ISSN: 2473-4901
v12 n6697

Shifting Skill Demands in the AI Job Market: A Comparative Text Mining Analysis of Job Postings from 2024 and 2025

Alan Peslak
arp14@psu.edu
Penn State University
Dunmore, PA 18512 USA

Wendy Ceccucci
ceccucci@qu.edu
Quinnipiac University
Hamden, CT, 06518, USA

D. Scott Hunsinger
hunsingerds@appstate.edu
Appalachian State University
Boone, NC 28607, USA

Abstract

This study examines how skill requirements in artificial intelligence (AI) and machine learning (ML) job postings have evolved between 2024 and 2025. Drawing on a dataset of 936 job descriptions collected from PostJobFree.com (436 from 2024 and 500 from 2025), we apply a structured text-mining approach to identify changes in employer demand. A regex-based pipeline was used to detect 246 skill terms organized into ten categories, with frequencies normalized to account for differences in sample size.

The analysis reveals broad and statistically significant shifts in employer expectations. Eight of the ten skill categories increased significantly over time, with Responsible AI emerging as the fastest-growing area. At the individual skill level, 16 of 21 selected technologies showed meaningful changes in prevalence. Demand for governance-related competencies—such as privacy, compliance, and regulatory knowledge rose sharply, while demand for certain legacy tools declined. The emergence of terms such as “agentic” suggests a transition toward more advanced AI system architectures.

Importantly, this study captures patterns in job posting language rather than actual hiring outcomes. As such, the findings reflect how employers signal demand, offering insight into how the AI labor market is being framed and communicated.

Keywords: artificial intelligence, machine learning, job market analysis, text mining, skill demand, responsible AI, workforce trends, Mann–Whitney U

 ISSN: 2473-4901
v12 n6697

Shifting Skill Demands in the AI Job Market: A Comparative Text Mining Analysis of Job Postings from 2024 and 2025

Alan Peslak, Wendy Ceccucci and D. Scott Hunsinger

1. INTRODUCTION

The labor market for artificial intelligence (AI) has changed considerably in recent years. This transformation has been driven not only by rapid advances in large language models but also by the increasing importance of the integration of AI technologies across a wide range of industries. What was once a specialized technical field has become embedded in numerous business functions, making it increasingly important to understand which skills employers value and how those job expectations continue to evolve.

Understanding these changes is important for multiple stakeholders. Students and job seekers must determine which technical and professional skills to develop, educators must continually adapt curricula to meet workforce needs, and policymakers must evaluate whether workforce development initiatives are keeping pace with technological advancement. Because hiring needs often evolve more quickly than formal labor statistics or educational programs, job postings provide a timely signal of workforce expectations by revealing the skills, technologies, and qualifications employers are actively seeking. As such, they offer a near real-time view of emerging labor market demand and changing employer priorities.

Prior research has examined technology labor markets using surveys, job board analytics, and publication data. More recently, text mining of online job postings has emerged as a practical and scalable approach for measuring employer skill demand (Ao et al., 2023; Papoutsoglou et al., 2021; Senger et al., 2024). By applying natural language processing and text mining techniques to large collections of job advertisements, researchers can systematically identify references to programming languages, software frameworks, AI tools, and professional competencies. These methods make it possible to quantify changes in employer demand over time and identify emerging trends in workforce expectations.

Rather than estimating the overall size or composition of the AI labor market, this study focuses on how employer demand for AI-related skills changed between 2024 and 2025. Job postings collected from PostJobFree.com (2025), an online employment aggregation platform, serve as a timely indicator of employer expectations during this period. The analysis compares the frequency of both broad AI skill categories and specific technologies across the two years to identify shifts in the competencies employers emphasize. In addition to descriptive comparisons, the study employs statistical hypothesis testing to determine whether observed differences in skill demand represent meaningful changes rather than random variation.

This study addresses two central questions:

- **RQ1:** Which categories of AI-related skills show statistically significant changes between 2024 and 2025?
- **RQ2:** Which specific skills increased or decreased significantly in employer demand?

2. LITERATURE REVIEW

2.1 Job Posting Analysis as Labor Market Data

The use of online job postings as a source of labor market insight has become increasingly common in recent years. (Rahhal, Kassou, Ghogho, 2024). Rahhal's survey paper reviewed recent research published between 2015 and 2022 that used data science analytics techniques on labor market data. Deming and Kahn (2018) analyzed 45 million job advertisements from Burning Glass Technologies. Their study grouped skill requirements into broad categories and showed how demand varied across firms and labor markets. Their keyword-matching approach, which identified skills through predefined term

 ISSN: 2473-4901
v12 n6697

lists, has become a widely adopted framework for subsequent research. Hershbein and Kahn (2018) used the same data source to demonstrate that the Great Recession accelerated structural increases in employer skill requirements, establishing that shifts in job posting language reflect genuine changes in labor demand rather than short-term fluctuations.

More recent research has built on these foundations by refining both data processing and analytical techniques. Tzimas et al. (2024), for example, proposed a framework for transforming raw job postings into usable labor market intelligence. Their framework covered data science techniques such as data cleaning, normalization, and deduplication. Their framework was applied to two cases. The first examined the Greek tourism labor market during the COVID-19 pandemic. Their analysis revealed notable shifts in job demand, required skills, and employment structures. In the second case they applied a data-driven ontology to extract skills from job postings using machine learning techniques. The findings highlight that the proposed methodology, utilizing NLP and machine-learning techniques instead of LLMs, can be applied to different labor market-analysis use cases and offer valuable insights for businesses and job seekers.

Purohit (2022) outlined a practical framework based on n-gram tokenization and keyword matching, which is a similar approach used in this study. Nguyen et al. (2024) explored the use of large language models for skill extraction. While these models can better capture context and handle complex phrasing, their findings suggest that traditional supervised approaches still outperform LLMs on benchmark tasks. Given these trade-offs, the present study adopts a middle-ground strategy: a regex-based method that emphasizes consistency, interpretability, and suitability for longitudinal comparison, even if it sacrifices some semantic nuance.

2.2 AI Workforce Trends

The AI job market is expanding rapidly across industries such as technology, healthcare, finance, and manufacturing. Data from UMD-LinkUp (2025) and the University of Maryland show a 68% increase in AI job postings in the United States, rising from 29,509 in 2022 to 49,577 by 2024.

Looking ahead, the World Economic Forum projects a major transformation in the global labor market by 2030, with 170 million new jobs created and 92 million displaced—a net gain of 78 million jobs. These shifts highlight the growing importance of workforce training and education reform. Alekseeva et al. (2021) provided one of the first aggregate analyses of AI skill demand in the United States, showing that the demand for AI-skilled labor steadily rose between 2010 and 2019 and that firms demanding AI skills pay a wage premium of approximately 5–10% even for non-AI positions within the same firm. Their work established online job postings as a valid proxy for AI technology adoption at the firm level.

The Federal Reserve Bank of Atlanta has produced ongoing analyses extending this line of research. Goldfarb et al. (2024) documented that AI skill demand accelerated sharply in 2024, with the share of online postings requiring AI skills reaching 1.62%, and that the composition of demand has shifted from ML-related skills toward broader AI-related skills, including generative AI. They also found that AI skill demand is spreading to a broader set of occupations, industries, and local labor markets.

At the international level, Green (2024) produced an OECD report finding that the skills most demanded in AI-exposed occupations are management and business skills rather than specialized AI competencies, and that the share of vacancies demanding emotional, cognitive, or digital skills has increased over time. The IMF (Brollo et al., 2025) analyzed Lightcast data across multiple countries and found that AI-user skills are now in higher demand than AI-developer skills, with approximately half of US AI-related postings mentioning only user-level skills. These findings provide important context for the present study's observation that soft skills and leadership represent the largest category of mentions in AI job postings.

AI is reshaping the workplace in complex ways. While it automates routine tasks, it also creates new roles requiring specialized expertise. Generative AI is expanding what workers can do, enabling roles like accounting clerks, nurses, and teaching assistants to handle more complex tasks, while helping professionals such as engineers and doctors work more efficiently.

Researchers such as Shen (2024) argue that the most impactful technologies are those that enhance

ISSN: 2473-4901
v12 n6697

human capabilities rather than replace them. Still, AI is contributing to a more polarized labor market, with rising skill requirements, widening income gaps, and higher compensation for AI specialists compared to lower-skilled workers. As automation replaces repetitive roles, workers without advanced skills face increasing risks of displacement.

These trends are evident across industries. In healthcare, AI supports diagnostics, drug discovery, and personalized medicine (Kolakowski, 2025). In finance, it is widely used for fraud detection, risk management, and investment analysis (Shen, 2024). In manufacturing, AI-driven automation improves efficiency and product quality. While AI is expected to create new job categories, demand for blue-collar roles may decline as knowledge-based work expands. Comunale and Manera (2024) note that occupations with high exposure to automation are especially vulnerable to displacement.

Asim and Ding (2025), using data from 500 professionals, found that AI adoption improves both engagement and job performance, though outcomes depend heavily on job complexity and individual expertise. Similarly, Han et al. (2025) showed that AI can enhance innovation by enabling employees to reshape their roles, particularly when supported by strong workplace environments.

2.3 Skills and Qualifications in AI Job Postings

Generative AI technologies, such as ChatGPT, are transforming industries and creating new types of professional roles that require both technical expertise and strategic thinking. Research by Jepperson (2025) finds growing demand for expertise in data analytics, STEM fields, and information technology. At the same time, organizations adopting AI are shifting toward more specialized and highly educated workforces. Hazan et al. (2024) found rising demand for roles in areas such as robotics, advanced analytics, and system integration.

Technical skills remain foundational. Programming languages like Python and R are widely expected, along with familiarity with machine learning frameworks such as TensorFlow and PyTorch. Data engineering has also become increasingly important, as professionals must manage, analyze, and interpret large datasets using tools like Power BI and Tableau.

However, technical ability alone is no longer sufficient. Employers are placing greater emphasis on soft skills, including critical thinking, communication, collaboration, and problem-solving (World Economic Forum, 2025). These skills are crucial for translating technical work into business value.

Educational pathways are also becoming more specialized. Many roles now require advanced degrees in computer science or AI-related fields, and new job titles, such as Generative AI Engineer or Computer Vision Engineer—reflect the increasing complexity of the field.

2.4 Responsible AI and Governance Skills

With the expansion of AI applications, there has been a growing emphasis on the ethical, legal, and governance aspects of AI systems. Oschinski et al. (2024) found that while technical skills remain important, a large share of in-demand capabilities in AI-related roles are nontechnical, including social and cognitive skills. This shift reflects the increasing complexity of deploying AI in real-world settings, where considerations such as fairness, accountability, and transparency play a critical role.

Institutional developments reinforce this trend. Frameworks such as the European Commission’s Skills Agenda and the ESCO classification system have begun to incorporate competencies related to AI ethics and governance. Regulatory initiatives, including the EU AI Act, further signal that organizations need to navigate a more structured and scrutinized environment when deploying AI technologies. The rapid growth in Responsible AI skills observed in this study appears less surprising and more indicative of a broader shift in how AI work is defined.

2.5 Methodological Approaches to Skill Comparison

Cross-temporal comparison of skill demand presents specific methodological challenges. Deming and Noray (2020) demonstrated that STEM skill demands change rapidly, with the value of specific technical skills declining as they become more widely held. Modestino et al. (2020) showed that employer skill

ISSN: 2473-4901
v12 n6697

requirements are sensitive to labor market conditions, with firms demanding more skills when workers are plentiful. These findings underscore the importance of normalization and statistical testing when comparing skill frequencies across time periods, as changes in posting volume, platform composition, and labor market tightness can all confound naive trend analysis.

The present study contributes to this literature by applying a regex-based text mining pipeline to a controlled cross-year comparison of AI job postings, supplemented by inferential statistical testing (Mann–Whitney U tests, two-proportion z-tests, and Poisson rate comparisons) that moves beyond the descriptive approaches employed in most prior job posting studies.

3.METHODOLOGY

3.1 Data Collection

The dataset for this study consists of AI-related job postings collected from PostJobFree.com. PostJobFree was selected because it offers many AI-related postings with complete descriptions and a consistent structure suitable for text mining. We limited the analysis to 2024 and 2025 to capture the most current shift in AI skill demand during the period of rapid generative AI adoption. Using two consecutive years also allows for a controlled comparison of changes in employer language over time. Additionally, PostJobFree was used because it provides full job descriptions and consistent access across both years, allowing direct comparison. It also does not have restrictions against scraping in their terms of service that other job boards such as Indeed have.

The PostJobFree 2024 sample includes 436 postings, while the 2025 sample includes 500. Each posting contains a full-text job description, which serves as the basis for analysis. While PostJobFree aggregates listings from multiple sources, it does not provide a fully controlled sample of the labor market. As a result, the findings should be interpreted as reflective of platform-level trends rather than a complete representation of the broader market.

During the cleaning process of the analysis, postings with missing descriptions were excluded and duplicate postings were identified within each year. The duplicates generally reflected multi-location listings. The primary analysis retains all postings to preserve the full demand signal, but a sensitivity analysis on deduplicated data (N = 853) is reported in Section 5.6 to confirm robustness. Table 1 summarizes the dataset characteristics.

3.2 Skill Taxonomy Development and Validation

A single taxonomy-driven extraction pipeline was implemented for analysis to ensure consistency across measures. The taxonomy was developed through a corpus-driven, iterative process based on recurring phrases found across the postings. The 246 extracted skill terms were consolidated into 10 categories using normalization of synonyms, abbreviation matching, and alignment with prior AI skill frameworks.

Stage 1: Corpus-driven identification. Every job posting in the dataset was reviewed to find recurring expressions that referred to AI methods, programming languages, machine-learning frameworks, analytics techniques, data-engineering technologies, cloud infrastructure, and professional competencies. Programming language refers to general coding and software development skills rather than specialized subdomains Those terms appearing multiple times and/or representing widely recognized AI skills were kept.

Stage 2: Consolidation and normalization. Synonymous abbreviations and full expressions such as ML/machine learning and LLM/large language models were mapped to canonical forms. Regex patterns were created to capture variants while ensuring consistent counting. Each skill was assigned to a single dominant category.

Stage 3: External validation and alignment. To ensure construct validity, the taxonomy was aligned with the ESCO classification system and prior job-posting research, serving as post hoc validation rather than as the generative source of the taxonomy.

ISSN: 2473-4901
v12 n6697

Subsequent extraction, frequency, and statistical analyses use this taxonomy as the foundational input.

Characteristic	**2024**	**2025**
Source	PostJobFree.com	PostJobFree.com
Postings (after cleaning)	436	500
File format	Excel (.xlsx)	CSV (.csv)

Table 1. Dataset Characteristics

Note. The 2024 workbook contains five supplementary sheets with pre-computed term frequencies and TF-IDF scores.

3.3 Analytical Pipeline

Python was utilized to conduct the analysis. The pipeline proceeds through the following six stages: data loading, skill taxonomy construction (10 categories with 246 terms), regex pattern compilation with word-boundary guards for short tokens, per-document skill counting, normalization to per-100-posting rates, and visualization/export. Skills mentioned multiple times in a single posting were counted (i.e., total mention frequency, not binary presence). Document frequency, the percentage of postings mentioning a skill at least once, was calculated separately.

Regex-based Extraction

The pattern-matching approach is unable to distinguish between required and optional skills, nor can it resolve semantic equivalence across synonymous terms. For example, “LLM” and “large language model” are counted separately. No differentiation was made concerning contextual usage, such as “experience with Python” vs. “familiarity with Python.” Those skills expressed through synonyms or phrasing not articulated by the taxonomy were missed. These limitations are inherent to lexical approaches and should be considered when interpreting results.

3.4 Statistical Analysis

To assess statistical significance, the following three complementary inferential tests, which move beyond descriptive comparisons, were applied.

Category-level: Mann–Whitney *U* tests

Per-document mention counts, such as the total number of skill mentions in a given category for each job posting, were compared across years using Mann–Whitney *U* tests, nonparametric tests appropriate for count data with non-normal distributions. Effect sizes were computed as $r = Z / \sqrt{N}$, where $N = 936$. Bonferroni correction was applied across the 10 tests ($\alpha = .005$).

Skill-level: Two-proportion *z*-tests

For 21 individual skills of substantive interest, two-proportion *z*-tests compared document frequency (the proportion of postings mentioning the skill) between 2024 and 2025. Benjamini–Hochberg (BH) false discovery rate correction was applied across all 21 tests.

Category-level: Poisson rate comparisons

To check for robustness on total mention counts, Poisson rate comparison tests (normal approximation) were conducted for each category, comparing the per-posting mention rate across years.

4. RESULTS

Results are presented in three layers: (1) Category-level changes in skill intensity, (2) Skill-level changes in prevalence, and (3) Descriptive trends in emerging and declining terms.

ISSN: 2473-4901
v12 n6697

4.1 Category-Level Findings (RQ1)

Table 2 displays normalized mention rates and inferential test results for the ten skill categories. Mann–Whitney *U* tests with Bonferroni correction found statistically significant increases in eight of the ten categories. Programming showed no significant change ($p = 1.00$ after correction), and Computer Vision did not reach significance after correction ($p = .039$, uncorrected). Poisson rate tests confirmed these results directionally, with eight categories reaching significance.

Category	Per100 ’24	Per100 ’25	Diff	% Chg	U	p	p(B)	r
Responsible AI	75.7	205.2	+129.5	+171.1	78,774	<.001	<.001	.268
Soft Skills and Leadership	285.1	385.8	+100.7	+35.3	88,904	<.001	<.001	.161
Cloud / MLOps	151.6	214.4	+62.8	+41.4	93,186	<.001	<.001	.140
AI Frameworks and Tools	57.1	92.6	+35.5	+62.1	90,672	<.001	<.001	.179
NLP / LLM	137.6	168.0	+30.4	+22.1	91,469	<.001	<.001	.158
ML Core	187.6	216.2	+28.6	+15.2	96,231	.001	.014	.105
Data Sci & Analytics	152.3	175.4	+23.1	+15.2	95,900	<.001	.007	.110
Data Engineering	31.9	40.4	+8.5	+26.7	99,168	<.001	.005	.114
Computer Vision	24.3	29.8	+5.5	+22.6	101,858	.004	.039	.094
Programming	152.8	147.0	−5.8	−3.8	102,979	.126	1.00	.050

Table 2. Category Mention Rates with Mann–Whitney U Test Results

Note. Per100 = mentions per 100 postings. p = uncorrected Mann–Whitney U p-value. p(B) = Bonferroni-corrected p-value (k = 10). r = effect size (Z / √N). Sorted by absolute difference.

Findings reflect changes in job posting language rather than direct measures of hiring outcomes.

Responsible AI exhibited the most dramatic growth (171.1% change), with the largest effect size among all categories ($r = .268$, $p < .001$). Its share of total mentions doubled from 6.0% to 12.3%. Soft Skills showed the largest absolute increase (+100.7 per 100 postings, $r = .161$, $p < .001$) and remained the dominant category in both years. Cloud / MLOps grew significantly ($r = .140$, $p < .001$), consistent with increasing emphasis on deployment infrastructure. AI Frameworks and Tools exhibited the second-largest effect size ($r = .179$, $p < .001$), driven by PyTorch and TensorFlow. Programming was the only category showing no statistically significant change ($p = 1.00$ after Bonferroni correction).

4.2 Individual Skill Findings (RQ2)

Table 3 provides two-proportion *z*-test results for 21 skills of substantive interest, testing whether the proportion of postings mentioning each skill significantly changed between 2024 and 2025. After Benjamini–Hochberg correction, 16 of 21 skills showed statistically significant changes in document frequency.

 ISSN: 2473-4901
v12 n6697

Skill	% ’24	% ’25	z	p	p(BH)	Sig
privacy	5.3	23.4	7.76	<.001	<.001	***
collaboration	16.5	37.0	7.01	<.001	<.001	***
GCP	4.1	16.2	5.99	<.001	<.001	***
Python	29.6	48.2	5.81	<.001	<.001	***
compliance	12.6	29.8	6.35	<.001	<.001	***
agentic	1.4	8.0	4.68	<.001	<.001	***
TensorFlow	15.4	27.2	4.38	<.001	<.001	***
PyTorch	17.0	29.0	4.34	<.001	<.001	***
R	26.4	15.0	−4.32	<.001	<.001	***
data science	20.9	33.2	4.22	<.001	<.001	***
regulatory	5.5	11.6	3.29	.001	.002	**
cross-functional	19.0	28.2	3.28	.001	.002	**
deep learning	17.0	24.4	2.79	.005	.009	**
Azure	14.4	21.4	2.75	.006	.009	**
AWS	16.5	23.4	2.62	.009	.012	*
strategy	20.4	27.6	2.56	.010	.014	*
leadership	28.2	35.4	2.35	.019	.023	*
generative AI	13.8	18.6	2.00	.046	.054	ns
SQL	11.0	14.6	1.63	.102	.113	ns
communication	55.0	52.0	−0.93	.351	.369	ns
machine learning	56.2	54.8	−0.43	.669	.669	ns

Table 3. Document Frequency with Two-Proportion z-Test Results (Selected Skills)

Note. % = percentage of postings mentioning skill at least once. p = uncorrected two-proportion z-test. p(BH) = Benjamini–Hochberg adjusted p-value (k = 21). *** p < .001, ** p < .01, * p < .05, ns = not significant after BH correction. Sorted by absolute z-value.

The largest effects were observed for Responsible AI terms: Privacy increased from 5.3% to 23.4% of postings (z = 7.76, p < .001), compliance from 12.6% to 29.8% (z = 6.35, p < .001), and regulatory from 5.5% to 11.6% (z = 3.29, p = .002). The term “agentic”—referring to autonomous, goal-directed AI agent architectures—showed a statistically significant increase from 1.4% to 8.0% of postings (z = 4.68, p < .001).

Python’s document frequency rose significantly from 29.6% to 48.2% (z = 5.81, p < .001), indicating that it is becoming a baseline expectation across a wider range of AI roles. R showed a significant decline from 26.4% to 15.0% (z = −4.32, p < .001). Machine learning and communication, despite being the two most frequently mentioned skills, showed no significant change in document frequency, suggesting they were already saturated in 2024.

Notably, the generic term “generative AI” did not reach significance after BH correction (p(BH) = .054), even as more specific LLM-related skills such as “agentic” and “fine-tuning” increased significant, which is consistent with a maturation of vocabulary from broad buzzwords to precise technical terms.

4.3 Emerging and Declining Skills

Beyond document frequency, normalized mention rate data (per 100 postings) identified the largest movers. Among gainers, compliance (+32.0), privacy (+29.3), and collaboration (+26.8) showed the largest absolute increases. Among decliners, R (−12.4), generative AI (−7.2), and cloud computing (−5.5) showed the largest absolute decreases.

ISSN: 2473-4901
v12 n6697

5. DISCUSSION

5.1 The Rise of Responsible AI (RQ1)

The most striking finding is the substantively large growth in Responsible AI mentions (r = .268, p < .001), which elevated the category from eighth to fourth in overall share. This shift is consistent with multiple converging forces: The proliferation of AI regulation worldwide (including the EU AI Act), increasing corporate governance requirements, and growing public scrutiny of AI systems. The specific terms driving this growth, compliance (z = 6.35), privacy (z = 7.76), and regulatory (z = 3.29), all reached statistical significance, suggesting that employers are seeking candidates who can navigate the legal and ethical dimensions of AI deployment, not merely build models.

5.2 Soft Skills as Differentiators

Soft Skills and Leadership registered the largest absolute gain and showed a significant overall increase (r = .161, p < .001). The significant growth in collaboration (z = 7.01, p < .001) and cross-functional skills (z = 3.28, p = .002) suggests that as AI matures, the ability to work across teams and communicate with non-technical stakeholders is increasingly valued. Importantly, this study measures demand signals in job posting language rather than actual hiring outcomes and thus reflects employer communication patterns rather than realized workforce composition.

5.3 Technical Skill Shifts

AI Frameworks and Tools showed the second-largest effect size (r = .179, p < .001) within the technical categories. PyTorch (z = 4.34) and TensorFlow (z = 4.38) reached significance, suggesting continued demand for deep learning frameworks despite the rise of higher-level APIs.

The term "agentic" showed a statistically significant increase (z = 4.68, p < .001), which may reflect the emergence of agent-based AI as a major architectural paradigm.

Programming was the only category showing no significant change (p = 1.00 after correction). Within the category, Python grew significantly (z = 5.81, p < .001), while R declined significantly (z = −4.32, p < .001), resulting in offsetting effects at the category level. This pattern is consistent with Python consolidating its position as the default language for AI work.

5.4 Cloud and Infrastructure Maturation

Cloud/MLOps increased significantly (r = .140, p < .001), with growth concentrated in specific platforms: GCP showed the largest individual increase (z = 5.99, p < .001), while AWS and Azure also reached significance. The sharp descriptive decline in cloud computing is consistent with a pattern in which employers now expect cloud proficiency as a baseline rather than a distinguishing qualification.

5.5 Limitations

Several limitations are acknowledged. First, the regex-based approach operates on surface-level text and cannot capture semantic nuance; a posting mentioning Python in a "nice-to-have" context is counted identically to one listing it as a core requirement. Regex cannot resolve semantic equivalence across synonymous terms, and contextual usage is not differentiated. Second, the taxonomy was manually curated and may underrepresent emerging tools. Third, PostJobFree.com aggregates listings from multiple sources and does not represent a controlled or stratified sample of the AI labor market; findings may reflect platform-specific biases. Fourth, while the statistical tests confirm that observed differences are unlikely due to chance, they cannot establish causal mechanisms underlying the shifts. Fifth, the two samples were collected independently, and differences in scraping methodology or temporal coverage could introduce mistakes. Finally, results should be interpreted as descriptive trends in employer-facing job posting language rather than as measures of actual hiring outcomes or workforce composition.

5.6 Sensitivity to Duplication

Effect sizes reported in this section are recalculated using the deduplicated dataset and therefore differ slightly from those reported in Table 2.

The datasets contained within-year exact duplicates (15 in 2024 (3.4%); 68 in 2025 (13.6%)), likely reflecting multi-location postings by the same employers. No exact cross-year duplicates were identified.

 ISSN: 2473-4901
v12 n6697

A sensitivity analysis removing all exact duplicates reduced the corpus from 936 to 853 postings (421 from 2024 and 432 from 2025).

All ten category-level directional findings were preserved, and eight of ten Mann–Whitney U significance levels were unchanged or equivalent. The two categories that shifted, ML Core and Computer Vision, both moved from marginally significant (Bonferroni-corrected $p = .014$ and $.039$, respectively) to non-significant, while the eight strongly significant categories, including the study's central findings regarding Responsible AI ($r = .268$, $p < .001$), Soft Skills and Leadership ($r = .163$, $p < .001$), and Cloud/MLOps ($r = .124$, $p = .003$), were fully preserved.

At the skill level, 16 of 21 two-proportion z-tests remained significant after deduplication, matching the count observed in the full dataset. The five that shifted (strategy, deep learning, AWS, Azure, and data science) all moved from marginally significant to non-significant, while core findings for privacy, compliance, collaboration, Python, R, agentic, and GCP remained highly significant (all $p < .001$).

These results indicate that within-year duplication does not materially affect the study's principal conclusions.

5.7 Recommendations for Future Research

Several extensions could strengthen this line of analysis. Semantic extraction methods, such as named entity recognition fine-tuned on job posting text, could supplement the regex pipeline. Longitudinal expansion to additional years would enable trend detection beyond a single comparison. Cross-platform validation using data from LinkedIn, Indeed, or Glassdoor would assess generalizability. Linking skill data with metadata such as job title, seniority level, and industry sector would enable richer segmentation. Document frequency, which provides a more robust measure of skill breadth than raw mention counts, should be further developed as a primary metric in future studies.

6. CONCLUSION

This study provides an empirical analysis of how technical and professional skill demands in the AI job market shifted between 2024 and 2025. Using 936 job descriptions, a 246-term skill taxonomy, and inferential statistical testing (Mann–Whitney U, two-proportion z-tests, and Poisson rate comparisons), the analysis identified several statistically validated trends: The rapid ascent of Responsible AI governance skills ($r = .268$), the growing importance of collaboration and cross-functional skills, the rise of agent-based AI terminology, and the significant decline of the R programming language. Programming was the only category to show no significant overall change, while 8 of 10 categories increased significantly after Bonferroni correction.

These findings indicate that the AI job market is shifting from broad technical literacy toward a more nuanced demand profile emphasizing governance, collaboration, deployment infrastructure, and specialized LLM competencies. For job seekers, the data highlight the growing value of compliance, privacy, and regulatory expertise alongside traditional ML skills. For educators, the results underscore the need to integrate responsible AI and cross-functional collaboration into AI curricula. The regex-based pipeline, while limited in semantic depth, provides a transparent and reproducible method for tracking skill demand at scale, and the addition of inferential testing strengthens confidence that the observed patterns reflect genuine shifts rather than sampling artifacts.

7. REFERENCES

Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. Labour Economics, 71, 102002. https://doi.org/10.1016/j.labeco.2021.102002

Ao, Z., Horváth, G., Sheng, C., Song, Y., & Sun, Y. (2023). Skill requirements in job advertisements: A comparison of skill-categorization methods based on wage regressions. *Information Processing & Management*, 60(2), 103185. https://doi.org/10.1016/j.ipm.2022.103185

Arendt, L., Gałecka-Burdziak, E., Núñez, F., Pater, R., & Usabiaga, C. (2023). Skills requirements across task-content groups in Poland: What online job offers tell us. Technological Forecasting and Social Change, 187, 122199.

ISSN: 2473-4901
v12 n6697

Asim, M. & Ding, W. (2025). AI Usage, Employee Engagement, and Work Performance: Examining the Roles of Job Complexity and AI Knowledge *Proceedings of the 2025 4th International Conference on Big Data and Digital Management*, 191, 294 – 303. pphttps://doi.org/10.2991/978-94-6463-710-6_34

Brollo, F., Cazzaniga, M., Cerutti, E., Ciminelli, G., Hansen, N.-J., Jaumotte, F., Khalil, M., Li, A., & Liu, Y. (2025). Bridging skill gaps for the future: New jobs creation in the AI age (IMF Staff Discussion Note SDN/2025/001). International Monetary Fund.

Comunale, M & Manera, A. (2024). The Economic Impacts and the Regulation of AI: A Review of the Academic Literature and Policy Actions. (n.p.): International Monetary Fund.

Deming, D. J., & Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. Journal of Labor Economics, 36(S1), S337–S369. https://doi.org/10.1086/694106

Deming, D. J., & Noray, K. (2020). Earnings dynamics, changing job skills, and STEM careers. The Quarterly Journal of Economics, 135(4), 1965–2005. https://doi.org/10.1093/qje/qjaa021

Goldfarb, A., Taska, B., & Teodoridis, F. (2024, October 15). Recent trends in the demand for AI skills. Federal Reserve Bank of Atlanta Macroblog. https://www.atlantafed.org/blogs/macroblog/2024/10/15/recent-trends-in-demand-for-ai-skills

Green, A. (2024). Artificial intelligence and the changing demand for skills in the labour market (OECD Artificial Intelligence Papers No. 14). OECD Publishing. https://doi.org/10.1787/88684e36-en

Han, X., Chen, F., Wang, H., & Xu, S. (2025). Unlocking innovation: Artificial intelligence usage and innovative behavior in the workplace. Social Behavior and Personality: An international Journal 53(3). https://doi.org/10.224.sbp.13851

Hazan, E., Madgavkar, A., Chui, M., Smit, S., Maor, D., Dandona, G., & Huyghues-Despointes, R. (2024). A new future of work: The race to deploy AI and raise skills in Europe and beyond. McKinsey Global Institute. Retrieved March 27, 2025 from https://www.mckinsey.com/mgi/our-research/a-new-future-of-work-the-race-to-deploy-ai-and-raise-skills-in-europe-and-beyond.

Hershbein, B., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. American Economic Review, 108(7), 1737–1772. https://doi.org/10.1257/aer.20161570

Jepperson, T. (2025). How Is Artificial Intelligence Impacting the Job Market? LINQTO. Retrieved March 30, 2026 from https://www.linqto.com/blog/how-is-artificial-intelligence-impacting-the-job-market/

Kolakowski, N. (2025). AI Job Market 2025: Trends and Opportunities Across the US. Dice. Retrieved March 30,2026 from https://www.dice.com/career-advice/artificial-intelligence-a-i-job-market-small-but-growing-in-many-states

Modestino, A. S., Shoag, D., & Ballance, J. (2020). Upskilling: Do employers demand greater skill when workers are plentiful? The Review of Economics and Statistics, 102(4), 793–805. https://doi.org/10.1162/rest_a_00835

Nguyen, K. C., Montariol, S., & Musat, C. (2024). Rethinking skill extraction in the job market domain using large language models. In Proceedings of the NLP for HR Workshop at EACL 2024 (pp. 18–28). Association for Computational Linguistics.

Oschinski, M., Crawford, A., & Wu, M. (2024). AI and the future of workforce training (CSET Report). Center for Security and Emerging Technology, Georgetown University. https://doi.org/10.51593/20240033

ISSN: 2473-4901
v12 n6697

Papoutsoglou, M., Mittas, N., Angelis, L., & Vakali, A. (2021). A survey on skill identification from online job ads*. IEEE Access*, 9, 118134–118153. https://doi.org/10.1109/ACCESS.2021.3106120

PostJobfoFree (2026). Retrieved Dec. 2025 from https://www.postjobfree.com/

Purohit, P. (2022). Technical framework for text mining of job advertisements. In D. Dennehy et al. (Eds.), Responsible AI and Analytics for an Ethical and Inclusive Digitized 13454, pp. 234–245. Springer.

Rahhal, I., Kassou, I., Ghogho. (2024). Data science for job market analysis: A survey on applications and techniques. *Expert Systems with Applications*, 251, https://doi.org/10.1016/j.future.2018.03.035.

Senger, E., Zhang, M., van der Goot, R., & Plank, B. (2024).Deep Learning-based Computational Job Market Analysis: A Survey on Skill Extraction and Classification from Job Postings. *Proceedings of the First Workshop on Natural Language Processing for Human Resources* (NLP4HR 2024), 1–15, St. Julian’s, Malta. Association for Computational Linguistics.

Shen, Y. (2024). Artificial Intelligence and the Economy . The Impact of Artificial Intelligence on the Job Market, Advances in Economics Management and Political Sciences, 92(1) 71-74. https://doi.org/10.54254/2754-1169/92/20231275

Tzimas, G.; Zotos, N.; Mourelatos, E.; Giotopoulos, K.C.; Zervas, P. (2024) From Data to Insight: Transforming Online Job Postings into Labor-Market Intelligence. *Information.* 15, 496. https://doi.org/ 10.3390/info15080496

World Economic Forum. (2025). Future of Jobs Report 2025. Retrieved March 30, 2026 from https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf