

A Systematic Review of MLOps Tools: Tool Adoption, Lifecycle Coverage, and Critical Insights

Zakkarija Micallef

Vrije Universiteit Amsterdam, The Netherlands

z.micallef@student.vu.nl

Abstract

Machine learning operations (MLOps) has become increasingly critical as more organisations move machine learning models into production. However, the growing landscape of MLOps solutions has introduced complexity for practitioners trying to select appropriate tools. To investigate how and why these tools are adopted in practice, this paper conducts a systematic review of academic literature focused on MLOps tools. The analysis maps tools to MLOps lifecycle stages to reveal their function, scope, and the challenges they are designed to address. It identifies usage trends and synthesises reported benefits and limitations based on practical implementations. The most commonly used components, according to the findings, are orchestration frameworks, data versioning, experiment tracking, and managed cloud platforms. No single tool covers the entire lifecycle, so researchers often combine multiple tools to build complete pipelines. This highlights the importance of interoperability across MLOps tools in real-world MLOps pipelines.

Keywords

Systematic Literature Review, MLOps, tool taxonomy, lifecycle mapping, tool adoption, MLflow, Kubeflow

1 Introduction

In recent years, AI has experienced a dramatic surge in popularity. As more companies deploy AI solutions, they quickly discover that specialised infrastructure is essential to support these new capabilities. However, many AI engineers and data scientists lack the software-deployment expertise of operations teams, especially when it comes to MLOps tools [14].

Traditionally, an operations team is in charge of tasks such as deployment and ongoing monitoring [16]. To reduce software time-to-value and create stronger collaboration between development and operations, software companies frequently use DevOps. DevOps is described as a culture that emphasises continuous collaboration throughout the software lifecycle. It involves practices like Continuous Integration, which involves frequent code merges, and Continuous Deployment, which automates the release process to ensure software remains deployable.

MLOps refers to the entire lifecycle of the machine-learning process, bridging the gap between data, development, and operations [22]. MLOps extends beyond applying DevOps principles to machine learning (ML). It involves continuous integration and continuous deployment automation for ML pipelines, orchestration of ML workflows, versioning of data, models, and code to ensure reproducibility, continuous training to keep models up to date, metadata tracking for experiment auditability, and continuous evaluation and performance monitoring.

Despite the growing interest in MLOps, existing reviews often remain high-level, focusing primarily on listing tools, comparing surface-level features, or distinguishing between open-source and proprietary solutions. However, there is limited synthesis of how these tools are actually used in practice [37]. Most studies fail to capture the practical experiences of teams deploying real-world ML systems. This gap in the literature limits the ability of researchers and practitioners to make informed choices based on implementation outcomes rather than tool specifications alone. In this paper, we explore the range of MLOps tools referenced in academic literature and examine their real-world applications. While other papers integrate general-purpose DevOps tools such as Docker, Kubernetes, Git, and Jenkins with ML-specific tools, our SLR (Systematic Literature Review) focuses exclusively on MLOps native tools that are developed specifically for machine learning workloads, excluding general-purpose tools. We explore how practitioners use these tools in their pipelines, why they select specific MLOps solutions, and the advantages and drawbacks they face in real-world scenarios. Through this analysis, we aim to provide a clear view of the current MLOps landscape.

2 Related Work

MLOps remains a vague umbrella term, with its scope and implications still ambiguous for both researchers and practitioners. To bring clarity, Kreuzberger *et al.* [22] conducted mixed-method research, where they combined a literature review, a tool review, and expert interviews to create a comprehensive description of MLOps that we adopt as our definition of MLOps in Section 1. Their work has become a *de facto* standard, widely cited by other papers [10] [45] [48]. While Kreuzberger *et al.* [22] identified principles, technical components, and roles of MLOps, we focus solely on the technical elements that recur across the literature: CI/CD pipelines for training and deployment, workflow orchestration, feature-store systems, model-training infrastructure, model registries, metadata stores, serving components, and monitoring tools. A synthesis of prior surveys shows that these components consistently form the backbone of MLOps solutions [37] [45] [48] [10] [32]:

- Data engineering
- Version control
- Hyperparameter tuning and experiment tracking
- CI/CD pipelines for training and deployment
- Workflow orchestration
- Model deployment/serving
- Automated testing and validation
- Continuous performance monitoring

Varon Maya's [26] argues that bringing DevOps principles into ML pushes organisations toward pipeline-like architectures: NVIDIA [27], Facebook [19], Spotify [1] and Google [18] each describe a

flow that runs from data collection and feature engineering through training, validation and model serving, occasionally adding feedback loops that trigger continuous retraining. This structure clarifies ownership, lets every stage use specialised tooling and aligns with CI/CD automation. However, higher-level concerns such as process orchestration and configuration management are largely unaddressed with this pipeline architecture, creating challenges for reproducibility and maintainability in complex ML systems. Our systematic literature review (SLR) adopts the same pipeline perspective, evaluating MLOps tools stage by stage.

In modern software development environments, multiple languages and libraries are combined. To determine which languages, frameworks, and runtimes each tool supports, Wazir *et al.* [45] looked at 22 open-source catalogues and research publications. Given the frequent combination of Python, Java, R, and other specialised libraries, their findings highlight the need for MLOps platforms to continue being language agnostic. Building on this, Ruf *et al.* [38] conducted interviews and real-world trials to produce a detailed feature matrix of API bindings and integration hooks offered by each candidate tool. They show that lacking support for languages like R or Java can hinder collaboration between operations and data science teams. Consequently, they recommend conducting several iterative selection rounds before beginning work, ensuring all departments agree on a unified toolchain.

In Recupito *et al.*'s review [37], thirteen prominent MLOps platforms were mapped to the stage of the MLOps pipeline they support: data management, model training, CI/CD, monitoring, and so on. Their comparison shows that no single product spans the entire lifecycle, so practitioners routinely assemble multi-tool pipelines. Multiple studies [37] [38] [45] highlight the importance of evaluating not just each tool's individual strengths but also how well they interoperate with other tools and the dependencies they introduce. Interestingly, nearly 50% of the sources in Recupito *et al.* [37] review are blog posts, underscoring how much MLOps expertise circulates through informal channels rather than peer-reviewed papers. While several recent papers embrace a multivocal approach that blends academic and practitioner sources, our decision to focus solely on academic papers means that some practitioner findings may be missed.

Our SLR compares tools by the benefits and limitations that paper authors explicitly report from their practical experience, in contrast to the majority of previous reviews that compare tools based on claimed features listed on the tool developer's websites or their official blogs. This method provides a more accurate understanding of the practical experience and shortcomings of MLOps tools.

3 Study Design

3.1 Research Goal

This study presents a clear landscape of MLOps tools through a SLR of academic papers. It examines the platforms and libraries adopted by practitioners and analyzes which stages each tool addresses to reveal its function, scope, and the challenges it is designed for. The review also investigates the factors influencing tool adoption, offering insights into the technologies most relevant to current MLOps practices and how they integrate into real-world workflows.

3.2 Research Questions

3.2.1 *RQ1: Which tools employed in MLOps workflows are most frequently reported in academic literature?* Tools with significant popularity frequently benefit from strong community support, extensive documentation, and rich integration possibilities, making them preferred candidates for further exploration. They give a great perspective on how many practitioners interact with MLOps. The goal of this research question is to determine which tools are most often used in the literature. We can identify which tools are most popular by looking at how frequently they are integrated.

3.2.2 *RQ2: Which stages of the MLOps lifecycle do these tools cover in the use cases reported in the academic literature?* ML pipelines are made up of several steps, ranging from data processing and model training to model deployment and continuous monitoring. MLOps tools rarely attempt to cover all of these stages comprehensively. They generally specialise in a single or subset of these stages. Analyzing which stages each tool addresses reveals its function, scope and challenges it is designed for. This understanding is crucial since, as mentioned before, engineers and researchers tend to combine multiple tools to get their final desired pipeline.

3.2.3 *RQ3: What are the reported benefits and limitations of the MLOps tool?* MLOps tools offer a variety of features designed to meet the diverse needs of the user. Following the analysis of existing literature and documented user experiences, we examine common trends in the advantages and limitations of various MLOps tools.

3.3 Pilot Study

Before starting the full review, we ran a pilot study with just the first ten papers from the Google Scholar results, following Kitchenham's SLR guidelines [21]. A pilot study is important to validate our methodology and ensure the consistency and reliability of selection and extraction before full deployment. We subjected the selected studies to the full review protocol. The pilot revealed that some inclusion and exclusion criteria were incorrect, so we refined the set of extraction fields by removing unnecessary items and adding missing elements. We adjusted the search string, updated the inclusion rules, and modified the extraction spreadsheet. The study design presented below is the final version after all of those tweaks were applied.

3.4 Initial search

In February 2025, we conducted our literature search on Google Scholar. We first identified relevant keywords and synonyms for the MLOps literature. The search-term selection process involved several iterations of testing various keywords in Google Scholar and assessing whether the top results aligned with MLOps tooling and lifecycle studies. The finalised search string is as follows:

Google Scholar Search String

```
("MLOps" OR "machine learning operations") AND ("tool"
OR "application" OR "framework" OR "platform" OR
"pipelines") AND ("comparison" OR "evaluation" OR
"benchmark" OR "analysis" OR "empirical")
```

175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253

255 The initial query returned 9 100 records; however, to keep the
 256 review scope manageable and respect our time constraints, we
 257 screened only the first ten pages of Google Scholar results and
 258 reviewed 96 papers.

260 3.5 Application of Selection Criteria

261 For a paper to qualify as a primary study, it must satisfy all inclusion
 262 criteria and none of the exclusion criteria.

263 Inclusion Criteria:

- 264 • I1: Papers that either analyse/compare MLOps tools or
 265 describe the development/implementation of MLOps tools,
 266 frameworks, or pipelines.
- 267 • I2 Papers published in the last five years to ensure relevance
 268 (i.e., 2020 and later)

270 Exclusion Criteria:

- 271 • E1: Papers only describing applications or projects using
 272 an MLOps tool without detailed tool analysis.
- 273 • E2: Papers discussing MLOps as a concept/architecture or
 274 highlighting the benefits of the MLOps culture.
- 275 • E3: The study must present original primary research; ex-
 276 cluded are secondary studies such as literature reviews or
 277 surveys.
- 278 • E4: Non-English literature.
- 279 • E5: The full text of the paper is not available for our insti-
 280 tution.

281 The selection criteria I2, E4, and E5 are standard criteria for SLRs.
 282 Criteria I1, E1, E2, and E3 ensure we exclude papers outside the
 283 scope of this review (those lacking original primary research or
 284 detailed analysis/development of MLOps tools) and thus focus only
 285 on studies that directly contribute to our research questions.

287 3.6 Snowballing

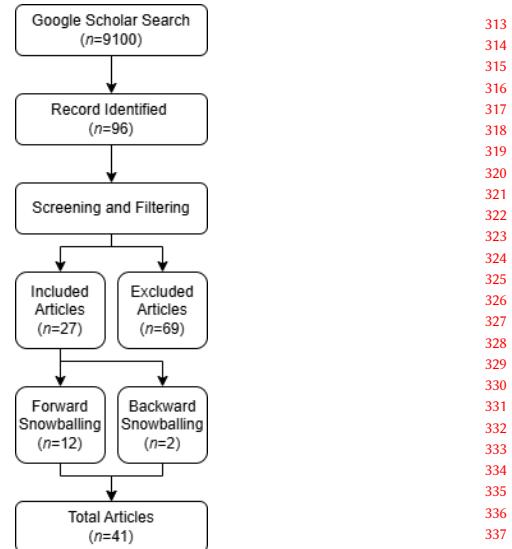
288 We applied both backward and forward snowballing to expand the
 289 results obtained through our initial search. Our snowballing process
 290 follows the approach suggested by Wohlin et al. [46]. To perform
 291 backward snowballing, we first reviewed the reference sections of
 292 our initial set of 96 papers by manually assessing titles to determine
 293 their relevance. If a title aligned with our research scope, we further
 294 examined the corresponding paper's abstract. Through this process,
 295 we identified four additional papers that met our inclusion criteria.
 296 On the other hand, during forward snowballing, we used Google
 297 Scholar to find studies that cite our initial set of papers. From this
 298 process, seven additional relevant papers were identified.

299 In both cases of forward and backward snowballing, we applied
 300 the same selection criteria and data extraction procedures (3.5, 3.7)
 301 as those used for the initial set of papers. We limited the snowballing
 302 process to a single round since further iterations yielded minimal
 303 additional relevant papers beyond those identified initially.

304 Figure 1 illustrates the complete selection process, from initial
 305 paper retrieval, through screening and inclusion, to snowballing.

307 3.7 Data Extraction

308 In this phase of our study, we performed a systematic analysis of the
 309 primary studies to extract data related to our research questions. As
 310 detailed in subsection 3.3, our initial pilot study helped us refine the
 311



313
 314
 315
 316
 317
 318
 319
 320
 321
 322
 323
 324
 325
 326
 327
 328
 329
 330
 331
 332
 333
 334
 335
 336
 337
 338
 339
 340
 341
 342
 343
 344
 345
 346
 347
 348
 349
 350
 351
 352
 353
 354
 355
 356
 357
 358
 359
 360
 361
 362
 363
 364
 365
 366
 367
 368
 369
 370
 371
 372
 373
 374
 375
 376
 377
 378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391

Figure 1: Flow diagram showing selection process record retrieval, screening, inclusion, and citation snowballing steps for the MLOps literature review

data extraction strategy. For each research question, we extracted a predetermined set of information from every primary research paper and recorded it in spreadsheet. With respect to RQ1, we identified tools that were actively implemented, excluding tools that were only mentioned. This distinction enabled us to evaluate the researchers' hands-on experiences with the tools. As for RQ2, we noted the tools' specific use cases mentioned in the paper. Finally, for RQ3, we gathered information on the reported benefits and limitations of each tool by analysing the justifications for their selection and the challenges encountered during their application. Table 1 lists the data extraction fields used.

Table 1: Data extraction fields

Field	Description
Name	Name of the MLOps tool
Uses	Practitioner use cases
Pipeline stage covered	Stage(s) of the MLOps lifecycle addressed
Paper	Citation of the academic study
Benefits	Reported benefits of the tool
Limitations	Reported limitations or challenges
Notes	Additional observations or context

3.8 Data Synthesis

3.8.1 *RQ1: Which tools employed in MLOps workflows are most frequently reported in academic literature?* In this stage of our study, we synthesise the data extracted from the selected primary research papers to address our research questions. To represent our findings for RQ1, we constructed Figure 2, a frequency graph that lists each

393 tool alongside the number of times it was used in the reviewed literature. This graph provides a quantitative measure of each MLOps
 394 tool's popularity, further details and interpretation are presented
 395 in the Results and Discussion sections.
 396

397 3.8.2 RQ2 To address RQ2, we categorised the tools based on
 398 their respective use cases, according to the MLOps components as
 399 defined by Najafabadi *et al.* [32]. In this section, we list the selected
 400 subset of the original components whilst adapting their names and
 401 responsibilities to our context. We have also introduced a new
 402 component, Visualisations, to better capture gaps in the existing
 403 architecture.
 404

405 Table 2 summarises the MLOps lifecycle categories that emerged
 406 from the primary studies discussed in this section.
 407

408 3.8.3 RQ3 To answer RQ3 we applied a descriptive synthesis,
 409 which means we looked for patterns in descriptions rather than
 410 combining numerical results. First, we copied every sentence that
 411 mentioned a benefit or limitation of a tool into a spreadsheet. We then
 412 grouped together statements that conveyed the same idea, allowing
 413 us to identify benefits and limitations that appeared across several
 414 studies, as well as those mentioned only once. Finally, we compiled
 415 a table for every tool that lists these combined benefits and limita-
 416 tions so readers can see at a glance what the literature agrees on
 417 for each tool. This analysis helped us capture the context of tool
 418 selection, including the rationale behind their adoption and the
 419 challenges associated with their use.
 420

3.9 Study Replicability

To ensure full replicability of our review, we have made a public Zenodo repository [28] with a spreadsheet that documents every step of the study. It comprises four sheets:

- (1) **Paper Selection** – the complete set of Google Scholar search results, showing whether each paper was included or excluded from our review and the exact criterion applied.
- (2) **Selected Papers** – all papers that passed screening, indicating whether they came from the primary search or from forward/backward snowballing, along with key metadata (author, publication year, and so on).
- (3) **Extraction** – the raw data taken from each selected paper: tool names, quoted passages, reported benefits or limitations, and possible dependencies or integrations.
- (4) **Synthesis** – a consolidated view that maps every extracted tool to the relevant component of the MLOps architecture as well as summaries of the common advantages and drawbacks. Key synthesis results are presented in this paper, while the full mapping is available in the supplementary material [28].

Researchers can follow the study design and the data in the provided sheets to reproduce our search, screening, extraction, and synthesis processes or to extend the analysis.

3.10 Threats to Validity

In this section, we follow the threat classification schemes for experiment validity described by Ampatzoglou *et al.* [3] and outline the threats that may affect the validity of our research.

Table 2: MLOps lifecycle categories observed in the primary studies

Component	Responsibility
Orchestrator	Provides system-wide orchestration and schedules multiple models while balancing throughput and latency.
Raw Data Store	Holds raw source data; needs specialised versioning tools because datasets exceed typical Git size limits.
Data Preprocessor	Transforms, cleans, and validates data before it becomes training input.
Dataset Repository	Stores and versions datasets; relies on large-file platforms.
Feature Store	Computes, stores, and serves reusable features with low latency.
Artefact Repository	Keeps packaged or containerised ML components that include a model.
ML Metadata Repository	Tracks training metadata for experiment tracking and model performance.
Code Repository	Versions source code, configuration files, and related artefacts.
Model Repository	Versions trained models together with basic metadata such as version numbers.
ML Training Pipeline (Online)	Automates continuous model training at runtime in production.
ML Experiment Pipeline (Offline)	Supports manual experimentation and training during design time.
MLOps User Interaction Manager	Enables interaction between the MLOps team and the platform.
ML Pipeline Editor	Builds, tests, and packages pipeline code into containers or similar environments.
Model Deployer	Deploys a trained model and its dependencies to production.
Model Evaluator	Measures and assesses model performance.
Runtime Model Monitor	Continuously watches serving performance and infrastructure metrics.
Visualisation	Presents dashboards and graphs for experiments, metrics, and system status.
End-to-End	Covers the full ML workflow from data ingestion to inference, though most tools still leave gaps.
Managed End-to-End	Provides an end-to-end pipeline fully managed by the platform, automating infrastructure, execution, and monitoring; often requires auxiliary services for complete coverage.

3.10.1 External validity External validity relates to the generalisability of our systematic literature review's findings. For our primary studies to accurately represent the MLOps field and draw correct conclusions, it is essential that these studies reflect the diverse MLOps landscape. One threat to generalisability is that the limitations reported in the literature may be based on older versions of MLOps tools, thereby inaccurately representing their current

451
 452
 453
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490
 491
 492
 493
 494
 495
 496
 497
 498
 499
 500
 501
 502
 503
 504
 505
 506
 507
 508
 509
 510
 511
 512
 513
 514
 515
 516
 517
 518
 519
 520
 521
 522
 523
 524
 525
 526
 527
 528
 529
 530

531 state. Such limitations might have already been addressed by the
 532 developer and thus may no longer hold.

533 Another threat comes from the sheer number of results returned
 534 by Google Scholar. Its proprietary relevance ranking does act as a
 535 partial mitigation by prioritizing influential papers first. Never-
 536 theless, we reviewed only the first ten result pages, covering 96
 537 papers, so some relevant studies may have been missed and this
 538 could introduce selection bias. A broader search by screening more
 539 Google Scholar pages would mitigate this limitation.

540
 541
 542 *3.10.2 Internal validity* Internal validity examines whether the
 543 study's design and execution provide a confident basis for linking
 544 causes with effects. A key threat in our review is the process of
 545 selecting appropriate papers and formulating an exhaustive set of
 546 search terms. Vital studies might be missed if the search terms are
 547 not thorough, which could introduce bias. In order to mitigate this
 548 threat, we pre-identified a number of key papers that needed to be
 549 on our final list of papers.

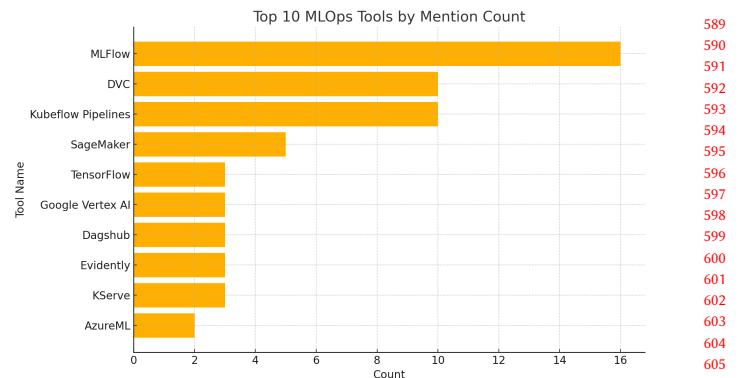
550 Another threat concerns our tool-component heatmap (Figure
 551 3). We can map only the capabilities that authors explicitly describe,
 552 so the heatmap may miss components that a tool supports but were
 553 not used in the included studies. When an author employs only
 554 a subset of the tools capabilities, any unmentioned components
 555 are left out, which makes the mapping non-exhaustive and might
 556 under-represent the true scope of certain tools.

557
 558
 559 *3.10.3 Construct validity* Construct validity concerns how well
 560 our measures and constructs align with the theoretical concepts we
 561 intend to study. A threat here is inherent in the source literature:
 562 many academic papers emphasise implementing MLOps pipelines
 563 or addressing broader issues rather than critically evaluating the
 564 tools themselves. This tendency often leads to detailed reporting on
 565 the benefits while underreporting limitations, which may result in
 566 conclusions that do not accurately capture the tool's effectiveness.

567
 568
 569 *3.10.4 Conclusion validity* Conclusion validity focuses on the ac-
 570 curacy of the deductions generated from our data analysis. Our
 571 conclusions, which are based on frequency counts, categorisation
 572 mappings, and feature evaluations, are meant to be rational. How-
 573 ever, verifying that the data analysis procedures are sound and
 574 executed correctly is crucial to confirming the validity of our find-
 575 ings. Finally, the extraction and synthesis process was not reviewed
 576 by a third party, which could introduce self-bias.

4 Results

581 In this section, we present the outcomes of our SLR structured
 582 by three research questions. Section 4.1 (RQ1) reports the most
 583 frequently mentioned MLOps tools (Figure 2); Section 4.2 (RQ2)
 584 examines tool capabilities across pipeline stages and categories
 585 (Figures 3 and 4); and Section 4.3 (RQ3) synthesizes reported
 586 benefits and limitations (Tables 3–10). Further discussion appears in
 587 Section 5.



588
 589
 590
 591
 592
 593
 594
 595
 596
 597
 598
 599
 600
 601
 602
 603
 604
 605
 606
 607
 608
 609
 610
 611
 612
 613
 614
 615
 616
 617
 618
 619
 620
 621
 622
 623
 624
 625
 626
 627
 628
 629
 630
 631
 632
 633
 634
 635
 636
 637
 638
 639
 640
 641
 642
 643
 644
 645
 646
 647
 648
 649
 650
 651
 652
 653
 654
 655
 656
 657
 658
 659
 660
 661
 662
 663
 664
 665
 666
 667
 668

Figure 2: Top 10 MLOps tools ranked by number of mentions in primary studies

4.1 RQ1: Which tools employed in MLOps workflows are most frequently reported in academic literature?

Our first research question aims to identify the most widely used MLOps tools based on evidence from academic studies. In order to determine the prevalent tools, Figure 2 presents a bar graph showcasing tool usage based on our primary data sources. As explained previously, each tool had to be both described in the literature and practically applied by the authors for it to be included in our study. A significant portion of the literature focused on authors researching multiple tools in order to implement an MLOps pipeline. In these studies, the literature review section typically describes various MLOps tools and evaluates them based on the features offered. In our analysis, we only listed the tools that were actually implemented in their final pipeline. As shown in Figure 2, MLflow emerged as the most commonly implemented tool, appearing 16 times, followed by DVC and Kubeflow Pipelines, each appearing 10 times. Further interpretation is provided in the discussion section.

4.2 RQ2: Which stages of the MLOps lifecycle do these tools cover in the use cases reported in the academic literature?

MLOps pipelines encompass a range of stages, including data ingestion, pre-processing, model development, training, validation, deployment, and ongoing monitoring. A heatmap (Figure 3) was created to visualise what component(s) each extracted MLOps tool fulfils, along with the frequency of each tool's implementation across the reviewed studies. Analysing the stages supported by each tool revealed common areas of focus among MLOps solutions. Overall, very few tools cover the entire pipeline, as most specialise in one or more phases. Since MLOps tools can address several components, authors may discuss only a subset. As a result, the figure may omit certain stages a tool could cover, making our assessment of functionality incomplete. Throughout this paper we use the terms "stage", "phase", and "component" interchangeably to refer to one functional step in that lifecycle.

Furthermore, Figure 4 presents a bar graph of MLOps tool counts grouped by broad categories. These categories are adapted from

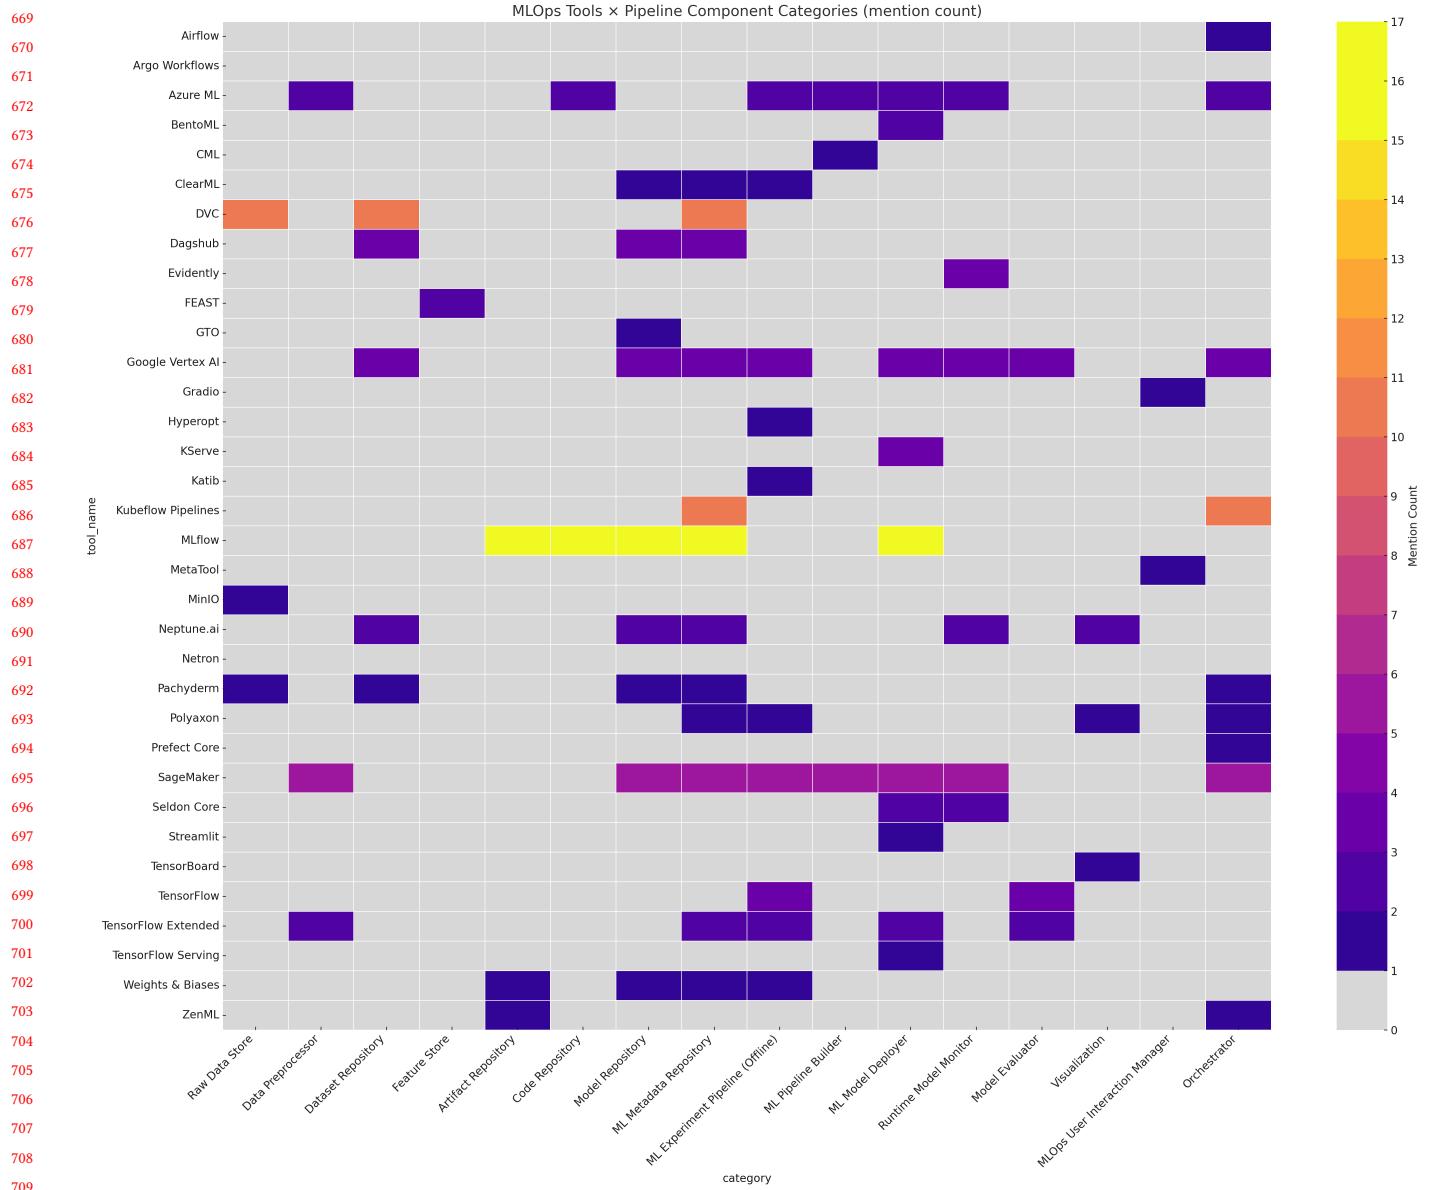


Figure 3: Heatmap illustrating MLOps tools mapped across different MLOps pipeline phases, showing which phase(s) each tool supports. Colour intensity represents the number of times each tool was implemented.

the taxonomy proposed by Najafabadi *et al.*, who originally defined six groups to classify MLOps components. To better reflect our corpus, we both added two new groups—End-to-End and Managed End-to-End—and removed unused groups such as Data Curation. The full set of categories is:

- **End-to-End**
- **Managed End-to-End**
- **Storage and Versioning**
- **Infrastructure and Supporting Services**
- **Inference**
- **ML Training**

This categorization is primarily intended to facilitate the organization and presentation of our findings in this paper. While individual tools may fall under multiple MLOps components, we assign each tool to a single category to provide a clear and high-level overview of the MLOps landscape.

4.3 RQ3: What are the reported benefits and limitations of the MLOps tools?

The following Tables 3–10 present a synthesis of the reported benefits and limitations of each MLOps tool identified in our review, thereby addressing Research Question 3.

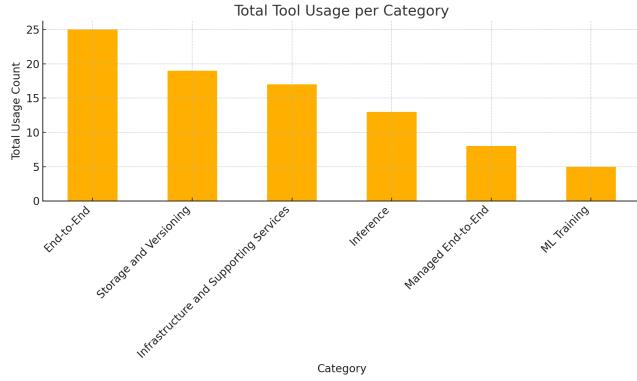


Figure 4: MLOps tool count grouped by category

These tables are grouped according to the categories defined in Section 4.2 (End-to-End, Managed End-to-End, Storage and Versioning, Infrastructure and Supporting Services, Inference, ML Training), which allows for straightforward comparison of the practical considerations involved in tool selection and implementation. Each table organises tools by one of these categories.

These summaries are drawn from direct quotations in the source literature, capturing the most common insights authors reported from academic projects or industrial deployments. The complete set of quotations, tools, and source papers is available in the accompanying supplementary sheets.

The tables are followed by Section 5, where we delve deeper into the implications of these findings.

4.3.1 Infrastructure and Supporting Services

Apache Airflow is an open-source platform for orchestrating production workflows and data pipelines [50].

Argo Workflow Argo Workflow is a container-native workflow engine for orchestrating parallel jobs on Kubernetes [30]. By combining Argo Events for webhook-triggered workflows and Argo Workflow for execution, developers can automate the full ML lifecycle in a reproducible, scalable, and hands-off manner. This includes training, evaluation, and deployment. While production environments can benefit from such pipelines for mature, production-ready teams, it is not recommended to incorporate such an environment in the early stages of ML services or businesses [30].

Kubeflow Pipelines handles managing and orchestrating containerised workloads [15] while taking care of model training, deployment, and coordination [24]. Since it runs on top of Kubernetes, it lets users codify preprocessing, training, and deployment steps in a single UI and execute many pipelines in parallel. Because every component ships as a Kubernetes resource, the same pipeline runs on-prem or in any managed Kubernetes service. It also has auto-scaling pods, notebook sessions provided by Jupyter, experiment tracking, hyperparameter tuning, and pluggable serving (Kubeflow's own KServe or other solutions such as Seldon Core), which all sit behind the central dashboard.

Prefect Core is a pipeline orchestrator with a modern Python API that keeps code readable and flexible when workflows are highly dynamic [11].

ZenML links the orchestration layer with an artefact store so that a stack can swap orchestration, or any other component, without rewriting user code [5].

CML enables CI/CD for ML projects. It wires continuous integration to ML experiments, tracking changes and auto-generating metric reports [35].

Gradio enables fast creation of interactive web interfaces for model demos and evaluation without touching HTML, CSS, or JavaScript [15].

TensorBoard is TensorFlow's official visualisation toolkit, offering interactive exploration of computation graphs, layer structures, and distributions of weights and biases through histograms. TensorBoard also provides real-time plotting of training metrics such as loss and accuracy over time [47].

4.3.2 Managed End-to-End

AzureML is a managed end-to-end service; users still provision the workspace, storage, networking, and container registry either interactively or through infrastructure-as-code [29].

AWS SageMaker is Amazon's cloud platform for building, training, tuning, and deploying models in a hosted production environment [34].

ClearML provides semi-automated experimental run tracking and auditing, capturing comprehensive metadata, parameters, and metrics across popular Data Science tools. It relies on a dedicated central web server underpinned by a MongoDB database, an Elastic-search deployment, and cloud object storage for artefact archiving, enabling end-to-end traceability and auditability of experiments [40].

4.3.3 End-to-End

MLflow is an end-to-end tool made of multiple modules [9] which was initially started as an open-source platform to provide comprehensive experiment tracking capabilities [25][31][5][44]. It has grown to cover most of the MLOps pipeline through the inclusion of the following components. The *MLflow Tracking* component records parameters, metrics, and artefacts for every run [40], ensuring a complete workflow history. *MLflow Projects* encapsulate data science code and dependencies, enabling reproducible execution across diverse environments. *MLflow pipeline* involves processing, cleaning, transforming, training models, and evaluating them, ensuring large-scale ML model deployment. *MLflow Models* then offers a standardised structure for packaging machine learning models and facilitates seamless deployment in various serving infrastructures [9]. Finally, *MLflow Registry* serves as a unified model store, providing versioning, annotation, stage transitions, and the promotion of specific model versions to "Production" for batch or real-time inference [31][39]. Autologging further streamlines experiments by automatically capturing parameters and metrics from supported ML libraries [5].

807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944

945 **Neptune.ai** positions itself as a complete platform that spans
 946 data versioning, testing, deployment, and monitoring [43].

947 **Pachyderm** offers GUI-based notebook services, experiment
 948 tracking, and hyperparameter tuning [23].

949 **Polyaxon** is an open-source, Kubernetes-native end-to-end MLOps
 950 platform focused on experiment tracking, visualisation, parallel
 951 executions and hyperparameter optimisation. Its main focus is
 952 parallelisable experiment tracking, where it attempts to provide
 953 computationally efficient mechanisms for running and interpreting
 954 parallel ML experiments at scale. While Polyaxon's core and exper-
 955 imentation tools are open source, its automation and management
 956 features are not [23].

957 **TensorFlow Extended (TFX)** offers a production-grade pipeline
 958 implementation for TensorFlow models [20].

959 **Google Vertex AI** is part of GCP, merging AutoML and AI Plat-
 960 form behind one API, client library, and UI [41][29]. The Workbench
 961 offers Jupyter notebooks [8] and AutoML streamlines training, ex-
 962 periment tracking, and metadata management. A Model Registry
 963 manages versions for online or batch prediction, while the Pipeline
 964 view visualises step dependencies, integrated logging, and moni-
 965 toring track performance over time.

966 4.3.4 ML Training

967 **Katib** is part of Kubeflow. It automates hyperparameter tuning
 968 across frameworks including TensorFlow, MXNet, PyTorch, and
 969 XGBoost [17].

970 4.3.5 Storage and Versioning

971 **Weights & Biases** is cited as an ML metadata repository, yet
 972 the sources offer no elaboration [5].

973 **Dagshub** hosts data repositories and models for collaborative
 974 versioning. Some authors store all pipeline outputs remotely in it
 975 [31][2].

976 **DVC** extends Git to manage large datasets and model files, en-
 977 abling version control of large data alongside code [31] [4] [25]
 978 [9] [36] [7][20][15]. Consequently, its Git-like workflow simplifies
 979 adoption for users familiar with Git [4]. DVC's pipeline feature
 980 modularizes data processing into stages with explicit input-output
 981 dependencies and parameterisation, ensuring reproducibility, au-
 982 tomation, and reusability across projects [4]. It also supports ex-
 983 periment tracking and storing metadata within Git, while actual
 984 data resides in DVC storage, with metadata files orchestrating data
 985 retrieval [35].

986 **Feast** provides both offline and online feature storage, keeping
 987 historical as well as live data in sync [44]. One paper mislabels it
 988 as a performance monitoring tool [47].

989 **GTO** is Iterative's GitOps-based model registry. It eliminates
 990 the need for separate servers or databases by leveraging Git and
 991 DVC repositories and provides streamlined support for promoting
 992 models through designated repository branches and stages [35].

993 **MinIO** is an S3-compatible object store that keeps artefacts and
 994 metadata from ZenML and MLflow but can host any data [5].

995 4.3.6 Inference

996 **Metaflow** is a framework for creating and executing data sci-
 997 ence workflows in local environments and scaling them to the cloud
 998 with ease [40].

999 **KServe** originates from the Kubeflow ecosystem and wraps
 1000 models as web services on Knative/Istio. This results in autoscaling,
 1001 isolation, and an event-driven path to downstream monitoring
 1002 [24][20][17].

1003 **BentoML** is a model serving tool that bundles a trained model
 1004 and its dependencies into a production-ready service [40][39].

1005 **Seldon Core** is a model deployment tool that exposes models
 1006 as web services [11] with robust support for Kubernetes [24].

1007 **Evidently** is an open-source model monitoring library that gen-
 1008 erates predefined monitoring reports with a few lines of code. It can
 1009 automatically detect drift in input data, target values, and predic-
 1010 tions, and supports multi-model dashboards by calculating metrics
 1011 across multiple models with a single monitor [2][44][24].

1012 **Streamlit** provides a quick route to deploy ML applications
 1013 through a simple Python interface [2].

1014 5 Discussion

1015 The following subsections are organised around our three research
 1016 questions and draw directly on the synthesis captured in Tables
 1017 1–8. Each benefit extracted from the studies is labelled B1, B2, ...
 1018 and each limitation as L1, L2 and so on. Thus, a reference such
 1019 as B31 sends you to the 31st benefit row, while L17 points to the
 1020 17th limitation. To examine the supporting evidence, consult the
 1021 corresponding rows in Tables 1–8, where the summarised finding
 1022 and the primary paper that reported it are.

1023 5.1 RQ1 – Which tools employed in MLOps 1024 workflows are most frequently reported in 1025 academic literature?

1026 Across the primary studies, four tools stand out: **MLflow**, **DVC**,
 1027 **Kubeflow Pipelines**, and **AWS SageMaker** as shown in Figure 2.
 1028 The first three are fully open-source, while AWS SageMaker is a
 1029 proprietary cloud service. This divide reflects a familiar compromise
 1030 in MLOps practice: practitioners prefer community-maintained
 1031 tools for tasks like experiment tracking and data versioning but
 1032 often turn to commercial platforms when they need infrastructure
 1033 that scales quickly and comes with operational support.

1034 Both MLflow and Kubeflow are classified as end-to-end plat-
 1035 forms, providing experiment tracking, model packaging, and de-
 1036 ployment in one bundle. However, their design philosophies diverge
 1037 significantly. MLflow focuses on accessibility: it is easy to use, has
 1038 a straightforward web UI, thorough and substantial documentation,
 1039 wide storage-backend support, and a lively community (B24–B30).
 1040 Kubeflow takes a different approach. It is rooted in Kubernetes and
 1041 focuses more on customisability with fine-grained control, modular
 1042 components, and automatic scaling (B3–B8). However, it demands
 1043 a tougher setup and steeper operational know-how (L3–L5). The
 1044 fact that both tools top the frequency chart suggests that while
 1045

1083 some users prefer easy-to-use tools and others value customisable
 1084 options, there is room for both approaches to succeed.

1085 DVC is a notable outlier among the most cited tools, as it tackles
 1086 only a single stage of the MLOps life cycle, data storage and ver-
 1087 sioning, and yet it appears almost as often as the full end-to-end
 1088 platforms. Synchronising large data artefacts with code has been
 1089 a historic pain point in ML. Authors repeatedly select DVC for its
 1090 easy-to-use Git-like interface (B45) as well as pipeline caching and
 1091 serverless architecture. A recurring pairing was the adoption of Git
 1092 for code with DVC for large data because of Git's size limits. Git LFS
 1093 was meant to alleviate this struggle, but unlike Git LFS, DVC needs
 1094 no extra server, a difference reviewers flag as decisive. L29–L30
 1095 confirm that none of the other surveyed tools offer a comparable
 1096 alternative.

1097 Another reason for MLflow's popularity is that, even though it is
 1098 open-source, it remains a mature, well-established platform trusted
 1099 by leading companies and backed by both a strong community and
 1100 Databricks (B26). Its popularity feeds a loop: more users attract more
 1101 contributors, the codebase improves, and the project becomes even
 1102 more appealing, which leads the community to further grow. The
 1103 same holds true for Kubeflow (B4). Some projects, however, are less
 1104 attractive because of their "uncertain vitality", meaning that their
 1105 long-term health is harder to predict; Polyaxon, for instance, relies
 1106 on a comparatively small contributor base (L27). Lower adoption
 1107 means fewer contributors, slow progress, and even less visibility;
 1108 a feedback loop in reverse that keeps new users away. In practice,
 1109 the presence of an active maintainer community weighs heavily
 1110 when choosing a tool for the long haul (B26, B4). MLflow does
 1111 not just entice developers who want a self-hosted solution, since
 1112 developers who prefer to skip infrastructure work can opt for a
 1113 managed MLflow provided by Databricks, which sidesteps the need
 1114 to run and maintain a separate tool (L16).

1115 The widespread use of AWS SageMaker highlights the ongoing
 1116 significance of managed services (Figure 2). Managed MLOps plat-
 1117 forms such as SageMaker remain popular because they include a
 1118 wide array of pre-integrated pipeline components, such as a fea-
 1119 ture store, model registry, and CI/CD templates (B17). Companies
 1120 accept the subscription fee and the inherent risks of vendor lock-in
 1121 because the alternatives present significant challenges and costs
 1122 (L10, L12, L9, L15). A managed service hides the heavy work of
 1123 provisioning and securing compute, storage, and networking, and it
 1124 scales on demand (B18, B15, B16, B23). The same attraction applies
 1125 to the managed stacks from Azure ML (B14) and Google Cloud
 1126 Vertex AI (B22). Organisations that are unwilling to rely on a third-
 1127 party cloud provider often choose to build an in-house pipeline
 1128 from open-source projects such as MLflow for experiment tracking,
 1129 Kubeflow for orchestration, and DVC for data versioning. How-
 1130 ever, this option demands far more integration effort and long-term
 1131 maintenance and may be better suited to larger teams and more
 1132 mature projects (L3, L29, L19, L16).

1134 5.2 RQ2 – Which stages of the MLOps lifecycle 1135 do these tools cover in the reported use 1136 cases?

1138 Mapping each tool to the taxonomy defined in Section 2 reveals
 1139 that no single solution addresses the entire ML lifecycle.. However,

1140 this is more due to what the papers mention than what the product
 1141 definitively offers. Kubeflow excels at orchestration (B4), MLflow
 1142 at experiment tracking (B28), while DVC and Feast handle data and
 1143 feature management, respectively (B45, B57).

1144 AWS SageMaker bundles a model registry, feature store, and
 1145 deployment tools, yet teams still turn to third-party services for
 1146 granular security and local runs (L10–L12). The component that
 1147 offers the least coverage is in feature stores and runtime monitoring;
 1148 outside of Feast and Evidently, non-managed options are almost
 1149 nonexistent (Tables 6 and 7). Consequently, multi-tool pipelines are
 1150 the norm, underpinning the importance of effective tool integration.

1151 While classifying tools, we found one component that did not fit
 1152 cleanly into Najafabadi *et al.*'s [32] component architecture: visu-
 1153 alisation dashboards such as TensorBoard, Weights & Biases, and
 1154 Gradio. The closest existing component was Runtime Monitoring,
 1155 which does not really capture their purpose, so we introduced a
 1156 dedicated visualisation component instead.

1157 Across all reviewed papers, the ML Metadata Repository (Experiment-
 1158 tracking) stage is addressed the most. This dominance is caused
 1159 by the ubiquity of MLflow, making it a default choice for an exper-
 1160 iment tracking tool that is open source. Managed solutions such
 1161 as SageMaker and Azure ML (B17, B16) further contribute to its
 1162 widespread adoption. Reliable metadata is essential for reproducibil-
 1163 ity, auditability, and model evaluation; consequently, practitioners
 1164 consistently prioritise this component.

1165 Yet orchestration tools sit a close second: they are the "glue" that
 1166 binds data prep, training, and deployment into a runnable pipeline,
 1167 so virtually every study that examines end-to-end workflows also
 1168 highlights orchestrators such as Kubeflow, Argo, or similar sched-
 1169 ulers.

1171 5.3 RQ3 – What are the claimed benefits and 1172 limitations of the MLOps tool?

1173 5.3.1 *Security and Managed Platform Trade-offs* MLflow lacks auto-
 1174 matic data versioning (L17), which explains why many studies pair
 1175 it with DVC (B45–B49). It is also missing role-based access control
 1176 (L19). In contrast, AWS SageMaker covers that gap through IAM in-
 1177 tegration (B19). A well-established trade-off emerges: open-source
 1178 tools need extra engineering to meet enterprise-grade security,
 1179 while cloud platforms shift that work to the provider at the cost
 1180 of vendor lock-in. Even then, the cloud does not guarantee full
 1181 coverage. For example, Google's Vertex AI still trails Azure ML and
 1182 SageMaker on security, user control, and governance despite tying
 1183 you to Google's stack.

1184 5.3.2 *Post-deployment Monitoring* Evidently is the only tool that
 1185 comes with a drift detection feature out of the box (B67–B71), and
 1186 it works only for tabular data (L37). The limited findings indicate
 1187 that post-deployment observability is still an emerging component
 1188 of the ecosystem.

1189 5.3.3 *Prerequisite Knowledge* A common limitation noted among
 1190 tools is the required prerequisite knowledge. Pachyderm requires
 1191 Helm and cloud-storage expertise (L23), while both Kubeflow and
 1192 Argo Workflows demand solid Kubernetes and containerisation
 1193 experience (L3, L7–L8). Several other services, including Feast's SDK
 1194 (L34), Weights & Biases' client code (L33), Neptune's Python API
 1195 (L20), and Pachyderm's Helm chart (L23), require significant expertise
 1196 to set up and maintain. This is a significant barrier for adoption, as
 1197 many organisations lack the internal expertise to support these tools.
 1198

(L22) and SageMaker's SDK (L11), call for advanced programming skills. These findings show that tool comparison should not focus solely on a tool's features and limitations, but it must also consider the skills a team possesses.

5.3.4 *Integration and Flexibility* Easy integration and language agnosticism are among the most frequently praised benefits. Orchestrators such as Kubeflow Pipelines are applauded for both their cloud-agnostic Kubernetes foundation and their seamless hooks into TensorFlow Extended (B3, B6), while managed stacks win favour largely because of the way they slot into their parent ecosystems—AWS SageMaker's tight coupling with IAM and the rest of the AWS suite being a prime example (B19). At the experiment-tracking layer, MLflow extends this integrative spirit through container-friendly, self-hosted deployments and pluggable back-end stores that work just as well with S3, Azure Blob, on-prem NFS, or SQL-compatible databases (B25, B30). Flexibility in storage backends is also demonstrated by DVC's remote options (B49) and MLflow's broad object-store support (B30). TFX is widely appreciated for its portability, with pipelines that can run seamlessly across multiple orchestrators rather than being locked to a single workflow engine (B41). Finally, inference services such as BentoML embrace framework diversity by supporting TensorFlow, PyTorch, Keras, XGBoost, and more out of the box (B61). Taken together, these examples show that the community consistently rewards tools that can drop into existing tech stacks without forcing a wholesale rewrite in a particular language, framework, or cloud. In contrast, the lack of flexibility is listed as a drawback by reviewers, specifically in TensorBoard and TFX which tie users to the TensorFlow stack (L39, L28).

5.3.5 *User Interfaces and Visualization* Finally, visual dashboards and UIs are widely appreciated. Kubeflow's central UI (B7), Vertex AI's pipeline view (B22), MLflow's experiment board (B24) and Pachyderm's web console (B35) are all reported as benefits, as they improve ease of use.

6 Conclusion

This review set out to identify which MLOps-native tools appear most frequently in academic work and to understand the reasons for their uptake. A structured Google Scholar search, followed by manual screening, filtered 96 papers to 41 selected studies which were extracted and synthesised in a systematic, structured manner.

Our review confirms a clear pattern in recent MLOps practice. MLflow, DVC, Kubeflow Pipelines, and AWS SageMaker appear most frequently in MLOps pipelines because each addresses a critical pain point: MLflow simplifies experiment tracking, DVC brings Git-style version control to data, Kubeflow coordinates cloud-agnostic workflows, and SageMaker lifts the infrastructure burden through a fully managed service. Their popularity is sustained either by lively open-source communities or by the deep resources of a major cloud provider.

Yet none of these platforms covers the entire lifecycle on its own. Researchers commonly assemble a mixed tool pipeline. Future work should move beyond cataloguing tools to evaluating tool interoperability and integration effort. This includes analyzing each tool's dependencies on external platforms and how easily it can be combined with others to build coherent workflows. A useful extension

would be a comparative table showing dependency relationships and open-source status, offering a clearer picture of ecosystem maturity and possible vendor lock-in. Benchmarking integration effort, metadata consistency, and runtime stability would offer the objective metrics needed to help practitioners choose interoperable end-to-end solutions.

References

- [1] [n. d.]. The Winding Road to Better Machine Learning Infrastructure Through Tensorflow Extended and Kubeflow | Spotify Engineering. <https://engineering.spotify.com/2019/12/the-winding-road-to-better-machine-learning-infrastructure-through-tensorflow-extended-and-kubeflow>
- [2] William Inouye Almeida. 2023. *Building an Automated MLOps Pipeline and Recommending an Open-Source Stack to Deploy a Machine Learning Application*. Master's thesis. Universidade do Porto (Portugal). <https://search.proquest.com/openview/3caeae32a5b14a5346907376fce17b338/1?pq-origsite=gscholar&cbl=2026366&diss=y>
- [3] Apostolos Ampatzoglou, Stamatia Bibi, Paris Avgeriou, Marijn Verbeek, and Alexander Chatzigeorgiou. 2019. Identifying, categorizing and mitigating threats to validity in software engineering secondary studies. *Information and Software Technology* 106 (2019), 201–230. <https://doi.org/10.1016/j.infsof.2018.10.006>
- [4] Vidushi Arora. 2024. Exploring real-world challenges in MLOps implementation: a case study approach to design effective data pipelines. (2024). <https://elib.uni-stuttgart.de/items/e6f46863-465d-4e88-861c-9dcabc746db>
- [5] Michal Bacigál. 2024. Design and Implementation of Machine Learning Operations. (Feb. 2024). <https://dspace.cvut.cz/handle/10467/113781> Accepted: 2024-02-09T23:53:17Z Publisher: České vysoké učení technické v Praze. Vypočetní a informační centrum.
- [6] Rahul Bagai, Ankit Masrani, Piyush Ranjan, Madhavi Najana, and Ankit Masrani. 2024. Implementing Continuous Integration and Deployment (CI/CD) for Machine Learning Models on AWS. *International Journal of Global Innovations and Solutions (IJGIS)* (May 2024). <https://doi.org/10.21428/e90189c8.9cb39c55> Publisher: The New World Foundation.
- [7] N. Bauman. 2022. Building a generalisable ML pipeline at ING. (2022). <https://repository.tudelft.nl/record/uuid:35c850eb-1d03-4185-a8c5-4469b2112327>
- [8] Ralph Bergmann, Felix Theusch, Paul Heisterkamp, and Narek Grigoryan. 2024. Comparative Analysis of Open-Source ML Pipeline Orchestration Platforms. (2024). https://www.researchgate.net/profile/Narek-Grigoryan-3/publication/382114154_Comparative_Analysis_of_Open-Source_ML_Pipeline_Orchestration_Platforms/links/668e31d6b15ba559074d9a4b/Comparative-Analysis-of-Open-Source-ML-Pipeline-Orchestration-Platforms.pdf
- [9] Anas Bodor, Meriem Hnida, and Daoudi Najima. 2023. From Development to Deployment: An Approach to MLOps Monitoring for Machine Learning Model Operationalization. In *2023 14th International Conference on Intelligent Systems: Theories and Applications (SITA)*. 1–7. <https://doi.org/10.1109/SITA60746.2023.10373733>
- [10] Anas Bodor, Meriem Hnida, and Daoudi Najima. 2023. MLOps: Overview of Current State and Future Directions. In *Innovations in Smart Cities Applications Volume 6*. Springer, Cham, 156–165. https://doi.org/10.1007/978-3-031-26852-6_14 ISSN: 2367-3389.
- [11] Antonio M. Burgueño-Romero, Cristóbal Barba-González, and José F. Aldana-Montes. 2025. Big Data-driven MLOps workflow for annual high-resolution land cover classification models. *Future Generation Computer Systems* 163 (Feb. 2025), 107499. <https://doi.org/10.1016/j.future.2024.107499>
- [12] Ji-hyun Cha, Heung-gyun Jeong, Seung-woo Han, Dong-chul Kim, Jung-hun Oh, Seok-hee Hwang, and Byeong-ju Park. 2023. Development of MLOps Platform Based on Power Source Analysis for Considering Manufacturing Environment Changes in Real-Time Processes. In *Human-Computer Interaction*. Springer, Cham, 224–236. https://doi.org/10.1007/978-3-031-35572-1_15 ISSN: 1611-3349.
- [13] Swati Choudhary. 2021. Kubernetes-Based Architecture For An On-premises Machine Learning Platform. (2021).
- [14] Thomas Davenport and Katie Malone. 2021. Deployment as a Critical Business Data Science Discipline. *Harvard Data Science Review* 3, 1 (feb 10 2021). <https://hdsr.mitpress.mit.edu/pub/2fu65ujf>.
- [15] Daniel Deutsch. 2023. *Machine learning operations – domain analysis, reference architecture, and example implementation / Author Daniel Deutsch, LL.B. (WU), LL.M. (WU)*. <http://epub.jku.at/obvulih8/8593075>
- [16] Christof Ebert, Gorka Gallardo, Josu Hernantes, and Nicolas Serrano. 2016. DevOps. *IEEE Software* 33, 3 (2016), 94–100. <https://doi.org/10.1109/MS.2016.68>
- [17] Kanwarpartap Singh Gill, Vatsala Anand, Rahul Chauhan, Ruchira Rawat, and Pao-Ann Hsiung. 2023. Utilization of Kubeflow for Deploying Machine Learning Models Across Several Cloud Providers. In *2023 3rd International Conference on Smart Generation Computing, Communication and Networking (SMARTGENCON)*. 1–7. <https://doi.org/10.1109/SMARTGENCON60755.2023.10442069>

1359 [18] Google Cloud Tech. 2020. Introduction to Kubeflow. <https://www.youtube.com/watch?v=cTZArDgbIWw> 1417

1360 [19] Kim Hazelwood, Sarah Bird, David Brooks, Soumith Chintala, Utku Diril, Dmytro 1418

1361 Dzhulgakov, Mohamed Fawzy, Bill Jia, Yangqing Jia, Aditya Kalro, James Law, 1419

1362 Kevin Lee, Jason Lu, Pieter Noordhuis, Misha Smtelyanskiy, Liang Xiong, and 1420

1363 Xiaodong Wang. 2018. Applied Machine Learning at Facebook: A Datacenter 1421

1364 Infrastructure Perspective. In *2018 IEEE International Symposium on High Performance 1422*

1365 Computer Architecture (HPCA). 620–629. <https://doi.org/10.1109/HPCA.2018.00059> ISSN: 2378-203X. 1423

1366 [20] Hannes Jämtner and Stefan Brynielsson. 2022. An Empirical Study on AI Work- 1424

1367 flow Automation for Positioning. (2022). 1425

1368 [21] Barbara Ann Kitchenham and Stuart Charters. 2007. *Guidelines for Per- 1426*

1369 forming Systematic Literature Reviews in Software Engineering. Technical Report 1427

1370 EBSE 2007-001. Keele University and Durham University Joint Report, 1428

1371 Keele, UK and Durham, UK. https://www.elsevier.com/__data/promis_misc/52544systematicreviewsguide.pdf 1429

1372 [22] Dominik Kreuzberger, Niklas Kühl, and Sebastian Hirschl. 2022. Machine 1430

1373 Learning Operations (MLOps): Overview, Definition, and Architecture. <https://doi.org/10.48550/arXiv.2205.02302> arXiv:2205.02302 [cs]. 1431

1374 [23] Anders Köhler. 2022. *Evaluation of MLOps Tools for Kubernetes : A Rudimentary 1432*

1375 Comparison Between Open Source Kubeflow, Pachyderm and Polyaxon. <https://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-488601> 1433

1376 [24] Yumo Luo. 2023. An Open-Source and Portable MLOps Pipeline for Continuous 1434

1377 Training and Continuous Deployment. (2023). 1435

1378 [25] Giulio Mallardi, Fabio Calefato, Luigi Quaranta, and Filippo Lanubile. 2024. 1436

1379 An MLOps Approach for Deploying Machine Learning Models in 1437

1380 Healthcare Systems. In *2024 IEEE International Conference on Bioinformatics and 1438*

1381 Biomedicine (BIBM). IEEE, 6832–6837. https://ieeexplore.ieee.org/abstract/document/10822603/?casa_token=5GBnjhwp4KAFFFFFF:EOEBUeWoY9ZBRsrZ5ij8AoSypmOyGTqTPEraM4kOitmHIRI-0s4zW9omjkMLOt65AuexQrDa3TwVGsQ 1439

1382 [26] Andres Felipe Varon Maya. [n. d.]. The State of MLOps. ([n. d.]). 1440

1383 [27] Rick Merritt. 2020. What is MLOps? <https://blogs.nvidia.com/blog/what-is-mlops/> 1441

1384 [28] Zakkaria Micallef. 2025. A Systematic Review of MLOps Tools: Practices, 1442

1385 Challenges, and Lessons Learned. <https://doi.org/10.5281/zenodo.15459745> 1443

1386 [29] Widad El Moutaouakal and Karim Bainâa. 2023. Comparative Experimentation of 1444

1387 MLOps Power on Microsoft Azure, Amazon Web Services, and Google Cloud 1445

1388 Platform. In *2023 IEEE 6th International Conference on Cloud Computing and 1446*

1389 Artificial Intelligence: Technologies and Applications (CloudTech). 1–8. <https://doi.org/10.1109/CloudTech58737.2023.10366138> 1447

1390 [30] Sasu Mäkinen. 2021. Designing an open-source cloud-native MLOps pipeline. 1448

1391 *University of Helsinki* (2021). <https://helda.helsinki.fi/server/api/core/bitstreams/d01f98ef-beef-4329-997d-3ebe70092590/content> 1449

1392 [31] Óscar A. Méndez, Jorge Camargo, and Hector Florez. 2025. Machine 1450

1393 Learning Operations Applied to Development and Model Provisioning. In *Applied 1451*

1394 Informatics, Hector Florez and Hernán Astudillo (Eds.). Vol. 2236. Springer Nature 1452

1395 Switzerland, Cham, 73–88. https://doi.org/10.1007/978-3-031-75144-8_6 Series 1453

1396 Title: Communications in Computer and Information Science. 1454

1397 [32] Faezeh Amou Najafabadi, Justus Bogner, Ilias Gerostathopoulos, and Patricia 1455

1398 Lago. 2024. An Analysis of MLOps Architectures: A Systematic Mapping Study. 1456

1399 Vol. 14889, 69–85. https://doi.org/10.1007/978-3-031-70797-1_5 arXiv:2406.19847 1457

1400 [33] Moses Openja, Forough Majidi, Fouts Khomh, Bhagya Chembakottu, and Heng 1458

1401 Li. 2022. Studying the Practices of Deploying Machine Learning Projects on 1459

1402 Docker. In *Proceedings of the 26th International Conference on Evaluation and 1460*

1403 Assessment in Software Engineering (EASE '22). Association for Computing 1461

1404 Machinery, New York, NY, USA, 190–200. <https://doi.org/10.1145/3530019.3530039> 1462

1405 [34] Alessandro Palladini. 2022. *Streamline machine learning projects to production 1463*

1406 using cutting-edge MLOps best practices on AWS. laurea. Politecnico di Torino. 1464

1407 <https://webthesis.biblio.polito.it/22607/> 1465

1408 [35] Productdock d.o.o., Nataša Radaković, Ivana Šenk, University of Novi Sad, Faculty 1466

1409 of Technical Sciences, Nina Romanić, and Productdock d.o.o. 2023. A MACHINE 1467

1410 LEARNING PIPELINE IMPLEMENTATION USING MLOPS AND GITOPS PRIN- 1468

1411 CIPLES. In *19th International Scientific Conference on Industrial Systems*. Faculty 1469

1412 of Technical Sciences, 94–99. https://doi.org/10.24867/IS-2023-T2.1-6_08141 1470

1413 [36] Katja-Mari Ratilainen. 2023. Adopting Machine Learning Pipeline in Existing 1471

1414 Environment. (2023). 1472

1415 [37] Gilberto Recupito, Fabiano Pecorelli, Gemma Catolino, Sergio Moreschini, 1473

1416 Dario Di Nucci, Fabio Palomba, and Damian A. Tamburri. 2022. A Multivo- 1474

1417 lutional Literature Review of MLOps Tools and Features. In *2022 48th Euromicro 1475*

1418 Conference on Software Engineering and Advanced Applications (SEAA). 84–91. 1476

1419 <https://doi.org/10.1109/SEAA56994.2022.00021> 1477

1420 [38] Philipp Ruf, Manav Madan, Christoph Reich, and Djaffar Ould-Abdeslam. 2021. 1478

1421 Demystifying MLOps and Presenting a Recipe for the Selection of Open-Source 1479

1422 Tools. *Applied Sciences* 11, 19 (Jan. 2021), 8861. <https://doi.org/10.3390/app11198861> Number: 19 Publisher: Multidisciplinary Digital Publishing Institute. 1480

1423 [39] Enrico Salvucci. 2021. MLOps-Standardizing the Machine Learning Workflow. 1481

1424 (2021). <https://amslaurea.unibo.it/id/eprint/23645/> 1482

1425 [40] Luca Scotton. 2021. Engineering framework for scalable machine learning 1483

1426 operations. (2021). <https://aaltodoc.aalto.fi/items/a1497a44-1c3a-46bf-b76a-c7cba635462c> 1484

1427 [41] Ladson Gomes Silva. 2022. A Review on How Machine Learning Operations 1485

1428 (MLOps) Are Changing the Landscape of Machine Learning Development for 1486

1429 Production. (2022). 1487

1430 [42] Afonso Rafael Carvalho Sousa. 2022. Orchestrator selection process for cloud- 1488

1431 native machine learning experimentation. (2022). 1489

1432 [43] Matteo Testi. 2024. Machine Learning Operations (MLOps) in Healthcare. (2024). 1490

1433 <https://www.iris.unicampus.it/handle/20.500.12610/83683> Publisher: Università 1491

1434 Campus Bio-Medico. 1492

1435 [44] T Vishwanbari and Sonali Agrawal. 2023. Integration of Open-Source Machine 1493

1436 Learning Operations Tools into a Single Framework. In *2023 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)*. 335–340. 1494

1437 <https://doi.org/10.1109/ICCCIS60361.2023.10425558> 1495

1438 [45] Samar Wazir, Gautam Siddharth Kashyap, and Parag Saxena. 2023. MLOps: A 1496

1439 Review. <https://doi.org/10.48550/arXiv.2308.10908> arXiv:2308.10908 [cs]. 1497

1440 [46] Claes Wohlin. 2014. Guidelines for snowballing in systematic literature studies 1498

1441 and a replication in software engineering. In *Proceedings of the 18th International 1499*

1442 Conference on Evaluation and Assessment in Software Engineering. ACM, London 1500

1443 England United Kingdom, 1–10. <https://doi.org/10.1145/2601248.2601268> 1501

1444 [47] Ting Chun Yau. 2023. *Investigate the challenges and opportunities of MLOps*. 1502

1445 <https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-324011> 1503

1446 [48] Mohammad Zarour, Hamza Alzabut, and Khalid T. Al-Sarayreh. 2025. MLOps 1504

1447 best practices, challenges and maturity models: A systematic literature review. 1505

1448 *Information and Software Technology* 183 (July 2025), 107733. <https://doi.org/10.1016/j.infsof.2025.107733> 1506

1449 [49] Yue Zhou, Yue Yu, and Bo Ding. 2020. Towards mllops: A case study of ml 1507

1450 pipeline platform. In *2020 International conference on artificial intelligence and 1508*

1451 computer engineering (ICAICE). IEEE, 494–500. https://ieeexplore.ieee.org/abstract/document/9361315/?casa_token=gVu7NAHT9ekAAAAAA:drRrkR3bXuCbJBAJNB3M95gMr9rn64LekVjhXlh2E8M5hgbYSzM5HaF7ysjw0VLIO3K-ilzXVPW6M8 1509

1452 [50] Iago Águila Cifuentes. 2023. *Design and Development of an MLOps Framework*. 1510

1453 Master's thesis. Universitat Politècnica de Catalunya. <https://upcommons.upc.edu/handle/2117/395348> Accepted: 2023-10-25T10:32:47Z. 1511

Table 3: Summary of MLOps Tools: Infrastructure and Supporting Service (Orchestrators)

Tool	Benefits	Limitations
Apache Airflow	<p>B1 Open source software licensed under Apache License 2.0 [50]</p> <p>B2 Provides an intuitive web interface for visualising and monitoring workflows [50]</p>	<p>L1 Installation and configuration can be complex in real-world environments [50]</p> <p>L2 The recommended installation method is complicated in real-world situations [50]</p>
Kubeflow Pipelines	<p>B3 Cloud-agnostic architecture enabling it to be executed on any cloud provider that supports Kubernetes [17] [24] [20] with distributions dedicated for major cloud providers [23].</p> <p>B4 More mature and widespread than its competitors and more specialised for ML when compared with Flyte, Apache Airflow, and so on [20].</p> <p>B5 Abstracts away the complexity of dynamically scaling workloads up or down via the Kubernetes engine. scale [13].</p> <p>B6 Full end-to-end MLOps solution (Kubeflow Notebooks, Kubeflow Pipelines, Katib) that is highly customisable, through either KServe or Seldon Core [23] and with seamless TFX integration [49].</p> <p>B7 Offers an easily accessible and configurable Kubeflow UI dashboard [23] [12].</p> <p>B8 Strong security (multi-user isolation) [23].</p>	<p>L3 Difficult to set up with a non-trivial learning curve. [23][42].</p> <p>L4 "Everything in one package" and hence can feel bloated compared with small, focused solutions aimed at solving specific pain points [42].</p> <p>L5 Out-of-date documentation for many Kubeflow features [23].</p>
ZenML	<p>B9 Simple local development (configured in three simple commands) [5].</p> <p>B10 Automatic deployment of a local MLflow tracking server [5].</p> <p>B11 Easily expandable to support different orchestration components [5].</p>	<p>L6 Incomplete user management features [5].</p>
Argo Workflow	<p>B12 Workflow definitions enable straightforward artefact storage and transfer between tasks.</p> <p>B13 Supports fully automated, reproducible, and scalable end-to-end model training, evaluation, and deployment without manual intervention.</p>	<p>L7 Requires developers to manage containerisation complexity.</p> <p>L8 Pipeline development and management demand in-depth knowledge of Kubernetes and associated tooling.</p>

1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633

Table 4: Summary of MLOps Tools: Managed End-to-End Platforms

Tool	Benefits	Limitations	
AzureML	<p>B14 Simple data imports and code-free drag-and-drop tools for data cleaning and transformation [29].</p> <p>B15 Quick and easy to set up for testing, orchestration, and robust security (network protection, RBAC) [29].</p> <p>B16 Fully complete security features. Only Microsoft Azure ML provides the ability for users to set up network and data protection policies as well as built-in RBAC features [29].</p>	<p>L9 Commercial product requiring an Azure subscription [29].</p>	1693 1694 1695 1696 1697 1698 1699 1700 1701 1702 1703 1704 1705 1706 1707 1708 1709 1710 1711 1712 1713 1714 1715 1716 1717 1718 1719 1720 1721 1722 1723 1724 1725 1726 1727 1728 1729 1730 1731 1732 1733 1734 1735 1736 1737 1738 1739 1740 1741 1742 1743 1744 1745 1746 1747 1748 1749 1750 1751 1752 1753 1754 1755 1756 1757 1758 1759 1760 1761 1762 1763 1764 1765 1766 1767 1768 1769 1770 1771
SageMaker	<p>B17 Robust and intuitive, offering feature parity with Azure ML [29]. Includes a good set of tools with its own model registry, a feature store and lineage tracking logic. Handles step creations and management as well as automates model deployment with CI/CD [34].</p> <p>B18 Facilitates easy scalable training and deployment that reduces operation overhead [6].</p> <p>B19 Seamless integration within the AWS ecosystem [34] such as the ability to integrate with AWS Single Sign-On (SSO) to manage identities and access [29].</p> <p>B20 Provides abstracted tools for model tuning [34].</p>	<p>L10 Requires a commercial licence [47].</p> <p>L11 Requires programming skills and has less comprehensive documentation [47].</p> <p>L12 If an organisation does not already use AWS IAM or SSO, additional adoption and migration work is required [29].</p>	1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 1674 1675 1676 1677 1678 1679 1680 1681 1682 1683 1684 1685 1686 1687 1688 1689 1690 1691 1692
Google Vertex AI	<p>B21 Offers AutoML, integrates Jupyter Notebook in Vertex AI Workbench and provides its own model registry and pipelines. Integrates monitoring and logging capabilities [41].</p> <p>B22 Easy to use central UI dashboard with visualised workflows where users can monitor the progress, review the execution history, and understand the dependencies between different components of the pipeline [8].</p> <p>B23 Flexible scalability and easy to use Kubernetes cluster with a choice of machine hardware [8].</p>	<p>L13 Dataset preparation tasks are considered the most "painful", with large multi-step coding tasks without a visual drag-and-drop interface when compared to Azure and AWS Sagemaker [29].</p> <p>L14 Lacks robust security control measures (IAM and workspace separation) [29].</p> <p>L15 Operates on a commercial, pay-as-you-go model.</p>	1736 1737 1738 1739 1740 1741 1742 1743 1744 1745 1746 1747 1748 1749 1750 1751 1752 1753 1754 1755 1756 1757 1758 1759 1760 1761 1762 1763 1764 1765 1766 1767 1768 1769 1770 1771

Table 5: Summary of MLOps Tools: End-to-End Platforms

Tool	Benefits	Limitations
MLflow	<p>B24 User-friendly with a UI providing visualisation capabilities for easy data interpretation and analysis, along with an intuitive web-based dashboard [31][36][25].</p> <p>B25 Completely open-source and self-hosting capabilities while natively supporting container deployments [5][40].</p> <p>B26 Strong vitality, since it is a well-established tool used by many world-renowned companies, backed by strong community support [5][11].</p> <p>B27 Offers comprehensive and extensive documentation [11].</p> <p>B28 Comprehensive Experiment Tracking with robust logging and visualisation of metrics and artefacts, which allows users to trace models to their training rounds via MLflow APIs/UI and to track algorithms, hyperparameters, dataset versions, and feature selections [31][4][49]. It is MLflow's key feature and module.</p> <p>B29 Simplifies performance evaluation with enhanced model performance monitoring. Moreover, it features autologging and seamlessly integrates with other monitoring tools for efficient model improvement [5][9][15][31].</p> <p>B30 Flexible data storage options with support for popular cloud storage services (Amazon S3, Azure Blob Storage, Google Cloud Storage) as well as on-premise or hybrid options (SFTP, NFS), local files, SQLAlchemy-compatible databases, or remote tracking servers [36].</p>	<p>L16 Requires a dedicated server or additional web service for collaboration, which complicates initial setup and maintenance [4][40].</p> <p>L17 Does not automatically reproduce data versions, necessitating manual intervention for consistency [36][4]. Relies on the user using other tools.</p> <p>L18 Lacks in-built alerting based on monitoring for insufficient resources [40].</p> <p>L19 Does not offer role-based access or user isolation, allowing unrestricted changes to experiments by any user. [5][36].</p> <p>L20 Lacks native collaboration features and automated deployment tools [36].</p> <p>L21 Certain advanced features or enterprise use cases require a commercial licence [47].</p>
Neptune	<p>B31 Easy setup with a clear guide [47] and integrates with tools like Google Colab, Git, and Docker [43].</p> <p>B32 Complete MLOps platform: monitoring, data versioning, and testing in one solution [43].</p> <p>B33 Open-source [47].</p> <p>B34 Can be used for monitoring [43].</p>	<p>L22 Requires basic programming skills [43].</p>
Pachyderm	<p>B35 Provides GUI-accessible services for notebook experimentation, experiment tracking, and hyperparameter optimisation [23].</p> <p>B36 Simple to parallelise experiments [23].</p>	<p>L23 Requires knowledge of, and manual configuration for, cloud storage and Helm deployments [23].</p> <p>L24 Entails a computationally heavy file system overhead compared with Kubeflow and Polyaxon [23].</p> <p>L25 Lacks "out-of-the-box" solutions.</p>
Polyaxon	<p>B37 Requires only basic Helm and Kubernetes operator knowledge [23].</p> <p>B38 Provides a meticulous audit trail that enhances reproducibility [23].</p>	<p>L26 Does not provide out-of-the-box scaling [24].</p> <p>L27 Uncertain open-source vitality due to a weaker community compared to Pachyderm and Kubeflow [23].</p>
TFX	<p>B39 Automates regular retraining, evaluation, and deployment [8].</p> <p>B40 Supports distributed computing for large workloads [8].</p> <p>B41 Portable and can be run on various orchestration platforms such as Apache Airflow [8].</p>	<p>L28 Requires TensorFlow, which might be limiting for non-TF use cases [20].</p>

Table 6: Summary of MLOps Tools: Storage & Versioning

Tool	Benefits	Limitations
Dagshub	<p>B42 Facilitates collaboration and versioning [31].</p> <p>B43 Circumvents GitHub's size limitations and integrates with DVC and MLflow [31][15].</p> <p>B44 Offers a free integrated MLflow server and unified storage for data and metadata [2][15].</p>	--
DVC	<p>B45 Easy to use as its DVC workflow is similar to Git's. This helps ease its adoption among users familiar with Git [4].</p> <p>B46 Modularises workflows into stages, ensuring reproducibility and reducing manual errors [4].</p> <p>B47 Pipeline data can be automatically pulled from a DagsHub repository so that the entire process can be run using only a single command, thus simplifying workflow execution [33].</p> <p>B48 Caching for unchanged stages saves time and minimises overhead [35].</p> <p>B49 Flexible storage options (e.g., Google Drive, HTTP, S3) without requiring a dedicated server or GitLFSserver [4].</p>	<p>L29 Requires manual updates for externally stored data [4].</p> <p>L30 File-level versioning can lead to extensive storage use in environments with frequent file changes and is not suited for versioning SQL databases [4].</p>
GTO	<p>B50 Functions as a GitOps-based model registry, removing the need for separate databases or servers by leveraging Git and DVC [35].</p> <p>B51 Supports promotion of models to specific stages, facilitating deployment across designated environments [35].</p> <p>B52 Seamlessly integrates with Iterative products such as DVC and CML [35].</p>	--
MinIO	<p>B53 S3-compatible storage solution optimised for AI workloads [5].</p> <p>B54 Used for storing artefacts and metadata generated by ZenML [43] and MLflow [4], and suitable for general data.</p>	--
Weights & Biases	<p>B55 Real-time experiment tracking with built-in hardware usage monitoring [5].</p> <p>B56 Free tier offers unlimited experiment runs and 100 GB of artefact storage. In addition, it includes dataset/model versioning, hyperparameter optimisation, report generation, and a model registry [5].</p>	<p>L31 Requires registration and a licence key for initial setup. The free plan is limited to personal projects [5].</p> <p>L32 Only the client application is open-source. The server-side infrastructure is proprietary, limiting self-hosting options [5].</p> <p>L33 Demands advanced programming knowledge for integration and usage [43][5].</p>
Feast	<p>B57 Manages historical and live data with offline/online feature storage [44].</p> <p>B58 Allows reuse of features across different ML projects [47].</p> <p>B59 Integrates with MLflow [44].</p> <p>B60 Open-source [47].</p>	<p>L34 Requires advanced programming skills [47].</p> <p>L35 Integration process is relatively complicated [47].</p>

Table 7: Summary of MLOps Tools: Inference (Model Serving)

Tool	Benefits	Limitations
BentoML	<p>B61 Open-source [40] and supports multiple frameworks (TF, PyTorch, Keras, XGBoost) [39].</p> <p>B62 Features automated micro-batching for better API performance and cloud-native deployment [40].</p>	<p>L36 Provides inference-only [40] and does not offer an automatic deployment of models, as KServe and Seldon Core do [24][20].</p>
KServe	<p>B63 Automatically wraps models as web services and is easily integrated with Kubeflow [24].</p> <p>B64 Provides scalable and isolated deployments [17] with support for HTTP/gRPC APIs [24].</p>	--
Seldon Core	<p>B65 Offers pre-packaged inference servers with robust Kubernetes support [11].</p> <p>B66 Supports advanced metrics tracking via Prometheus [11].</p>	--
Evidently	<p>B67 Provides model monitoring for data, target, and prediction drifts [2][44].</p> <p>B68 Predefined monitoring reports can be generated with few lines of code [2].</p> <p>B69 Enables a single monitor to calculate metrics across multiple models [24].</p> <p>B70 Integrates with Prometheus and Grafana for interactive reports, scheduled tests, and result logging [44].</p> <p>B71 Automatically logs results as artefacts in MLflow [44].</p>	<p>L37 Supports tabular data only. For image or text data, consider Alibi Detect [24].</p>

2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186

Table 8: Summary of MLOps Tools: Inference (Model Deployment)

Tool	Benefits	Limitations
Metaflow	<p>B72 Provides a framework for creating and executing data science workflows locally and scaling to the cloud with ease [40].</p> <p>B73 Rigorous checkpoint system enables great tracking and logging [40].</p> <p>B74 Fully integrated with AWS for automatic resource management and native parallelisation via AWS Batch [40].</p>	--
Streamlit	<p>B75 Simple interface for deploying ML applications to the cloud [2].</p> <p>B76 Convenient GitHub integration streamlines deployment workflows [2].</p>	<p>L38 1 GB limit for public application deployments [2].</p>
CML	<p>B77 Manages ML experiments and tracks modifications automatically [35].</p> <p>B78 Generates comprehensive reports with essential metrics and plots [35].</p> <p>B79 Reports are created and displayed directly in pull request comments, enhancing collaboration and review efficiency [35].</p>	--

Table 9: Summary of MLOps Tools: Infrastructure and Supporting Services

Tool	Benefits	Limitations
Gradio	<p>B80 Enables creation of interactive web interfaces for model evaluation, demonstration, and deployment without HTML/CSS/JS knowledge [15].</p>	--
TensorBoard	<p>B81 Visualisation toolkit for model graphs, weight/bias histograms, and training metrics [47].</p> <p>B82 Provides full exploration and visualisation functionality [47].</p> <p>B83 Open-source and integrates with multiple tools and applications [47].</p>	<p>L39 Requires familiarity with TensorBoard tooling and community support for effective use [47].</p>

2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323

Table 10: Summary of MLOps Tools: ML Training

2325	2383
2326	2384
2327	2385
2328	2386
2329	2387
2330	2388
2331	2389
2332	2390
2333	2391
2334	2392
2335	2393
2336	2394
2337	2395
2338	2396
2339	2397
2340	2398
2341	2399
2342	2400
2343	2401
2344	2402
2345	2403
2346	2404
2347	2405
2348	2406
2349	2407
2350	2408
2351	2409
2352	2410
2353	2411
2354	2412
2355	2413
2356	2414
2357	2415
2358	2416
2359	2417
2360	2418
2361	2419
2362	2420
2363	2421
2364	2422
2365	2423
2366	2424
2367	2425
2368	2426
2369	2427
2370	2428
2371	2429
2372	2430
2373	2431
2374	2432
2375	2433
2376	2434
2377	2435
2378	2436
2379	2437
2380	2438
2381	2439
2382	2440
2383	2441
2384	2442
2385	2443
2386	2444
2387	2445
2388	2446
2389	2447
2390	2448
2391	2449
2392	2450
2393	2451
2394	2452
2395	2453
2396	2454
2397	2455
2398	2456
2399	2457
2400	2458
2401	2459
2402	2460
2403	2461
2404	2462