Full Paper - The Hidden Labor Market Titles of AI Engineering Education

Chibuzor Okocha, University of Florida

Chibuzor Okocha is a Ph.D. student in Computer Science at the University of Florida. His research focuses on Artificial Intelligence in audio and speech processing, with a specific interest in leveraging large audio language models (LALMs) for reasoning tasks and audio question answering. He is particularly interested in the application of these multimodal technologies to enhance engineering education and accessibility.

Atayliya Natasha Irving, University of Florida

Atayliya N. Irving is a PhD Candidate at the University of Florida (UF) in the Engineering Education department. She earned her Bachelor's degree in Computer Science from Jackson State University. The scope of her interests is to explore the experiences of marginalized students and faculty in higher learning to broaden participation in computing. She specifically looks at the experiences and motivations of women in computing-related PhD programs, with the overall goal of recreating these experiences for the next generation of computer scientists. Upon graduating, Atayliya aspires to pursue a career that broadens interest and representation in the STEMM+ C fields.

Idalis Villanueva Alarcón, University of Florida

Dr. Idalis Villanueva Alarcón is Chair and Professor in the Department of Engineering Education in the Herbert Wertheim College of Engineering at the University of Florida. A presidential awardee (PECASE), she has led multiple pioneering efforts in engineering education including multimodal methods in engineering education using sensor technologies and biophysiological tools, hidden curriculum, mentoring, active learning, professional identity, among others. She is a renowned national and international leader in engineering education research and practice earning her multiple accolades and honors through professional organizations such as the National Academy of Engineering, IEEE, and ASEE. She integrates her technical experiences as a Chemical Engineering, Biological Engineer, Analytical Cell Biologist, and Engineering Education Researcher to tackle complex engineering education problems across the learner life span.

Dr. Curtis R. Taylor, University of Florida

Dr. Curtis Taylor, Ph.D. is an Associate Professor in the Departments of Engineering Education and Mechanical and Aerospace Engineering at the University of Florida.

Christan Grant Ph.D., University of Florida

Dr. Christan Grant is the Arnold and Lisa Goldberg Rising Star Associate Professor in the Department of Computer & Information Science & Engineering at the University of Florida, where he directs the UF Data Studio research group. His work focuses on end-to-end data and AI systems—how data are collected, labeled, modeled, and visualized—in domains ranging from civic and health data to intelligent infrastructure, with support from the National Science Foundation, federal agencies, and industry partners. At UF he teaches data engineering and introductory data science and has previously developed courses in databases, operating systems, and text analytics that foreground real-world datasets and collaborations reflecting contemporary AI and data science careers. He earned both his Ph.D. in Computer Science and B.S. in Computer Engineering from the University of Florida and has led community-building efforts such as the Broadening Participation in Data Mining workshop and a human-centered data science seminar series. In this paper, his interest lies in how AI-related roles across higher education, edtech, and industry encode teaching work in their titles and descriptions, and how that "hidden" labor market can inform more transparent, career-relevant preparation for engineering students and instructors.

The Hidden Labor Market for AI in Engineering Education

Abstract

Artificial intelligence (AI) is reshaping how learning happens in engineering, yet many roles that carry out core pedagogical work in AI are not labeled as "teaching." Across universities, educational technology (edtech) companies, startups, and large technology firms, labor-market titles such as developer educator, AI educator, learning experience designer, and AI product researcher describe roles that encompass education. These roles include designing learning materials, facilitating training, documenting practices, and evaluating learning outcomes.

This study examines how AI engineering education-related job titles are described and operationalized, the roles and responsibilities associated with these titles, and their implications for preparing engineering education graduate students and instructors. A stratified corpus of public job postings from higher education, edtech, and industry teams that support engineering or STEM learning was analyzed using a directed content analysis approach.

The analysis focused on identifying patterns in: (a) pedagogical actions (e.g., design, facilitate, assess, document); (b) required instructional artifacts (e.g., curricula, labs, tutorials, assessment plans); (c) intended learner audiences; (d) delivery contexts and modalities; (e) assessment and evaluation practices; and (f) references to learning science frameworks. Content patterns were examined to identify areas of convergence and divergence in competency expectations using a combination of large language model (LLM)-assisted evaluation and human-led analysis.

Findings highlight how pedagogical work is distributed across sectors and embedded within a range of AI-related roles, often without being explicitly labeled as teaching. These results provide insight into the competencies, artifacts, and practices associated with AI-related educational work and offer guidance for aligning engineering education curricula with emerging workforce expectations.

Introduction

Artificial intelligence (AI) is reshaping how learning occurs in engineering and other STEM fields [1, 2]. Beyond influencing course content and instructional tools, AI is transforming where and by whom pedagogical work is carried out [3]. Teaching and learning about AI now extend across universities, edtech companies, industry teams, and government workforce initiatives [4]. In these settings, many roles responsible for designing and supporting learning are not labeled as “teaching” positions, yet they perform core instructional functions essential to preparing engineers for AI-enabled practice [5, 6].

Across sectors, job titles such as developer educator, learning experience designer, enablement engineer, and responsible AI program manager describe roles that routinely involve designing curricula and learning materials, facilitating workshops and trainings, documenting practices, and evaluating learning outcomes [7, 8]. These roles often support engineering students, educators, and professional engineers as they learn to use, evaluate, or govern AI systems [9, 10]. However, because this pedagogical work is distributed across product teams, research groups, and organizational support structures, it is largely absent from how teaching is traditionally conceptualized in engineering education . This gap matters for engineering education programs. As departments seek to prepare students for careers shaped by AI, they must understand not only technical skill requirements but also the pedagogical competencies embedded in emerging roles.

Prior research in engineering education and AI in education has focused primarily on classroom applications, instructional technologies, and student outcomes [11, 12]. Far less attention has been given to how pedagogy is articulated in the labor market, how teaching-like work is embedded in non-academic roles, or how expectations for pedagogical expertise differ across academic, industry, and edtech contexts. Without this insight, curricular design risks misalignment with the professional practices graduates are likely to encounter [13, 14]. In this study, we examine how pedagogy is described and operationalized in AI-related job postings that support engineering or STEM learning. We analyze a stratified corpus of public job advertisements from higher education, edtech, and industry using directed content analysis. Our analytic framework captures pedagogical actions (e.g., design, facilitate, assess, document), instructional artifacts, learner audiences, delivery contexts, assessment language, and references to learning science and responsible AI practices.

While prior work has documented the growing presence of AI in education and the emergence of new roles, less attention has been paid to how pedagogical work itself is being restructured within these roles. In industry and edtech contexts, teaching is often not organized as a discrete activity but is instead embedded within products, platforms, and organizational workflows. Drawing on patterns observed in job postings including the prevalence of documentation, platform-based learning environments, and analytics-driven evaluation this study interprets how pedagogical work is being transformed in practice. Specifically, we conceptualize three overlapping trajectories: pedagogy becomes technologized, embedded within tools and systems; productized, as learning is packaged into scalable artifacts such as tutorials and curricula; and metricized, as learning is evaluated through analytics, experimentation, and performance indicators. These conceptual distinctions extend beyond description of job requirements to explain how and why pedagogical work takes its current form across sectors.

Problem Statement

Although AI-related job postings are abundant, there is no systematic way to identify roles that combine AI with pedagogy and engineering education, making it difficult for academic programs to align their curricula with workforce trends and for people to find related jobs in this sector.

Purpose of the Study

The purpose of this work is to systematically identify and characterize roles at the intersection of artificial intelligence and pedagogy within engineering and STEM contexts. By analysing a stratified set of job postings from higher education, edtech, industry and government, we seek to uncover how teaching-like work is embedded in contemporary AI job titles and descriptions, to surface the competencies implied by these postings, and to assess how well current engineering education programs align with emerging labor–market expectations. In doing so, we provide evidence to inform curriculum design, career advising and micro-credential development for AI-enabled engineering education.

Research Questions

This study addresses the following research questions:

> **RQ1:** To what extent do large language models converge with one another and with a human-curated corpus in identifying AI-related roles that involve pedagogical work in engineering education?
>
> **RQ2:** What qualifications, skills, and labor-market characteristics define AI-related roles that support engineering and STEM education across sectors?

Background and Literature Review

The Evolving Workforce of AI Jobs

In recent years, the AI related workforce has expanded rapidly across almost every industry [15]. Alongside these advancements, employers are not only creating new AI-specific roles, AI is being infused into many other job sectors. This surge is driven by increased technological advancements, which in turn raise expectations for the workforce to continuously upskill [16]. It is anticipated that AI will completely transform the job market, with a prediction of a 70 percent change of skills needed by 2030 [17]. AI competencies have become a core demand for many jobs and the workforce is expected to adapt to this change.

Educational institutions are struggling to keep up with the pace of AI innovation and the demands of the evolving workforce [18]. There is a growing curricular gap between what students are taught and what industry now expects and AI is exacerbating this. AI technologies have moved from only being in specialized research labs to widespread use, yet no standardized approach to integrating AI into curricula exists [18]. Some universities have began adding AI courses or modules, but these efforts are often impromptu or ad-hoc. As a result engineering education curricula lags behind the industries development and expectations for AI competencies. A major aspect of this challenge is academia's slow pace in change [19].

This means that many graduates enter the workforce without formal AI preparation. Engineering education then has the pressure to update its content and pedagogy to address the realities of the current AI integrated job market. Another obstacle is that AI roles have become increasingly interdisciplinary, blurring the line between traditional fields and as a result, AI jobs are no longer confined to the tech sector [19]. Employers now seek people who in addition to understanding AI techniques, can also apply them to specific domain related contexts. This means that students need to be prepared on how to qualify and find jobs that may appear under different titles across a plethora of industries. Although an engineering background does equip students with a versatile problem-solving and technical skillset, students need to market these skills for AI application [20] [21].

In this, to better prepare students for the AI-driven workforce, engineering education curricula could incorporate both technical applications of AI and domain specific skills, with students also being exposed to AI opportunities across these different sectors and domains. Analyzing current job postings can then illuminate what skills and keywords students should focus on when searching for jobs, which can in turn inform future curriculum design and professional development so students can better navigate the evolving job market.

Prior Research on Job Analysis Postings

Researchers have looked at job postings analysis as a data-driven method to scope job market needs. By mining online job as for required skills and qualifications, one is able to obtain insight into what employers are actually seeking [20]. This approach has emerged as a hypothesized valuable complement to traditional curriculum design, which typically relies on accreditor perspectives [20]. The literature generally confirms that analyzing job postings is an effective method to keep track of new competency requirements, especially in the fast evolving job market [21]. Unlike one-time surveys or general course guidelines, the job posting analysis approach captures the shifts in employer demands as they occur. Prior studies applying this method have revealed important patterns relevant to AI and engineering education. Fleming et al. (2024) conducted a large-scale analysis of over 26,000 engineering job postings to identify the skills gap between engineering education graduates and industry expectations [22]. They found that certain professional competencies are underemphasized in academic settings [20] [21]. This research demonstrates that job posting analytics can pinpoint which skills are in high demand. These studies validate the approach of analyzing job postings to keep educational outcomes aligned with the workforce needs.

Gap in the Literature

Despite prior work, a clear gap remains for the need for an up-to-date analysis on AI related jobs, with a focus on job titles [23]. In a rapidly evolving job market, students and educators need to know which jobs and keywords to target. By analyzing AI engineering education jobs, we can gain a better understanding of the current landscape. Many existing studies either examine broad engineering roles [22] or specific subfields such as data science [20]. However, the AI job market of today spans many sectors, and evolves at an alarming rate [24]. The current landscape calls for a targeted examination of the skills and qualifications that define “AI jobs” across industries and how they align with the skills students acquire [25, 26]. As AI applications are increasingly interdisciplinary, engineering education can leverage this to provide a foundation applicable across a range of contexts [12, 27]. In many instances, it is unclear what combination of skills is most valuable, so bridging this gap is crucial for preparing competitive graduates for the job market [28, 29].

Recent graduates report feeling unprepared and uncertain about AI's role in their careers [18, 23], confirming the need for academia to catch up with industry's adoption of AI. By conducting an updated job-posting analysis centered on AI-related positions, we aim to shed light on the current demands of the AI-driven workforce as they pertain to engineering education graduates. This research would identify the key technical proficiencies and valued skills that employers list when listing AI-infused roles. The findings also could directly inform engineering educators on which emerging competencies to integrate into curricula or co-curricular. This focused analysis addresses the interdisciplinary nature of modern AI employment by mapping it against the interdisciplinary skill set of engineering education. It responds to the urgent call for academic programs to better align with the AI landscape, ensuring that graduates are not only technically capable but also equipped to apply AI responsibly and effectively in whichever domain their career takes them [18]. This is especially vital for preparing the next generation of engineers and platforms, where a future of AI competency is expected to be a baseline requirement across the board.

Data and Methodology

Data Sources

The corpus analyzed in this study consists of 50 public job advertisements from higher education institutions, education technology companies, and industry teams with explicit ties to engineering or STEM learning. A PRISMA-style workflow was employed to systematically identify AI-related job postings relevant to engineering education. Job postings were collected from publicly accessible platforms, including LinkedIn, Google Jobs, and Indeed, yielding an initial set of N = 65 postings.

The search strategy was developed through an iterative, researcher-informed process. The first and second authors, current engineering education graduate students with backgrounds in computer science, generated and refined search terms based on domain knowledge of artificial intelligence and education-related roles. While "AI in Education Roles" served as a primary search phrase, related variants (e.g., combinations of "artificial intelligence," "machine learning," and "education") were also explored to assess consistency in retrieved postings.

After removal of duplicates, N = 60 unique postings remained and were screened based on job title and summary. Full-text review of postings was conducted using predefined inclusion and exclusion criteria.

Inclusion and Exclusion Criteria

Job postings were included in the analytic corpus if they met all of the following criteria:

1. *Explicit AI requirement:* The posting references at least one artificial intelligence–related skill, tool, or methodology (e.g., machine learning, deep learning, natural language processing, computer vision, generative AI, or AI model development, deployment, or evaluation).
2. *Engineering education relevance:* Required or preferred qualifications indicate alignment with engineering or STEM education, learning design, instructional support, or educational enablement.
3. *Workforce-oriented role:* The position represents a full-time or long-term professional role rather than an internship, fellowship, or short-term training opportunity.
4. *U.S.-based position:* The role is located in or explicitly tied to the United States.
5. *Public availability:* The posting is publicly accessible at the time of data collection.

Postings were excluded from the analysis if they met any of the following criteria:

1. *No explicit AI content:* The posting does not reference AI-related skills, tools, or methods.
2. *Lack of engineering education relevance:* The role does not involve engineering or STEM education, learning design, or educational support activities.
3. *Training-only offerings:* The posting describes standalone training programs, bootcamps, certifications, or workshops rather than a professional role.
4. *Insufficient role information:* The posting lacks adequate detail regarding responsibilities, required skills, or qualifications.

These criteria were applied during both initial title screening and full-text review to ensure consistency across selection stages.

Sixty-five jobs were identified and with the use of the above criterion, N=50 job postings were included in the final analysis. Only positions requiring explicit AI-related skills and engineering-aligned competencies were retained. This process is shown in the PRISMA-style diagram depicted in Fig. 1.

As shown above, the selection criteria required that each advertisement (i) mention AI, machine learning, or related technologies in the description and (ii) describe responsibilities related to teaching, training, documentation or curriculum design. Duplicate postings and roles without clear educational components were removed. Because the data are publicly available job ads, no human subjects were involved. Additionally, we did not anonymize any identifiable organizational names when reporting examples. Table 1 below shows examples from the validated job corpus, including qualifications and skills summarized to highlight pedagogical and AI-related responsibilities across organizations.

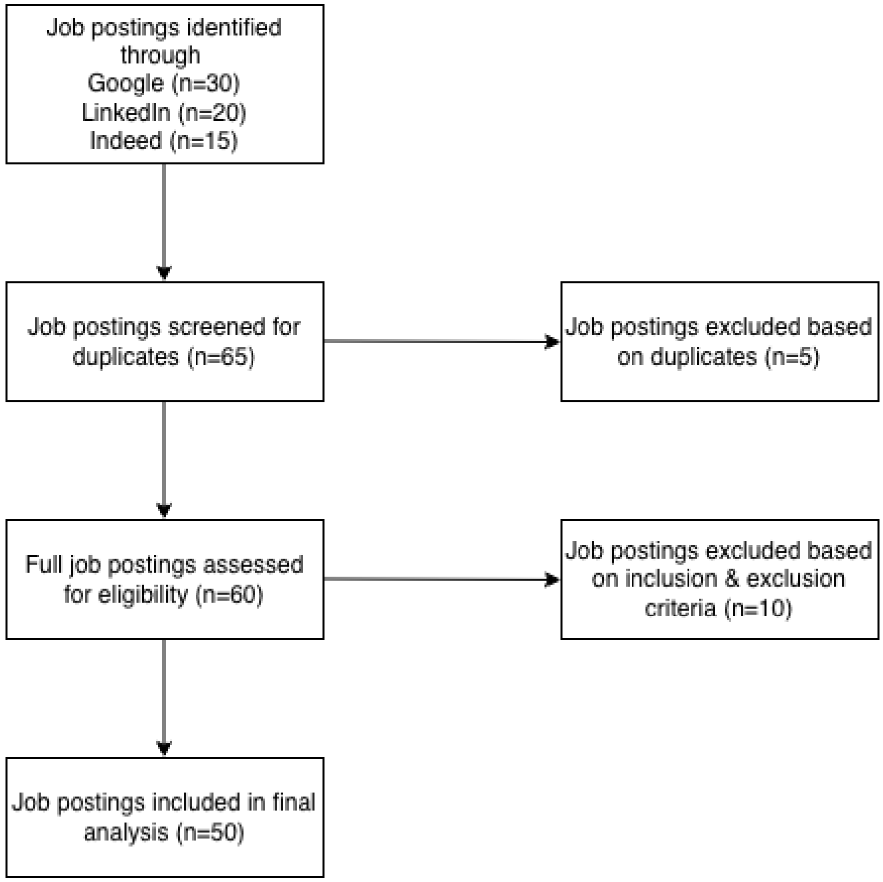

Figure 1: PRISMA-style flow diagram for job posting selection.

Table 1: Illustrative examples of AI–Education roles from the validated job corpus.

Company	Job Title	Representative Qualifications	Representative Skills
Apple	Machine Learning Educator	• Degree in computer science, engineering, or related field • Prior experience teaching or training technical audiences • Industry experience with machine learning systems	• Designing technical learning materials • Explaining ML concepts to non-expert audiences • Developer education and enablement
OpenAI	Customer Education Community Manager, Champion Network	• Extensive experience in education program design or enablement • Experience supporting technical products or platforms • Strong communication and coordination background	• Designing and scaling AI education programs • Community-based learning and engagement • Translating AI capabilities into learning resources

Company	Job Title	Representative Qualifications	Representative Skills
Google	Senior Instructor, AI and Machine Learning	• Advanced degree or equivalent experience in a technical field • Prior instructional or training experience • Experience with AI or machine learning tools	• Curriculum development for AI topics • Technical instruction and assessment • Learner-focused explanation of AI systems
Google	Associate Director, AI Fluency and Education	• Senior-level experience in education or workforce development • Experience leading cross-functional teams • Background in AI or emerging technologies	• Strategic design of AI education initiatives • Organizational AI literacy and fluency • Stakeholder communication and alignment
Duolingo	Learning Engineer	• Degree in computer science, engineering, or related field • Experience working at the intersection of learning and technology • Familiarity with data-driven experimentation	• Designing learning systems and content • Applying ML to educational products • Evaluation of learning outcomes
Intuit	Teaching and Learning Expert, AI Enablement	• Experience in instructional design or corporate learning • Familiarity with AI-enabled products • Background in professional training or enablement	• Developing AI training programs • Facilitating workforce learning • Creating instructional documentation
Oracle	Instructional Designer, Development Enablement	• Degree in education, instructional design, or related field • Experience supporting technical or developer audiences • Knowledge of enterprise software environments	• Instructional design for technical products • Learning content development • Enablement and documentation

Large Language Models Used in the Analysis

Five instruction-tuned large language models were used in this study to support structured evaluation, information extraction, and normalization tasks. The models were selected to represent a range of architectures, training lineages, and organizational origins, allowing us to assess the robustness of LLM-assisted analysis across diverse systems rather than relying on a single model. All models were used through a consistent instruction format and fixed output schemas to ensure comparability.

The models included in the analysis were: Google Gemma 3 27B Instruct, Meta Llama 3.1 8B Instruct, Mistral 7B Instruct v0.3, AllenAI OLMo 2 7B Instruct, and Qwen 2.5 7B Instruct. Each model is publicly available and optimized for instruction-following tasks, making them suitable for structured text transformation and rubric-based evaluation.

Although the corpus size (N = 50) is within the range of manual qualitative analysis, LLM-assisted evaluation was employed to support structured consistency and comparative robustness rather than to replace human interpretation. Specifically, the use of multiple instruction-tuned models enabled (1) standardized application of a fixed analytic rubric across all postings, reducing variability that may arise from single-coder subjectivity; (2) cross-model comparison, allowing us to assess whether patterns in pedagogical content were stable across models with different architectures and training data; and (3) transparent, schema-based extraction and normalization of skills and qualifications, improving traceability between raw text and analytic categories.

Importantly, LLMs were used only for constrained transformation and scoring tasks, while all higher-order interpretation, pattern identification, and synthesis were conducted by the research team. The relatively small dataset enabled close human oversight of all outputs, ensuring accuracy while leveraging LLMs to enhance consistency and reproducibility. In this sense, LLM-assisted analysis functioned as a complementary methodological tool that strengthened analytic rigor rather than as a substitute for human qualitative analysis.

These models were used for three primary purposes. First, they were used as evaluators to score job postings along predefined rubric dimensions related to pedagogy, learner audience, AI relevance, responsible AI content, cross-sector context, and overall alignment with AI-enabled educational work. Second, they were used to extract structured information such as skills and qualifications verbatim from job descriptions. Third, they were used to support rule-based canonicalization by mapping extracted phrases to standardized categories under explicit normalization constraints.

The use of multiple models served two methodological goals. The first was to reduce dependence on the idiosyncrasies or biases of any single model. The second was to enable comparison between LLM-assisted outputs and human-led selection analysis across the same set of fifty job postings that were selected through manual screening. Human analysis was treated as the interpretive reference point, while LLM outputs were examined for alignment, divergence, and consistency across models. [30]

The selected models differ in scale, training data composition, and governance orientation. For example, Gemma and Llama represent models developed by large technology companies with broad instruction tuning, while Mistral and Qwen reflect models designed for efficiency and multilingual or cross-domain instruction-following. OLMo was included to incorporate an openly developed model with an explicit emphasis on transparency and research use. Together, this set of models allowed us to examine whether observed patterns in pedagogical content were stable across different model families.

Analytic Design

We employed a qualitative analytic design combining LLM-assisted structured coding with human-led selection. The purpose of this hybrid approach was to support systematic extraction and normalization of job posting content while preserving interpretive depth through human analysis. The LLM was used to perform constrained, schema-based transformations of job posting text. All higher-order interpretation, pattern identification, and synthesis were conducted by the research team.

LLM-Assisted Structured Coding

The first analytic stage used LLM-assisted structured coding to extract and normalize skills and qualifications from job postings. A single instruction prompt with a fixed output schema was used to ensure consistency across inputs. The prompt directed the model to map raw skill and qualification phrases to canonical categories using explicit normalization rules, or to create a new general category when no rule applied. Outputs were restricted to valid JSON. All outputs were reviewed for completeness and formatting accuracy prior to analysis. This procedure supported transparency and traceability by linking canonical categories to raw phrases. Representative mappings are shown in Tables 2 and 3. The full normalization rules are provided in the Appendix A.

Data Extraction Prompts

To extract and normalize data from the job postings, we used three targeted prompts:

- Qualification extraction: Extract qualifications verbatim (see Prompt A.2).
- Skill normalization: Map raw skills to canonical categories (see Prompt A.3).
- Qualification canonicalization: Normalize degrees and years of experience (see Prompt A.4).

Table 2: Representative skill normalization mappings used to standardize raw skill phrases into canonical categories.

Raw skill phrases (examples)	Canonical skill category
machine learning; ML; deep learning; AI/ML	AI/ML Fundamentals
large language models; LLMs; GenAI; generative AI	LLMs / Generative AI
prompt engineering; prompt design; prompt crafting	Prompt Engineering
curriculum design; instructional design; learning experience design	Instructional / Curriculum Design
workshops; training delivery; facilitation	Facilitation / Training Delivery
documentation; tutorial writing; guides	Technical Documentation / Tutorials
communication; presentation; stakeholder engagement	Communication and Stakeholder Collaboration
A/B testing; analytics; learning outcomes measurement	Evaluation / Analytics
LMS; Canvas; Moodle; learning management system	Learning Platforms (LMS)
Python; programming; coding	Programming / Coding Skills
data analysis; data visualization; analytics	Data Analysis and Visualization
research; conducting research; research experience	Research Experience

Table 3: Representative qualification normalization mappings used to standardize degrees and experience requirements, including domain-context experience banding.

Raw qualification phrases (examples)	Canonical qualification category
bachelor's degree; undergraduate degree; BS; BA	Bachelor's Degree
master's degree; MS; MA; graduate degree	Master's Degree
PhD; doctorate; doctoral degree	Doctoral Degree
minimum of 1 year of AI/ML experience; 1 year of AI experience	1–3 Years AI/ML Experience
3+ years of AI/ML experience; 4 years of ML experience	4–6 Years AI/ML Experience
7+ years of AI/ML experience; 8 years of AI experience	7+ Years AI/ML Experience
minimum of 1 year teaching; 1 year instructional experience	1–3 Years Teaching Experience
3+ years teaching; 4 years teaching or instruction	4–6 Years Teaching Experience
7+ years teaching; 8+ years instructional experience	7+ Years Teaching Experience
minimum of 1 year in EdTech; 1 year learning technology	1–3 Years EdTech Experience
3+ years in EdTech; 4 years learning technology	4–6 Years EdTech Experience
7+ years in EdTech; 8+ years learning technology	7+ Years EdTech Experience
minimum of 1 year in higher education	1–3 Years Higher Education Experience

Human-Led Selection

Following structured coding, the research team conducted human-led selection analysis to examine how pedagogical work was represented across roles.Researchers analyzed coded outputs alongside original job posting text to identify patterns in pedagogical verbs, artifacts, learner audiences, modalities, assessment practices, and references to responsible AI documentation. Interview data from practitioners across academia, edtech, and industry were used to contextualize patterns observed in postings and to surface instances where job descriptions and enacted practices diverged.

Analytic Framework and Coding Scheme

We developed an analytic framework to characterize pedagogical work embedded in AI-related job postings. The framework served two complementary purposes: (1) guiding human-led selection analysis and (2) structuring large language model–based evaluation, extraction, and normalization. Rather than functioning as an inductive codebook, the framework operationalizes theoretically informed dimensions of pedagogy, AI relevance, and cross-sector educational work.

The framework consists of six dimensions, each rated on a five-point ordinal scale (1 indicating absence and 5 indicating strong presence): The six rubric dimensions were selected to reflect the study's central concern: whether a job posting meaningfully combined AI-related work with pedagogical or educational labor. The dimensions therefore capture both the presence of educational work and the organizational context in which that work appears. Pedagogical indicators, learner/audience focus, and AI relevance were used to determine whether a posting aligned with the study's core inclusion logic. Responsible AI/governance was included because AI-related educational roles may involve explaining, documenting, or supporting ethical and responsible use of AI systems. Cross-sector context captured whether the role was situated in higher education, edtech, industry, or related organizational settings. Overall fit provided a holistic judgment of whether the role represented the type of AI-related pedagogical work examined in this study.

To support consistency across postings, each dimension was rated on a five-point ordinal scale. A score of 1 indicated that the dimension was absent or not meaningfully represented. A score of 3 indicated moderate or implicit evidence. A score of 5 indicated strong, explicit, and central evidence in the posting. Scores of 2 and 4 represented intermediate levels of evidence. As shown in Table 4, each dimension of the analytic rubric captures a distinct aspect of AI-related pedagogical work.

To extract specific skills from the text prior to normalization, we used a dedicated skill-extraction prompt (see Prompt A.5).

Sector Inference and Salary Parsing Rules

Sector and salary variables were derived from job posting metadata and text using transparent, rule-based procedures. Sector was inferred from employer and posting text using keyword cues. Postings containing higher education indicators (e.g., "university," "college," "campus," "community college") were labeled as Higher Ed. Postings containing education technology indicators (e.g., "edtech," "learning platform," "LMS," "Canvas," "Moodle," "MOOC") were labeled as EdTech. Remaining postings were labeled as Industry. This heuristic classification was applied consistently across the corpus and is reported as inferred sector.

Salary information was extracted from the salary field when present and standardized to annual US dollar values for descriptive reporting. When salary text included explicit pay periodicity (e.g., per year, per hour, per month, per week), values were converted to annual equivalents using fixed multipliers (hourly multiplied by 2,080 hours per year; weekly multiplied by 52; monthly multiplied by 12). When salary ranges were provided, minimum and maximum values were extracted and the annual midpoint was computed as the average of the annualized minimum and maximum. Salary strings that referenced benefits without explicit pay values (e.g., retirement plans) or that provided compensation policies without numeric values were treated as not disclosing salary. Reported salary summaries use only postings with extracted numeric compensation values.

Results

This study analyzed a validated corpus of fifty ($N = 50$) job postings to characterize the emerging labor market for AI-related roles that support engineering education. Results are presented in three stages. First, we provide descriptive context by examining job titles and

Table 4: Analytic rubric for evaluating AI-related pedagogical work in job postings

Dimension	What it measures	Low score example	Mid score example	High score example
Pedagogical indicators	Evidence that the role involves teaching, training, instructional design, facilitation, documentation, or assessment	No teaching, training, or learning-related responsibilities	Some mention of training, onboarding, or documentation	Clear responsibility for curriculum design, workshops, learning materials, or assessment
Learner/audience focus	Whether the posting identifies people who are expected to learn from or use the educational work	No clear learner or user audience	General users, clients, or stakeholders are mentioned	Specific learners such as students, educators, engineers, developers, or professional trainees are named
AI relevance	Degree to which the role explicitly involves AI, ML, LLMs, generative AI, or AI-enabled tools	No clear AI-related content	AI is mentioned generally or as a preferred skill	AI/ML, LLMs, generative AI, or AI tools are central to the role
Responsible AI/-governance	Evidence of ethics, fairness, transparency, safety, explainability, accountability, or governance	No responsible AI or ethics language	General mention of responsible technology, compliance, or quality	Explicit responsibility for responsible AI, fairness, safety, transparency, or governance documentation
Cross-sector context	Whether the role sits within an educational, edtech, industry, or workforce-learning context relevant to engineering/STEM learning	No clear education, workforce, or sector relevance	Some connection to training, product support, or organizational learning	Clear alignment with higher education, edtech, industry enablement, developer education, or STEM workforce learning
Overall fit	Holistic assessment of whether the posting integrates AI-related work with pedagogical or educational responsibilities	AI role with no meaningful educational component, or education role with no AI component	Some overlap between AI and education, but one component is weak	Strong integration of AI content and pedagogical work

compensation patterns in the corpus. Second, we evaluate the degree of agreement between multiple Large Language Models (LLMs) and a human-curated job set to assess the reliability of the LLM-assisted analytic approach (RQ1). Third, we characterize the qualifications and skills associated with these roles using the validated corpus (RQ2).

Most Frequent Job Titles in the Corpus

Table 5 presents the ten most frequent job titles in the corpus, reported verbatim. Job titles varied substantially across postings, with most appearing only once. Only a small number of titles recurred across multiple postings, reflecting the heterogeneous and evolving nature of AI-Education roles.

Despite this variation, recurring titles such as Instructional Designer, Developer Advocate, AI Educator, and AI Tutor indicate that pedagogical work is formally recognized in some roles, even as many positions embed educational responsibilities within technically oriented or product-facing titles. This dispersion of titles underscores the challenge of identifying AI-related pedagogical work using title-based searches alone.

Table 5: Frequent job titles (verbatim) in the corpus ($N = 50$).

Job Title (Verbatim)	Count
Instructional Designer	2
AI Educator	2
AI Tutor	2
AI/ML Program Analyst	1
Technical Instructor	1
Information Systems Instructor	1
Senior Human-AI Collaboration Research Engineer	1
Software Engineer, AI/ML Education	1

Compensation Disclosure by Sector

To provide additional labor-market context, we examined compensation disclosure across inferred sectors. Table 6 summarizes salary transparency and annualized median salary midpoints for postings that included compensation information.

Salary disclosure rates varied by sector. Higher education postings were more likely to disclose compensation than industry postings, while edtech postings exhibited the highest disclosure rate overall. Among postings with disclosed salaries, median annualized midpoints were highest in edtech roles, followed by industry roles, with higher education roles reporting lower median compensation. These differences align with broader sectoral compensation patterns and provide additional insight into how pedagogical AI labor is valued across contexts.

Table 6: Salary disclosure and annualized median salary midpoint (USD) by inferred sector.

Sector	N	Salary Disclosed (n)	Disclosure Rate	Median Annual Midpoint (USD)
EdTech	5	4	0.80	150,025
Higher Ed	14	10	0.71	97,725
Industry	31	18	0.58	125,725

RQ1: Agreement Between LLM Evaluators and Human-Selected Roles

We next evaluated whether LLM-assisted methods provide a reliable approach for identifying pedagogically oriented AI roles. Five instruction-tuned models were used to score each job posting using a shared rubric capturing pedagogical relevance, learner focus, AI content, responsible AI considerations, and cross-sector context. The human-curated corpus served as an interpretive reference point rather than a strict ground-truth label.

Table 7 reports mean rubric scores, score dispersion, and classification rates for each model. Classification rate represents the proportion of postings that met or exceeded the predefined inclusion threshold. Across models, classification rates ranged from 70% to 94%, indicating that most postings were consistently identified as AI-Education roles. Mean scores were tightly clustered, and standard deviations were modest, suggesting limited ambiguity in model evaluations.

Table 7: Summary statistics of LLM evaluators. Classification rate indicates the percentage of jobs meeting the threshold (≥ 3.5) for inclusion as an AI-Education role.

Model	Mean Score	Std. Dev.	Classification Rate (%)
Gemma	3.67	0.57	70
Llama	3.77	0.87	74
Mistral	4.07	0.60	88
OLMo	4.27	0.51	94
Qwen	3.99	0.56	86

To further assess consistency, we computed pairwise Pearson correlation coefficients between model scores for all job postings (Table 8). Correlations ranged from $r = 0.61$ to $r = 0.85$, indicating strong positive agreement across models with distinct architectures and training lineages. No single model exhibited systematic divergence in job ranking.

Table 8: Pairwise Pearson correlations between model scores. Strong positive correlations indicate high inter-model agreement in the evaluation of job postings.

	Gemma	Llama	Mistral	OLMo	Qwen
Gemma	1.00	0.83	0.70	0.75	0.85
Llama	0.83	1.00	0.73	0.73	0.84
Mistral	0.70	0.73	1.00	0.61	0.72
OLMo	0.75	0.73	0.61	1.00	0.69
Qwen	0.85	0.84	0.72	0.69	1.00

Together, these results demonstrate strong convergence among LLM evaluators and alignment with the human-selected corpus. This convergence supports the use of LLM-assisted evaluation as a stable and scalable method for identifying pedagogically oriented AI roles in engineering education. All subsequent analyses draw on this validated set of job postings.

RQ2: Competency Profile of AI-Education Roles

Using the validated corpus, we examined the qualifications and skills employers emphasize to characterize the competency profile of AI-Education roles.

Qualifications

Formal education requirements were common across postings. As shown in Figure 2, a bachelor's degree was the most frequently cited baseline requirement, appearing in 52% ($n = 26$) of postings. Advanced degrees were also prevalent, with 26% ($n = 13$) of roles explicitly requiring a master's or doctoral degree. These patterns suggest that employers view AI-related educational roles as requiring substantial disciplinary and theoretical preparation.

Experience requirements reflected a hybrid professional profile. Many postings specified technical experience related to AI or machine learning alongside pedagogical experience such as teaching or instructional design. The most frequently cited experience bands were four to six years and seven or more years, indicating that these roles are typically targeted toward experienced professionals rather than entry-level candidates.

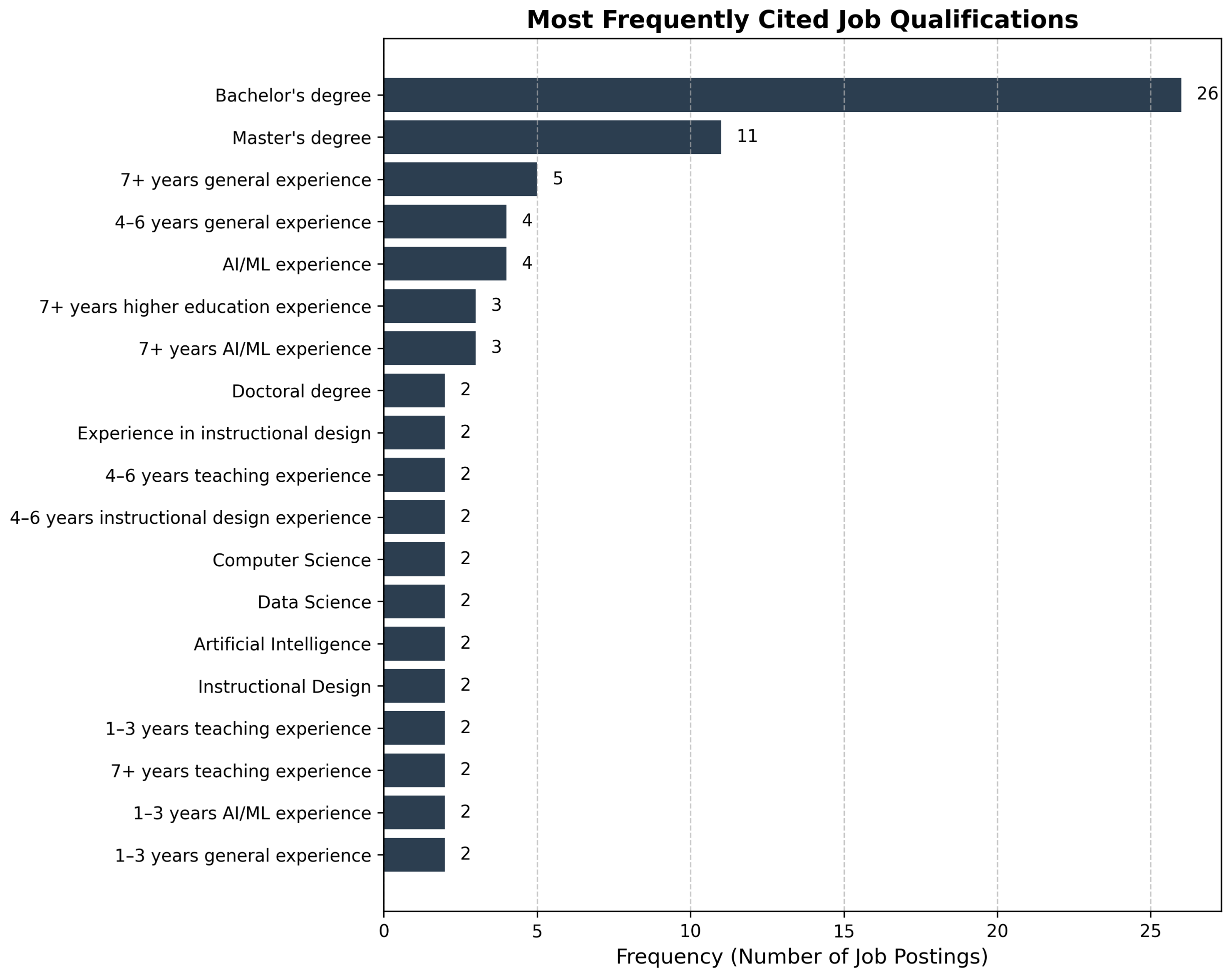

Figure 2: Frequency of canonical qualifications in AI-Education job postings ($N = 50$).

Skills

Skill analysis revealed a strong emphasis on communication and pedagogical expertise. As shown in Figure 3, communication and stakeholder collaboration was the most frequently cited skill category, appearing in 56% ($n = 28$) of postings. Instructional and curriculum design and facilitation, or training delivery, were also prominent.

Technical skills related to AI were present but emphasized applied rather than infrastructural knowledge. AI or machine learning fundamentals, large language models or generative AI, and prompt engineering appeared more frequently than low-level programming or systems-oriented competencies. This pattern suggests that the primary function of these roles is to enable learning, interpretation, and adoption of AI tools rather than to develop AI systems from first principles.

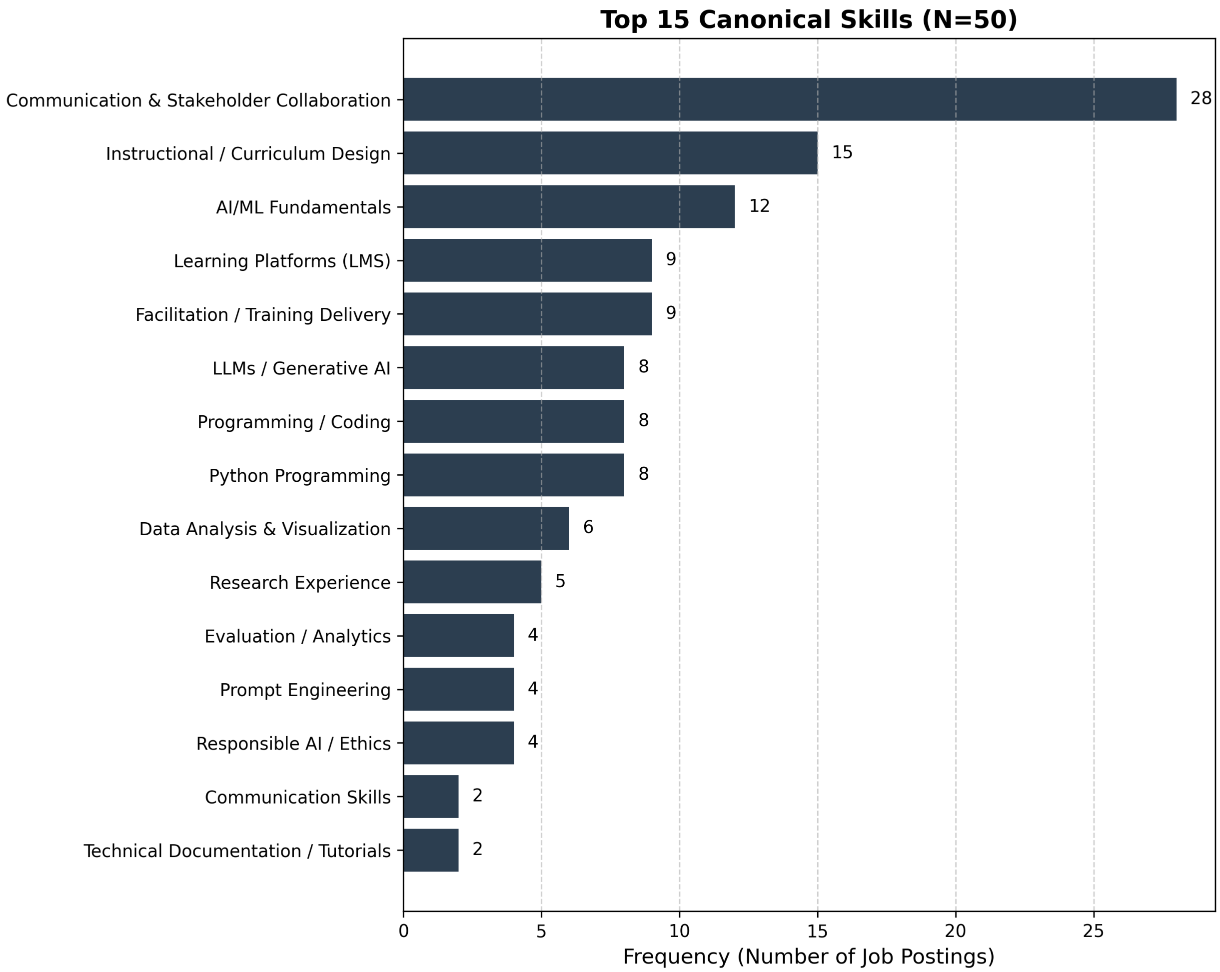

Figure 3: Top canonical skills identified in the corpus.

Taken together, the results establish both the methodological viability of LLM-assisted analysis and the emergence of a distinct competency profile for AI-Education roles. The strong agreement between models and human selection provides confidence in the observed qualification and skill patterns. In the following section, we interpret these findings in relation to engineering education theory and practice, considering how they reframe teaching as a distributed activity across academic, industry, and product contexts, and we discuss implications for curriculum design, workforce preparation, and future research.

Discussion

This study examined how pedagogical work is embedded in AI-related job postings supporting engineering and STEM education, and whether large language models can reliably identify such roles at scale. By combining sectoral analysis, compensation signals, and LLM-assisted evaluation, the findings extend existing engineering education research in two important ways: first, by reframing teaching as a distributed, cross-sector form of labor, and second, by illuminating how this pedagogical labor is differentially valued across institutional contexts. Collectively, these findings point to the broadening of engineering education work beyond traditional roles, with implications for how graduates conceptualize their professional trajectories. The diversification of positions may shape how graduates construct their professional identities and navigate career pathways. At the same time, uneven access to such opportunities raises important questions about equity in who can enter and benefit from this evolving job market. With this, the evolving job market warrants continued attention in both research and practice.

Pedagogy as Distributed, Stratified, and Transformed Labor

Across higher education, edtech, and industry postings, pedagogical responsibilities were consistently present in roles that did not explicitly reference teaching in their titles. These responsibilities included curriculum and content design, facilitation of learning experiences, documentation of complex AI systems, and evaluation of learner or user outcomes. The recurrence of these functions across sectors supports the notion that teaching-related work in AI is increasingly embedded within non-traditional educational roles. [31–33].

Importantly, however, this pedagogical labor is not only distributed but also stratified. While higher education roles tended to formalize pedagogy through instructional titles and degree requirements, industry and edtech roles embedded similar teaching functions within enablement, advocacy, or product-facing positions. This stratification suggests that pedagogy operates under different institutional logics: as formal instruction in higher education, as user enablement and adoption in industry, and as platform-mediated learning in edtech.

Building on this stratification, the findings further indicate that pedagogical labor in AI-related roles is not merely relocated across sectors but fundamentally transformed in how it is enacted. Three overlapping trajectories of transformation are evident across the corpus. First, pedagogy is technologized, with instructional work embedded in AI systems, platforms, and tools. In this form, teaching occurs through interfaces, documentation systems, learning management platforms, and AI-enabled environments that structure how users interact with knowledge. Second, pedagogy is productized, with educational activities restructured into scalable assets such as reusable curricula, tutorials, developer documentation, and training modules designed for distribution across large, diverse user bases. Third, pedagogy is metricized, where learning is increasingly evaluated through analytics, A/B testing, and performance indicators tied to engagement, adoption, or system outcomes. [34, 35].

These transformations reflect a broader shift from teaching as an interpersonal, context-dependent activity toward teaching as a system-level function shaped by scalability, standardization, and organizational performance objectives. In contrast to traditional classroom instruction, where success may be defined through student understanding or interaction, AI-related pedagogical roles often require demonstrating impact through measurable outputs such as user retention, product adoption, or learning analytics.

For engineering education, this reconceptualization has important implications. It challenges static conceptions of teaching as confined to academic roles and instead positions pedagogy as a transferable and economically valuable form of labor that operates across technical, organizational, and product contexts. It also suggests that preparing students for AI-enabled careers requires not only foundational pedagogical knowledge, but also the ability to design learning systems, create scalable educational artifacts, and interpret data-driven measures of learning effectiveness. Recognizing pedagogy as distributed, stratified, and transformed labor, therefore, expands how engineering education defines teaching competence and aligns curricular preparation with the realities of the evolving AI workforce.

Sectoral Differences in Valuation and Compensation

Salary disclosure and compensation patterns further reveal how pedagogical AI labor is valued across sectors. Industry roles had the highest median compensation, followed by edtech, while higher education roles offered substantially lower median salary midpoints. At the same time, higher education postings were more likely to disclose salary information than industry postings, reflecting differing transparency norms rather than differences in labor demand.

These patterns suggest a tension between the *visibility* and *valuation* of pedagogical work. In higher education, teaching-related labor is explicit and transparent but comparatively undervalued in monetary terms. In industry, similar pedagogical competencies command higher compensation but are often obscured under technical or managerial titles, and salary information is less consistently disclosed. Edtech occupies an intermediate position, blending instructional clarity with market-oriented compensation structures.

This differential valuation has implications for workforce mobility and equity. Engineering graduates with pedagogical expertise may encounter structurally different incentives depending on sector, which may in turn shape career trajectories, retention, and professional identity. Recognizing pedagogy as economically valuable labor rather than ancillary “soft skill”—is therefore critical for understanding AI-enabled engineering careers.

The Hybrid Competency Profile Revisited

The hybrid competency profile identified in the Results section gains additional significance when viewed through a sectoral lens. Across all sectors, communication and stakeholder collaboration emerged as the most frequently cited skill, surpassing narrowly defined technical competencies. Instructional design and facilitation were also consistently emphasized, while technical requirements focused on applied AI knowledge such as large language models and generative AI rather than systems-level development.

These patterns suggest that AI-Education roles function primarily as translation and enablement positions. Professionals in these roles are expected to bridge technical systems and human understanding explaining, contextualizing, and governing AI rather than solely building it. The fact that these competencies recur across sectors reinforces their centrality to contemporary engineering practice.

For engineering education programs, this finding strengthens the case for integrating pedagogical design, communication, and responsible AI practices into core curricula [36, 37]. These competencies are not peripheral but foundational to how AI is taught, adopted, and governed in real-world contexts.

Implications for Engineering Education and Career Preparation

Taken together, the results point to several implications for engineering education. First, curricular models that treat pedagogy as the exclusive domain of future faculty may be misaligned with the labor market [8]. Pedagogical competencies appear to be widely demanded across AI-related roles, including those in industry and edtech.

Second, the prominence of pedagogical artifacts such as tutorials, curricula, notebooks, and evaluation plans suggests the value of portfolio-based assessment [38, 39]. Engineering programs may better support student career readiness by designing assignments and capstones that produce artifacts recognizable to employers across sectors [12].

Third, the observed sectoral differences in compensation and role framing highlight the importance of career advising that makes these distinctions explicit [40, 41]. Students should be equipped not only with technical skills, but with an understanding of how similar competencies are framed, valued, and rewarded differently across institutional contexts [36, 42].

Methodological Contributions and the Role of LLMs

This study demonstrates that multiple large language models can converge in identifying pedagogically oriented AI roles when guided by a transparent and theoretically grounded rubric. Rather than replacing human judgement, the LLMs functioned as scalable evaluators whose agreement provided evidence of construct stability.

The use of multiple models mitigated single-model bias and strengthened confidence in classification outcomes [43]. This approach illustrates how LLMs can be responsibly incorporated into mixed-methods engineering education research, particularly for analyzing dynamic labor market data where manual coding alone may be impractical [43, 44].

More broadly, the findings suggest that LLM-assisted analysis can serve as a methodological bridge between qualitative theory and large-scale empirical observation, enabling engineering education researchers to study emerging professional practices as they evolve [45, 46].

Limitations

There are a few limitations to consider when interpreting these findings. First, the corpus consisted of fifty U.S.-based, publicly accessible job postings. While appropriate for exploratory analysis, this scope may not fully capture the breadth and diversity of AI education and pedagogy-related roles, particularly those emerging in international contexts, smaller organizations, or internal hiring pipelines that are not publicly advertised. In addition, duplicate job postings across platforms were not included in the final sample, nor were the few duplicates identified as a result of job postings collected by two separate researchers. When identical or substantially overlapping postings were identified, only one was retained, yielding a final sample of 50 unique job postings. This factor may have reduced the visibility of cross-platform overlap in hiring practices. Consequently, the observed distribution of postings across the three job venues may underrepresent the extent to which employers disseminate identical positions through multiple sites.

Second, although multiple large language models were employed to mitigate model-specific bias, the analysis did not incorporate formal human annotation of rubric-based scores. The absence of human-coded validation limits the ability to assess inter-rater reliability and consensus across evaluators. Future work would integrate domain experts in engineering education and AI to support more nuanced classification and strengthen the credibility of the findings. [47]

Future Work

This study opens several avenues for research and methodological advancement. One primary direction involves expanding the analysis to include human annotation alongside current computational approaches. Incorporating expert human annotators—particularly those with backgrounds in engineering education and AI—would enable more nuanced interpretations by disambiguating expectations that may not be fully captured by keyword-based methods [48]. Building on this, future work could explore consensus-building mechanisms, such as inter-rater reliability methods or Delphi-style approaches, to ensure consistency in the classification of AI-related skills. Such efforts would support the development of a robust skills taxonomy, serving as a foundation for consensus-driven frameworks in workforce trend analysis.

To complement these computational findings, a critical next step involves a qualitative investigation into stakeholder expectations. Future research should include semi-structured interviews and focus groups with diverse stakeholders, including students, employers, recruiters, and faculty. This qualitative layer will provide deeper insight into how different groups perceive AI-integrated roles and identify the greatest gaps between current engineering curricula and industry needs. Understanding these varied perspectives is essential for validating the ”hidden” requirements identified in job postings and for uncovering latent expectations that are rarely formalized in text.

Finally, another direction is to develop an interactive evaluation dashboard to visualize AI workforce demands and their alignment with engineering education curricula. Such a tool would enable stakeholders to explore trends in AI skill requirements across sectors and job types in real-time. By integrating the dashboard with continuously updated job-posting data, researchers and administrators can support data-informed curriculum improvements and pedagogical decision-making [49]. This mapping provides actionable guidance for designing the interdisciplinary experiences necessary to prepare students for the evolving AI-integrated labor market [50]

Conclusion

This study provides a systematic analysis of how pedagogical work is embedded in AI-related job postings supporting engineering education, and evaluates the feasibility of using large language models to identify such roles at scale. By analyzing job postings across sectors using LLM-assisted evaluation, this work provides new insights into the evolving AI workforce and its implications for engineering education.

The findings highlight that pedagogical labor associated with AI and engineering education is not confined to traditional academic settings, but is instead distributed across industry, nonprofit, and many institutional contexts. Moreover, the study demonstrates that while AI-related pedagogical roles are increasingly prevalent, their value varies by sector. This variability stresses the need for engineering education programs to prepare students not only with technical AI competencies, but also with transferable pedagogical, communicative, and interdisciplinary skills that align with emerging workforce demands. In addition, this work illustrates the potential of large language models as tools for workforce analysis in engineering education research. Together, these findings position job posting analysis as a valuable lens for informing curriculum design, pedagogical preparation, and broader discussions about the future of AI-integrated engineering education.

Overall, this study advances understanding of the intersection between artificial intelligence, pedagogy, and engineering education. As AI continues to reshape professional roles across sectors, systematic approaches that connect workforce signals to educational practice will be increasingly critical for designing curricula and pedagogies that prepare engineers for both technical and educational responsibilities in an evolving landscape.

A Prompt Engineering Archive

This appendix contains the full text of the prompts used in the data processing and evaluation pipelines. Python variable assignments and wrapper syntax have been removed to show the raw text input provided to the Large Language Models.

Prompt A.1: Skill Normalization Prompt

```
SYSTEM  You are a meticulous data curator for a labor-market study on AI in Education and
    Engineering Education roles.
Your task is to normalize raw skills into standardized canonical skill categories.

Instructions:
1) For each raw skill phrase, assign ONE canonical category.
2) Canonical labels must be short (26 words), Title Case, and reusable.
3) If two phrases mean the same thing, map them to the same canonical category.
4) If a phrase does not fit any common category, create a new general category.
5) Return STRICT JSON only. No markdown. No extra text.
```

Prompt A.2: Zero-Shot Evaluation Prompt

```
SYSTEM You are an expert evaluator of cross-sector job postings related to Artificial
    Intelligence (AI) and Pedagogy in Engineering Education.

Evaluate the JOB POSTING below to determine how strongly it aligns with AI-related
    pedagogical work  roles that combine AI/ML with teaching, learning design, training,
    documentation, or educational enablement.

JOB POSTING:
{job_posting}

Rate each criterion from 1 (very poor / not present) to 5 (strongly evident):
1. Pedagogical Indicators  Mentions of teaching, designing, facilitating, assessing,
    documenting, or iterating learning activities, materials, or experiences (e.g.,
    curriculum, tutorials, training, workshops, rubrics, or educational guides).
2. Learner / Audience Focus  Clear target learners such as students, educators,
    professional engineers, developers, or broader learning communities.
3. AI Relevance  Mentions of AI/ML, LLMs, GenAI, deep learning, or integration of AI tools
    in learning, training, or documentation contexts.
4. Responsible AI or Governance  References to ethics, fairness, transparency, safety,
    explainability, or responsible AI training and documentation.
5. Cross-Sector Context  Fit with academia, edtech, industry, or government roles that
    blend AI with learning, enablement, or curriculum development.
6. Overall Fit  How well this job represents the type of AI-pedagogical role described in
    the abstract (AI-related teaching, documentation, or educational enablement work).

Respond STRICTLY in valid JSON (no explanations, markdown, or extra text) with the
    following keys:
"score_pedagogy", "score_audience", "score_ai_relevance", "score_responsible_ai",
"score_cross_sector_context", and "score_overall_fit".
```

Prompt A.3: Job Qualification Extraction Prompt

```
SYSTEM You are an expert data analyst. Your task is to extract ALL job qualifications
    mentioned in the provided text.

Extract ONLY the qualifications as they appear in the original text - do NOT rephrase,
    normalize, or group them. Keep them exactly as written.

Qualifications may include:
- Educational requirements (degrees, certifications)
- Years of experience requirements
- Specific experience in domains (AI, ML, education, etc.)
- Skills or knowledge areas mentioned as qualifications
- Any other formal requirements listed

Return a JSON array of strings, where each string is a qualification extracted verbatim
    from the text.

JOB QUALIFICATION TEXT:
{qualification_text}

Respond STRICTLY in valid JSON (no markdown, comments, or extra text).
Example:
[
  "bachelor's degree in Education, Instructional Design Technology...",
  "Minimum of 1 year of demonstrated experience...",
  "master's degree in a relevant field"
]

If no qualifications are found, return an empty array: []
```

Prompt A.4: Qualification Canonicalization Prompt

```
SYSTEM You are a meticulous data curator for a labor-market study on AI in Education and
    Engineering Education roles.
Your task is to normalize the given raw qualifications (degrees and experience
    requirements) into standardized canonical categories, using domain context where
    relevant.

NORMALIZATION RULES FOR QUALIFICATIONS (DEGREES):
- "bachelor's degree", "undergraduate degree", "BS", "BA" -> "Bachelor's degree"
- "master's degree", "MS", "MA", "graduate degree" -> "Master's degree"
- "PhD", "doctorate", "doctoral degree" -> "Doctoral degree"

EXPERIENCE REQUIREMENTS (by domain -- use context to determine):
AI/ML Domain:
- "1 year", "minimum of 1 year" of AI/ML experience -> "13 years AI/ML experience"
- "3+ years", "3 years", "4 years" of AI/ML experience -> "46 years AI/ML experience"
- "7+ years", "8 years" of AI/ML experience -> "7+ years AI/ML experience"

[... rules for Teaching, EdTech, Instructional Design, Higher Ed, General repeated here as
    per original text ...]

CRITICAL INSTRUCTIONS:
1. For each raw qualification, assign it to ONE canonical category.
2. If a phrase contains multiple distinct requirements, split it into multiple canonical
    categories.
3. If a phrase does not fit the rules above, CREATE a new canonical category that is:
   - Short (2-6 words)
   - Title Case
   - General and reusable
4. Ignore filler words and ellipses ("...").
5. Keep the response format STRICTLY as specified below.

INPUT DATA:
{raw_data}
```

Prompt A.5: Skill Extraction Prompt

```
SYSTEM You are an expert data analyst. Your task is to extract ALL required skills
    mentioned in the provided text.

Extract ONLY the skills as they appear in the original text - do NOT rephrase, normalize,
    or group them. Keep them exactly as written.

Skills may include:
- Technical skills (programming languages, tools, frameworks)
- Domain knowledge (AI, ML, education, etc.)
- Soft skills (communication, collaboration, leadership)
- Tool proficiency (software, platforms)
- Any other abilities or competencies mentioned

Return a JSON array of strings, where each string is a skill extracted verbatim from the
    text.

REQUIRED SKILLS TEXT:
{skills_text}

Respond STRICTLY in valid JSON (no markdown, comments, or extra text).
Example:
[
  "Python proficiency",
  "Experience with deep learning frameworks",
  "Strong communication and presentation skills",
  "Ability to explain complex technical concepts"
]

If no skills are found, return an empty array: []
```

Prompt A.6: Full Skill Normalization Rules

```
SYSTEM You are a meticulous data curator for a labor-market study on AI in Education and
    Engineering Education roles.

Your task is to normalize raw skills into standardized canonical skill categories.

NORMALIZATION RULES FOR SKILLS:
- "machine learning", "ML", "deep learning", "AI/ML" -> "AI/ML Fundamentals"
- "large language models", "LLMs", "GenAI", "generative AI" -> "LLMs / Generative AI"
- "prompt engineering", "prompt design", "prompt crafting" -> "Prompt Engineering"
- "curriculum design", "instructional design", "learning experience design" ->
    "Instructional / Curriculum Design"
- "workshops", "training delivery", "facilitation" -> "Facilitation / Training Delivery"
- "documentation", "tutorial writing", "guides" -> "Technical Documentation / Tutorials"
- "AI ethics", "bias", "fairness", "responsible AI" -> "Responsible AI / Ethics"
- "communication", "presentation", "stakeholder engagement" -> "Communication and
    Stakeholder Collaboration"
- "LMS", "Canvas", "Moodle", "learning management system" -> "Learning Platforms (LMS)"
- "A/B testing", "analytics", "learning outcomes measurement" -> "Evaluation / Analytics"
- "Python", "programming", "coding" -> "Programming / Coding Skills"
- "data analysis", "data visualization", "analytics" -> "Data Analysis and Visualization"
- "research", "conducting research", "research experience" -> "Research Experience"

CRITICAL INSTRUCTIONS:
1. For each raw skill phrase, assign it to ONE canonical category.
2. If a phrase contains multiple distinct skills, split it into multiple canonical
    categories.
3. If a phrase does not fit the rules above, create a new canonical category that is
    short, Title Case, and reusable.
4. Ignore filler words and ellipses.
```

Prompt A.7: Full Qualification Normalization Rules

```
SYSTEM You are a meticulous data curator for a labor-market study on AI in Education and
    Engineering Education roles.

Your task is to normalize raw qualifications into standardized canonical categories.

NORMALIZATION RULES FOR QUALIFICATIONS (DEGREES):
- "bachelor’s degree", "undergraduate degree", "BS", "BA" -> "Bachelor’s Degree"
- "master’s degree", "MS", "MA", "graduate degree" -> "Master’s Degree"
- "PhD", "doctorate", "doctoral degree" -> "Doctoral Degree"

EXPERIENCE REQUIREMENTS (by domain, use context to determine):

AI/ML Domain:
- "1 year", "minimum of 1 year" of AI/ML experience -> "1--3 Years AI/ML Experience"
- "3+ years", "3 years", "4 years" of AI/ML experience -> "4--6 Years AI/ML Experience"
- "7+ years", "8 years" of AI/ML experience -> "7+ Years AI/ML Experience"

Teaching/Instruction Domain:
- "1 year", "minimum of 1 year" in teaching or instruction -> "1--3 Years Teaching
    Experience"
- "3+ years", "3 years", "4 years" in teaching or instruction -> "4--6 Years Teaching
    Experience"
- "7+ years", "8 years" in teaching or instruction -> "7+ Years Teaching Experience"

EdTech/Learning Technology Domain:
- "1 year", "minimum of 1 year" in EdTech -> "1--3 Years EdTech Experience"
- "3+ years", "3 years", "4 years" in EdTech -> "4--6 Years EdTech Experience"
- "7+ years", "8+ years" in EdTech -> "7+ Years EdTech Experience"

Instructional Design Domain:
- "1 year", "minimum of 1 year" in instructional design -> "1--3 Years Instructional
    Design Experience"
- "3+ years", "3 years", "4 years" in instructional design -> "4--6 Years Instructional
    Design Experience"
- "7+ years", "8+ years" in instructional design -> "7+ Years Instructional Design
    Experience"

Higher Education Domain:
- "1 year", "minimum of 1 year" in higher education -> "1--3 Years Higher Education
    Experience"
- "3+ years", "3 years", "4 years" in higher education -> "4--6 Years Higher Education
    Experience"
- "7+ years", "8+ years" in higher education -> "7+ Years Higher Education Experience"

General/Unspecified Domain:
- "1 year", "minimum of 1 year" (general) -> "1--3 Years General Experience"
- "3+ years", "3 years", "4 years" (general) -> "4--6 Years General Experience"
- "7+ years", "8+ years" (general) -> "7+ Years General Experience"

CRITICAL INSTRUCTIONS:
1. For each raw qualification, assign it to ONE canonical category.
2. If a phrase contains multiple distinct requirements, split it into multiple canonical
    categories.
3. If a phrase does not fit the rules above, create a new canonical category that is
    short, Title Case, and reusable.
4. Ignore filler words and ellipses.
```

References

[1] B. K. Nagaraj, A. Kalaivani, S. Begum, S. Akila, H. K. Sachdev, *et al.*, "The emerging role of artificial intelligence in stem higher education: A critical review," *International Research Journal of Multidisciplinary Technovation*, vol. 5, no. 5, pp. 1–19, 2023.

[2] Y. Yang, W. Sun, D. Sun, and S. Z. Salas-Pilco, "Navigating the ai-enhanced stem education landscape: a decade of insights, trends, and opportunities," *Research in Science & Technological Education*, vol. 43, no. 3, pp. 693–717, 2025.

[3] A. S. George, "Preparing students for an ai-driven world: Rethinking curriculum and pedagogy in the age of artificial intelligence," *Partners Universal Innovative Research Publication*, vol. 1, no. 2, pp. 112–136, 2023.

[4] F. Pedro, M. Subosa, A. Rivas, and P. Valverde, "Artificial intelligence in education: Challenges and opportunities for sustainable development," 2019.

[5] R. Luckin and M. Cukurova, "Designing educational technologies in the age of ai: A learning sciences-driven approach," *British Journal of Educational Technology*, vol. 50, no. 6, pp. 2824–2838, 2019.

[6] N. A. of Sciences, Medicine, D. of Behavioral, S. Sciences, B. on Science Education, and C. on Educator Capacity Building in K-12 Engineering Education, *Building capacity for teaching engineering in K-12 education*. National Academies Press, 2020.

[7] P. Napier and T. Wada, "Bridging the skills gap: The design education landscape with ai and experiential learning," *Design Management Journal*, vol. 19, no. 1, pp. 19–30, 2024.

[8] A. Padovano and M. Cardamone, "Towards human-ai collaboration in the competency-based curriculum development process: The case of industrial engineering and management education," *Computers and Education: Artificial Intelligence*, vol. 7, p. 100256, 2024.

[9] L. Cao, "Ai science and engineering: A new field," *IEEE Intelligent Systems*, vol. 37, no. 1, pp. 3–13, 2022.

[10] C. Krishnamoorthy and S. Rajeev, *Artificial intelligence and expert systems for engineers*. CRC press, 2018.

[11] C. Liu, G.-C. Wang, and H.-F. Wang, "The application of artificial intelligence in engineering education: A systematic review," *IEEE Access*, 2025.

[12] G. D. Furman, "Enhancing engineering education: The role of artificial intelligence in personalizing learning and outcomes," in *2024 4th International Conference on Big Data Engineering and Education (BDEE)*, pp. 61–65, IEEE, 2024.

[13] L. Ainsworth, *Rigorous curriculum design: How to create curricular units of study that align standards, instruction, and assessment*. Lead+ Learn Press, 2011.

[14] C. Crosthwaite, I. Cameron, P. Lant, and J. Litster, "Balancing curriculum processes and content in a project centred curriculum: in pursuit of graduate attributes," *Chemical Engineering Research and Design*, vol. 84, no. 7, pp. 619–628, 2006.

[15] C. Machovec, M. J. Rieley, and E. Rolen, "Incorporating ai impacts in bls employment projections: Occupational case studies," *Monthly Labor Review*, 2025.

[16] S. Duke, "Ai is changing work — the time is now for strategic upskilling," *World Economic Forum*, 2025.

[17] S. Galeano, N. Hodge, and A. Ruder, "Measuring employer demand for ai skills by educational requirements," *Workforce Currents*, 2025.

[18] G. Alnahdi, J. P. Martin, and R. Maloy, "Ai and engineering careers: Recent graduates' outlook on opportunities and challenges," *International Journal of STEM Education*, vol. 12, no. 1, pp. 1–18, 2025.

[19] S. M. Fain, "To employers, ai skills aren't just for tech majors anymore," *The Hechinger Report*, 2023.

[20] J. Vogt, P. Weber, and U.-D. Ehlers, "Development of a job advertisement analysis for assessing data science competencies," *Data Science Journal*, vol. 22, no. 1, p. 33, 2023.

[21] E. da Silveira, C. Silva, and P. Santos, "Identifying required competencies for data scientists: A job postings analysis," in *Proceedings of the IEEE International Conference on Big Data*, pp. 4543–4552, IEEE, 2020.

[22] G. C. Fleming, J. Smith, and R. Miller, "What engineering employers want: An analysis of technical and professional skills in engineering job advertisements," *Journal of Engineering Education*, vol. 113, no. 2, pp. 251–279, 2024.

[23] S. R. Thomson, B. A. Pickard-Jones, S. Baines, and P. C. Otermans, "The impact of ai on education and careers: What do students think?," *Frontiers in artificial intelligence*, vol. 7, p. 1457299, 2024.

[24] L. Liu, "The future of labor markets: The evolution of ai and changing human roles," *Geographical Research Bulletin*, vol. 2, pp. 238–246, 2023.

[25] R. Baldwin, *The globotics upheaval: Globalization, robotics, and the future of work*. Oxford University Press, 2019.

[26] D. M. West, *The future of work: Robots, AI, and automation*. Bloomsbury Publishing USA, 2018.

[27] M. Alghazo, V. Ahmed, and Z. Bahroun, "Exploring the applications of artificial intelligence in mechanical engineering education," in *Frontiers in Education*, vol. 9, p. 1492308, Frontiers Media SA, 2025.

[28] N. R. Aljohani, A. Aslam, A. O. Khadidos, and S.-U. Hassan, "Bridging the skill gap between the acquired university curriculum and the requirements of the job market: A data-driven analysis of scientific literature," *Journal of Innovation & Knowledge*, vol. 7, no. 3, p. 100190, 2022.

[29] S. Zeidan and M. Bishnoi, "An effective framework for bridging the gap between industry and academia," *International Journal on Emerging Technologies*, vol. 11, no. 3, pp. 454–461, 2020.

[30] J. Gu, X. Jiang, Z. Shi, H. Tan, X. Zhai, C. Xu, W. Li, Y. Shen, S. Ma, H. Liu, S. Wang, K. Zhang, Y. Wang, W. Gao, L. Ni, and J. Guo, "A survey on llm-as-a-judge," *The Innovation*, vol. 6, 2025.

[31] R. Luckin, *Enhancing Learning and Teaching with Technology: What the Research Says*. ERIC, 2018.

[32] H. S. Beetham, H. Beetham, and R. Sharpe, *Rethinking pedagogy for a digital age*. routledge London, 2007.

[33] A. E. McFarlane, "Devices and desires: Competing visions of a good education in the digital age," *British Journal of Educational Technology*, vol. 50, no. 3, pp. 1125–1136, 2019.

[34] C. J. Okocha, J. J. Blanchard, G. J. Kim, and M. Adachi, "Wip: Navigating unfamiliar waters: Enhancing intercultural understanding and academic self-efficacy among computer science students in japan," in *2024 IEEE Frontiers in Education Conference (FIE)*, pp. 1–5, IEEE, 2024.

[35] J. W. C. Chibuzor Joseph Okocha, Gloria J. Kim and Y. K. Yoon, "Microelectronics research and global competencies: Unpacking research abroad experiences of engineering students," in *2024 ASEE Annual Conference & Exposition*, no. 10.18260/1-2–47777, (Portland, Oregon), ASEE Conferences, June 2024. https://peer.asee.org/47777.

[36] D. Santandreu Calonge and M. Aman Shah, "Moocs, graduate skills gaps, and employability: A qualitative systematic review of the literature," *International review of research in open and distributed learning*, vol. 17, no. 5, pp. 67–90, 2016.

[37] J. Cañavate, E. Martínez-Marroquín, and X. Colom, "Engineering a sustainable future through the integration of generative ai in engineering education," *Sustainability*, vol. 17, no. 7, p. 3201, 2025.

[38] R. Lam, *Portfolio assessment for the teaching and learning of writing*. Springer, 2018.

[39] D. J. Cole, *Portfolios across the curriculum and beyond*. Corwin Press, 2000.

[40] A. Ramaswami, G. F. Dreher, R. Bretz, and C. Wiethoff, "Gender, mentoring, and career success: The importance of organizational context," *Personnel Psychology*, vol. 63, no. 2, pp. 385–405, 2010.

[41] C. K. Chan and L. N. Hedden, "The role of discernment and modulation in enacting occupational values: How career advising professionals navigate tensions with clients," *Academy of Management Journal*, vol. 66, no. 1, pp. 276–305, 2023.

[42] V. N. Gordon, “The 3-i process: A career-advising framework,” *NACADA Journal*, vol. 39, no. 2, pp. 64–71, 2019.

[43] J. P. Wiebe, R. Khan, S. Burns, and J. D. Slotta, “Qualitative research in the age of llms: A human-in-the-loop approach to hybrid thematic analysis,” in *Proceedings of the 19th International Conference of the Learning Sciences-ICLS 2025, pp. 1123-1131*, International Society of the Learning Sciences, 2025.

[44] X. Ding, M. Gopannagari, K. Sun, A. Tao, D. L. Zhao, S. Varadhan, B. L. B. Hardy, D. Dalpiaz, C. Vogiatzis, L. Angrave, *et al.*, “Evaluation of llms and other machine learning methods in the analysis of qualitative survey responses for accessible engineering education research,” in *2024 ASEE Annual Conference & Exposition*, 2024.

[45] R. H. Tai, L. R. Bentley, X. Xia, J. M. Sitt, S. C. Fankhauser, A. M. Chicas-Mosier, and B. G. Monteith, “An examination of the use of large language models to aid analysis of textual data,” *International Journal of Qualitative Methods*, vol. 23, p. 16094069241231168, 2024.

[46] A. S. Hayes, ““conversing” with qualitative data: Enhancing qualitative research through large language models (llms),” *International Journal of Qualitative Methods*, vol. 24, p. 16094069251322346, 2025.

[47] X. Wang, H. Kim, S. Rahman, K. Mitra, and Z. Miao, “Human-llm collaborative annotation through effective verification of llm labels,” in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pp. 303:1–303:21, Association for Computing Machinery, 2024.

[48] A. Szymanski, N. Ziems, H. A. Eicher-Miller, T. J.-J. Li, M. Jiang, and R. A. Metoyer, “Limitations of the llm-as-a-judge approach for evaluating llm outputs in expert knowledge tasks,” in *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI)*, Association for Computing Machinery, 2025.

[49] S. Luna, A. Lopes, A. Bahabry, and A. Akundi, “Trends of systems engineering job postings and their implications for curriculum development,” *ASEE Annual Conference & Exposition*, 2022.

[50] M. Alhamadi, “Challenges, strategies and adaptations on interactive dashboards,” in *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, Association for Computing Machinery, 2020.