOps challenges to DevOps, this paper showcases the challenges which are found in both practices. Practitioners can look at this to get an overview of challenges and try to reduce these challenges on an organisational level. Tool vendors can look at the technical challenges mentioned to provide possible solutions in their latest tools. The findings contribute to models or frameworks concerning technology adoption and implementation challenges. In general, most interviewees stated that they believed my research about this phenomenon is very relevant, and hence a copy of this would guide them further.

Lastly, the implications extend to the realm of ethics and responsible AI. As ML models become more deeply integrated into business operations, organizations must prioritize the development of ethical guidelines and practices for AI development and deployment. This includes addressing issues of model bias, ensuring model interpretability and explainability, and establishing mechanisms for ongoing monitoring and auditing of ML models in production.

Limitations and future research

Qualitative research can provide valuable insights and an in-depth understanding of the complex phenomenon (Gioia et al., 2013), and this research methodology was chosen carefully, and steps were taken to mitigate some limitations, as mentioned in the quality criteria section.

However, there are still more limitations to be addressed.

Firstly, due to the newness of the MLOps concept, there is a lack of academic research articles focusing on the challenges associated with MLOps implementation. This led the review to be shallow and failed to provide a comprehensive overview of the current state of knowledge in the field. Secondly, this study included cross-industry interviewees, which allowed the results to be combined across industries. Given that industry type influences technology adoption decision-making (Tornatzky & Klein, 1982), future research could focus on MLOps challenges within a specific industry. As a third point, although we achieved saturation with our codes, the relatively small sample size of twelve respondents may make this study less generalisable. A quantitative approach involving a larger number of respondents could enhance the robustness of the study's findings and make them more generalisable. Finally, as we have conducted the semi-structured interviews at a single point in time, in this research we are limited to insights from the historical experience of the interviewees. We therefore cannot determine the evolution of MLOps in organizations. Future work could overcome this limitation by utilizing longitudinal qualitative interviews (Hermanowicz, 2013) and therefore arrive at challenges that change during the implementation.

The academic community has concentrated on machine learning model development and benchmarking but not on operating complicated machine learning systems in real-world applications. In the real world, we find data scientists and ML engineers are struggling to adopt MLOps. So, in this research, we interviewed ML practitioners from different industries and consolidated the results. Firstly, future studies may perform a comparative analysis across different industries or regions. Secondly, they can conduct a broader quantitative study (for example, based on a survey). This would make it possible to derive quantifiable conclusions, identify the scope and gravity of each difficulty, and pinpoint the associated area of need for further study. Thirdly, all semi-structured interviews were conducted, keeping the focus only on MLOps challenges. Future studies can extend our typology (in Fig. 2). They could also find the resolutions to solve these challenges. Lastly, future studies can adopt and further explore the additional findings "Prerequisites to implementing MLOps" and "Benefits of MLOps implementation" mentioned earlier.

Conclusion

Applications of machine learning (ML) have sparked numerous new technological developments in both industry and academia in recent years. But there are still a lot of unanswered problems regarding the practical applications of machine learning (Shollo et al., 2022), and a lot of ML projects don't make it to production (Shankar et al., 2022). Although many organisations have not adopted it, MLOps, an extension of DevOps, appears to be the answer for releasing ML-based software to production.

In this paper, we aim to answer the research question "What are the challenges organizations face while implementing MLOps?" by conducting a structured literature review and a semi-structured interview with 12 ML experts working in different organizations. We have identified a typology of challenges (shown in Fig. 2), which is a Type 1 theory according to Gregor (2006). This typology is the primary contribution of our paper, and future studies can extend it while finding resolutions to these challenges.

Our results show that there are challenges associated with using MLOps in organisations. Finding individuals with the data engineering and machine learning engineering expertise needed to build MLOps pipelines is one of the main obstacles. It is difficult to keep talent in the market due to high demand. The reluctance of certain team members to adopt new tools and technology could potentially impede the effective use of MLOps pipelines. As a result, companies need to make sure that the teams working on ML applications understand and make use of MLOps.

Organisational structure and decision-making procedures also aid in MLOps. Long approval procedures could restrict engineers' independence and cause delays in implementation. One of the ongoing challenges of DevOps is that of MLOps teams having various aims and limited capabilities, which leads to problems with collaboration and communication. The management of massive volumes of data within pipelines, scaling the ML infrastructure as data and models grow, the scarcity of clean, production data, integration problems with current IT systems, and challenges scaling and managing the underlying infrastructure are a few of the technical challenges that have been identified.

The intricacy of continuous integration, reliance on external teams (a

problem acknowledged by DevOps), and the necessity of balancing priorities and optimisation efforts are examples of operational problems. Getting financial allocations and communicating to management the commercial value of MLOps are the final business-related issues identified in this research.

CRediT authorship contribution statement

Chintan Amrit: Writing – review & editing, Writing – original draft, Conceptualization. **Ashwini Kolar Narayanappa:** Writing – original draft, Conceptualization.

Appendix A

Here we present the inclusion and exclusion criteria for the structured literature review.

Table 3
Inclusion criteria.

Criteria	Inclusion Criteria
IC 1	Research articles and Conference Papers
IC 2	Studies that address Machine Learning Operations (MLOps) in general
IC 3	Studies that identify challenges associated with MLOps
IC 4	Studies that included AIOps challenges

Table 4
Exclusion criteria.

Criteria	Exclusion Criteria
EC 1	The study that was not published in English
EC 2	Studies that talk about the building and application of ML models
EC 3	Studies that do not allow access to its content
EC 4	Papers that did not have relevance to the research question

Appendix B

Here we present the protocol for the systematic literature review.

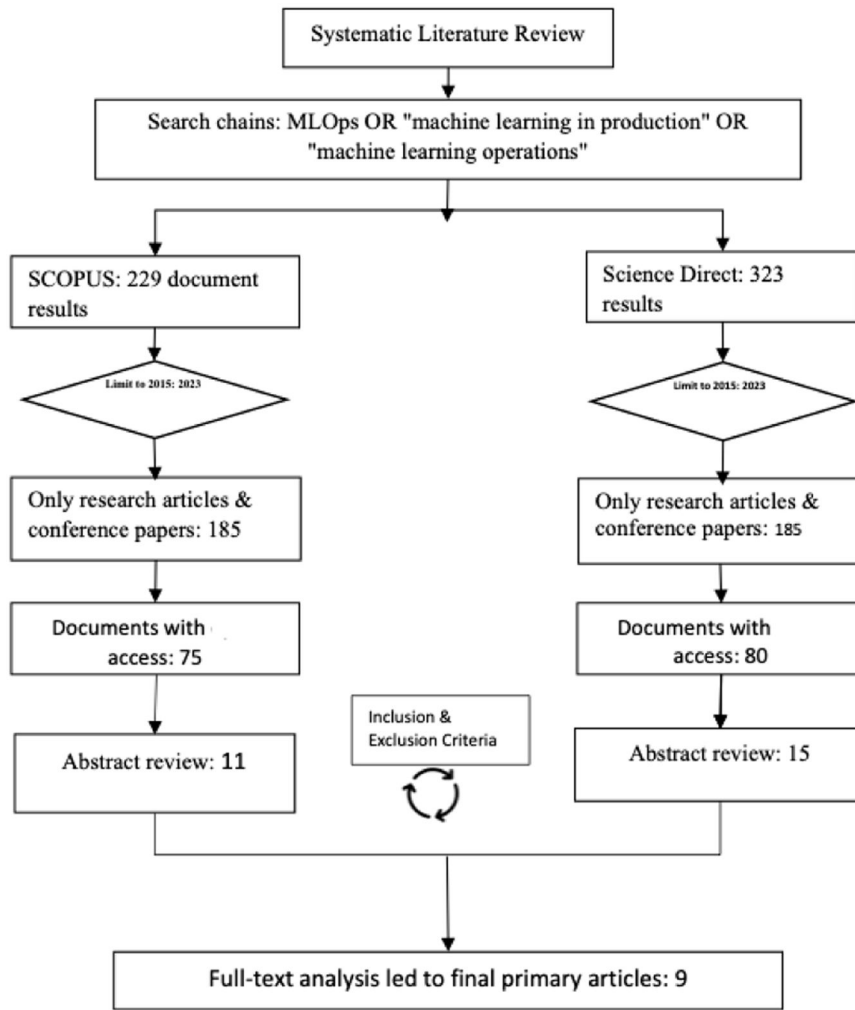


Fig. 3. Systematic literature review selection protocol.

Appendix C

Table 5
Participant profile.

Participant	Participant Job title	ML Experience (Years)	Country	Domain
1	MLOps Lead Engineer	9	Netherlands	Publishing
2	ML and DS Manager	6	India	Tech
3	VP of Data Science	9	India	EdTech
4	AI and ML Manager	4	Norway	Software provider
5	ML team Manager	7	Denmark	Software provider
6	Data Scientist	4	Netherlands	Software provider
7	ML Engineer	4,5	Netherlands	Consulting
8	Data Scientist	9	Netherlands	Bank
9	Machine learning consultant	5	Netherlands	Consulting
10	Product Owner for ML	3	Netherlands	Insurance
11	MLOps Engineer	2	Netherlands	Insurance
12	AI architect	4	Netherlands	Startup

Appendix D




Fig. 4. Interviews and themes ratio through a Sankey diagram.

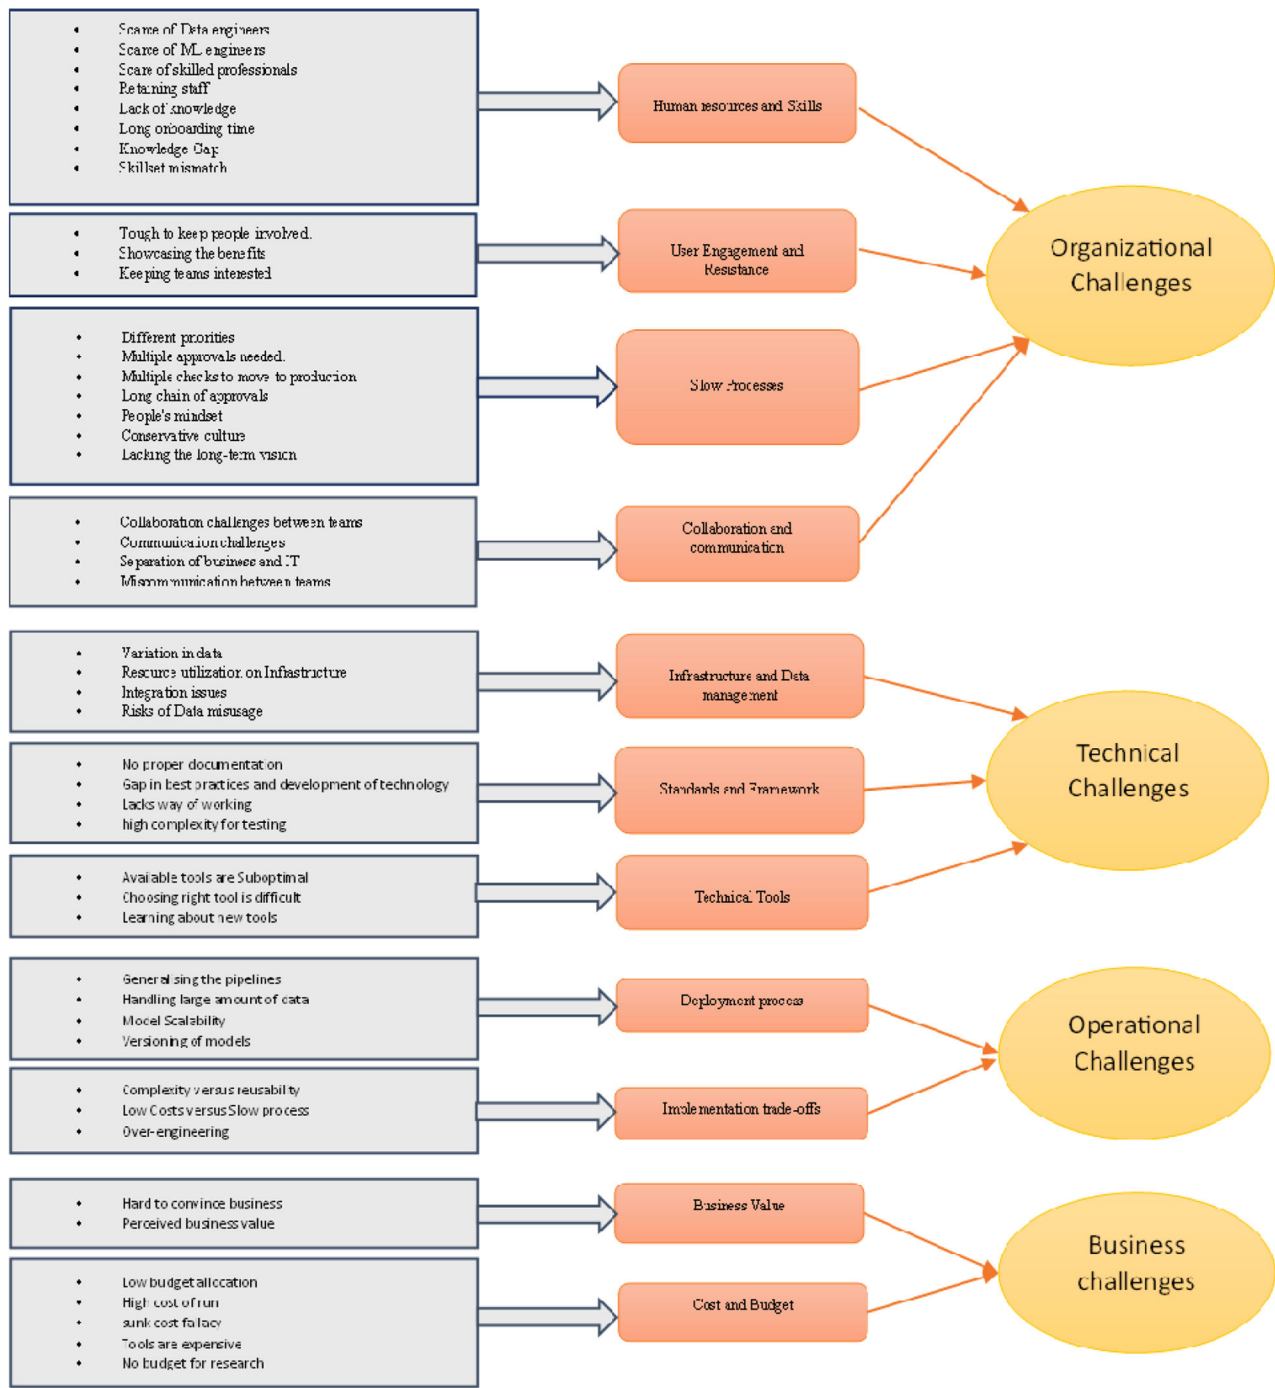


Fig. 5. The Coding Schema (grey boxes represent first level codes, the orange boxes represent the second level codes and the yellow ovals on the right represent the third level codes).

Quotations	First Order Concepts	Second Order Themes	Aggregate Dimensions
And the number of available data engineers in the market is far lower than the demand. Data engineers are almost always understaffed and overworked.	Scarcity of data engineers	Human Resources and Skillset	
We have a real scarcity of good machine learning engineers and DevOps engineers.	Scarcity of ML engineers		
When I talk about the people front, I would say the talent shortage is there for skilled professionals on the entire ML setup because we need that interdisciplinary skill set of math or statistics with coding, with business.	Scarcity of skilled professionals		
The other situation is we have found lots of resource problems and employee turnover. What we have planned didn't work because most of our intelligent engineers left. This market is volatile because there is high demand.	Retaining staff		
What I realized is that a lot of people, especially from risk, but maybe also some dev or OPS engineers don't really understand how machine learning works.	Lack of knowledge	User Engagement and Resistance	Organisational Challenges
This machine-learning world is new. So some level of skill set mismatch will be there in typical software you will find these skills, but on the ML side it will be lesser.	Skillset mismatch		
I would say the biggest challenge was keeping the teams that were used with and for involved.	Tough to keep people motivated		
So we spent a lot of time involving different teams and making them feel engaged in the process and because our biggest challenge was not making the model or making the pipeline, but actually getting people to just think about some of it and making people see the benefits.	Showcasing the benefits		
Because along with our pipeline came all the new learning to do. It could be a challenge for some people in the teams to adapt to these new tools. Keeping them engaged, interested, and enthusiastic about what we wanted to achieve and what was necessary for them. To help us achieve it, I think I would say our biggest challenge.	Keeping teams interested	Slow Processes	
Very often it's not one team that manages this whole thing, and that means that if you for example have a data engineer that works with you from a different team, they're also working on seven other different projects. And they need to manage all of those, and they all have expectations. And that means that they always don't always have time to help you out with whatever you need at that time and that can cause friction.	Different priorities		
With large enterprises, everything goes very slow because there's a lot of people that need to sign off on picking a tool. So, if you want to choose something it can take a lot of time.	Multiple approvals needed		
By far the biggest difference is all the checks. We have a model validation part where you must describe what your model does. There's a different team that looks at if your model is biased. And in my previous experience with smaller companies, the steps were not that much. The idea is cautious.	Multiple checks to move to production		
We see within ING that risk is always saying you cannot do it. It's production data. I think this took us months to get an OK and they allowed us to have that training data on that machine.	Long chain of approvals	Collaboration & Communication	
But in this case, they were a lot more hesitant at first as the mentality is if it isn't broken, don't fix it.	People's mindset		
It's difficult to collaborate in the sense that we say, we need to do exploratory data analysis I don't know how long it will take. I don't know what data quality is. I'm not sure about the next steps.	Collaboration challenges between teams		
We need to check the data, we need to run the test, deploy it, and update the dashboards. So sometimes a little bit difficult to communicate the effort and time needed.	Communication challenges		
You get a lot of miscommunications between those teams and also a lack of accountability or responsibility. Because it's a lot easier to say 'Yeah, there's a problem. That's not my problem. That's their problem. And you just start going in circles.	Miscommunication between teams	Infrastructure & data management	
I think the big challenge we have remaining is the enormous variation in data.	Variation in data		
And then the second one is to ensure the right infrastructure management is there. Deploying and managing these ML models can become, very challenging if you're a small setup because these licensed versions would cost a lot, and effectively managing becomes important.	Resource utilisation on Infrastructure		
I had a really hard time doing ML, OPS part. They provide subset of the things you would like and the rest do not. And this makes it more difficult to attach to your working environment. I would say it's quite a basic thing and this is also quite an expensive and professional tool. But even these things are sometimes difficult. So basically, there are integration challenges with all these tools.	Integration issues		
The first challenge is that we are basically trying to build a solution that depends on open-source projects like Kubernetes and others.	No proper documentation	Standards and Framework	Technical Challenges
These projects are normally nurtured by Google, and they are not properly documented.	Lacks way of working		
We have two different teams where we have real dedicated OPS people that help us with some OPS work. But in practice, we're kind of in a gray area right now where they don't have that much experience. So that is something we do within our own team. And to be honest that's still something we must negotiate a little bit within ING.	High complexity for testing		
If you have like regular software, I'm used to writing unit tests. For machine learning, these tests are a little bit more complex because it's not deterministic.	Available tools are suboptimal		
There was Sage Maker, which is ML/OPS framework or what you want to call it. I didn't want to use it because I thought it was not built for data scientists who wanted to get something out in production. It was quite slow there are a variety of tools, multiple tools which the industry has or is within the company. Also, there are many tools so there's no standard every team uses, and it keeps changing. Us then there is an environment with the packages, and related ones. You know we use some things which are not suitable for the ML, OPS. So things like that also will be a challenge.	Choosing right tool is difficult	Technical Tools	
What really is required is at the beginning itself. There has to be some agreement between the ML, OPS engineer and the data scientist to build the models so that this friction will be lesser.	Generalising the pipelines		
ML pipelines, which generates this models. Fitting all this into one picture I think was probably one of the most challenging things.	Handling large amount of data		
How do you handle large amounts of data in ML, OPS environment right? This means the models require high-quality clean process data to train and produce. And that's when we talk about the whole garbage in garbage out analogy: what if I do not have the cleaner data then we cannot expect high output of it.	Model Scalability		
As the size of the data increases and the model it increases, it can become very difficult to scale up the ML infra.	Versioning of models	Deployment process	Operational challenges
Reasoning things or maintaining the integrity of that model, so keeping a track of these model versions with their code and the data and the parameters is very challenging.	Complexity versus reusability		
Another trade off we have a complexity versus reusability, we want to make it as easy and effortless as possible to deploy new machine-learning services, that, we need to generalize a lot of the back-end code and a lot of the infrastructure and so on. So it was super easy to just copy and the plan is a new file, so deploy machine learning to solve any problem kind of and that, but that also creates a lot of complexity because you have to have this general perspective in their background the whole time. So that's the trade off.	Low Costs versus Slow process		
Not every company has that amount of money. You always have to make a trade off if you go to open source, the cost will be low, but you spend more time developing.	Hard to convince business		
I mean the business just wants to see the results to be honest. And if they are already getting the results prior to the ML/OPS, then it's hard to say that you really need ML/OPS too to improve on the results. It's hard to sell without putting a number on it.	Resistance from business	Cost & Budget Constraints	Business challenges
Another problem in business guys are only looking for latency and accuracy. So sometimes when the quantized model's accuracy and latency are varied, it's difficult to explain to the Business owners.	Perceived business value		
There are multiple products like SageMaker GCP which provide generic solutions. Sometimes our team can adopt it, but the costing will be so much higher.	High cost of run		
Then for the tools, there were a lot of people adamant about using the tools that they already had in-house. It's kind of a sunk cost fallacy. So, we pay a lot of money to get this Azure service for example. So we have to use it even though there were other, better suited tools available.	Sunk cost fallacy		
That's very deliberate because the cost of research phase would be too high.	No budget for research		

Fig. 6. An impression of the Codebook.

References

- Abrahamsson, P., Jedlitschka, A., Nguyen Duc, A., Felderer, M., Amasaki, S., & Mikkonen, T. (2016). DevOps adoption benefits and challenges in practice: A case study. *Product-Focused software process improvement* (pp. 590–597). Springer International Publishing AG. https://doi.org/10.1007/978-3-319-49094-6_44
- Alla, S., & Adari, S. K. (2020). What is MLOps?. *Beginning MLOps with MLFlow* (pp. 79–124). Apress L. P.. https://doi.org/10.1007/978-1-4842-6549-9_3
- Aubert, B. A., Barki, H., Patry, M., & Roy, V. (2008). A multi-level, multi-theory perspective of information technology implementation. *Information Systems Journal*, 18(1), 45–72. <https://doi.org/10.1111/j.1365-2575.2007.00279.x>
- Awa, H. O., Ukoha, O., & Igwe, S. R. (2017). Revisiting technology-organization-environment (T-O-E) theory for enriched applicability. *The Bottom Line*, 30(1), 2–22. <https://doi.org/10.1108/BL-12-2016-0044>
- Baier, L., Jöhren, F., & Seebacher, S. (2019). Challenges in the deployment and operation of machine learning in practice. In , I. *ECIS*. <https://doi.org/10.5445/ir/1000095028>
- Bansal, P., Smith, W. K., & Vaara, E. (2018). New ways of seeing through qualitative research. *Academy of Management Journal*, 61(4). <https://hal.archives-ouvertes.fr/hal-02312197>
- Beitin, B. (2012). *Interview and sampling: How many and whom The SAGE handbook of interview research: The complexity of the craft*, 243–253.
- Bunker, D., Kautz, K., & Anhtuan, A. (2008). An exploration of information systems adoption: Tools and skills as cultural artefacts - the case of a management information system. *Journal of Information Technology*, 23(2), 71–78. <https://doi.org/10.1057/palgrave.jit.2000134>
- Calefato, F., Lanubile, F., & Quaranta, L. (2022). *A preliminary investigation of MLOps practices in github*. Ithaca: Cornell University Library. <https://doi.org/10.1145/3544902.3546636>. arXiv.org.Retrieved from Publicly Available Content Database <https://search.proquest.com/docview/2718002010>.
- Cao, L. (2017). Data science: Challenges and directions. *Communications of the ACM*, 60(8), 59–68. <https://doi.org/10.1145/3015456>
- Cardoso Silva, L., Rezende Zagatti, F., Silva Sette, B., Nildaimon dos Santos Silva, L., Lucredio, D., Furtado Silva, D., & de Medeiros Caseli, H. (2020). Benchmarking machine learning solutions in production. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)* (pp. 626–633). IEEE. <https://ieeexplore.ieee.org/document/9356298>
- Choudhary, F., & Krensky, P. (2022). *Hype cycle for data science and machine learning 2022*. June 29. Gartner <https://www.gartner.com/>.
- Conway, M. E. (1968). How do committees invent. *Datamation*, 14(4), 28–31.
- Ebert, C., Gallardo, G., Hernantes, J., & Serrano, N. (2016). DevOps. *IEEE Software*, 33(3), 94–100. <https://doi.org/10.1109/MS.2016.68>
- Diaz-De-Arcaya, J., Torre-Bastida, A. I., Zárate, G., Minón, R., & Almeida, A. (2023). A joint study of the challenges, opportunities, and roadmap of MLOps and AIOps: A systematic survey. *ACM Computing Surveys*, 56(4), 1–30.
- Erich, F. M. A., Amrit, C., & Daneva, M. (2017). A qualitative study of DevOps usage in practice. *Journal of Software: Evolution and Process*, 29(6). <https://doi.org/10.1002/smr.1885>. n/a.
- Flaounas, I. (2017). Beyond the technical challenges for deploying machine learning solutions in a software company. [10.48550/arxiv.1708.02363](https://arxiv.org/abs/1708.02363).
- Gangwar, H., Date, H., & Ramaswamy, R. (2015). Understanding determinants of cloud computing adoption using an integrated TAM-TOE model. *Journal of Enterprise Information Management*, 28(1), 107–130. <https://doi.org/10.1108/JEIM-08-2013-0065>
- Garg, S., Pundir, P., Rathee, G., Gupta, P. K., Garg, S., & Ahlawat, S. (2022). On continuous integration /continuous delivery for automated deployment of machine learning models using MLOps. In *2021 IEEE fourth international conference on artificial intelligence and knowledge engineering (AIKE)* (pp. 25–28). IEEE.
- Gaskin, J., Berente, N., Lyytinen, K., & Rose, G. (2018). Innovation among different classes of software development organizations. *Information Systems Journal*, 28(5), 849–878. <https://doi.org/10.1111/isj.12171>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research. *Organizational Research Methods*, 16(1), 15–31. <https://doi.org/10.1177/1094428112452151>
- Granlund, T., Kopponen, A., Stirbu, V., Myllyaho, L., & Mikkonen, T. (2021). MLOps challenges in multi-organization setup: experiences from two real-world cases. [10.48550/arxiv.2103.08937](https://arxiv.org/abs/2103.08937).
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 611–642.
- Hermanowicz, J. C. (2013). The longitudinal qualitative interview. *Qualitative Sociology*, 36, 189–208.
- Hui Kang, Le, M., & Shu, Tao (2016). Container and microservice driven design for cloud infrastructure DevOps. In *2016 IEEE International Conference on Cloud Engineering (IC2E)* (pp. 202–211). IEEE. <https://ieeexplore.ieee.org/document/7484185>.
- John, M. M., Olsson, H. H., & Bosch, J. (2021). Towards MLOps: A framework and maturity model. In *2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (pp. 1–8). IEEE. <https://ieeexplore.ieee.org/document/9582569>.
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science (American Association for the Advancement of Science)*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Joseph, D., Ng, K., Koh, C., & Ang, S. (2007). Turnover of information technology professionals: A narrative review, meta-analytic structural equation modeling, and model development. *MIS Quarterly*, 31(3), 547–577. <https://doi.org/10.2307/25148807>

- Khan, M. S., Khan, A. W., Khan, F., Khan, M. A., & Whangbo, T. K. (2022). Critical challenges to adopt DevOps culture in software organizations: A systematic review. *IEEE Access*, 10, 14339–14349. <https://doi.org/10.1109/ACCESS.2022.3145970>
- Kim, H., & Kankanhalli, A. (2009). Investigating user resistance to information systems implementation: A Status Quo Bias perspective. *MIS Quarterly*, 33(3), 567–582. <https://doi.org/10.2307/20650309>
- Kitchenham, B. A. (2004). *Procedures for performing systematic reviews*, 33. Keele, UK: Keele University, 2004.
- Klaus, T., & Blanton, J. E. (2010). User resistance determinants and the psychological contract in enterprise system implementations. *European Journal of Information Systems*, 19(6), 625–636. <https://doi.org/10.1057/ejis.2010.39>
- Kreuzberger, D., Kühl, N., & Hirschl, S. (2023). Machine Learning Operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11, 31866–31879. <https://doi.org/10.1109/ACCESS.2023.3262138>
- Krey, M., Kabbout, A., Osmani, L., & Saliji, A. (2022). DevOps adoption: Challenges & barriers. In *Paper presented at the 55th Hawaii International Conference on System Sciences (HICSS), Virtual, 3-7 January 2022* (pp. 7297–7309).
- Lapointe, L., & Rivard, S. (2005). A multilevel model of resistance to information technology implementation. *MIS Quarterly*, 29(3), 461–491. <https://doi.org/10.2307/25148692>
- Lassenius, C., Dingsoyr, T., & Paasivaara, M. (2015). DevOps: A definition and perceived adoption impediments. *Agile processes in software engineering and extreme programming* (pp. 166–177). Springer International Publishing AG. https://doi.org/10.1007/978-3-319-18612-2_14
- Leite, L., Rocha, C., Kon, F., Milojicic, D., & Meirelles, P. (2020). A survey of DevOps concepts and challenges. *ACM Computing Surveys*, 52(6), 1–35. <https://doi.org/10.1145/3359981>
- Lima, A., Monteiro, L., & Furtado, A. (2022). MLOps: Practices, maturity models, roles, tools, and challenges – A systematic literature review. In *Paper presented at the Proceedings of the 24th International Conference on Enterprise Information Systems - Volume 2: ICEIS* (pp. 308–320). <https://doi.org/10.5220/0010997300003179>
- Liu, Y., Ling, Z., Huo, B., Wang, B., Chen, T., & Mouine, E. (2020). Building a platform for machine learning operations from open source frameworks. *IFAC PapersOnLine*, 53(5), 704–709. <https://doi.org/10.1016/j.ifacol.2021.04.161>
- Lwakatere, L. E., Crnkovic, I., & Bosch, J. (2020). DevOps for AI - Challenges in development of AI-enabled applications. In *2020 international conference on software, telecommunications and computer networks (SoftCOM)* (pp. 1–6). IEEE. <https://ieeexplore.ieee.org/document/9238323>
- Lwakatere, L. E., Kilamo, T., Karvonen, T., Sauvola, T., Heikkilä, V., Ikonen, J., Kuvaja, P., Mikkonen, T., Oivo, M., & Lassenius, C. (2019). DevOps in practice: A multiple case study of five companies. *Information and Software Technology*, 114, 217–230. <https://doi.org/10.1016/j.infsof.2019.06.010>
- Macarthy, R. W., & Bass, J. M. (2020). An empirical taxonomy of DevOps in practice. In *Paper presented at the 2020 46th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (pp. 221–228).
- Makinen, S., Skogstrom, 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)* (pp. 109–112). <https://ieeexplore.ieee.org/document/9474355>
- Matsui, B. M. A., & Goya, D. H. (2022). MLOps: Five steps to guide its effective implementation. In *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI* (pp. 33–34). <https://ieeexplore.ieee.org/document/9796456>
- Maya, V., & Felipe, A. (2021). *The State of MLOps* <http://hdl.handle.net/1992/51495>.
- Mucha, T., Ma, S., & Abhari, K. (2022). Beyond MLOps: The lifecycle of machine learning-based solutions. In *AMCIS 2022 Proceedings*. <https://amcis2022.aisconference.org/>
- Painoli, G., & Datrika, V. (2021). Artificial intelligence in business-benefits and challenges. *Turkish Online Journal of Qualitative Inquiry*, 12, 1377–1388.
- Paleyes, A., Urma, R., & Lawrence, N. D. (2022). Challenges in deploying machine learning: A survey of case studies. *ACM Computing Surveys*, 55(6), 1–29. <https://doi.org/10.1145/3533378>
- Qu, S. Q., & Dumay, J. (2011). The qualitative research interview. *Qualitative Research in Accounting and Management*, 8(3), 238–264. <https://doi.org/10.1108/11766091111162070>
- Ruf, P., Madan, M., Reich, C., & Ould-Abdeslam, D. (2021). Demystifying MLOps and presenting a recipe for the selection of open-source tools. *Applied Sciences*, 11(19), 8861. <https://doi.org/10.3390/app11198861>
- Rzig, D. E., Hassan, F., & Kessentini, M. (2022). An empirical study on ML DevOps adoption trends, efforts, and benefits analysis. *Information and Software Technology*, 152, Article 107037. <https://doi.org/10.1016/j.infsof.2022.107037>
- Schröder, T., & Schulz, M. (2022). Monitoring machine learning models: A categorization of challenges and methods. *Data Science and Management*, 5(3), 105–116. <https://doi.org/10.1016/j.dsm.2022.07.004>
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J., & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Paper presented at the Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2* (pp. 2503–2511).
- Serban, A., Blom, K. v. d., Hoos, H. H., & Visser, J. M. W. (2020). Adoption and effects of software engineering best practices in machine learning. In *Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)* (pp. 1–12). <https://hdl.handle.net/1887/3307600>
- Shankar, S., Garcia, R., Hellerstein, J. M., & Parameswaran, A. G. (2022). Operationalizing machine learning: An interview study. *arXiv preprint 10.48550/arxiv.2209.09125*.
- Shollo, A., Hopf, K., Thiess, T., & Müller, O. (2022). Shifting ML value creation mechanisms: A process model of ML value creation. *The Journal of Strategic Information Systems*, 31(3), Article 101734. <https://doi.org/10.1016/j.jsis.2022.101734>
- Sonnenwald, D. H., Maglaughlin, K. L., & Whitton, M. C. (2001). Using innovation diffusion theory to guide collaboration technology evaluation: Work in progress. In *Proceedings Tenth IEEE International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises. WET ICE* (pp. 114–119). IEEE. <https://ieeexplore.ieee.org/document/953399>
- Subramanya, R., Sierla, S., & Vyatkin, V. (2022). From DevOps to MLOps: Overview and application to electricity market forecasting. *Applied Sciences*, 12(19), 9851. <https://doi.org/10.3390/app12199851>
- Symeonidis, G., Nerantzis, E., Kazakis, A., & Papakostas, G. A. (2022). MLOps - Definitions, tools and challenges. In *2022 IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC)* (pp. 0453–0460). IEEE.
- Mboweni, T., Masombuka, T., & Dongmo, C. (2022). A systematic review of machine learning DevOps. In *Paper presented at the - 2022 International Conference on Electrical, Computer and Energy Technologies (ICECET)* (pp. 1–6). <https://doi.org/10.1109/ICECET55527.2022.9872968>
- Tamburri, D. (2020). A. Sustainable MLOps: Trends and challenges. In *2020 22nd international symposium on symbolic and numeric algorithms for scientific computing (SYNAS)* (pp. 17–23). IEEE. <https://ieeexplore.ieee.org/document/9356947>
- Tanzil, M. H., Sarker, M., Uddin, G., & Iqbal, A. (2023). A mixed method study of DevOps challenges. *Information and Software Technology*, 161, Article 107244. <https://doi.org/10.1016/j.infsof.2023.107244>
- Testi, M., Ballabio, M., Frontoni, E., Iannello, G., Moccia, S., Soda, P., & Vessio, G. (2022). MLOps: A taxonomy and a methodology. *IEEE Access*, 10, 63606–63618. <https://doi.org/10.1109/ACCESS.2022.3181730>
- Tornatzky, L. G., & Klein, K. J. (1982). Innovation characteristics and innovation adoption-implementation: A meta-analysis of findings. *IEEE Transactions on Engineering Management*, EM-29(1), 28–45. <https://doi.org/10.1109/TEM.1982.6447463>
- VB Staff. (2019). *Why do 87% of data science projects never make it into production?*
- Wohlin, C. (2014). Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th international conference on evaluation and assessment in software engineering* (pp. 1–10).
- Zhang, Z., Nandhakumar, J., Hummel, J., & Waardenburg, L. (2020). Addressing the key challenges of developing machine learning AI systems for knowledge-intensive work. *MIS Quarterly Executive*, 19(4), 221–238. <https://doi.org/10.17705/2msqe.00035>