AI Education in a Mirror: Challenges Faced by Academic and Industry Experts

Mahir Akgun(✉) and Hadi Hosseini

College of Information Sciences and Technology, Pennsylvania State University, University Park, PA 16802, USA
makgun@psu.edu, hadi@psu.edu

Abstract. As Artificial Intelligence (AI) technologies continue to evolve, the gap between academic AI education and real-world industry challenges remains an important area of investigation. This study provides preliminary insights into challenges AI professionals encounter in both academia and industry, based on semi-structured interviews with 14 AI experts—eight from industry and six from academia. We identify key challenges related to data quality and availability, model scalability, practical constraints, user behavior, and explainability. While both groups experience data and model adaptation difficulties, industry professionals more frequently highlight deployment constraints, resource limitations, and external dependencies, whereas academics emphasize theoretical adaptation and standardization issues. These exploratory findings suggest that AI curricula could better integrate real-world complexities and interdisciplinary learning, while recognizing the broader educational goals of building foundational and ethical reasoning skills.

Keywords: AI Education · Expert Challenges

1 Introduction

As Artificial Intelligence (AI) becomes increasingly integral to industries and everyday life, ensuring that the next generation of AI professionals is well-equipped with both theoretical knowledge and practical skills is essential. Undergraduate AI education, therefore, plays a crucial role in preparing students to meet the demands of this rapidly evolving field and addressing the critical gaps that exist between academic training and industry requirements.

Considerable efforts have been made in AI curriculum development. Since 2018, the AI4K12 Initiative has been creating national guidelines for AI education in K-12 schools, focusing on the '5 Big Ideas in AI' [1]. These guidelines outline the essential AI concepts and skills students should master at each grade level, providing a framework for curriculum developers and standards writers. Similarly, the ACM/IEEE-CS/AAAI's CS2023 guidelines emphasize core AI topics, ethical considerations, and interdisciplinary applications [2]. However, the

A. I. Cristea et al. (Eds.): AIED 2025, LNAI 15878, pp. 336–348, 2025.
https://doi.org/10.1007/978-3-031-98417-4_24

fast-paced advancements in AI technology pose considerable challenges in maintaining a comprehensive and relevant curriculum.

To ensure that undergraduate AI education not only keeps pace with technological advancements but also meets ever-changing industry needs, it is essential to understand the real-world challenges faced by AI professionals. Real-world challenges– involving data scarcity, unrealistic assumptions, and stakeholder constraints – provide critical learning opportunities for students, enabling them to develop a deeper understanding of the complexities and nuances of AI work. By identifying and analyzing these characteristics, educators can design undergraduate curricula that more accurately reflect the realities of the field, thus bridging the gap between academic learning and industry practice. This approach not only ensures that students are better prepared to tackle multifaceted challenges in their careers but also promotes the development of essential skills such as critical thinking, problem-solving, and ethical decision-making.

This study offers an exploratory examination of the challenges faced by AI experts in academia and industry, providing preliminary insights into potential gaps in AI education. Through semi-structured interviews with a limited but diverse group of experts, we identify emerging patterns and propose potential directions for better aligning AI education with real-world complexities, while acknowledging broader educational missions beyond immediate industry requirements.

2 Background

2.1 AI Education Challenges

Research in AI education has identified several key challenges in preparing students for the rapidly evolving AI landscape. These challenges span from pedagogical concerns to practical implementation issues and industry alignment.

Students often struggle with translating theoretical AI concepts into practical problem-solving skills [3,14]. These difficulties manifest in several ways: misunderstanding mathematical foundations, challenges in debugging AI models, and applying algorithmic decision-making in real-world contexts [3]. Particularly in machine learning, students exhibit misconceptions about model behavior, such as overfitting, generalization errors, and hyperparameter tuning [14], as well as fundamental misunderstandings of model-data relationships [15].

Pedagogical strategies such as scaffolding, active learning, and interdisciplinary integration have been proposed to alleviate these issues [3]. These methods aim to equip students not only with technical proficiency in AI and ML but also with a broader, contextual understanding of how AI systems operate in diverse, real-world environments [3,14].

2.2 Industry-Academia Gaps in AI Education

The disconnect between AI education and industry expectations has been a recurring theme in computing education research. Studies have analyzed how AI graduates often lack exposure to real-world deployment challenges, such as data drift, model monitoring, scalability, and ethical considerations [4,12]. Research in this area suggests that AI curricula should integrate interdisciplinary perspectives, including regulatory compliance, human-centered AI design, and ethical AI principles, to prepare students for diverse career paths [12]. Additionally, prior studies highlight the need for experiential learning, where students work on industry-relevant AI problems that reflect the complexity of real-world applications [6]. Approaches such as university-industry collaborations, capstone projects, and AI competitions have been explored as potential solutions to bridge this gap and better align academic AI training with professional requirements [5,7].

3 Method

This study employed semi-structured interviews with two distinct groups of AI experts from academia and industry. Participants answered a series of pre-defined questionnaires and open-ended questions aimed at uncovering the most challenging problems they face in their work. The responses were analyzed qualitatively using inductive content analysis, which facilitated the identification of key themes and labels characterizing common AI challenges. This analysis provided a foundation for comparing the challenges encountered by professionals and faculty.

3.1 Participants

To identify experts, we considered individuals with a degree in Computer Science, Information Sciences, Engineering, or related fields, and at least ten years of relevant experience. The criterion of ten years was chosen based on the widely accepted notion that extensive experience, often described as "10 years or 10,000 h of deliberate practice", is indicative of expertise in a given field [10].

This study brought together fourteen experts from both industry and academia to provide comprehensive insights into the challenges facing AI development and education. The participant pool included eight AI industry practitioners working across diverse sectors (P1-P8) and six tenured faculty members with significant experience in AI research and teaching (A1-A6).

The industry participants represent a diverse range of organizations, including streaming services, social media platforms, e-commerce companies, pharmaceutical technology firms, and supply chain optimization companies. These practitioners offer perspectives from the front lines of applied AI development and implementation.

The academic experts, all tenured faculty members, teach both graduate and undergraduate courses while conducting research across various AI-related areas, including language models, data mining, machine learning, multi-agent systems, and computational social science.

This diverse group of experts provides a well-rounded view of the practical challenges and educational considerations in the rapidly evolving field of artificial intelligence.

3.2 Research Design

The study is structured to elicit the major challenges from experts in the field through semi-structured interviews lasting about 60 min. Semi-structured interviews involve a verbal interchange where the interviewer asks prepared questions while allowing the conversation to unfold naturally. This format enables participants to delve into issues they find important, providing richer and more nuanced insights.

We chose semi-structured interviews with open-ended questions for several reasons. Open-ended questions allow respondents to express their thoughts and experiences in their own words, leading to more detailed and contextually rich data. This is particularly important when exploring complex and subjective topics like the challenges faced by AI experts. The open-ended nature of our questions facilitates a conversational flow, encouraging AI professionals to narrate their experiences and reflect on various aspects of their work. This approach aligns with our goal of understanding the diverse and multifaceted challenges encountered by AI professionals and academics.

The selection of the interview questions was driven by the goal of understanding not only the technical challenges faced by AI professionals but also the contextual factors that make certain problems particularly difficult or unique. The questions were designed to prompt participants to reflect on both specific experiences and the broader characteristics that distinguish routine tasks from particularly complex ones. This approach allowed us to explore the multi-dimensional nature of challenges in AI and identify patterns that might not be immediately apparent in more straightforward inquiries. The questions posed in the interviews were:

- *Can you tell us about two-three most interesting or most challenging problems/cases you encountered in the past in your career?*
- *Why do you think these cases/problems were especially interesting or challenging?*
- *Are there any characteristics of these cases that are common with respect to the challenge? If yes, what are they?*
- *Are there any other characteristics that can be used to define "challenging/tough" problems/cases?*
- *What makes these tough/challenging problems/cases different from typical/routine problems/cases?*

3.3 Data Analysis

The analysis focused on the challenges that experts have when developing and deploying AI solutions. The line-by-line reading was used as the analytical process of separating the transcribed data into constituent qualitative elements, but we also concentrated on portions of the data that were qualitatively meaningful units for signifying the challenges we aimed to identify in the study. A meaningful unit may be a line, a sentence, a paragraph, or any other entity, so we did not use a single entity as a unit of analysis in this study.

Interview recordings were transcribed verbatim and analyzed using a systematic four-phase coding process to identify key themes in AI practitioners' challenges. In the first phase, two researchers independently conducted line-by-line inductive coding of the transcripts, remaining open to emergent themes without predefined categories. The second phase involved comparing and rationalizing codes, where researchers identified instances of different terminology representing similar concepts and developed a unified coding vocabulary. In the third phase, researchers resolved all coding discrepancies through detailed discussion, achieving 100% inter-rater reliability. Finally, in the fourth phase, related codes were collaboratively grouped into broader thematic categories that captured the key challenges reported by participants.

4 Results

The analysis of interviews with AI experts from both industry and academia revealed a comprehensive set of characteristics that define the challenges they encounter in their work. These characteristics were categorized into distinct themes based on the qualitative data. Each theme represents a specific aspect of the challenges faced by AI professionals when developing and deploying AI solutions (see Table 1).

4.1 Identified Themes of AI Challenges

This section presents the findings from the study, categorized into five overarching themes, each encapsulating specific challenges identified through participant interviews. These themes highlight the difficulties faced by AI practitioners in both academia and industry, illustrating how challenges manifest across different professional environments.

Data-Related Challenges. Data issues were frequently highlighted by experts, encompassing various aspects such as data quality, availability, and imbalance. These challenges directly impact the performance and reliability of AI models. Experts often face situations where the dataset has a significantly uneven distribution of classes, making it difficult to train effective machine learning models because the model tends to be biased towards the majority class. One participant described this challenge:

> "...problems related to fraud,..., are very challenging because the data is heavily imbalanced and you don't know what kind of fraud you would face in the future."

In many enterprise settings, the user base may be small, resulting in a lack of adequate data. This scarcity makes it hard to fine-tune algorithms and make accurate predictions, which is particularly challenging when trying to deliver personalized or precise outputs. Additionally, the absence of high-quality, domain-specific data necessary for building robust models was a recurring issue. This problem is exacerbated when there is no ground truth data or domain expertise

Table 1. Codes and their descriptions categorized by themes

Codes	Description
Theme 1: Data-Related Challenges	
Imbalanced data	Situations where the dataset has a significantly uneven distribution of classes affecting model performance.
Lack of good data	Issues arising from the absence of high-quality domain-specific data necessary for robust model building.
Limited data	Challenges related to the availability of insufficient data to train models effectively.
Low-quality feedback	Scenarios where feedback from users is inconsistent or not representative of actual performance.
Theme 2: Model Adaptation and Scalability	
Difficulty in detecting new kinds of incidents	The challenge of identifying novel or evolving incidents that deviate from historical patterns.
Handling unpredictable situations and novel contexts	This label covers scenarios where AI models encounter unforeseen behaviors or unfamiliar environments.
Problems involving risk	Scenarios where decisions have significant potential consequences such as financial trading.
Scalability issues	The difficulty of scaling AI solutions from small-scale implementations to larger populations or settings.
Irregular and variable data structures	Dealing with irregular and variable data structures where relationships and connections between data points can vary greatly.
Overcoming unrealistic theoretical assumptions	The need to eliminate or adjust theoretical assumptions not feasible in real-world applications.
Domain knowledge gaps	This label highlights difficulties AI practitioners face when they lack expertise in the specific domain where a model is applied.
Theme 3: Practical Constraints and External Factors	
Internal data influenced by external factors	Situations where model accuracy is affected by external variables beyond the control of the dataset.
Constraints defined by stakeholders	Limitations and requirements set by various stakeholders that influence AI system development.
Constraints shaped by practical settings	Practical limitations encountered in real-world environments differing from theoretical research settings.
Resource and infrastructure constraints	This label encompasses limitations related to computational power, workforce availability, and financial resources.
Theme 4: User Behavior and Interaction	
Constraints defined by user action	Situations where user behavior introduces constraints that must be considered in model development.
Making incorrect assumptions about users	Scenarios where models are built based on incorrect assumptions about user behavior.
Challenges in understanding and measuring user impact	This label highlights the difficulty of predicting user responses to AI system outputs and defining appropriate long-term success metrics.
Theme 5: Trust, Explainability, and Communication	
Explainability and trust building	The necessity of making AI models transparent and understandable to build trust among stakeholders.
Gap of understanding	Communication barriers between technical and non-technical stakeholders leading to misaligned expectations.
Overcoming domain expertise resistance	Challenges in overcoming resistance from domain experts who may distrust AI models.

to guide the model development and validation process. Furthermore, feedback from users that is inconsistent or not representative of actual system performance can mislead model improvement efforts, making it difficult for experts to assess the true efficacy of their systems.

Model Adaptation and Scalability. Adapting AI models to dynamic environments and ensuring their scalability emerged as a prominent concern among experts. Many AI models struggle to generalize beyond their training data, particularly when encountering new or unforeseen data patterns. One participant highlighted this issue:

> "... not all scams follow the same patterns. Learning from the past doesn't always help because newer fraud methods constantly emerge."

Experts highlighted the need for models to handle novel contexts and unexpected scenarios, where predefined rules or past experiences do not always provide sufficient guidance. In high-risk applications, such as finance or healthcare, the consequences of incorrect predictions are significant, making robust and fail-safe model design crucial. Moreover, scalability remains a persistent challenge, with many AI models failing to perform optimally when deployed at larger scales due to increasing computational costs and data variability. One expert noted that models often fail to scale effectively due to resource limitations and unpredictable system behavior:

> "During deployment, software behaves unpredictably, both in terms of input variations and the ways models interact with real-world environments."

Additional difficulties arise from processing non-standardized data structures that demand flexible and adaptive algorithms. Finally, experts noted that theoretical assumptions often fail to align with real-world applications, leading to models that do not adequately reflect operational constraints. A lack of domain-specific knowledge further exacerbates these challenges, as AI solutions require contextual expertise to be effectively integrated into specialized fields.

Practical Constraints and External Factors. Numerous constraints and external factors influenced the work of AI experts, highlighting the importance of considering real-world limitations in AI development. Various stakeholders set limitations and requirements that influence AI system development, often leading to conflicting expectations from senior executives, project managers, and end-users, complicating the development process. Stakeholder expectations play a crucial role in shaping AI development:

> "People want solutions quickly, but the requirements keep evolving, making it difficult to define a stable approach."

Practical limitations encountered in real-world environments, such as unpredictable factors like traffic, weather, and operational constraints, differ from controlled or theoretical research settings. These constraints must be considered during model development to ensure applicability and effectiveness. Organizations face a broad range of limitations, such as shortages of people, money, and services, which necessitates efficient allocation of limited resources to maximize impact. Developing AI solutions that can operate on limited hardware resources is particularly relevant for organizations with limited budgets, posing a significant challenge for experts who need to ensure their models are both effective and resource-efficient. Model accuracy can also be affected by external variables beyond the control of the dataset, such as economic changes, seasonal trends, or competitor actions. These factors introduce variability that impacts performance, making it challenging for experts to maintain model accuracy.

User Behavior and Interaction. The unpredictable nature of user behavior and interaction with AI systems posed significant challenges for the experts. User behavior introduces constraints that must be considered in model development. For instance, users interacting with a chatbot in unexpected ways can require the system to handle off-topic or irrelevant queries effectively to maintain user engagement. Diverse and unpredictable user responses to system outputs can impact engagement and satisfaction, necessitating robust designs that can accommodate this unpredictability. Measuring the long-term impact of AI on user engagement and decision-making is another complex issue, as traditional evaluation metrics may not fully capture the evolving nature of AI-human interactions.

Models built on incorrect assumptions about user behavior often fail to meet actual needs and preferences, resulting in less effective models. Limited user research can lead to these incorrect assumptions, making it essential for experts to gather comprehensive and accurate user data:

> "...after deployment, we realized our assumptions were flawed, leading to unexpected failures."

Trust, Explainability, and Communication. Building trust and ensuring clear communication between technical and non-technical stakeholders were crucial challenges for AI experts. Making AI models transparent and understandable is essential for building trust among stakeholders. Explainable AI solutions help stakeholders understand how models arrive at their decisions, increasing their willingness to adopt and rely on these systems. However, achieving this transparency can be challenging, especially when dealing with complex models.

Communication barriers between technical and non-technical stakeholders can lead to misaligned expectations and solutions that do not fully address the intended issues. Ensuring effective communication and understanding is key to overcoming these barriers and aligning objectives:

> "One of the biggest challenges is interacting with non-technical stakeholders who struggle to articulate their problems in ways that AI researchers can interpret."

Additionally, professionals in specialized fields may resist AI-driven solutions due to concerns about reliability and lack of domain expertise in AI implementations.

4.2 Comparison of Challenges Between AI Professionals and Academics

Our analysis revealed both shared challenges and distinctive concerns between industry practitioners and academic researchers. While both groups identified data-related issues (imbalanced datasets, limited data availability, and lack of quality domain-specific data) as fundamental challenges, they emphasized different aspects of AI development (see Table 2).

These findings highlight a fundamental difference in focus: industry professionals emphasize deployment constraints, resource limitations, and external

Table 2. Comparison of challenges between industry professionals and academic researchers

Industry Professionals	Academic Researchers
Low-quality user feedback hindering model refinement	Bridging theoretical research with practical applications
Detecting novel patterns in rapidly evolving scenarios	Overcoming unrealistic theoretical assumptions in real-world contexts
Stakeholder-defined constraints limiting development options	Building explainability and trust with domain experts
Resource and infrastructure limitations affecting deployment	Adapting models for resource-constrained environments
External factors continuously influencing internal data	Domain knowledge gaps hampering effective implementation
Unpredictable user behavior affecting model effectiveness	Challenges in translating research findings to practical systems
Communication gaps with non-technical stakeholders	Resistance from domain experts to AI-based solutions

dependencies, whereas academics prioritize theoretical adaptation and addressing the gap between idealized models and practical implementation. These complementary perspectives suggest opportunities for AI curricula to better integrate real-world complexities while maintaining strong theoretical foundations.

5 Discussion

The findings from our study provide valuable insights into the common challenges faced by AI experts. Below, we provide recommendations for educators to enhance AI curricula.

5.1 Aligning AI Curricula with Real-World Challenges

Traditional AI curricula emphasize algorithmic foundations, statistical modeling, and theoretical underpinnings, yet our study highlights significant real-world challenges that students may not encounter in a classroom setting. Issues such as imbalanced data, low-quality feedback, and resource constraints require AI practitioners to develop problem-solving skills beyond algorithmic implementation. Prior research has emphasized the importance of integrating pedagogical approaches that allow students to engage with real-world problems, gain hands-on experience and develop practical solutions. Allen et al. [3] highlight the importance of aligning teaching strategies with difficulties students in AI courses face when learning threshold concepts. Their study suggests best practices for teaching AI, including the use of practical examples and problem-based learning. Similarly, Sulmont et al. [14] identify design decisions and model evaluation as the challenging aspects of AI education–areas that closely align with the challenges reported by industry professionals in our study. Furthermore, recent work by Skripchuk et al. [13] found that students often struggle when handling open-ended, real-world datasets, frequently making mistakes during data preprocessing and feature engineering stages. These findings reinforce the importance of preparing students to navigate the inherent messiness and unpredictability of real-world AI development environments.

5.2 Bridging the Industry-Academia Divide in AI Education

The distinct priorities of industry and academia identified in our study reflect a significant educational opportunity. Paleyes et al. [12] similarly found that operational and deployment challenges often receive insufficient attention in AI education, despite being critical barriers in industry practice.

Multiple approaches can bridge this divide, including experiential learning opportunities through university-industry collaborations. Structured industry engagements–such as co-developed capstone projects, internships, and project-based coursework with external stakeholders–provide students with early exposure to the practical realities of AI system development. Additionally, emphasizing interdisciplinary training that combines AI coursework with ethics and domain-specific expertise can help students develop the flexibility needed to succeed across diverse AI applications.

These bridging strategies should not aim to simply prioritize industry needs over theoretical foundations, but rather create complementary learning experiences that value both perspectives. The goal should be developing adaptable practitioners who understand fundamental principles while navigating the constraints, stakeholder dynamics, and resource limitations that characterize professional AI development.

From this broader view, academic training serves not only to prepare students for specific professional roles, but also to foster foundational capabilities that support lifelong learning and critical engagement with evolving technologies. Denning [8] cautions against reducing computing education to algorithmic manipulation or symbolic problem-solving. He argues for a broader framing that treats computing as a professional practice, emphasizing principles such as computation, communication, coordination, automation, evaluation, and design. This perspective reinforces the importance of equipping students with reflective and transferable skills, which remain essential in an AI landscape marked by rapid technical and ethical change.

At the same time, AI curricula can benefit from a closer integration with practice-informed challenges. As Fincher and Petre [9] explain, educators often engage with expert practice not to replicate it wholesale, but to enhance students' conceptual understanding and better support their transition into professional environments. This perspective reinforces the idea that academic curricula can selectively incorporate industry insights while maintaining their broader educational mission. In this sense, our study contributes to AI education by offering one perspective – grounded in practitioner experience – on how academic programs might evolve to address real-world complexity without compromising their foundational commitments.

5.3 Integrating Real-World Problem Solving into AI Education

To bridge the gap between academic learning and industry requirements in AI education, integrating project-based learning (PBL) can be highly effective. PBL allows students to engage with real-world problems, thereby developing deeper and more usable knowledge [11]. Many of the skills essential for successful AI practice–such as critically evaluating models, making design decisions under uncertainty, and adapting solutions to dynamic conditions–are inherently higher-order cognitive tasks. Unlike traditional lecture-based courses, PBL provides an authentic context where students must actively apply theoretical knowledge, confront messy, open-ended challenges, and iteratively refine their solutions. Sulmont et al. [14] also argue that students find higher-order machine learning tasks–such as model evaluation and design decision-making–especially challenging, further highlighting the need for project-based learning or similar learning approaches. By structuring AI education around authentic, complex projects, students can develop the analytical, evaluative, and design skills required to navigate real-world AI challenges.

By incorporating projects that simulate these challenges, students can gain hands-on experience and develop practical solutions. For instance, projects could involve designing AI systems that operate efficiently under limited computational resources, which mirrors the constraints often faced by nonprofits and other resource-limited organizations. This real-world application ensures that students understand the importance of optimizing algorithms and systems for environments where high-performance computing resources are not available.

Moreover, PBL encourages collaboration and social interaction, essential components in understanding and overcoming challenges related to user behavior and interaction with AI systems. As noted by Krajcik and Shin [11], social interactions in PBL environments help students construct shared understanding and engage in disciplinary practices. Projects that require students to work together to solve complex AI problems can mirror the collaborative nature of industry work, preparing them to navigate constraints defined by diverse stakeholders.

Engaging students in authentic tasks through PBL also helps them address the unpredictability of user reactions and the need for explainability in AI. Projects could involve developing AI systems for specific user groups, followed by testing and refining these systems based on user feedback. This iterative process not only enhances technical skills but also builds an appreciation for user-centered design and the ethical implications of AI technologies.

In addition to project-based work, incorporating structured failure analysis exercises into AI coursework could further strengthen students' critical thinking. By analyzing real-world cases where AI systems failed due to deployment challenges, ethical oversights, or model drift, students can develop a deeper understanding of the complex factors influencing AI system success in practice.

5.4 Limitations

Several limitations should be considered when interpreting our findings. The majority of the participants were affiliated with US-based institutions and industries. This geographical concentration means our results may not fully represent the global diversity of AI education and practice. Our relatively small sample size ($n = 14$) further suggests these findings should be considered exploratory rather than definitive, highlighting areas that warrant further investigation through broader and cross-cultural studies.

The semi-structured interview methodology, while yielding rich qualitative data, introduces potential variability in response depth and is susceptible to self-reporting biases, as participants may emphasize certain challenges based on personal experiences or perceptions. Future work could triangulate these findings with other data sources, such as surveys or observational studies, to provide a more comprehensive understanding of the challenges faced by AI professionals and academics.

6 Conclusion

Our study explored the challenges that AI professionals face in both industry and academia, highlighting key gaps between current AI education practices and the realities of professional AI development and deployment. In light of the study findings, we propose several strategies that could strengthen AI undergraduate education:

- **Integrate Real-World Data Complexity into Coursework:** Incorporate projects and assignments using noisy, imbalanced, or incomplete datasets to expose students to practical data challenges.
- **Introduce Failure Analysis Exercises:** Embed structured analyses of real-world AI system failures into coursework to cultivate critical reflection on operational risks and system limitations.
- **Foster Interdisciplinary Collaboration:** Design learning experiences that involve working with domain experts and non-technical stakeholders, reflecting the cross-disciplinary nature of modern AI deployments.
- **Model User Behavior Variability in Design Projects:** Encourage students to anticipate and design for diverse, unpredictable user behaviors and evolving system requirements.
- **Promote Experiential Learning Opportunities:** Expand internships, industry-sponsored projects, and university-industry collaborations to offer students direct exposure to real-world AI development environments.
- **Strengthen Capstone Project Requirements:** Encourage capstone projects that simulate realistic resource constraints, dynamic conditions, and stakeholder negotiation processes.

These suggestions aim to complement the existing strengths of AI programs by better aligning technical instruction with the complexities and uncertainties encountered in real-world practice. By enhancing experiential, operational, and interdisciplinary training, AI curricula can foster a new generation of professionals who are technically adept, operationally resilient, and ethically aware.

References

1. The artificial intelligence (AI) for k-12 initiative (ai4k12) (2024). https://ai4k12.org/. Accessed 18 Jul 2024

2. ACM, IEEE-CS, AAAI: Cs2023: ACM/IEEE-CS/AAAI computer science curricula. Tech. rep., ACM, IEEE-CS, and AAAI (2024). https://csed.acm.org/, endorsed: ACM – January 18th 2024, IEEE-CS – January 22nd, 2024, AAAI – February 22nd 2024
3. Allen, B., McGough, A.S., Devlin, M.: Toward a framework for teaching artificial intelligence to a higher education audience. ACM Trans. Comput. Educ. (TOCE) **22**(2), 1–29 (2021)
4. Baier, L., Jöhren, F., Seebacher, S.: Challenges in the deployment and operation of machine learning in practice. In: ECIS, vol. 1 (2019)
5. Bublin, M., Schefer-Wenzl, S., Miladinović, I.: Educating AI software engineers: challenges and opportunities. In: Auer, M.E., Hortsch, H., Michler, O., Köhler, T. (eds.) Mobility for Smart Cities and Regional Development - Challenges for Higher Education: Proceedings of the 24th International Conference on Interactive Collaborative Learning (ICL2021), Volume 2, pp. 241–251. Springer International Publishing, Cham (2022). https://doi.org/10.1007/978-3-030-93907-6_26
6. Chenoweth, S., Linos, P.K.: Teaching machine learning as part of agile software engineering. IEEE Trans. Educ. **67**(3), 377–386 (2024). https://doi.org/10.1109/TE.2023.3337343
7. DeNero, J., Klein, D.: Teaching introductory artificial intelligence with Pac-Man. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 24, pp. 1885–1889 (2010)
8. Denning, P.J.: The profession of it beyond computational thinking. Commun. ACM **52**(6), 28–30 (2009)
9. Fincher, S., Petre, M.: Mapping the territory. In: Fincher, S., Petre, M. (eds.) Computer Science Education Research, pp. 1–8. Routledge (2004)
10. Gladwell, M.: Outliers: The Story of Success. Little, Brown (2008)
11. Krajcik, J.S., Shin, N.: Project-based learning. In: Sawyer, R.K. (ed.) The Cambridge Handbook of the Learning Sciences, pp. 72–92. Cambridge University Press, third edn. (2022)
12. Paleyes, A., Urma, R.G., Lawrence, N.D.: Challenges in deploying machine learning: a survey of case studies. ACM Comput. Surv. **55**(6), 1–29 (2022)
13. Skripchuk, J., Shi, Y., Price, T.: Identifying common errors in open-ended machine learning projects. In: Proceedings of the 53rd ACM Technical Symposium on Computer Science Education, vol. 1, pp. 216–222 (2022)
14. Sulmont, E., Patitsas, E., Cooperstock, J.R.: Can you teach me to machine learn? In: Proceedings of the 50th ACM Technical Symposium on Computer Science Education, pp. 948–954 (2019)
15. Zimmermann, R.M., Allin, S., Zhang, L.: Common errors in machine learning projects: a second look. In: Proceedings of the 23rd Koli Calling International Conference on Computing Education Research, pp. 1–12 (2023)