This article has been accepted for publication in IEEE Transactions on Technology and Society. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTS.2025.3567143

AI Ethics and Governance in the Job Market: Trends, Skills, and Sectoral Demand

Lucas Wiese, *IEEE Student Member*, Sonali Subbu Rathinam, Matthias Oschinski, Bryan DeWitt, and Daniel S. Schiff*, *IEEE Member*

***Abstract*—Demand for an AI-literate workforce has surged, in large part to counter a growing skills gap. Meanwhile, expertise in ethical and governance dimensions of AI is increasingly deemed crucial to handle various organizational, regulatory, and social concerns. However, the focus of AI literacy efforts to date has been primarily technical. This paper helps close this gap by providing the first large-scale analysis of AI ethics and governance skills sought by employers in the labor market. Drawing on more than four million job postings for AI-related professions over the years 2018-2023, we provide an empirically-grounded characterization of AI ethics and governance competencies and perform associated descriptive analyses. We find that professionals with AI ethics and governance competencies are requested by employers to hold diverse skill sets, covering technical, managerial, and regulatory domains, though the two professions remain distinct. Moreover, the demand for expertise in these domains has grown rapidly, both in absolute terms and as a proportion of AI-related job postings. More than 100,000 professionals with expertise in AI ethics and governance are now requested annually, with the concentration highest in the financial and information sectors. These findings can help individuals, employers, and institutions of higher education better design job requirements, educational programs, and individual learning pathways, closing the career competency gap in AI ethics and governance.**

***Index Terms*—AI literacy, AI ethics, AI governance, workforce development, labor market, career competencies**

I. INTRODUCTION

THE demand for artificial intelligence (AI)-literate workers has grown substantially in recent years, with significant efforts by policymakers, industry, and scholars to cultivate a workforce capable of realizing AI's role in innovation and provision of goods and services. Calls to prepare workers emanate from local to international scales, with initiatives led by government, educational institutions, and the private sector. For example, U.S. initiatives like AI4K12 and AI4ALL, Finland's Elements of AI, and China's AI+X program aim to cultivate a literate workforce through primary, secondary, and tertiary education, certificates, and retraining initiatives [1], [2], [3], [4], [5]. Initiatives are aimed not only at future AI developers with advanced graduate education, but also at workers who may interact with or utilize AI, and to members of the general public [6]. Government policies like the US National AI Initiative Act [7] are now providing funding in the billions of dollars, as an extension of prior STEM-focused workforce development strategies. These calls to invest in and develop an AI workforce signal a concerning gap regarding the current workforce's capacity and the desired skills needed to maintain an innovative yet equitable AI workforce. Extending from a decades-long concern about STEM skills competitiveness [8], this mismatch is broadly described as the *AI skills gap*.

As defined by Wu and Lin [9, p. 1], "this gap refers to the disparity between the growing demand for AI expertise and the available talent pool—spanning both individual workforce capabilities and national technological capacities." Thus, the AI skills gap has both domestic and global invocations, serving as a hidden catalyst that exacerbates economic disparity and social inequality. On this logic, those who hold the skills, resources, and knowledge to shape and use AI will continue to innovate, rendering it less feasible for other communities and countries to catch up, causing the gap to expand at an unprecedented rate compared to previous technologies.

An underlying skills mismatch in the labor market has long been recognized [10], [11], especially in STEM. However, the skills mismatch in AI is particularly unique; the large investment in AI is hypothesized to play a heightened role in inducing labor disruption and displacement. Thus, it may be that AI advancement both silos AI jobs and disrupts others. The scale of change, the requirements for upskilling or reskilling, and the pace of change raise the stakes of the skills gap from concerning to potentially economically and socially existential [12], [13], [14], [15]. Heightened power concentration, withdrawn economic development ladders, and uncertainty about even the feasibility of keeping up with the pace of innovation threatened to exacerbate the distance between the haves and have-nots.

Received 14 February 2025. Wiese, Rathinam, Oschinski, and Schiff contributed equally and are co-first authors. (*Corresponding author: Daniel S. Schiff). This work was in part supported by the U.S. National Science Foundation under award #2134667.

Lucas Wiese (lwiese@purdue.edu), Bryan DeWitt (bdewitt@purdue.edu), and Daniel S. Schiff (dschiff@purdue.edu) are with Purdue University, West Lafayette, IN 47907 USA. Sonali Subbu Rathinam (ss4608@georgetown.edu) and Matthias Oschinski (matthias.oschinski@georgetown.edu) are with the Center for Security and Emerging Technology, Georgetown University, Washington, DC 20001 USA.

A. Beyond Technical AI: The Need for Socio-Ethical and Governance Competencies

To mitigate the AI skills gap, stakeholders have thus proposed a wide array of strategies: increased international collaboration, social dialogue, equitable policy frameworks, and the potential of educational reform and institutional acceptance of AI education in preparing workers across disciplines [16], [17], [18], [19]. Yet, a critical element of this effort is the recognition that the AI skills gap is not manifested just through the (uneven) distribution of technical expertise. Instead, desired skills in AI literacy also consist of the capacity to navigate complex socio-ethical issues, advanced critical thinking and creative reasoning, and interpersonal and interdisciplinary collaboration skills, as well as more traditional technical expertise [20], [21], [22], [23]. Associated AI ethics and AI governance skills gaps are the focus of this paper.

Importantly, the aims of AI literacy—or competency—initiatives also go beyond depositing necessary job skills for a given workplace [21], [22]. They include training students and future professionals to be independent agents capable of their own learning and personal and professional growth [24], [25]. Insights from ethics education, professional ethics, AI literacy, and business ethics are all pertinent. For instance, [20] recommends that AI literacy efforts address how professionals understand the socio-legal-technical impacts of their work, practice collaboration with each other, and engage in self-reflection over professional practices. Sidra and Mason [25] similarly highlight the unique metacognitive skills required for effective work in human-AI collaboration, corresponding with accounts of day-to-day problems for AI ethicists [26], [27]. AI ethics and governance professionals also do not solve problems in a detached, clean environment; moral decision-making is subject to cognitive biases and heuristics and beset by organizational pressures [28], and these dimensions of "what it means to be an AI ethicist" are among many of importance to consider for AI deployment in organizational contexts.

Moreover, while these skill sets and competencies are indeed essential for the *technical design and development* of AI systems, they are equally essential for their *responsible adoption, implementation, monitoring, and governance*, including at the level of executives [29]. These efforts are likely to require efforts of (at least) tens of thousands of individuals over many years, and across societal sectors. For instance, to help manage AI's implications for safety, ethics, and human rights, recent years have seen landmark governance efforts like the European Union AI Act [30], or the U.S. NIST AI Risk Management Framework [31] which include proposals for independent auditing, conformity assessment, safety testing, organizational governance, and risk management.

Recent policy and recommendations emanating globally, such as from the since-revoked 2023 U.S. Executive Order on Safe, Secure, and Trustworthy AI [32] include: requirements for public sector agencies to recruit chief AI officers and advisory councils; the public posting of AI inventories and evaluation results; the development of responsible AI maturity models for private organizations; and consortia of interested professionals who participate in organizations, such as the Ethical AI Governance Group and Responsible AI Institute, potentially receiving certification in AI ethics or governance skills [33], [34]. Meanwhile, among consortiums and public policy, an industry has emerged to address responsible AI, trustworthy or ethical AI, and AI governance within private and public sector organizations [35], [36]. That is, a consistent feature of AI policy and innovation is the need for professional practitioners who can design and maintain frameworks for responsible AI within and across organizations, ranging from hospitals to schools to manufacturing floors.

Thus, while efforts to build the technical AI workforce, emphasizing skills in computer science, software engineering, and data science remain dominant [6], this paper highlights the need to understand and prepare the workforce with respect to competence in AI governance and AI ethics. A growing number of initiatives do exist, including undergraduate and graduate courses [37], graduate degree programs [38], professional certifications [39], and trainings offered by civil society groups. Another body of research discusses the prospects for teaching AI ethics and governance to future practitioners [40], [41], [42], aligned with a growing body of scholarship addresses the practice of AI ethics within organizations and associated challenges [35], [43], [44]. Further, a portion of research has sought to specifically understand the competencies for AI ethics practitioners in particular, a necessary step and literature to which this study most directly contributes [45], [46], [47].

B. Addressing Foundational Knowledge Gaps in AI Ethics and Governance Workforce Development

Despite growing recognition of the importance of AI ethics and governance competencies, major gaps remain in scholars' and decision-makers' understanding of how these fields are developing. Key uncertainties include: 1) the scale and distribution of demand for AI ethics and governance skills across sectors; 2) the degree of alignment or mismatch between employer needs and available talent; and 3) the design of effective educational programs, career pathways, and instructional strategies to prepare future practitioners. For instance, to what extent are skills in related fields, like privacy, compliance, or risk management, transferable to these new AI ethics and governance roles? Further, should organizations or educators concentrate AI ethics and governance skills within dedicated specialist roles, or should these skills be fostered across all technical and AI-engaged professionals? This paper takes an initial step toward addressing these foundational issues by conceptualizing these issues, analyzing large-scale labor market data, and offering empirical insights into the emerging composition of AI ethics and governance professions.

In short, calls for safe, trustworthy, and responsible AI governance will require attention comparable to that of training of future data scientists and technical AI practitioners. Yet we have very little empirical understanding of these professionals, their roles, and which skills are demanded of them or truly needed of them. This paper aims to help close these gaps by focusing on the composition of skills demanded by AI ethics and governance jobs. Through our analysis, it is important to note that our findings are drawn from empirical evidence in *job*

postings, but there remains a larger gap between the realities of these jobs and the market's conception of them. This may further reflect a limited understanding or even normatively undesirable neglect of key skills and roles surrounding ethics. This paper is a first step in connecting prior work [26], [27], [48] on AI ethics to these system-level trends.

Specifically, we examine skills related to either AI ethics or AI governance in job posting data from 2018-2023 compiled by Lightcast.[2] In this descriptive analysis, we consider the following three research questions (RQ):

- **RQ1**: What core competencies characterize AI ethics and governance roles as reflected in the labor market? Specifically, what skill subcategories—including technical, managerial, regulatory, and socio-ethical—are most desired for professionals in AI ethics and governance-aligned job postings?
- **RQ2**: How has the demand for professionals with AI ethics and governance competencies changed from 2018 to 2023?
- **RQ3**: How does the demand for professionals with AI ethics and governance competencies vary across industry sectors?

Next, we present our conceptual framework that outlines the assumptions guiding our research and findings of the AI (ethics and governance) skills gap. This framework situates the role of individual learners, employing organizations, and higher education institutions as some key stakeholders in addressing the skills gap. Following, we present our methods which describe our use of job posting data to determine competencies and the analytical strategies used on this data to investigate job postings for AI ethics and AI governance roles from 2018 to 2023. Our results section highlights key findings across our three research questions, from understanding core competencies to trends over time and sectors. Finally, we discuss the implications of our findings as the first comprehensive analysis on the demand of AI ethics and AI governance using real-time job posting data. As our results further operationalize AI ethics and AI governance as concrete professional domains, and draw distinctions between them, our study has practical implications for educational design, policy development, and professional organizations who are defining these fields moving forward.

II. Conceptual Framework

We approach our investigation under the assumptions that (1) individual learners, (2) employers, and (3) higher-education institutions have a shared responsibility to effectively contribute to a skilled AI ethics and AI governance workforce. As a corollary, none of these actors alone can be solely responsible for driving an AI-literate workforce: the uncertainties surrounding these skills, regulatory and socioeconomic developments, and the change of pace and AI technology make planning especially difficult. Under these conditions, it is unrealistic to expect an individual to fully navigate these dynamic conditions, to expect employers to be fully equipped or incentivized to upskill their employees, or to expect higher-education institutions to become prescient workforce development factories. Understanding the interaction between these three actors (while not comprehensive) is critical to close the gap in the availability and empowerment of AI ethics and AI governance professionals.

The interactions between these three actors are additionally influenced by efforts made by the government and regulatory agencies, at both national and local levels. As depicted in Fig. 1, these government and regulatory bodies are foundational to shaping, informing, and encouraging the competencies and capabilities that the labor market demands. For instance, government agencies are motivated to establish national priorities, such as economic growth and national security, and build public trust in AI for a fair and competitive market. Government agencies also strategically engage in agenda-setting for AI regulations and policy. This may include establishing standards and guidelines for AI development, or funding and investing in STEM education initiatives or AI research at large. In total, individual learners, higher educational institutions, and employers are influenced by the strategic actions emerging from the public sector. Finally, these dynamics are likely affected by other actors as well not depicted here, like media, civil society, and the general public.

The framework presented in Fig. 1 represents an operational understanding of the role of these three actors in addressing the aforementioned AI skills gap [16]. It builds on the existing *Competency-Capability Chasm* (CCC) as developed by Ward *et al*. [49], [50], which represents the distance between the capabilities taught from higher education institutions and the competencies sought in the job market. That is, the CCC "exists between what is taught and what is sought" [50, p. 898]. We expand this understanding to the AI context by situating individuals, employers, and institutions within the chasm to conceptualize responsible actors, institutions, and mechanisms that might be activated to address the AI skills gap.

For each actor, we present examples of key motivations potentially driving their behavior in this ecosystem as well as strategic actions they might take. For instance, individual learners may be motivated both by the interest in securing career opportunities in AI ethics or governance roles, as well as promoting their job security in an uncertain labor market generally. Meanwhile, institutions of higher education are motivated by ensuring their graduates are prepared for these roles, as well as supporting organizational goals like student recruitment and stable funding. Employers may likewise be motivated by the need to comply with regulations and promote trust of consumers or the public.

[2] Lightcast is a labor market analysis company that collects job postings from multiple sources, and provides insight into the labor market.

This article has been accepted for publication in IEEE Transactions on Technology and Society. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTS.2025.3567143

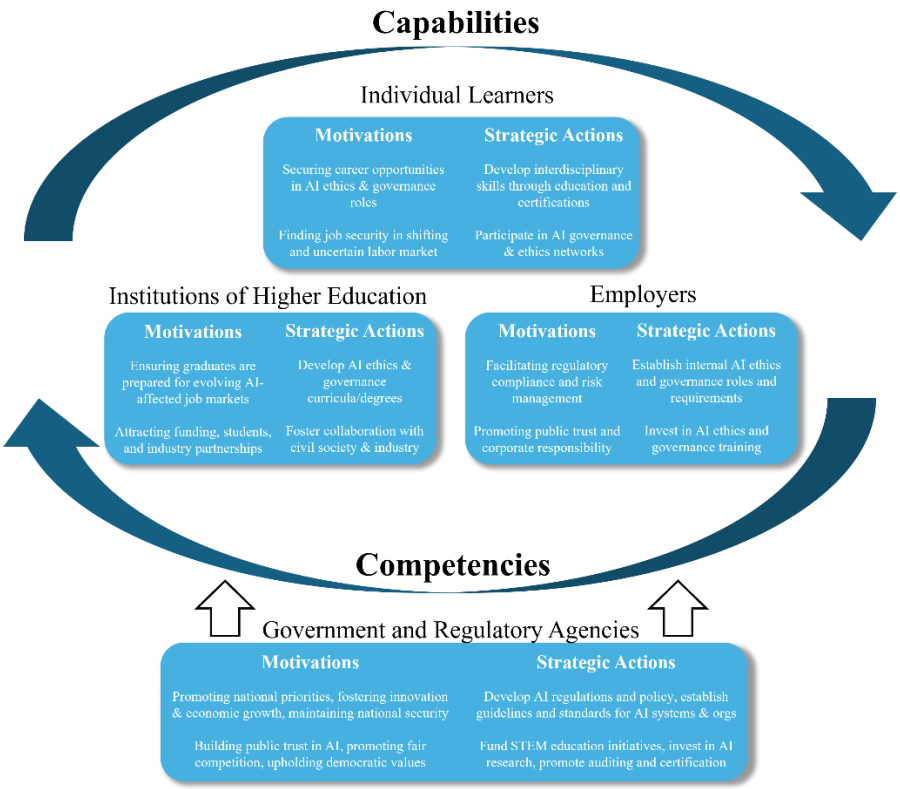

Fig 1. Adapted Competency-Capability Chasm from [50] to account for individuals, employers, higher education institutions, and the public sector. The framework represents how shared commitment between individuals, employers, and higher education institutions is necessary to close the capability-competency chasm for AI ethics and governance.

These motivations suggest ways in which the actors' decisions and actions are entangled. For instance, universities that remain unaware of needs in industry may fail to promote the relevant skills needed in their programs. Meanwhile, employers which present a narrow vision of AI ethics and governance roles may lead to the creation of educational programs that are unattractive for universities and students. Along these lines, Fig. 1 presents examples of suggested strategic actions that can help support all three actors. In turn, this conceptualization affords scholars the ability to map the findings of the present study onto the broader literature, and helps future researchers and practitioners identify pathways to close the gap further.

A. The Nature of Instructional Design and Aligning Individuals, Employers, and Institutions

For this vision to be realized, it is important to understand the traditional approach universities take with instructional design, including when developing new educational programs or degrees. This includes how institutions respond to individual learners and draw ideas from employer needs (or not). Instructional designers in particular must apply their expertise, experience, and understand the context to solve complex problems in educational design, which are often novel [51]. That is, designing for AI ethics and AI governance education requires synthesizing the requirements and constraints for each actor in this model.

However, conventional approaches, while valuable, have focused primarily on the start of a learner's journey (what skills or competencies they bring when beginning a program) and typically design forward from there to scaffold their learning [52]. Unfortunately, this strategy often fails to account for how the learning journey will culminate into professional applications down the road; that is whether the articulated learning goals are reasonable, useful, and sufficient. This problem is not easily solved by backwards design approaches either, but we argue it requires deeper synthesis and collaboration.

For instance, appropriate learning paths must not only account for the sought-after skills articulated as important in the labor market; they should also account for learners' motivations, needs, and expectations [53]. Moreover, learners' motivations are themselves often inextricably linked with the job market–they want the education that will afford them the flexibility, opportunity, impact, and compensation that they desire [54], [55]. Consequently, this framework's effectiveness depends on recognizing what each stakeholder brings to the table dynamically, rather than from fixed assumptions. Importantly, this does not assume that employers' articulation of what the job market does or should require is the ground truth. Rather, employers may also have gaps in their knowledge

and understanding, and may be overly focused on short-term goals or attached to prior ways of thinking. As argued in [49], [50], the CCC is a space where stakeholders should constantly assess, reassess, and realign workforce development and educational efforts.

When addressing the CCC, a few strategic options may be prudent. First, individual learners must do the necessary reflection and articulate their desires, trajectories, and interests. This is no doubt informed by both the evolving demands of the labor market and the messaging from higher education institutions that construct curricula and opportunities for students. Second, employers must maintain a culture and atmosphere of continuous learning, growth, and professional development [56]. With this, they can maintain an adaptive workforce that responds to the fast paced AI landscape, without outsourcing employee learning entirely to individuals or educational institutions [57], a growing approach which risks increased employee turnover and skill stagnation.

Helpfully, employers can leverage programs like (in the US) the Section 127 Educational Assistance Program for up to $5,250 in annual expenses per employee for professional development [58]. Third, higher education institutions must evolve beyond traditional instructional design practices and foster foundational capabilities, respond to evolving market-aligned competencies, and respect the individual's experiences and motivations. This might include innovative approaches like non-credit-to-credit or earned-admission pathways that provide flexible degree trajectories [59], [60], [61]. Yet, universities are also in a position to help identify and shape the skill demands and future roles of students, not merely react to the market, especially in dynamic settings like AI, or in arguably underdeveloped settings like ethics.

These shared commitments, as we propose in Fig. 1, can thus assist with more robustly closing the AI skills gap. Next, we will examine the context surrounding AI ethics and AI governance skills in particular.

B. Specifying the AI Ethics and AI Governance Domains

Designing and developing AI systems comes with a plethora of social and ethical considerations to protect human life and societal welfare. As such, ethical principles and practices help to guide responsible AI, and governance mechanisms must exist to connect and align AI systems with human values [62], [63]. Across settings and discipline, there are practices that professionals can implement at each stage of the AI development lifecycle, policies that organizations can enforce, and national policy to govern the AI industry and AI systems at large [64] [25]. While a full review of relevant AI ethics and AI governance strategies is outside the scope of this paper, it suffices to say that accomplishing these widely demanded goals of safe and responsible AI requires unique socio-technical skill sets that are important to target.

Responsible AI practitioners may range from technical researchers, computer scientists, and engineers, to policy experts and managerial professionals. Across these occupations, [65] present a helpful taxonomy of different skills thought necessary across technical and non-technical professionals including: technical research capacities, interdisciplinary communication, analytical systems thinking, policy integration strategies, and business or human resources expertise [65]. Extending the complexity of this unique, multidisciplinary domain [26] argues that an AI ethicist should possess a unique capacity to "apply and communicate ethical principles in the context of artificial intelligence within a corporate structure" which may require special skills like bravery [26, p. 4]. Moreover, the AI ethicist must be able to hold conflicting personal and professional identities; As identified by [27], the motivations of an AI ethicist may be in conflict with those of their employing institution—whether because of financial, social, or philosophical differences.

While more explicit research has been done on the professional capacities of the AI ethicist, we argue that this understanding applies to AI governance professionals as well. AI governance professionals are needed to understand technical components of the AI development lifecycle and integrate legal, regulatory, and compliance frameworks seamlessly, minimizing the likelihood of harm while maximizing the capacity and innovation of AI systems. And while many of these dimensions are discussed in policy and educational discourse, and present in general AI literacy frameworks, it remains unclear exactly which sets of skills are important or prioritized for AI ethics and governance professionals.

As such, we next detail the methods we used to conduct a large-scale job posting analysis to study the sought-after skills for these types of professionals. We emphasize again that the current state of demands from employers do not fully capture the ground truth about these professions or their future. However, they provide important signals about how AI ethics and governance skills and professions are emerging; what skills are imagined to be important; to what extent these professionals are unique compared to professionals in related roles; and to what extent AI ethics and governance goals are concentrated or distributed.

III. METHODS

A. Rationale and Data Collection

To ascertain skill requirements in AI ethics and AI governance occupations, we leverage real-time labor market information from Lightcast, particularly their job posting data. Lightcast collects job postings data by scraping over 65,000 websites globally, including company career pages, national and local job boards, as well as data from job aggregators [66]. The use of job posting data, and Lightcast data in particular, is increasingly common in labor market studies [67], [68], [69] and has been used for purposes such as analyzing skills demand [70], [71], assessing emerging job markets [72], [73] and measuring the effects of technological change [74], [75]. Closely related to the current context, these data have been used to assess trends in digital literacy [76] and competencies in the green economy [77]. These data have also proven valuable for analyzing the evolution of AI-related job postings and skill requirements [78], [79].

This article has been accepted for publication in IEEE Transactions on Technology and Society. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTS.2025.3567143

Fig. 2 illustrates our methodological approach. In the first step, we follow [80] in using keywords that designate an occupation as AI-related. The initial dataset includes U.S. job postings with associated metadata such as company names, skills, industry taxonomies [81] and occupational taxonomies [82]. Lightcast removes duplicates and we only include job postings that contain skills mapped to Lightcast's publicly-available skills taxonomy.

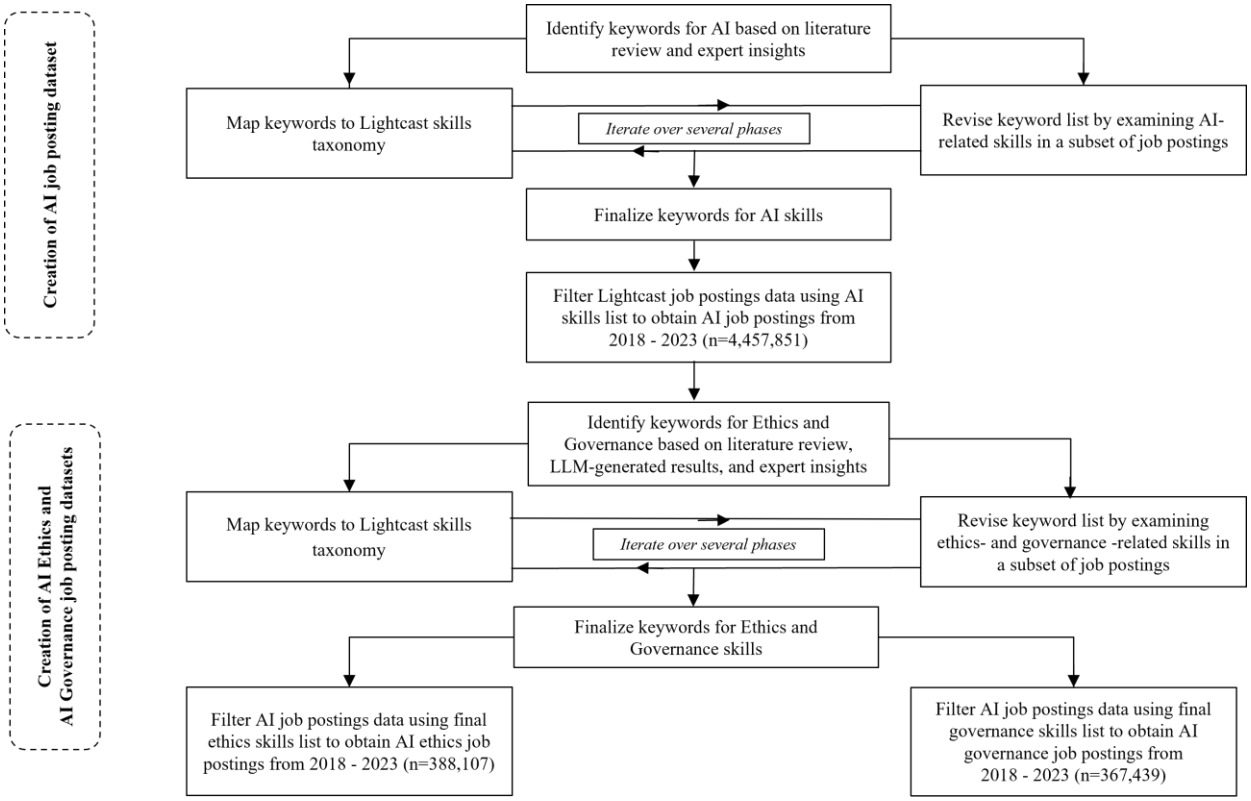

Fig. 2. Determining AI ethics and AI governance occupations

We used a total of 72 AI (skills) keywords to extract the primary AI-related dataset, all available in Lightcast's formal typology of terms [83]. To come up with our list of AI keywords that we crosswalk with Lightcast's available terms, we drew on prior taxonomic research from [79] and augment this list by examining other AI-related keywords articulated in the Lightcast skills taxonomy. We iterated this list over several rounds as outlined as best practice for textual analysis in [84], [85] and depicted in Fig. 2. This process involved examining a subset of identified job postings and adjusting keywords iteratively to ensure greater completeness and accuracy of coverage.

Due to contractual limitations with our data provider, we cannot share the proprietary set of final keywords for the AI, ethics, and governance datasets. However, Table 1 presents examples of overarching categories synthesized from the 72 skills keywords which were used to identify the initial AI-related job postings. Finally, note that all keywords must definitionally be part of the formal Lightcast typology to enable extraction based on those keywords. The resulting dataset comprises 4,457,851 job postings for the years spanning 2018 to 2023, covering the demand for AI-related occupations.

TABLE 1
EXAMPLE CATEGORIES OF AI-RELATED SKILLS KEYWORDS

Artificial intelligence	Automate	Computer vision
Data mining	Deep learning	DevOps
Human-computer interaction	Knowledge engineering	Machine learning
Natural language	Neural net	Prompt engineering
Reinforcement learning	Robotic	Speech recognition
Supervised learning	Unsupervised learning	

Next, we conducted a second stage of iterative data collection to pinpoint job postings specifically associated with AI ethics and AI governance, respectively, using a refined set of 30 ethics and 58 governance keywords. This process was similarly informed by literature and expert sources but also augmented by large language model support. To begin, we used OpenAI's GPT-4 large language model to brainstorm keywords, consistent with the strategy used by [80]. For example, terms associated with AI ethics included: fairness, accountability,

transparency, privacy, bias, mitigation, and ethical decision-making. For AI governance, terms included: policy development, compliance, risk management, standards development, oversight, and stakeholder engagement. Next, we expanded this keyword set based on additional keywords from previous studies on AI ethics and AI governance concepts [86], [87], [88]. After expert review and several rounds of iteration, similar to the process employed for the superset of AI job postings, we identified final sets of (skills) keywords for AI ethics and governance jobs.

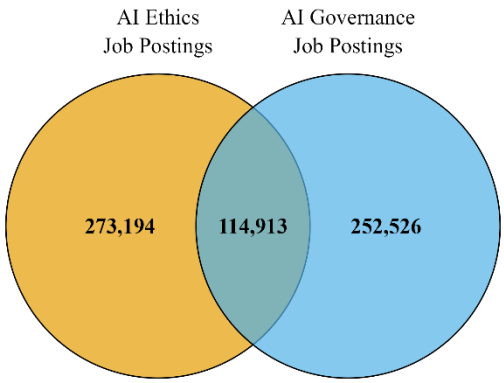

Fig. 3. Frequency of AI ethics and AI governance job postings

The results provided us with an AI ethics (job posting) dataset of *n*=388,107, an AI governance (job posting) dataset of *n*=367,439, and a dataset of job postings which contain *both* AI ethics and AI governance terms of *n*=114,913. See Fig. 3 for a Venn diagram of these datasets and their overlap.

We describe these individuals as *AI ethics or governance professionals*, though this should be understood to include individuals with a variety of roles and both greater or lesser degrees of focus on AI ethics and governance. For instance, a data scientist or risk management professional might devote much, or only some, of their job to AI ethics or governance. However, importantly, all individuals are ostensibly working in AI-related roles, and the jobs they apply to require skills like 'privacy law' or 'fairness.' So, while most of these professionals would not have job titles like "AI Governance Lead" or "AI Ethics Manager," they are working at the intersection of AI and ethics or governance, at least to some degree.

B. Analysis

Our analysis procedure was designed to address each of our RQs. To answer RQ1, regarding the core competencies required of AI ethics and AI governance-aligned professionals, we conducted a *skill analysis using location quotients* (LQ) and a *job title frequency analysis*. As such, first, we calculated the LQ for ethics- and governance-related skills across job postings.

The LQ is typically used to measure the concentration of a specific industry or occupation in a region to its concentration in a broader geography as a measure of specialization [89]. In this context, it measures the relative representation (or importance) of a given skill in AI ethics or governance job postings compared to that skill's overall representation (or importance) in the *entire* job posting dataset. For instance, the entire job posting dataset for 2023 has approximately 39 million job postings. Skills with an LQ greater than 1 are considered highly represented in the 388,107 AI ethics job postings or 367,439 AI governance job postings, respectively, meaning these skills are especially likely to be *particular emphases of AI ethics or governance jobs* rather than just generally important skills (i.e., they are core competencies). Note that, for ease of interpretation, we aggregated individual skills to a pre-defined set of higher-level skill subcategories, discussed in more detail later.

Next, we identified the most frequent *jobs or job titles* associated with AI ethics and AI governance skills by calculating the share of total AI ethics or governance postings in 2023 as a percentage of all job postings for that job title. Here, we drew on a standardized set of job titles provided by Lightcast, given the diversity of informal job titles used in the marketplace. This analysis effectively highlights the top types of occupations—or roles—that employers are targeting when they ask candidates to have AI ethics or governance skills when applying for a position.

To answer RQ2, regarding the evolution of demand from 2018 to 2023, we conducted a trend analysis by examining how the frequency of job postings changed over this time period. Specifically, we analyzed the amount of AI, AI ethics, and AI governance job postings by year, which allows us to glean insight into growth patterns, relative frequencies, and changing employer priorities over time.

To answer RQ3, regarding how these skills vary across economic sectors, we analyzed the percentage distribution of AI ethics and governance skills, broken out by industry. We calculated the prevalence of AI ethics and governance job postings for each year from 2018 to 2023 as a percentage of all job postings (not just AI-related) in those industries. The results offer a sectoral perspective on the demand for these competencies.

IV. Results

A. RQ1: Core Competencies of AI Ethics and AI Governance Professionals

As a first step, we determined the *core competencies* required of AI ethics and AI governance professionals using the concept of the location quotient (LQ). Using Lightcast's built-in skill subcategories attached to a set of 39M online job postings from 2023, we calculated the LQ for AI ethics and AI governance competencies as described in (1) and (2). In this context, a LQ greater than 1 indicates that a given skill is more concentrated within a specific domain (i.e., AI ethics or AI governance occupations) relative to its prevalence across the full set of job postings. This suggests that the skill is disproportionately important for that domain and can be considered a core competency. Here, all skills are aggregated into higher-level Lightcast-provided *skill subcategories* to simplify analysis and interpretation.

This article has been accepted for publication in IEEE Transactions on Technology and Society. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTS.2025.3567143

$$\frac{\frac{k_{AIEthics}}{n_{AIEthics}}}{\frac{P_k}{N}} = LQ \quad (1)$$

$$\frac{\frac{k_{AIGovernance}}{n_{AIGovernance}}}{\frac{P_k}{N}} = LQ \quad (2)$$

Where

LQ = location quotient
k = number of job postings for skill subcategory
n = number of job postings for domain
P = number of total postings for subcategory k
N = total number of job postings

Table 2 presents skill subcategories that are most sought after in AI ethics job postings, relative to all job postings. For instance, the Lightcast-provided subcategory 'Regulation and Legal Compliance' encompasses 'Data Ethics,' 'Knowledge of Data Privacy Laws,' and 'Ethical Standards and Conduct.' Importantly, note that the core competencies demanded of these professionals need not be limited to (obviously) ethics- or governance-focused skills; more general skills are certainly also relevant. The diverse range of skill subcategories identified suggests that desired AI ethics competencies cut across multiple disciplines, with prominent skills in domains like technology, strategy, research, management, policy, and regulation.

TABLE 2
TOP SKILL SUBCATEGORIES IN AI ETHICS JOB POSTINGS

Business Intelligence	Business Leadership
Business Management	Business Strategy
Cybersecurity	Education Software and Technology
Financial Modeling	Financial Regulation
General Science and Research	Intelligence Collection and Analysis
Medical Science and Research	Policy Analysis, Research and Development
Regulation and Legal Compliance	Risk Management

In contrast, in AI governance job postings, we see less science, research, and technical skills required and more project management-related skills. Table 3 showcases the most sought-after skill subcategories for AI governance job postings. Here, we see desired business-specific skills like analysis, communication, leadership, operations, and management as well as field-specific skills such as regulatory, legal, and data-related skills.

TABLE 3
TOP SKILL SUBCATEGORIES IN AI GOVERNANCE JOB POSTINGS

Business Analysis	Business Communications
Business Leadership	Business Management
Business Operations	Data Management
Environment and Resource Management	Environmental Regulations
Financial Management	Financial Regulation
Human Resource Management and Planning	People Management
Policy Analysis, Research and Development	Product Development
Program Management	Project Management
Regulation and Legal Compliance	Risk Management

Overall, while the skill compositions that the job market is seeking for professionals in AI ethics- and AI governance-aligned roles differ, both are interdisciplinary and multifunctional in nature.

To elaborate on this understanding, we next looked at the specific *occupations* and *job titles* which seek AI ethics and governance skills most. Table 4 showcases the top five standard occupational classification (SOC) codes where AI ethics job postings are common, subset to 2023, as a portion of all job postings. The SOC codes come from a frequently used taxonomy, linked to the Lightcast dataset, enabling us to view the concentration of AI ethics job postings by (standardized) occupation.

TABLE 4
TOP SOC CODES WITH AI ETHICS JOB POSTINGS

Job Title (SOC code)	Percent of Postings
Computer and Information Research Scientists	4.95%
Computer and Information Systems Managers	2.70%
Data Scientists	2.69%
Database Architects	2.51%
Financial Specialists, All Other	2.27%

As shown, job postings for Computer and Information Research Scientists in 2023 most frequently cited AI ethics-related skills in their requirements, with nearly 5% of all postings for this occupation including a focus on AI ethics. This concentration is followed by postings for Computer and Information Systems Managers, Data Scientists, and Database Architects, each with similar shares in AI ethics postings. Beyond the prevalence of AI ethics skills in computing-related roles, we found that over 2% of job postings for Financial Specialists were related to AI ethics. This is suggestive of the importance of AI ethics professionals in this specific sector, one which is highly regulated.

TABLE 5
TOP SOC CODES WITH AI GOVERNANCE JOB POSTINGS

Job Title (SOC code)	Percent of Postings
Financial Specialists, All Other	4.95%
Financial Risk Specialists	2.70%
Computer and Information Systems Managers	2.69%
Statisticians	2.51%
Data Scientists	2.27%

This article has been accepted for publication in IEEE Transactions on Technology and Society. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTS.2025.3567143

Similarly, for AI governance job postings, Table 5 showcases the top five SOC codes with AI governance job postings. As seen, job postings for "Financial Specialists" had the highest share of AI governance skills requested. Then, a similar share of Financial Risk Specialist, Computer and Information Systems Manager, Statistician, and Data Scientist job postings reflected a desire for AI governance-competent professionals. Currently, these concentrations suggest close alignment between the financial sector and AI governance skills, and between computer professionals and AI ethics skills. Importantly, however, these 'top' five example occupations only reflect a fraction of the tens of thousands of total job postings calling for AI ethics and governance skills, as seen in Fig. 4.

Overall, the findings suggest that ethical and governance-related skills are becoming core requirements for established occupations, such as data scientists, machine learning engineers, and financial specialists. Concurrently, however, we also observe the emergence of new, dedicated roles in these areas. For example, McKinsey & Company describe the growing trend of organizations hiring dedicated AI governance and AI ethics specialists to complement their AI workforce strategy, along with transforming expectations for data science and machine learning roles [90]. Along these lines, a consortium of leading global corporations identifies AI ethics and responsible AI skills as the overall most critical new skill domain for information and communication technology jobs [91], increasingly relevant to 100% of jobs analyzed in their study.

Together, these developments offer an initial map of the diverse literacies that may become increasingly important for jobseekers and institutions of higher education seeking to prepare the future AI workforce.

B. RQ2: Evolution of AI Ethics and AI Governance Skills

The demand for skilled AI, AI ethics, and AI governance jobs has grown from 2018 to 2022, with a noticeable dip in job postings in 2023, potentially related to broader labor market trends (i.e., sweeping layoffs in the technology sector). In Fig. 4, we present a bar graph that depicts the total number of each type of job posting, which provides a relative sense of the prevalence of ethics and governance roles versus AI roles more generally. Note the definition of skills associated with each type of role remains consistent across these years in our analysis.

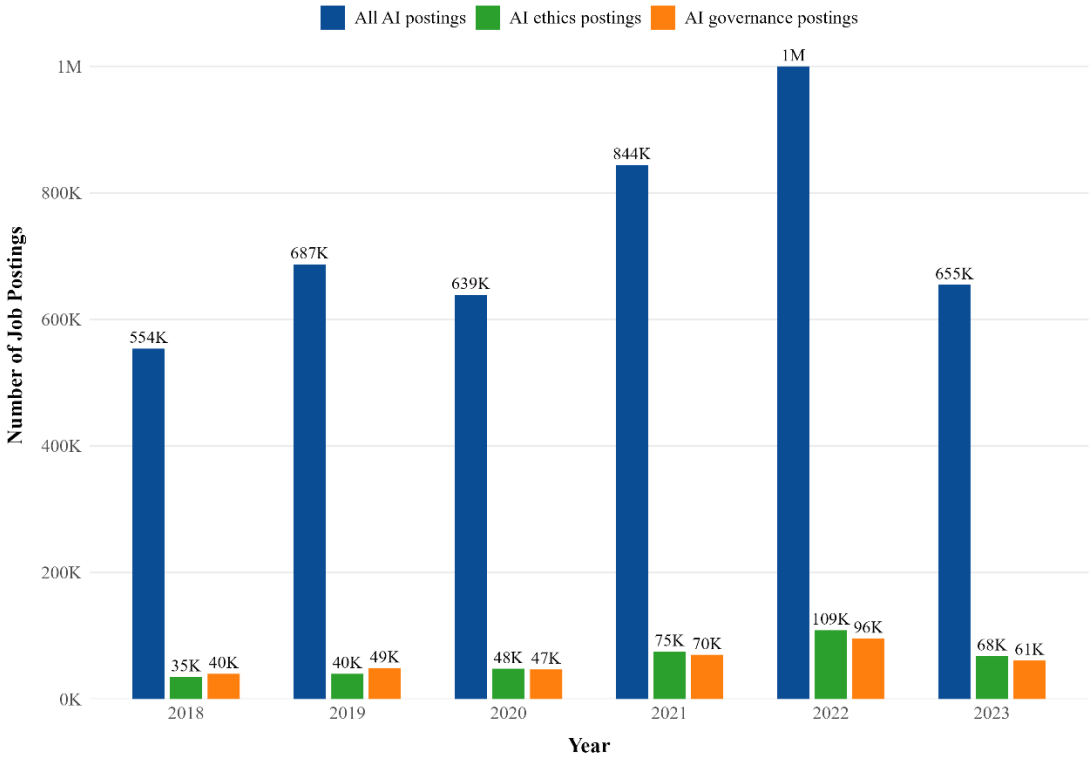

Fig. 4. Frequency of AI, AI ethics, and AI governance job postings over time, from 2018 to 2023

As seen, job postings in AI nearly doubled from 2018 to 2022. Then, in 2023, numbers returned almost to 2018 levels (albeit, with a higher proportion of AI ethics and governance jobs). The 2024 AI Index Report [92] finds a similar dip in AI labor demand in 2023, as the authors found that "every AI-related skill cluster tracked by Lightcast had a decrease in market share, with the exception of generative AI, which grew by more than a factor of 10" [92, p. 224]. Notwithstanding 2023's decline, AI ethics and AI governance job postings do not seem as affected by this dip, and it will be interesting to observe whether this trend continues. Instead, the relative demand for professionals with these skills is broadly increasing, both as an absolute number compared to earlier years, and as a proportion of AI-related jobs. For instance, while only around 6% of AI job postings featured ethics and governance skills in 2018, around 10% did only five years later.

C. RQ3: Skills Across Sectors

Lastly, to address our third RQ, we analyzed the distribution of AI ethics and AI governance job postings across industries, focusing on the top four industries. We used the standard North American Industrial Classification System (NAICS) codes to label industries and to subset our analyses. Presented in Table 6 and Table 7, we see that the "Finance and Insurance" industry has the highest demand for AI ethics and AI governance skills, followed closely by jobs in the "Information" sector.

For both AI ethics and AI governance skills, the percentage share out of all job postings in each industry steadily increased from 2018 to 2022, but again, in 2023 seemed to stagnate. As an example of the significance of the change: the prevalence of these skills has grown from about a quarter of a percentage point to a one-half of a percentage point, roughly doubling over only a five-year period.

TABLE 6
INDUSTRIES WITH HIGHEST SHARE OF AI ETHICS JOB POSTINGS

Industry	2018	2019	2020	2021	2022	2023
Finance and Insurance	0.26%	0.41%	0.44%	0.62%	0.66%	0.66%
Information	0.35%	0.44%	0.52%	0.57%	0.65%	0.63%
Manufacturing	0.24%	0.29%	0.36%	0.35%	0.46%	0.40%
Professional, Scientific and Technical Services	0.24%	0.33%	0.41%	0.49%	0.71%	0.54%

TABLE 7
INDUSTRIES WITH HIGHEST SHARE OF AI GOVERNANCE JOB POSTINGS

Industry	2018	2019	2020	2021	2022	2023
Finance and Insurance	0.53%	0.75%	0.69%	0.80%	1.04%	0.83%
Information	0.29%	0.35%	0.44%	0.49%	0.51%	0.53%
Manufacturing	0.25%	0.26%	0.31%	0.30%	0.35%	0.30%
Professional, Scientific and Technical Services	0.27%	0.30%	0.35%	0.45%	0.53%	0.49%

Future research will be necessary to evaluate the stability of these patterns, especially in light of other economic shocks and cycles. Nevertheless, we see that the same industries are represented, and in the same order: Finance and insurance, Information, Manufacturing, and Professional, Scientific, and Technical Services. These sectors are amongst the leaders in seeking AI ethics and governance skills, though again, tens of thousands of positions are requested in other sectors, ranging from healthcare to education to public administration.

V. DISCUSSION AND CONCLUSION

As demands for AI literacy in the general public and demands for an AI-skilled workforce have grown, an ecosystem has emerged to better understand and prepare for these changes. While there has been ample focus on technical AI literacy, AI ethics and governance skills are increasingly deemed critical to address a wide range of social, regulatory, and organizational needs. Meanwhile, employers and higher education institutions are only beginning to make sense of the skills required for these roles, and stakeholders broadly have only a modest knowledge of the scale of labor demand.

Adapting ideas from the Competency-Capability Chasm, this paper aims to provide insight to employers, educational institutions, and individuals pursuing these professional goals, along with scholars and policymakers. Drawing on more than four million job postings related to AI over the years 2018-2023, we adduce an empirically-grounded characterization of AI ethics and AI governance as professional domains, as articulated by job postings—that is, from the perspective of employers. This definition encompasses both core specialists in these fields, as well as other adjacent professionals for whom these skills are important. For example, occupations as diverse as statisticians to financial risk specialists now require these AI ethics or AI governance skills.

Moreover, we find that professionals with AI ethics or AI governance competencies—whether in technical, regulatory, or business domains—require largely distinct yet overlapping skill sets. Relatedly, while we find some overlap between AI ethics and AI governance job postings, the majority (as depicted in Fig. 2) belong to one domain alone. Next, we find clear evidence of growing demand for both AI ethics and governance roles since 2018, in absolute numbers and as a percentage of all AI job postings. There is modest evidence of a shift in relative demand as well; while AI governance skills were initially more prominent, AI ethics skills have slightly surpassed them in importance as of 2021. Either way, the overall growth trend suggests increasing organizational recognition of these roles' and skills' importance.

Finally, we find the greatest demand for AI ethics or AI governance-skilled professionals within the financial and information sectors, though the demands for these professionals are broader and extend to many sectors beyond the four we highlight in our analysis. As of the early 2020s, more than 100,000 job postings annually in the AI industry in the US require AI ethics or governance skills.

The study provides the first comprehensive analysis of demand for AI ethics and AI governance-aligned roles using comprehensive job posting data, drawing on empirical evidence of how organizations conceive of these professional domains. It also contributes by demonstrating the growth of these fields as distinct professional domains, and by articulating the specific competencies (not limited to ethics or governance skills construed narrowly) requested for individuals in job postings.

However, while our study reviews the market's demands of AI ethics and governance roles, these should be carefully interpreted and not conflated with the complex realities and deeper normative challenges that demand the work of AI ethicists and governance professionals in the first place. As [27] highlights, the AI ethicist may be required to work in industries that have potentially opposing interests, objectives, and epistemological perspectives. For instance, an AI ethicist may have to critically evaluate their organization and recommend actions that conflict with the economic interests of a firm, or, they may have to apply or evaluate practices using approaches not already formalized, understood, or valued within a firm. For instance, we might expect ethicists to have particular skills in normative ethics, systems thinking, sociotechnical thinking, sociology, and much more, including even atypical skills like 'bravery' [26] which are unlikely to show up in job postings.

On the one hand, it is understandable that job postings and skills taxonomies might not list all skills down to a maximal level of detail, which may be difficult and unrealistic, especially given the novelty of some of these job roles. In line with the Competency-Capability Chasm, it may be up to educational institutions to articulate all the micro skills that feed into "AI ethics" or "AI governance," allowing employers to rely on this careful training. On the other hand, this gap could suggest that broader work is needed. This could include better elucidating the deeper sets of skills required, educating employers on the *need* for certain unfamiliar skills and roles, and even advancing societal, normative, or regulatory efforts to promote associated activities or conceptions of those activities. In short, while the job market's current demands constitute an important signal, they do not constitute a complete answer to how AI ethics and governance skills, needs, and job roles will or should evolve.

While more work is needed, our paper helps identify how the market thinks about these professionals, and helps surface gaps between industry, educational institutions, and broader normative goals. As a starting point, we encourage educational programs to not only consider the desired qualifications of professionals from job posting data like those evaluated here, but also consider the ethnographic reports and other taxonomic, qualitative, and conceptual work discussing AI ethics and governance in industry. Educational decision makers also have an important role in being future-oriented: they should think about how skills and roles will evolve—or should evolve—in five or ten years.

These findings offer actionable guidance for educational programs, professional certifications, and other community-serving initiatives aimed at education and career development. The identified skills, and associated overlaps, may be useful for helping guide competency-based educational programs as well as helping employers efficiently structure AI ethics and governance job descriptions and job functions. And critically, for individuals, understanding required competencies and labor market demand can help guide their career pathways, whether individuals are new to the labor market or transitioning.

Our research is limited by several factors. Data sharing agreements limit the full set of details we are able to present here. Further, while our methodology for defining AI, ethics, and governance roles are based on the literature, iterative research design, and the available taxonomies, it necessarily involved subjective decisions along the way and remains a snapshot of an evolving process. Significantly more research is needed, including longitudinal studies of labor market demand and supply, improved definitions of requisite skills and competencies, studies of career trajectories, and evaluations of educational programs' effectiveness in supporting individuals. Next, our research speaks most directly to the US context, where the data are drawn from. Other regions may have different trends with respect to AI, AI ethics, and AI governance: for instance, the European Union might have greater need for these competencies and professionals given its heightened regulatory focus, whereas low-income countries with modest AI industries might have less of a focus. Questions about the definition and composition of these skills and roles, including nuances like focuses on risk management versus ethics, are worthy of future study. Finally, an important limitation of our analysis is that it does not fully differentiate between core AI ethics and governance professionals and professionals in more peripheral roles. Future work should systematically analyze associated differences, including how responsibilities, skill compositions, and organizational positioning vary between core and peripheral roles.

Along these lines, we hope this analysis encourages future investigations into opportunities for growth as well as potential misalignment between the skills articulated by employers in job postings and the complex realities faced by AI ethics and governance professionals. Our findings suggest that these roles are diverse in nature and growing in prevalence and importance. Looking ahead, additional research will be essential to support concerted efforts by educators, employers, regulators, and practitioners to better align training with societal goals and professional practices, ultimately enabling more effective AI ethics and governance work.

REFERENCES

[1] A. Klein, "Bipartisan Bill in Congress Seeks to Help Schools Teach AI Literacy," *Education Week*, Dec. 21, 2023. https://www.edweek.org/technology/bipartisan-bill-in-congress-seeks-to-help-schools-teach-ai-literacy/2023/12

[2] Google, "How we're improving AI literacy in young people," Google. https://blog.google/technology/families/improving-ai-literacy-in-young-people/

[3] N. DeCario and O. Etzioni, "America needs AI literacy now," Allen Institute for AI, Dec. 2021. https://pnw.ai/article/america-needs-ai-literacy-now/72515409

[4] T. Roos, "Elements of AI has introduced one million people to the basics of artificial intelligence | University of Helsinki." https://www.helsinki.fi/en/news/artificial-intelligence/elements-ai-has-introduced-one-million-people-basics-artificial-intelligence

[5] C. Wei, "Nanjing University pioneers AI education system," Mar. 20, 2024. https://www.chinadaily.com.cn/a/202403/20/WS65fa3a4ea31082fc043bda45.html

[6] D. Schiff, "Education for AI, not AI for Education: The Role of Education and Ethics in National AI Policy Strategies," *Int J Artif Intell Educ*, vol. 32, no. 3, pp. 527–563, Sep. 2022, doi: 10.1007/s40593-021-00270-2.

[7] M. Heinrich, *Artificial Intelligence Initiative Act*. 2019. https://www.congress.gov/bill/116th-congress/senate-bill/1558

[8] National Commission on Excellence in Education, "A nation at risk: The imperative for educational reform," *The Elementary School Journal*, vol. 84, no. 2, pp. 113–130, 1983.

[9] N. Wu, "'Restrict foreigners, not robots': Partisan responses to automation threat," *Economics & Politics*, vol. 35, no. 2, pp. 505–528, 2023, doi: 10.1111/ecpo.12225.

[10] M. J. Handel, "Skills Mismatch in the Labor Market," *Annual Review of Sociology*, vol. 29, pp. 135–165, 2003, doi: 10.1146/annurev.soc.29.010202.100030.

[11] J. Handelsman and M. Smith, "STEM for all," *The White House Blog*, 2016.

[12] All Party Parliamentary Group on AI (APPG AI) Secretariat, United Kingdom, "The New Frontier: Artificial Intelligence at Work," All Party Parliamentary Group on AI (APPG AI) Secretariat, United Kingdom, London, UK, Nov. 2021.

[13] D. M. W. and J. R. Allen, "How artificial intelligence is transforming the world," Brookings. https://www.brookings.edu/research/how-artificial-intelligence-is-transforming-the-world/

[14] D. Gehlhaus and S. Mutis, "The U.S. AI Workforce," Center for Security and Emerging Technology, Washington, D.C., Jan. 2021. https://live-cset-georgetown.pantheonsite.io/research/the-u-s-ai-workforce/

[15] S. Bankins, A. C. Ocampo, M. Marrone, S. L. D. Restubog, and S. E. Woo, "A multilevel review of artificial intelligence in organizations: Implications for organizational behavior research and practice," *Journal of Organizational Behavior*, doi: 10.1002/job.2735.

[16] Y. Wu and C. Lin, "The Hidden Multiplier: Unraveling the True Cost of the Global AI Skills Gap," *IEEE Technology and Society Magazine*, vol. 43, no. 4, pp. 15–23, Dec. 2024, doi: 10.1109/MTS.2024.3505618.

[17] J. Southworth *et al.*, "Developing a model for *AI Across the curriculum*: Transforming the higher education landscape via innovation in AI literacy," *Computers and Education: Artificial Intelligence*, vol. 4, p. 100127, Jan. 2023, doi: 10.1016/j.caeai.2023.100127.

[18] X. Fan, "Artificial Intelligence's Impact on Employment: Challenges, Potential Consumers, and Policy Responses Through Automation and Workforce Rehabilitating," *Lecture Notes in Education Psychology and Public Media*, vol. 73, pp. 50–55, Nov. 2024, doi: 10.54254/2753-7048/73/20241040.

[19] P. Cardon, C. Fleischmann, M. Logemann, J. Heidewald, J. Aritz, and S. Swartz, "Competencies Needed by Business Professionals in the AI Age: Character and Communication Lead the Way," *Business and Professional Communication Quarterly*, Nov. 2023, doi: 10.1177/23294906231208166.

[20] T. K. F. Chiu, Z. Ahmad, M. Ismailov, and I. T. Sanusi, "What are artificial intelligence literacy and competency? A comprehensive framework to support them," *Computers and Education Open*, vol. 6, p. 100171, Jun. 2024, doi: 10.1016/j.caeo.2024.100171.

[21] D. Long and B. Magerko, "What is AI Literacy? Competencies and Design Considerations," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, in CHI '20. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–16. doi: 10.1145/3313831.3376727.

[22] A. Bewersdorff, M. Hornberger, C. Nerdel, and D. S. Schiff, "AI advocates and cautious critics: How AI attitudes, AI interest, use of AI, and AI literacy build university students' AI self-efficacy," *Computers and Education: Artificial Intelligence*, vol. 8, p. 100340, Jun. 2025, doi: 10.1016/j.caeai.2024.100340.

[23] M. Hornberger, A. Bewersdorff, and C. Nerdel, "What do university students know about Artificial Intelligence? Development and validation of an AI literacy test," *Computers and Education: Artificial Intelligence*, vol. 5, p. 100165, Jan. 2023, doi: 10.1016/j.caeai.2023.100165.

[24] G. Biesta, "Good education in an age of measurement: on the need to reconnect with the question of purpose in education," *Educ Asse Eval Acc*, vol. 21, no. 1, pp. 33–46, Feb. 2009, doi: 10.1007/s11092-008-9064-9.

[25] S. Sidra and C. Mason, "Reconceptualizing AI Literacy: The Importance of Metacognitive Thinking in an Artificial Intelligence (AI)-Enabled Workforce," in *2024 IEEE Conference on Artificial Intelligence (CAI)*, Jun. 2024, pp. 1181–1186. doi: 10.1109/CAI59869.2024.00211.

[26] O. Gambelin, "Brave: what it means to be an AI Ethicist," *AI Ethics*, vol. 1, no. 1, pp. 87–91, Feb. 2021, doi: 10.1007/s43681-020-00020-5.

[27] D. Burema, "The challenges of being an in-house AI ethicist and how to overcome them," *Journal of Responsible Innovation*, vol. 12, no. 1, p. 2445322, Dec. 2025, doi: 10.1080/23299460.2024.2445322.

[28] G. Gigerenzer, "Moral Satisficing: Rethinking Moral Behavior as Bounded Rationality," *Topics in Cognitive Science*, vol. 2, no. 3, pp. 528–554, 2010, doi: 10.1111/j.1756-8765.2010.01094.x.

[29] M. Pinski, T. Hofmann, and A. Benlian, "AI Literacy for the top management: An upper echelons perspective on corporate AI orientation and implementation ability," *Electronic Markets*, vol. 34, no. 1, Art. no. 1, Apr. 2024, doi: 10.1007/s12525-024-00707-1.

[30] European Commission, *Artificial Intelligence Act: Regulation (EU) 2024/1689*. 2024. http://data.europa.eu/eli/reg/2024/1689/oj/eng

[31] NIST, "AI Risk Management Framework: AI RMF (1.0)," National Institute of Standards and Technology, Gaithersburg, MD, error: NIST AI 100-1, 2023. doi: 10.6028/NIST.AI.100-1.

[32] White House, "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," White House, Office of Science and Technology Policy, Washington, D.C., Oct. 2023.

[33] RAII, "Responsible AI Institute," Responsible AI. https://www.responsible.ai/

[34] EAIGG, "Ethical AI Governance Group," EAIGG. https://www.eaigg.org

[35] S. Kelley, "Employee Perceptions of the Effective Adoption of AI Principles," *J Bus Ethics*, Feb. 2022, doi: 10.1007/s10551-022-05051-y.

[36] D. S. Schiff, S. Kelley, and J. Camacho Ibáñez, "The emergence of artificial intelligence ethics auditing," *Big Data & Society*, vol. 11, no. 4, p. 20539517241299732, 2025, doi: 10.1177/20539517241299732.

[37] L. Tuovinen and A. Rohunen, "Teaching AI Ethics to Engineering Students: Reflections on Syllabus Design and Teaching Methods," in *CEUR Workshop Proceedings*, K.-J. S. Koskinen J. Rantanen M. M. ,. Tuikka A. M. ,. Knaapi-Junnila S., Ed., CEUR-WS, 2021, pp. 19–33.

[38] Evan Castillo, "Purdue Launches Online Master of Artificial Intelligence Degree," *BestColleges*, Mar. 06, 2024. https://www.bestcolleges.com/news/purdue-online-master-artificial-intelligence-degree/

[39] IEEE, "IEEE CertifAIEd™," IEEE Standards Association. https://standards.ieee.org/products-programs/icap/ieee-certifaied/

[40] V. Bogina, A. Hartman, T. Kuflik, and A. Shulner-Tal, "Educating Software and AI Stakeholders About Algorithmic Fairness, Accountability, Transparency and Ethics," *Int J Artif Intell Educ*, Apr. 2021, doi: 10.1007/s40593-021-00248-0.

[41] J. Borenstein and A. Howard, "Emerging challenges in AI and the need for AI ethics education," *AI Ethics*, vol. 1, no. 1, pp. 61–65, Feb. 2021, doi: 10.1007/s43681-020-00002-7.

[42] University of Pittsburgh, "DINS and CS faculty get grant from Google to enhance ethical competence in curriculum." https://www.dins.pitt.edu/news/dins-and-cs-faculty-get-grant-google-enhance-ethical-competence-curriculum

[43] G. Auld, A. Casovan, A. Clarke, and B. Faveri, "Governing AI through ethical standards: learning from the experiences of other private governance initiatives," *Journal of European Public Policy*, vol. 0, no. 0, pp. 1–23, Aug. 2022, doi: 10.1080/13501763.2022.2099449.

[44] CSIRO, "Catalogue of RAI Maturity Models," Software Systems. https://research.csiro.au/ss/science/projects/responsible-ai-pattern-catalogue/rai-maturity-model/

[45] B. Walsh, B. Dalton, S. Forsyth, and T. Yeh, "Literacy and STEM Teachers Adapt AI Ethics Curriculum," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 13, Art. no. 13, Jun. 2023, doi: 10.1609/aaai.v37i13.26906.

[46] L. Floridi, "Who is an AI Ethicist? An Empirical Study of Expertise, Skills, and Profiles to Build a Competency Framework," Aug. 2024.

[47] T. Turja, T. Särkikoski, P. Koistinen, and H. Melin, "Basic human needs and robotization: How to make deployment of robots worthwhile for everyone?," *Technology in Society*, vol. 68, p. 101917, Feb. 2022, doi: 10.1016/j.techsoc.2022.101917.

[48] M. Z. Cocchiaro, J. Morley, C. Novelli, E. Panai, A. Tartaro, and L. Floridi, "Who is an AI Ethicist? An Empirical Study of Expertise, Skills, and Profiles to Build a Competency Framework," *Social Science Research Network, Rochester, NY*: 4891907.

[49] R. Ward *et al.*, "Towards a 21st Century Personalised Learning Skills Taxonomy," in *2021 IEEE Global Engineering Education Conference (EDUCON)*, Apr. 2021, pp. 344–354. doi: 10.1109/EDUCON46332.2021.9453883.

[50] R. Ward, S. Grant, M. W. Larsen, and K. Giovacchini, "The Universal Micro-Credential Framework: The Role of Badges, Micro-Credentials, Skills Profiling, and Design Patterns in Developing Interdisciplinary Learning and Assessment Paths for Computing Education," *IEEE Transactions on Education*, vol. 67, no. 6, pp. 897–906, Dec. 2024, doi: 10.1109/TE.2024.3486016.

[51] P. A. Ertmer *et al.*, "How instructional design experts use knowledge and experience to solve ill-structured problems," *Performance Improvement Quarterly*, vol. 21, no. 1, pp. 17–42, 2008, doi: 10.1002/piq.20013.

[52] W. Dick and L. Carey, *The systematic design of instruction*. Glenview, Ill: Scott, Foresman and Co, 1978.

[53] J. Fulgencio and T. I. Asino, "Conducting a learner analysis," *Design for Learning: Principles, Processes, and Praxis. EdTech Books. https://edtechbooks. org/-TgiD*, 2021.

[54] Eduventures Research Staff, "What We Learned in 2024: The Top 5 Takeaways From Our Analysts."

[55] Risepoint, "Voice of the Online Learner," Risepoint, 2024. https://risepoint.com/wp-content/uploads/2024/09/VOL-2024.pdf

[56] M. J. Ahsan, "Cultivating a culture of learning: the role of leadership in fostering lifelong development," *TLO*, Aug. 2024, doi: 10.1108/TLO-03-2024-0099.

[57] A. Muduli, "Workforce agility: a review of literature.," *IUP Journal of Management Research*, vol. 12, no. 3, 2013.

[58] IRS, "Frequently asked questions about educational assistance programs | Internal Revenue Service." https://www.irs.gov/newsroom/frequently-asked-questions-about-educational-assistance-programs

[59] ASU, "ASU Earned Admission," Home | ASU Earned Admission. https://ea.asu.edu/

[60] USU, "USU Earned Admission," Utah State University. https://www.usu.edu/admissions/apply/earned-admission

[61] Purdue Graduate Council, "Non-credit to Credit Policy and Guidelines for Graduate Degrees at Purdue University," Purdue University, 18–40a, 2018.

[62] Y.-C. Chen and M. Ahn, "Governing AI Systems for Public Values: Design Principles and a Process Framework," in *The Oxford Handbook of AI Governance*, J. B. Bullock, Y.-C. Chen, J. Himmelreich, V. M. Hudson, A. Korinek, M. M. Young, and B. Zhang, Eds., Oxford University Press, 2024, p. 0. doi: 10.1093/oxfordhb/9780197579329.013.31.

[63] M. Sadek, E. Kallina, T. Bohné, C. Mougenot, R. A. Calvo, and S. Cave, "Challenges of responsible AI in practice: scoping review and recommended actions," *AI & Soc*, Feb. 2024, doi: 10.1007/s00146-024-01880-9.

[64] Q. Lu, L. Zhu, X. Xu, and J. Whittle, "Responsible-AI-by-Design: a Pattern Collection for Designing Responsible AI Systems," Sep. 20, 2022, *arXiv*: arXiv:2203.00905. doi: 10.48550/arXiv.2203.00905.

[65] S. Rismani and Aj. Moon, "What does it mean to be a responsible AI practitioner: An ontology of roles and skills," in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, in AIES '23. New York, NY, USA: Association for Computing Machinery, Aug. 2023, pp. 584–595. doi: 10.1145/3600211.3604702.

[66] Lightcast, "Job Posting Analytics (JPA) Methodology | Lightcast Knowledge Base."

https://kb.lightcast.io/en/articles/6957446-job-posting-analytics-jpa-methodology

[67] H. Choi and I. Marinescu, "Data for labor market concentration using Lightcast (formerly Burning Glass Technologies)," *Data in Brief*, vol. 55, p. 110647, Aug. 2024, doi: 10.1016/j.dib.2024.110647.

[68] B. Hershbein and L. B. Kahn, "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," National Bureau of Economic Research, 2017.

[69] D. J. Deming and K. L. Noray, "STEM Careers and the Changing Skill Requirements of Work," National Bureau of Economic Research, 2018. https://www.nber.org/papers/w25065

[70] J. Fuller, S. Hansen, T. Ramdas, and R. Sadun, "The demand for executive skills," Centre for Economic Performance, No.1797, 2021.

[71] S. Hemelt, B. Hershbein, S. Martin, and K. Stange, "College Majors and Skills- Evidence from the Universe of Online Job Ads.pdf," National Bureau of Economic Research, 2021. https://www.nber.org/papers/w29605

[72] W. Markow, Hughes, and A. Bundy, "The new foundational skills of the digital economy: developing the professionals of the future" *Business-Higher Education Forum*, 2018,

[73] L. Jose Gonzalez-Gomez, S. Margarita Hernandez-Munoz, A. Borja, J. Daniel Azofeifa, J. Noguez, and P. Caratozzolo, "Analyzing Natural Language Processing Techniques to Extract Meaningful Information on Skills Acquisition From Textual Content," *IEEE Access*, vol. 12, pp. 139742–139757, 2024, doi: 10.1109/ACCESS.2024.3465409.

[74] D. Acemoglu, D. Autor, J. Hazell, and P. Restrepo, "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, vol. 40, no. S1, pp. S293–S340, Apr. 2022, doi: 10.1086/718327.

[75] P. Bastos, K. Stapleton, D. Taglioni, and H. Y. Wei, "Firm networks and global technology diffusion," World Bank Group, 2024.

[76] A. Bergson-Shilcock, R. Taylor, and N. Hodge, "Closing the Digital Skill Divide- The Payoff for Workers, Business, and the Economy," Federal Reserve Bank of Atlanta, 2023.

[77] C. Larsen *et al.*, Eds., "The Green Transition in Labour Markets Across Europe: an Analysis of Job Postings Data.," in *Pathways of Greening Labour Markets: Opportunities and Challenges for Regional and Local Labour Market Observation in Europe and Beyond*, Nomos Verlagsgesellschaft mbH & Co. KG, 2023. doi: 10.5771/9783957104236.

[78] E. Beckett, "Demand for AI Skills Continues Climbing," Lightcast. https://lightcast.io/resources/blog/demand-for-ai-skills-continues-climbing

[79] D. Gehlhaus and S. Mutis, "The U.S. AI Workforce," Center for Security and Emerging Technology, Washington, D.C., Jan. 2021.

[80] T. Eloundou, S. Manning, P. Mishkin, and D. Rock, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Aug. 21, 2023, doi: 10.48550/arXiv.2303.10130.

[81] US Census Bureau, "Economic Census: NAICS Codes & Understanding Industry Classification Systems," Census.gov.

[82] U.S. Bureau of Labor Statistics, "Standard Occupational Classification (SOC) system," 2018. https://www.bls.gov/soc

[83] Lightcast, "Job postings data." https://lightcast.io/products/data/overview

[84] J. Fereday and E. Muir-Cochrane, "Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development," *International Journal of Qualitative Methods*, vol. 5, no. 1, pp. 80–92, Mar. 2006, doi: 10.1177/160940690600500107.

[85] G. C. Banks, H. M. Woznyj, R. S. Wesslen, and R. L. Ross, "A Review of Best Practice Recommendations for Text Analysis in R (and a User-Friendly App)," *J Bus Psychol*, vol. 33, no. 4, pp. 445–459, Aug. 2018, doi: 10.1007/s10869-017-9528-3.

[86] M. Prates, P. Avelar, and L. C. Lamb, "On Quantifying and Understanding the Role of Ethics in AI Research: A Historical Account of Flagship Conferences and Journals," Sep. 25, 2018, http://arxiv.org/abs/1809.08328

[87] J. Fjeld, N. Achten, H. Hilligoss, A. Nagy, and M. Srikumar, "Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI," *SSRN Journal*, 2020, doi: 10.2139/ssrn.3518482.

[88] D. S. Schiff, "Framing contestation and public influence on policymakers: evidence from US artificial intelligence policy discourse," *Policy and Society*, p. puae007, Apr. 2024, doi: 10.1093/polsoc/puae007.

[89] A. M. Isserman, "The Location Quotient Approach to Estimating Regional Economic Impacts," *Journal of the American Institute of Planners*, vol. 43, no. 1, pp. 33–41, Jan. 1977, doi: 10.1080/01944367708977758.

[90] A. Singla, A. Sukharevsky, L. Yee, M. Chui, and B. Hall, "The State of AI: How Organizations Are Rewiring To Capture Value," McKinsey & Company, Mar. 2025.

[91] Cisco, "The Transformational Opportunity of AI on ICT Jobs," AI-Enabled ICT Workforce Consortium, 2024.

[92] N. Maslej *et al.*, "The AI Index 2024 Annual Report," AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Apr. 2024.

ACKNOWLEDGMENT

The authors are grateful to Lightcast for their willingness to share data with the research team. Thank you to the Governance and Responsible AI Lab (GRAIL) at Purdue University and the Center for Security and Emerging Technology (CSET) at Georgetown University for support and feedback, especially Veronica Jade Kinoshita.