2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #50927

A Multi-Method Comparative Analysis of Data Science Education and Industry Practice: Integrating MLOps and Operational Readiness into University Curriculum

Sakhi Aggrawal, Purdue Polytechnic Institute, Purdue University – West Lafayette

Sakhi Aggrawal is a Graduate Research Fellow in Computer and Information Technology department at Purdue University. She completed her master's degree in Business Analytics from Imperial College London and bachelor's degree in Computer and Information Technology and Organizational Leadership from Purdue University. She worked in industry for several years with her latest jobs being as project manager at Google and Microsoft. Her current research focuses on integrating project management processes in undergraduate education. Her main goal is to understand how work management and product development practices widely used in industry can be modified and adapted to streamline undergraduate STEM education.

Revathi Kadari, Purdue Polytechnic Institute, Purdue University – West Lafayette

A Multi-Method Comparative Analysis of Data Science Education and Industry Practice: Integrating MLOps and Operational Readiness into University Curriculum

1. Introduction

Data science (DS) and machine learning (ML) have become foundational to decision-making and automation across a wide range of industries [1]. As organizations increasingly deploy data-driven systems at scale, expectations for DS and ML professionals have expanded beyond model development to encompass operational responsibilities for deployment, monitoring, maintenance, and collaboration in production environments [2]. These expectations are reflected in emerging industry roles such as ML engineer and MLOps (Machine Learning Operations) specialist, which emphasize operational and engineering competencies alongside analytical expertise. "MLOps" is an umbrella term used for the development of ML-enabled systems, the principles of which are "rooted in Software Engineering and inspired by DevOps" [3]. In response to workforce demand, universities in the United States have rapidly expanded undergraduate and graduate programs in data science and machine learning. These programs have been successful in providing students with strong foundations in statistics, programming, and predictive modeling [1]. However, concerns have emerged regarding graduates' readiness to work with production-level data and ML systems [1]. This exploratory study examines the emerging gap in educational practice through the lens of MLOps and operational readiness. By focusing on competencies beyond modelling, the study aims to contribute to ongoing discussions in data science education on how curricula can better align with contemporary professional practice.

2. Literature Review

The academic discipline of data science has its roots in statistics, computer science, and domain-specific analytics. Early curricular models emphasized mathematical foundations, statistical inference, and algorithmic development, reflecting both the field's historical origins and the structure of traditional academic departments [4], [5], [6]. As a result, Data Science and Machine Learning programs have tended to prioritize theoretical understanding and individual analytical proficiency [1], [7]. Over the past decade, industry practice has evolved rapidly. Machine learning systems are now commonly embedded within larger software and data infrastructures, requiring coordinated workflows for data ingestion, model training, validation, deployment, and monitoring [8], [9]. These workflows have given rise to Machine Learning Operations (MLOps), a set of practices and tools designed to support the reliable and scalable operation of ML systems in production [10]. While MLOps has become increasingly central in industry, its integration into formal university curricula remains uneven [11].

Educational research has begun to document this misalignment, noting that many programs emphasize modeling outcomes without sufficient attention to operational context [11]. This background motivates a closer examination of how DS and ML curricula represent the full data and model lifecycle, and whether students are being adequately prepared for the realities of professional practice [12], [13]. Early research on data science education has identified a core set of competencies commonly addressed in university curricula, including statistical reasoning, programming, data management, and machine learning algorithms [1]. These competencies align well with traditional academic assessment structures and are supported by established pedagogical approaches. However, scholars have begun to surface that these frameworks often underrepresent skills related to system integration, collaboration, and lifecycle management [11]. However, the research focused on understanding the gap is limited and is not comprehensive [14], [15]. Additionally, the impact of this gap on industry is unknown within the existing academic literature and is primarily captured by only the companies and employers [16].

Recent literature has begun to explore MLOps as a distinct domain of practice, emphasizing its role in addressing technical debt, reliability, and governance in ML systems [17]. While professional training programs and industry-led certifications increasingly incorporate MLOps concepts, their presence in accredited degree programs remains limited and expensive [12], [18]. Educational researchers have argued that introducing operational practices through project-based learning and industry-aligned tools may help bridge this gap, but systematic evidence on implementation strategies and outcomes is still sparse [19]. Collectively, the literature points to a consistent pattern that data science education excels at developing analytical and modeling skills but provides limited exposure to the operational realities of deploying and maintaining ML systems. This synthesis highlights the need for empirical studies that examine curricula, instructional practices, and stakeholder perspectives in an integrated manner.

To address this gap, this exploratory study takes multiple perspectives of hiring managers and new graduates for industry readiness, along with conducting gap analysis of job descriptions of relevant roles across industries and comparing to the syllabi from the Data Science and Machine Learning programs. Specifically, the research questions this study attempts to answer: *How do academic approaches to teaching data science and machine learning align with industry expectations for deployable, maintainable, and ethical data and model systems?* To address this, the study adopts a multi-method design that integrates survey data from new graduates and hiring managers, along with a systematic analysis of instructional artifacts and job postings.

3. Methods

3.1 Study Design

This exploratory study employs a multi-method triangulation design to assess the alignment between academic preparation and industry requirements. By integrating distinct data streams, subjective perceptions from surveys and objective data from artifact analysis, the study aims to minimize the bias inherent in any single data source.

The design utilizes a convergent parallel approach, where data from three distinct sources is collected simultaneously and compared: (1) *Industry Perspective:* Assessing the expectations and satisfaction of hiring managers. (2) *Learner Perspective:* Evaluating the confidence and reported gaps of recent graduates. (3) *Curricular & Market Artifacts:* Systematically comparing the content of university syllabi against the requirements of active job postings. This triangulation allows the research to identify gaps that are not limited to the subjective experiences of graduates but are structurally present in educational materials and demanded by the market.

3.2. Context and Participants

To evaluate the alignment between academic preparation and industry needs, this study engaged two primary participant groups: (1) *The Hiring Manager* cohort (n=7) consisted of industry professionals responsible for recruiting and managing data science and machine learning talent. Participants were selected based on their experience interviewing or supervising entry-level roles in the field. (2) *The Early-Career Graduate* cohort (n=6) comprised recent graduates currently employed in or seeking Data Science (DS) or Machine Learning (ML) roles. Demographic data indicated a gender distribution of 5 male and 1 female participants. In terms of professional experience, one participant was a new graduate (0 years), one reported less than 1 year of experience, and four reported between 1 and 3 years of relevant industry experience.

3.2 Data Collection

The study utilized a multi-method approach to gather primary perception data and secondary documentary evidence.

1. *Surveys:* Primary data was collected using two Qualtrics surveys developed to capture the perspectives of the respective participant groups. Refer to Table 1 and 2 for specific questions.
2. *Job Description Sourcing:* To represent current market demand, job postings were collected from major employment portals (e.g., LinkedIn, Company Careers pages). The search strategy targeted *entry-level* roles with titles restricted to ‘Machine Learning Engineer’ or ‘MLOps Engineer’. Five distinct postings were selected from diverse industries including Defense, Healthcare, Finance, and Semiconductor Manufacturing to

ensure the identified skills were transferable across sectors rather than domain-specific. Refer to Appendix A for job listings used in the study.

3. *Syllabus Selection:* To represent the academic landscape, a purposive sample of Data Science and Machine Learning curricula was gathered from U.S.-based institutions. The selection criteria prioritized diversity in program type, including traditional Master of Science degrees, online professional programs, and undergraduate minors from both public and private research universities. Course descriptions, learning outcomes, and publicly available syllabi were retrieved from university registrars and department websites. A total of five university courses were reviewed for this study. Refer to Appendix B for syllabi collected.

Table 1. Survey Questions for Hiring Managers

Question	Type
How prepared are the graduates regarding MLOps and deployment (e.g., containerization, cloud services, API endpoints)?	5-point Likert Scale: (1: Terrible, 5: Excellent)
How proficient are the graduates in using version control (Git) and collaboration tools in a team environment?	5-point Likert Scale: (1: Extremely incompetent, 5: Extremely competent)
Based on interview discussions, how realistic and complex are the projects typically presented by the graduates in the Data Science & ML program?	5-point Likert Scale: (1: Purely academic, 5: Extremely industry-relevant)
Compared to other entry-level candidates you interview, how would you rate the overall competitiveness of graduates from the Data Science & ML program?	5-point Likert Scale: (1: Terrible, 5: Excellent)
If you could add any courses/modules to this program to improve graduate readiness, what would they focus on?	Open text
How much training time is required on average, for entry-level graduates to be fully productive in their work?	Open text

Table 2. Survey Questions for Early-Career Graduates

Question	Type
Gender	Male/Female/Non-binary/Prefer not to say
Are you currently employed? If yes, how long are you employed in an ML-specific role?	Never/Less than 1 year/1-3 years
Rate your level of confidence using industry-standard tools (e.g., Git, Docker, cloud services like AWS/Azure/GCP, MLflow).	5-point Likert Scale: (1: Very low confidence, 5: Very high confidence)

How realistic were the data sizes and messiness of the datasets used in assignments compared to what you expect in a professional setting?	5-point Likert Scale: (1: Not realistic, 5: Very realistic)
How well did the school projects cover the entire ML life cycle from data acquisition to model deployment and monitoring?	5-point Likert Scale: (1: Not well at all, 5: Extremely well)
How confident are you that the curriculum has prepared you for an entry-level Data Scientist or ML Engineer role?	5-point Likert Scale: (1: Very unprepared, 5: Highly prepared)
Based on your job search or internships, what are the skills and knowledge gaps between this curriculum and current industry expectations? List the highest or most important missing skill	Open text
How much time does it take to start a job to get ready to be productive in your work. In other words, how much training time is required on average?	Open text

3.3 Data Analysis

The analysis employed a convergent approach to integrate quantitative survey ratings with qualitative content analysis. Quantitative data from the Likert-scale survey questions were analyzed using descriptive statistics to identify central tendencies regarding graduate preparedness (refer to Appendix C for statistics tables). Due to the sample size, this statistical view provided a baseline sentiment which was then contextualized by the qualitative data.

Qualitative data including open-ended survey responses, course syllabi, and job descriptions underwent comparative content analysis. Two independent coders analyzed the job descriptions to extract explicit technical and professional requirements. Simultaneously, they reviewed the university syllabi to catalog the learning outcomes, course topics, and project deliverables explicitly taught in the sampled programs. Then the two were cross-mapped to categorize competencies into three distinct groups: (1) *Aligned Competencies* (skills present in both datasets), (2) *Curricular Gaps* (high-demand industry skills absent in syllabi), and (3) *Curricular Surplus* (academic topics emphasized in coursework but absent from entry-level job descriptions). Discrepancies in coding were resolved through discussion to achieve consensus, resulting in a finalized alignment matrix that visualizes the relationship between academic focus and industry reality (refer to Appendix D for codebooks).

3.4 Ethics and Trustworthiness

To ensure ethical integrity, all participants were provided with a clear statement of purpose and provided informed consent prior to participation. No personally identifiable information (PII) was collected to ensure the anonymity of respondents. To ensure the trustworthiness and validity of the findings, the study relied on data triangulation and inter-coder consensus. Triangulation was achieved by cross-verifying the subjective perceptions of hiring managers against the objective requirements found in job descriptions. Furthermore, the use of two independent coders for the artifact analysis reduced individual bias in interpreting the syllabi and job requirements, ensuring that the identified gaps were based on consistent evidence rather than a single researcher's interpretation.

4. Results

4.1 Stakeholder Perceptions (Survey Results)

The survey data highlights a distinct disparity in perceived readiness between the hiring managers and the graduates themselves, particularly regarding operational technical skills. Hiring managers (n=7) consistently rated graduate preparedness in operational domains as low. When asked to evaluate graduates on MLOps and deployment skills (e.g., containerization, cloud services), 71% (n=5) of respondents rated them as ‘Poor’, with one respondent rating them as ‘Terrible’. Only one manager rated graduates as ‘Average’, and none selected ‘Good’ or ‘Excellent’. A similar trend appeared regarding version control (Git) and collaboration tools. The majority of managers (71%) classified graduates as ‘Incompetent’ (either ‘Extremely’ or ‘Somewhat’) in a team environment. When assessing the nature of academic projects presented by candidates, the consensus was that they were only ‘Somewhat industry-relevant’. Qualitative feedback from managers suggested that to improve readiness, programs should focus on *“MLOps and scalable deployment environments”* and *“real-world data scenarios*”. Notably, hiring managers estimated that the time required for an entry-level graduate to become fully productive ranges from 6 months to 1 year.

Recent graduates (n=6) reported mixed confidence levels regarding their technical training. While two respondents expressed ‘Mostly high’ confidence, a significant portion (66%) reported ‘Very low’ to ‘Medium’ confidence in using industry-standard tools like Docker, AWS/Azure, and MLflow. Regarding the realism of their education, graduates largely felt their academic datasets did not match professional complexity. Responses regarding data realism were predominantly ‘Not realistic’ or ‘Moderately realistic’, with no respondents selecting ‘Very realistic’. Furthermore, when asked how well their school projects covered the entire ML life cycle (from acquisition to monitoring), results were polarized; while one student felt it was covered ‘Very well’, others rated coverage as ‘Slightly well’ or ‘Not well at all’. In terms of onboarding, graduates were generally more optimistic than managers, with most estimating their time to productivity between 2 and 6 months.

4.2 Comparative Content Analysis (Artifacts)

Analysis of job descriptions and university syllabi revealed a systemic gap between the competencies demanded by entry-level roles and those explicitly prioritized in academic curricula. The analysis of entry-level Machine Learning Engineer and MLOps job postings identified a strong emphasis on the post-model operational lifecycle. Across the sampled industries, job descriptions consistently listed the following required competencies: (1) *Deployment & Containerization:* roles explicitly required experience with 'Docker' for packaging applications and deploying to production environments, (2) *Continuous Integration/Delivery (CI/CD):* employers demanded the application of software engineering best practices, specifically 'CI/CD pipelines' and 'automated ML pipelines', (3) *Monitoring & Observability:* a recurring requirement was the ability to *"maintain and extend monitoring solutions"* to track model performance, drift, and latency in real-time, and (4) *Cloud Infrastructure:* proficiency with cloud platforms (AWS, Azure) and infrastructure-as-code tools (Terraform) was frequently listed as a core qualification.

In contrast, the content analysis of Data Science and ML syllabi from the sampled U.S. universities showed a predominant focus on model development and algorithmic theory. The reviewed curricula emphasized: (1) *Core Modeling:* strong coverage of regression, classification, clustering, and neural networks, and (2) *Data Analysis:* focus on data management, preparation, and visualization using tools like Pandas and Spark. However, the analysis found a near-total absence of operational learning outcomes. Specifically, syllabi did not explicitly list modules on 'model deployment', 'production workflows', or 'MLOps'. While some courses covered tools like TensorFlow, they focused on implementation and tuning rather than the 'production pipeline' or 'monitoring' aspects required by industry.

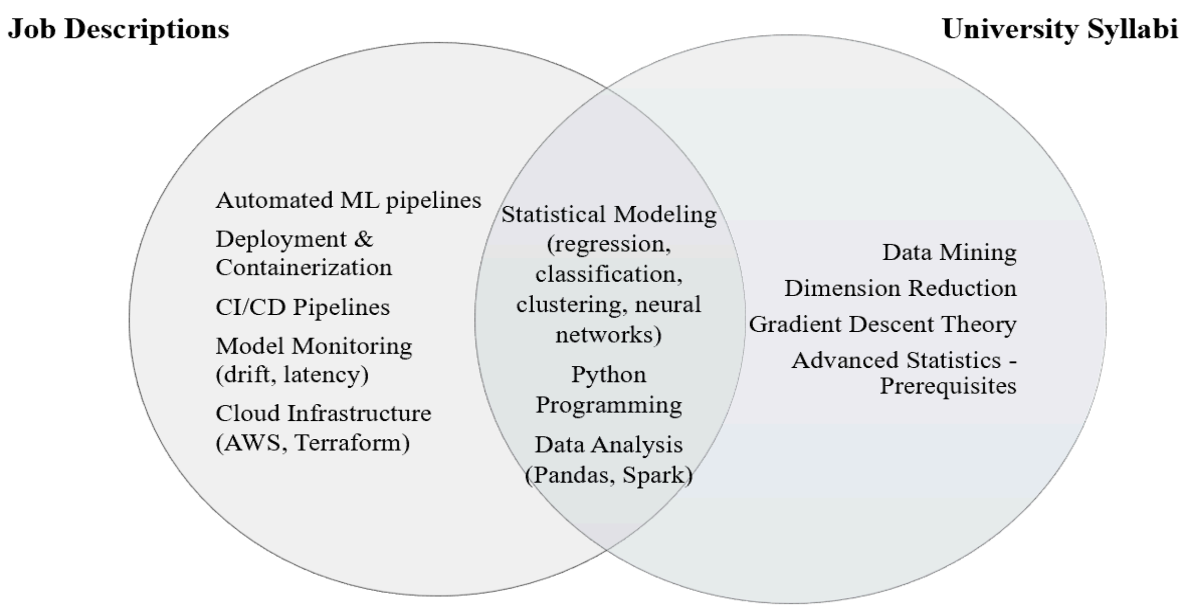

Figure 1. Comparative Analysis between Job Descriptions and University Syllabi

As shown in Figure 1, comparison of these two datasets reveals a distinct *operational gap* where academic programs successfully align with industry needs regarding 'Statistical Modeling' and

‘Basic Programming (Python)’, they diverge significantly in *Engineering & Operations* where critical skills such as ‘CI/CD’, ‘Model Monitoring’, and ‘Cloud Operations’ were present in nearly all industry job descriptions but absent from the sampled university syllabi.

5. Discussion and Implications

5.1 Discussion

This study sought to determine how academic approaches to teaching data science align with industry expectations for deployable systems. The triangulation of hiring manager perceptions, graduate confidence, and artifact analysis confirms a structural disconnect: while university curricula successfully build *analytical capacity* (modeling and statistics), they largely fail to develop *operational capability* (deployment and lifecycle management).

The most significant finding is the fundamental divergence in how ‘Data Science’ is defined by academia versus industry. The syllabus analysis reveals that university programs predominantly operate under a *Model-Centric* paradigm, where the primary learning outcome is the optimization of algorithmic performance (e.g., accuracy, F1 score). Conversely, the job description analysis and hiring manager feedback reflect a *System-Centric* paradigm, where the model is merely one component of a larger, engineered system requiring reliability, scalability, and monitoring. This misalignment explains the *Technical Proficiency Gap* observed in the results. Skills like Docker, CI/CD, and version control are not merely missing tools in the curriculum; they are artifacts of an engineering discipline that is currently absent from many Data Science programs. As evidenced by the hiring manager survey, this omission results in graduates who are viewed as ‘Incompetent’ in team-based collaborative environments, despite likely possessing strong theoretical knowledge.

The study also highlights a critical gap in project realism regarding the data lifecycle. Graduates reported that their academic projects largely failed to simulate the *messiness* and *scale* of professional data environments. By relying on pre-processed, static datasets (often provided by instructors or repositories like Kaggle), academic programs inadvertently shield students from the complexities of the full ML lifecycle - specifically data acquisition, data pre-processing, and post-deployment monitoring. This creates a disconnect where graduates are proficient in tuning models on static CSV files but unprepared for the latency, drift, and integration challenges inherent in production data streams. Notably, similar patterns of misalignment between academic preparation and professional demands have been documented in adjacent computing disciplines, where structured project management and self-regulated learning processes have been shown to improve the realism and rigor of student research experiences [20-22].

This identified skills gap imposes a tangible economic burden on the industry. Hiring managers estimated a training period of 6 to 12 months for new graduates to reach full productivity. This extended training period indicates that employers are effectively taking on the role of educators to fill critical skill gaps. Since these operational skills are now standard requirements in job descriptions, they arguably belong within the core academic curriculum rather than relying on on-the-job training. The discrepancy between graduate expectations (2–6 months) and manager reality (6–12 months) further indicates that students are entering the workforce unaware of the depth of their own operational deficiencies.

5.2 Implications

The findings of this study suggest the need for a multi-faceted approach to bridge the gap between education and practice.

1. *For University Curriculum Designers:* Curricula must evolve to treat Machine Learning as an engineering discipline rather than solely a mathematical one. This implies moving beyond ‘Algorithms’ courses to include ‘ML Systems’ modules that cover version control, containerization, and continuous deployment. Capstone projects should be restructured to require end-to-end delivery, grading not just on model accuracy, but on the reproducibility and deployability of the solution [23]. Additionally, AI-assisted tools can be embedded into existing curricula to provide low tech, personalized and scalable learning experience regarding new topics [24].
2. *For Students and Graduates:* Students must recognize that model development is only a fraction of the professional role. To improve employability, they should independently seek exposure to MLOps tools (e.g., Docker, Git, Cloud Platforms) and prioritize portfolio projects that demonstrate full-lifecycle management rather than just notebook-based analysis. Additionally, developing reflective and critical thinking practices alongside technical skills has been shown to strengthen graduates' ability to navigate complex professional environments [25].
3. *For Employers and Industry:* The industry should actively partner with academic institutions to provide real-world problem sets and anonymized, messy datasets that reflect actual business challenges [26]. Furthermore, internship programs should be explicitly designed to expose students to production workflows, helping to mitigate the training burden prior to full-time employment [27]. Research on structured mentoring models in engineering education has demonstrated that such partnerships can effectively accelerate the development of both technical and professional competencies in students [28].

6. Conclusion, Limitations and Future Work

As Data Science matures from an experimental discipline into a core industrial function, the educational requirements for entry-level professionals are shifting. This study utilized a multi-method triangulation design to assess whether current university curricula align with this evolution. The synthesis of hiring manager perceptions, early-career graduate feedback, and artifact analysis reveals a consistent, structural disconnect: while academic programs excel at producing graduates proficient in algorithmic theory and statistical analysis, the industry increasingly demands professionals capable of end-to-end lifecycle management. The data indicates that this disconnect is driven by conflicting priorities. Job descriptions and hiring managers prioritize operational competencies such as containerization, CI/CD, and monitoring that ensure models function reliably in production. Conversely, the reviewed university syllabi remain rooted in a modeling-focused approach, often treating deployment as a secondary concern. This misalignment results in a significant training period of 6 to 12 months, during which employers must train graduates in the fundamental engineering practices required for their roles. To align with professional reality, Data Science education must integrate MLOps not as an elective specialization, but as a foundational pillar of the curriculum.

The primary limitation of this study is the small sample size of both the survey cohorts and the analyzed artifacts. This small sample prevents statistical generalization of the findings and suggests the results should be interpreted as exploratory indicators of a broader trend rather than definitive population metrics. Additionally, the syllabus and job description analysis were restricted to the United States market, potentially excluding global perspectives on Data Science training. Finally, the purposive sampling of syllabi may not fully capture the diversity of elective courses available to students, focusing instead on core program requirements.

Future iterations of this research will aim to significantly expand the participant pool and artifact library to achieve statistical significance. This includes analyzing a larger dataset of job descriptions to segment requirements by industry and reviewing a wider array of university curricula to identify outlier programs that may successfully integrate MLOps. A critical next step involves the design and implementation of a pilot ‘ML Systems’ curriculum module. By measuring student confidence and competency before and after such an intervention, researchers can provide empirical evidence on effective pedagogical strategies for introducing operational concepts to students. Additionally, future work should investigate how these findings extend to other dimensions of professional readiness in computing education where there are analogous education-industry alignment challenges.

References

[1] ACM Data Science Task Force, *Computing competencies for Undergraduate Data Science Curricula*, Jan. 2021. doi:10.1145/3453538

[2] P. Y. Simard, S. Amershi, D. M. Chickering, A. E. Pelton, S. Ghorashi, C. Meek, G. Ramos, J. Suh, J. Verwey, M. Wang, and J. Wernsing, "Machine teaching: A new paradigm for Building Machine Learning Systems," *arXiv.*, 2017. doi:10.48550/arXiv.1707.06742

[3] L. E. Lwakatare, I. Crnkovic, and J. Bosch, "DevOps for AI – challenges in development of AI-enabled applications," in *Proc. IEEE Conf.*, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9238323/ (accessed Jan. 19, 2026)

[4] J. C. Adams, "Creating a Balanced Data Science Program," *Proc. 51st ACM Tech. Symp. Comput. Sci. Educ.*, Feb. 2020. doi:10.1145/3328778

[5] S. C. Hicks and R. A. Irizarry, "A guide to teaching data science," *Am. Stat.*, vol. 72, no. 4, pp. 382–391, Oct. 2018. doi:10.1080/00031305.2017.1356747

[6] I. Song and Y. Zhu, "Big Data and data science: What should we teach?" *Expert Syst.*, vol. 33, no. 4, pp. 364–373, Oct. 2015. doi:10.1111/exsy.12130

[7] S. Vinay, "Data scientist competencies and skill assessment: A comprehensive framework.," *Int. J. Data Sci.Technol.,* vol. 1, no. 1, pp. 1544–1660, 2024. [Online]. Available: https://iaeme.com/Home/issue/IJDST?Volume=1&Issue=1 (accessed Feb. 15, 2026)

[8] S. Jones, "The explosive growth of Mlops Talent in the United States: An in-depth analysis," People in AI, https://www.peopleinai.com/blog/the-explosive-growth-of-mlops-talent-in-the-united-states-an-in-depth-analysis (accessed Feb. 15, 2026)

[9] G. Ramesh et al., "A comprehensive review on scaling machine learning workflows using cloud technologies and DevOps," *IEEE Access*, 2025. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=11126113

[10] L. Berberi, V. Kozlov, G. Nguyen, J. Sáinz-Pardo Díaz, A. Calatrava, G. Moltó, V. Tran, and Á. López García, "Machine learning operations landscape: Platforms and tools," *Artif. Intell. Rev.*, vol. 58, no. 6, p. 167, 2025. doi:10.1007/s10462-025-11164-3

[11] F. Lanubile, S. Martínez-Fernández, and L. Quaranta, "Teaching Mlops in higher education through project-based learning," *arXiv.org*, 2023. doi:10.48550/arXiv.2302.01048

[12] F. Lanubile, S. Martínez-Fernández, and L. Quaranta, "Training future machine learning engineers: A project-based course on MLOps," *IEEE Softw.*, vol. 41, no. 2, pp. 60–67, Mar. 2024. doi:10.1109/ms.2023.3310768

[13] G. Nadzinski, B. Gerazov, S. Zlatinov, T. Kartalov, M. M. Dimitrovska, H. Gjoreski, R. Chavdarov, Z. Kokolanski, I. Atanasov, J. Horstmann, U. Sterle, and M. Gams, "Data science and machine learning teaching practices with focus on vocational education and training," *Inform. Educ.*, vol. 22, no. 4, pp. 671–690, 2023. doi:10.15388/infedu.2023.28

[14] E. Breck, S. Cai, E. Nielsen, M. Salib, and D. Sculley, “The ML test score: A rubric for ML production readiness and technical debt reduction,” in *Proc. IEEE Int. Conf. Big Data*, 2017. doi:10.1109/BigData.2017.8258038

[15] A. Lavin, C. M. Gilligan-Lee, A. Visnjic, S. Ganju, D. Newman, S. Ganguly, D. Lange, A. G. Baydin, A. Sharma, A. Gibson, S. Zheng, E. P. Xing, C. Mattmann, J. Parr, and Y. Gal, “Technology readiness levels for machine learning systems,” *Nat. Commun.*, vol. 13, no. 1, p. 6039, 2022. doi:10.1038/s41467-022-33128-9

[16] 2020 State of Enterprise Machine Learning, https://www.coriniumintelligence.com/hubfs/0284%20CDAO%20FS/Algorithmia_2020_State_of_Enterprise_ML.pdf?hsLang=en-gb (accessed Jan. 19, 2026)

[17] D. Kreuzberger, N. Kühl, and S. Hirschl, “Machine Learning Operations (MLOps): Overview, definition, and architecture,” *IEEE Access*, vol. 11, pp. 31866–31879, 2023. doi:10.1109/access.2023.3262138

[18] “Global mlops jobs 2025: Hiring trends and career outlook,” EliteRecruitments, https://eliterecruitments.com/the-rise-of-mlops-jobs-global-hiring-trends-and-future-outlook/ (accessed Jan. 19, 2026).

[19] M. D. Beckman, M. Çetinkaya-Rundel, N. J. Horton, C. W. Rundel, A. J. Sullivan, and M. Tackett, “Implementing version control with git and GitHub as a learning objective in statistics and data science courses,” *J. Stat. Data Sci. Educ.*, vol. 29, suppl. 1, Jan. 2021. doi: 10.1080/10691898.2020.1848485

[20] A. J. Magana, T. Amuah, S. Aggrawal, and D. A. Patel, "Teamwork dynamics in the context of large-size software development courses," *Int. J. STEM Educ.*, vol. 10, no. 1, p. 57, 2023. doi: 10.1186/s40594-023-00441-2

[21] S. Aggrawal and A. J. Magana, "Teamwork conflict management training and conflict resolution practice via large language models," *Future Internet*, vol. 16, no. 5, p. 177, 2024. doi: 10.3390/fi16050177

[22] S. Aggrawal and A. J. Magana, "Undergraduate student experience with research facilitated by project management and self-regulated learning processes," in *2023 ASEE Annu. Conf. Expo.*, 2023.

[23] C. Polpanumas *et al.*, "AI Builders: Teaching Thai students to build end-to-end machine learning projects online," in *2021 IEEE Int. Conf. Eng., Technol. Educ. (TALE)*, 2021, pp. 565–572. doi: 10.1109/TALE52509.2021.9678620

[24] V. Tudor, S. Aggrawal, and C. Böhm, "Enhancing agile project management education using AI via ChatGPT in computer science undergraduate courses," in *Proc. EdMedia*, 2025, pp. 482–491.

[25] A. Jaiswal, S. Aggrawal, G. Nanda, A. J. Magana, and M. D. Ward, "Beyond data: Nurturing critical thinking through reflective practices in a data science learning community," *Cogent Educ.*, vol. 12, no. 1, p. 2581415, 2025.

[26] S. Aggrawal and P. J. Thomas, "Investigating the industry perceptions and use of AI tools in project management: Implications for educating future engineers," in *2024 ASEE Annu. Conf. Expo.*, 2024.
[27] A. Abdoli, "Integrating industry AI/ML practices into academia: Towards bridging the industry/academia gap," in *Proc. 57th ACM Tech. Symp. Comput. Sci. Educ. V. 2*, 2026, pp. 1215–1216.
[28] S. Aggrawal and A. Jaiswal, "Enhancing research and mentoring skills in engineering education through a mentoring triad model," in *2024 IEEE Frontiers Educ. Conf. (FIE)*, 2024, pp. 1–7.

Appendix A

Job Descriptions

Appendix A1

General Atomics - Machine Learning Ops Engineer

Job Summary

General Atomics (GA), and its affiliated companies, is one of the world's leading resources for high-technology systems development ranging from the nuclear fuel cycle to remotely piloted aircraft, airborne sensors, and advanced electric, electronic, wireless and laser technologies.

From concept-to-deployment, General Atomics North Point Defense, Inc. (GA-NPD), a division of General Atomics Integrated Intelligence, Inc. (GA-Intelligence), provides AI/ML-based autonomous signal processing and data dissemination solutions providing real-time actionable intelligence supporting tactical and national mission priorities. At GA-NPD, we take a tailored approach meeting our customers' unique intelligence needs.

We are seeking an MLOps engineer who will streamline the end-to-end machine learning lifecycle, from research, development and experimentation to deployment and monitoring in production environments. This role demands applying software engineering best practices, such as continuous integration (CI) and continuous delivery (CD), to machine learning systems, ensuring ML models are not just developed but are also scalable, reliable, and continuously perform well in real-world applications.

Our team of experts work closely with the end-user in the development and implementation of a defense solution meeting platform and/or site-specific requirements. We pride ourselves as a trusted Defense Industry partner and deliver top-notch services far exceeding typical industry standards.

DUTIES AND RESPONSIBILITIES:

- Package ML models in containers, i.e. Docker, and deploy to production environments.
- Design and implement ML pipelines for data ingestion, training, evaluation, and deployment.
- Setup and maintain model monitoring and logging of deployed models to track performance metrics like accuracy, latency, and resource utilization.

- Collaborate with a diverse team including data scientists to transition models from research to production, software engineers to integrate ML models into broader application architectures, and system engineers to maximize hardware resources (cpu, fpga, gpu) to optimize performance.

We recognize and appreciate the value and contributions of individuals with diverse backgrounds and experiences and welcome all qualified individuals to apply.

Job Qualifications

- Typically requires a bachelors degree in computer science, engineering, mathematics, or a related technical discipline from an accredited institution. May substitute equivalent machine learning engineer experience in lieu of education.
- Strong proficiency in Python. Experience with other languages like C++ is also valuable.
- Understanding of machine learning principles and frameworks like PyTorch (preferred), TensorFlow, etc.
- Practical experience with Docker for deployment and packaging applications.
- Experience with optimizers such as TensorRT, onnx, and openVino.
- Proficient with Linux command line environment.
- Ability to obtain and maintain DoD Security Clearance is required.

Job Category

Engineering

Experience Level

Entry-Level (0-2 years)

Source:
https://sjobs.brassring.com/TGnewUI/Search/home/HomeWithPreLoad?partnerid=25539&siteid=5313&PageType=JobDetails&jobid=5191474#jobDetails=5191474_5313

Appendix A2

N1 Health - Machine Learning Engineer

Job Summary

N1 Health is an AI Platform company that helps healthcare organizations prioritize and maximize patient interactions. They are seeking a skilled Machine Learning Engineer to work with data scientists and engineers to scale ML model development tooling and enhance model performance.

Responsibilities:

- Work closely with our machine learning and data science teams to understand their requirements and contribute to the development of scalable ML model frameworks.
- Pair with data scientists to troubleshoot model performance issues in terms of both latency and prediction quality.
- Prototype and demo new machine learning tooling functionality for data scientists
- Design, implement, and maintain infrastructure for ML model testing and deployment.
- Leverage cloud platforms and containerization technologies to ensure scalability and flexibility.
- Implement automation processes for seamless integration of ML models into the development pipeline.
- Maintain and extend monitoring solutions to track the performance of ML models in real-time.
- Ensure compliance with industry standards and regulations related to data science and ML.
- Lead the model infrastructure components and prioritization on our product roadmap

Qualifications:

Required:

- 2+ Years of Machine Learning Engineering or MLOps experience
- Bachelor's or Master's degree in Computer Science, Mathematics, or related field.
- Proven proficiency in utilizing MLOps tools (SageMaker, MLflow, etc.) to deploy, monitor, and manage models in production environments.
- Strong proficiency in Python and SQL, with plenty of familiarity using popular libraries for machine learning (e.g. scikit-learn, XGBoost, LightGBM, PyTorch) and data manipulation (e.g. Pandas, NumPy, Polars, DuckDB, Dask).

- Experience applying software engineering best practices to both greenfield and brownfield development (e.g. testing, CI/CD, containerization, observability)
- Excellent technical communication and collaboration skills, with a passion for being at the helm of challenging problems in a fast-paced environment.

Preferred:

- Prior experience working with healthcare data, predictive models, or clinical workflows is a strong plus.
- Comfortable working in cloud-based environments, with familiarity in AWS services and modern ML tooling.
- Hands-on experience with Infrastructure-as-Code (Terraform), automation, and scripting languages such as Python and Bash.
- Strong skills in data visualization and dashboarding tools (Apache Superset, Grafana) for communicating insights effectively.
- Understanding of CI/CD/CT pipelines, version control (Git), and MLOps best practices.
- Familiarity with workflow orchestration frameworks like Airflow or Prefect for managing data and ML pipelines.
- Experience in Natural Language Processing, leveraging Python libraries such as NLTK, spaCy, or Hugging Face Transformers.
- Exposure to the Go programming language for performance-oriented components.
- Knowledge of containerization technologies like Docker for building, packaging, and deploying scalable ML applications.

Company:

N1 Health is the applied AI platform that drives measurable business results for healthcare organizations. Founded in 2017, the company is headquartered in Boston, Massachusetts, USA, with a team of 51-200 employees. The company is currently Growth Stage.

Source:
https://www.linkedin.com/jobs/view/machine-learning-engineer-entry-level-at-jobright-ai-4272420220/?skipRedirect=true

Appendix A3

JP Morgan - Machine Learning Engineer

Job Description

The mission of the Digital Intelligence team is to leverage large-scale computation, substantial datasets, and machine learning to enhance our most crucial customer products across a broad spectrum. We manage a wide array of impactful products and services, reaching countless customers and households. We place high value on our customers' direct feedback and employ a truly agile methodology to incorporate improvements that enhance user experiences.

As a Machine Learning Engineer in the Digital Intelligence team, you will leverage large-scale computation, substantial datasets, and machine learning to enhance our most crucial customer products across a broad spectrum. Your expertise will promote the team to solve complex relevance and ranking problems, develop scalable features, and manage ML Ops for building innovative systems that benefit our customers across all lines of business.

Job Responsibilities:

- Oversee the analysis of complex datasets to inform decisions on real-world applications.
- Lead the development and implementation of models and algorithms to enhance existing systems, processes, and products.
- Supervise data analysis activities and ensure effective visualizations are provided.
- Ensure the writing and deployment of software code in production systems is efficient and meets standards.
- Anticipate risks associated with machine learning solutions and prediction/classification systems and strategize mitigation.
- Communicate complex issues clearly and credibly to senior management and stakeholders.
- Foster a transparent cross-functional work environment and influence peers and team members to uphold these standards.

Source:
https://jpmc.fa.oraclecloud.com/hcmUI/CandidateExperience/en/sites/CX_1001/job/210665979

Appendix A4

Analog Devices - Associate Machine Learning Operations (MLOps) Engineer

Analog Devices, Inc. (NASDAQ: ADI) is a global semiconductor leader that bridges the physical and digital worlds to enable breakthroughs at the Intelligent Edge. ADI combines analog, digital, and software technologies into solutions that help drive advancements in digitized factories, mobility, and digital healthcare, combat climate change, and reliably connect humans and the world. With revenue of more than $9 billion in FY24 and approximately 24,000 people globally, ADI ensures today's innovators stay Ahead of What's Possible™. Learn more at www.analog.com and on LinkedIn and Twitter (X).

This role is an entry-level position within the global XOps team - which includes MLOps, LLMOps, AgentOps, and DevOps - whose mission is to deliver a world-class AI/ML developer experience for our software engineers and data scientists. You will join a mission-driven interdisciplinary team that spans data science, software engineering, product management, cloud architecture, and security expertise. A team that believes in a culture founded on trust, mutual respect, growth mindsets, and an obsession for building extraordinary products with extraordinary people.

Role Summary

As an entry-level MLOps Engineer, you will start building your foundational skills while learning how core processes and tools support technical success in machine learning operations. You will independently design and optimize complete systems, resolve technical issues via systematic analysis, and apply industry best practices and advanced methodologies for continuous improvement. You'll help develop major ML/AI operational features that span multiple aspects of the ML/AI developer experience— from infrastructure to pipelines, deployment, monitoring, governance, and cost/risk optimization.

Key Responsibilities

Operational Excellence

- Foster and contribute to a culture of operational excellence: high-performance, mission-focused, interdisciplinary collaboration, trust, and shared growth.
- Help drive proactive capability and process enhancements to ensure enduring value creation, analytic compounding interest, and operational maturity of the ML/AI platform.
- Help build resilient cloud-based ML/AI operational capabilities that advance our system attributes: learnability, flexibility, extensibility, interoperability, and scalability.

ML/AI Cloud Operations & Engineering

- Assist in setting up cloud resources (e.g., EC2 instances, S3 buckets, and SageMaker environments) to support the lifecycle of ML models and services.
- Learn and apply foundational concepts of cloud architecture under the guidance of senior engineers.
- Document configuration steps and contribute to maintaining infrastructure scripts for scalability and reliability.
- Support efforts to monitor cloud resources and ML workflows by setting up basic monitoring tools or dashboards.
- Collaborate with the team to ensure compliance by following pre-defined frameworks and flagging any issues or inconsistencies.
- Gain exposure to infrastructure lifecycle management concepts, such as drift detection and provisioning.
- Assist in testing and validating ML pipelines by running test cases and capturing results for review.
- Contribute to quality assurance efforts by creating detailed documentation of testing processes and outcomes.
- Learn how GenAI/LLM-based proofs-of-concept are evaluated and assist with basic tasks, such as setting up testing environments.
- Gain hands-on experience with Kubernetes by helping manage clusters and working on tasks like deploying sample ML workflows.
- Learn about workflow orchestration tools (e.g., Argo, Kubeflow) by assisting in their setup and testing under supervision.
- Support the creation and governance of simple data pipelines using tools like Airflow.

Required Skills & Experience

- Basic understanding of the machine learning lifecycle (e.g., data preprocessing, model training, evaluation).
- Familiarity with cloud-based services (e.g., AWS, Azure, or Google Cloud).
- Exposure to infrastructure-as-code tools (e.g., Terraform, AWS CDK) and workflow orchestration tools (e.g., Airflow or Kubeflow) is a plus.
- Experience with programming languages such as Python or Bash.
- Strong communication and documentation skills.

Source:

https://analogdevices.wd1.myworkdayjobs.com/en-US/External/details/Associate-Machine-Learning-Operations--MLOps--Engineer_R258871?q=machine%20learning

Appendix A5

Cyfersoft - Junior MLOps Engineer

Job Description:

We are looking for a motivated Junior MLOps Engineer to join our MLOps & Data Engineering team. In this entry-level role, you will support the deployment, monitoring, and maintenance of machine learning models in production environments. You will work closely with senior engineers and data scientists to build scalable and reliable ML pipelines, ensuring smooth model lifecycle management.

Key Responsibilities:

- Assist in building and maintaining automated ML pipelines for training, deployment, and monitoring.
- Support the integration of machine learning models into production systems.
- Monitor model performance, detect issues, and collaborate on troubleshooting.
- Help manage infrastructure for ML workflows using cloud platforms and container orchestration tools.
- Work with data engineering teams to ensure data quality and availability for ML pipelines.
- Participate in code reviews, testing, and documentation of MLOps processes.
- Learn and apply best practices in model versioning, reproducibility, and CI/CD for ML.
- Stay up to date with emerging MLOps tools and technologies.

Requirements:

- Bachelor's degree in Computer Science, Engineering, Data Science, or related field, or equivalent practical experience.
- Basic understanding of machine learning concepts and software engineering principles.
- Familiarity with programming languages such as Python or Bash.
- Exposure to cloud platforms (AWS, Azure, GCP) and container technologies like Docker and Kubernetes is a plus.
- Knowledge of CI/CD pipelines and version control systems (e.g., Git).
- Strong problem-solving skills and eagerness to learn.
- Good communication and teamwork abilities.

Source:
https://careers.cyfersoft.com/jobs/junior-mlops-engineer

Appendix B

University Syllabi

Appendix B1

University of Minnesota – M.S. in Data Science Curriculum

- Includes core machine learning and data analysis courses
- No explicit listing of model deployment, MLOps, production workflows, or operationalization topics in the published curriculum.

Source: U of Minnesota Data Science MS curriculum (course list only) such as Machine Learning, Data Mining, Deep Learning, etc. - no deployment parts shown
https://cse.umn.edu/datascience/ms-plan-b-curriculum

Appendix B2

Boston University Online MS in Data Science (Curriculum Sequence)

- Lists foundational and advanced ML coursework
- Curriculum items shown do not explicitly include deployment, MLOps pipelines, or production monitoring topics in the published sequence.

Source: BU MSDS curriculum overview — core and elective modules, but no explicit MLOps/deployment classes listed in summary
https://www.bu.edu/cds-faculty/programs-admissions/online-msds/curriculum/

Appendix B3

University of Chicago – MS in Applied Data Science

- Focuses on ML I & II, predictive modeling, deep learning and analytics
- No explicit course or module listed on deployment, MLOps, pipeline automation, or production workflows.

Source: UChicago Applied Data Science curriculum overview (machine learning courses emphasize theory/applications, not deployment/ops)
https://applieddatascience.psd.uchicago.edu/academics/curriculum/

Appendix B4

University of Colorado Boulder – Machine Learning Course Listings

- Provides introduction to ML, unsupervised algorithms, and deep learning
- No indication this set of courses includes model deployment, production engineering, or operationalization curriculum.

Source: UColorado Boulder Machine Learning course offerings overview (intro and specialized ML classes)
https://www.colorado.edu/program/data-science/big-opportunities-in-big-data-Landing-Page#accordion-279954036-1

Appendix B5

Rice University – Data Science Tools and Models Course

- Covers tools like TensorFlow/Spark and machine learning implementation
- Syllabus does not include model deployment, production pipelines, or MLOps concepts.

Source: Course syllabus overview showing focus on tools and model implementation — not ML in production environments https://ga.rice.edu/programs-study/courses/dsci/

DSCI 302 - INTRODUCTION TO DATA SCIENCE TOOLS AND MODELS

Short Title: DATA SCIENCE TOOLS AND MODELS
Department: Data Science
Grade Mode: Standard Letter
Course Type: Lecture
Credit Hours: 3
Restrictions: Enrollment is limited to students with a major in E, E, E or Sport Analytics. Enrollment is limited to students with a minor in E, E, E or Data Science. Enrollment is limited to Undergraduate, Undergraduate Professional or Visiting Undergraduate level students. Students in the Computer Science department may not enroll.
Course Level: Undergraduate Upper-Level

Prerequisite(s): COMP 140 or DSCI 101
Description: This course introduces key concepts in data management, preparation, and modeling and provides students with hands-on experience in performing these tasks using modern tools, including relational databases, pandas, and Spark. Models covered include kNearest Neighbors, linear regression and gradient descent. For registration purposes, DSCI 101 or COMP 140 is a required prerequisite for this course. With instructor permission, students who have experience with the Python programming language may be allowed to special register for this course. Note that these students may be required to demonstrate proficiency with Python. Priority for this course is given to students enrolled in the data science minor or sport analytics major. Other students may be permitted to enroll at the discretion of the instructor. Mutually Exclusive: Cannot register for DSCI 302 if student has credit for COMP 330.

DSCI 303 - MACHINE LEARNING FOR DATA SCIENCE

Short Title: MACHINE LEARNING FOR DS
Department: Data Science
Grade Mode: Standard Letter
Course Type: Lecture
Credit Hours: 3
Restrictions: Enrollment is limited to Undergraduate, Undergraduate Professional or Visiting Undergraduate level students.
Course Level: Undergraduate Upper-Level
Prerequisite(s): DSCI 301 or STAT 310 or STAT 280 or STAT 305 or ELEC 303 or PSYC 339 or SOCI 382 or SOSC 302 or BIOE 439 or ECON 307
Description: This course is a practical introduction to machine learning, emphasizing when and how to apply techniques and how to interpret results. Topics covered include regression, classification, dimension reduction, clustering, decision trees, ensemble learning, and neural networks. Recommended Prerequisite(s): STAT 315 or DSCI 101 or COMP 140 and DSCI 302 (may be taken concurrently) or COMP 330 or COMP 430 Mutually Exclusive: Cannot register for DSCI 303 if student has credit for ELEC 478.

Appendix C

Descriptive Statistics

Appendix C1

Hiring Manager Perceptions (n=7)

			Ratings				
Question	**Mean**	**SD**	**1**	**2**	**3**	**4**	**5**
Q1: Preparedness regarding MLOps & Deployment *(Scale: 1=Terrible, 5=Excellent)*	2.00	0.58	1	5	1	0	0
Q2: Proficiency in Version Control (Git) & Collaboration *(Scale: 1=Extremely Incompetent, 5=Extremely Competent)*	2.43	1.13	1	4	0	2	0
Q3: Realism & Complexity of Academic Projects *(Scale: 1=Purely Academic, 5=Extremely Industry-Relevant)*	3.00	0.00	0	0	7	0	0
Q4: Overall Competitiveness of Graduates *(Scale: 1=Terrible, 5=Excellent)*	3.29	0.49	0	0	5	2	0

Appendix C2

Early Career Graduates Perceptions (n=6)

			Ratings				
Question	Mean	SD	1	2	3	4	5
Q1: Confidence using Industry-Standard Tools *(Scale: 1=Very Low, 5=Very High)*	2.50	1.38	2	1	1	2	0
Q2: Realism of Data Sizes & Messiness *(Scale: 1=Not Realistic, 5=Very Realistic)*	2.17	0.98	2	1	3	0	0
Q3: Coverage of Entire ML Life Cycle *(Scale: 1=Not Well at All, 5=Extremely Well)*	2.50	1.05	1	2	2	1	0
Q4: Confidence in Overall Preparation for Role *(Scale: 1=Very Unprepared, 5=Highly Prepared)*	3.83	1.17	0	1	1	2	2

Appendix D

Codebook

Appendix D1

Qualitative Codebook: Thematic Analysis of Open-Text (n=13)

This codebook synthesizes the open-text responses from both groups regarding "Missing Skills" and "Suggested Improvements."

Theme	Code	Description	Example Quotes (Data Source)
Operational & Engineering Gaps	DEPLOY_ENV	Mentions of deployment, containers, environments, or latency.	*"ML Ops and scalable deployment environment"* (Manager) *"Deployment Understanding and Tools"* (Graduate)
	E2E_FLOW	References to the full lifecycle, end-to-end pipelines, or integration.	*"Creating an end to end flow"* (Graduate) *"ML life-cycle management"* (Manager)
	TOOLING	Specific mentions of Git, Version Control, or big data tools.	*"Significant focus on version control and collaboration"* (Manager) *"Visualizations... dealing with Big Data"* (Graduate)
Business & Contextual Gaps	BIZ_ACUMEN	Needs for understanding business value,	*"Data science -> business integration"* (Manager)

		storytelling, or integration.	*“Visualizations... Storytelling”* (Graduate)
	REALISM	Need for messy data, real-world scenarios, or practical implementation.	*“Real-world data scenarios... messiness of datasets”* (Manager) *“Practical Part, Implementation matters”* (Graduate)
Curriculum vs. Reality	THEORY_FOCUS	Feedback that school focused on modeling/theory over practice.	*“School projects focused mainly on modeling"* (Graduate) *"Purely academic”* (Manager Survey Anchor)
Onboarding Burden	RAMP_UP	Estimates of time required to become productive.	*“6-12 months”* (Manager) *“Everything else I learnt via personal projects”* (Graduate)

Appendix D2

Job Descriptions & Syllabi Codebook (n=10)

Category	Keywords/Codes Searched	Found in Job Descriptions?	Found in Syllabi?
Model Operations	*Deployment, Containerization, Docker, Kubernetes, Production*	**High Frequency** (‘Package ML models in containers’)	**Absent/Low** (Only mentions of ‘implementation’)
Engineering Practices	*CI/CD, Pipelines, Git, Version Control, Testing, Automation*	**High Frequency** (‘Automated ML pipelines’, ‘CI/CD’)	**Absent** (Focus on ‘Algorithms’, ‘Derivations’)

Monitoring	*Drift, Latency, Observability, Maintenance, Logging*	**Medium Frequency** (‘Track performance metrics’, ‘Drift’)	**Absent**
Core Modeling	*Regression, Classification, Neural Networks, Pandas, Scikit*	**High Frequency**	**High Frequency** (This confirms alignment in this specific area)