

DOI: https://doi.org/10.48009/2_iis_2026_106

UpSkill-Scanner: An open-source LLM pipeline for curriculum alignment with labor market demand

Adam Patterson, *Amazon*¹, adapattk@amazon.com

Abstract

Academic departments across all disciplines face a persistent challenge of aligning curricula with evolving labor market demands. Yet systematic, data-driven approaches to measuring this alignment remain inaccessible to most programs, constrained by proprietary platforms, high licensing costs, and discipline-specific methodologies. This paper introduces UpSkill-Scanner, an open-source, discipline-agnostic Python pipeline that leverages web scraping and large language model (LLM) based skill extraction to analyze entry-level job postings for any field. The pipeline scrapes listings from major job boards, extracts structured skill data using Anthropic's Claude Haiku via the Batch API, normalizes synonyms through LLM-based semantic grouping, and produces easy to understand visualizations, all at a cost of approximately \$1.00 per full analysis. To demonstrate the tool's capabilities, it is applied to entry-level economics positions across six job titles and 51 metropolitan areas spanning five U.S. regions. Results reveal that Excel and SQL dominate technical skill demand, communication and analytical thinking lead soft skill requirements, and regional variation, while modest, carries practical implications for region-based curricular decisions. Critically, the pipeline requires only changing input parameters (job titles and locations) to replicate this analysis for any discipline: computer science, nursing, business, or any other field. UpSkill-Scanner is freely available, enabling any academic department or individual to conduct data-driven curriculum assessment without institutional data subscriptions.

Keywords: curriculum alignment, skills gap, higher education, large language models, labor market analysis

Introduction

Every academic department faces a fundamental question: are educators teaching what graduates actually need? The accelerating pace of technological change, the growing demand for data-driven decision-making, and the shifting composition of entry-level roles have made this question more urgent and more difficult to answer than ever before. Whether in economics, computer science, nursing, business, or the liberal arts, the gap between what curricula emphasize and what employers seek has become a central concern for higher education administrators, accrediting bodies, and faculty alike.

Efforts to measure this curriculum–workforce gap have relied primarily on proprietary labor market analytics platforms such as Lightcast (formerly Burning Glass), employer surveys, and alumni outcome data. While valuable, these approaches share common limitations that proprietary platforms require

¹ This work does not relate to the author's position

expensive institutional subscriptions (Onodipe et al., 2025), surveys suffer from response bias and temporal lag, and alumni data capture past labor markets rather than current demand. The result is that systematic, timely, discipline-specific skill analysis remains out of reach for some academic departments, particularly smaller programs without dedicated assessment budgets. These tools are also not easily available to individual students seeking advice on what skills to learn or courses to take.

This paper addresses that gap by introducing UpSkill-Scanner, an open-source Python package that automates the collection, extraction, normalization, and visualization of skill data from online job postings. The tool is discipline-agnostic by design, as any department can conduct a full analysis by simply changing the job titles and geographic parameters. The pipeline leverages large language model (LLM) based extraction, specifically, Anthropic's Claude Haiku via the Batch API, to identify and categorize skills from job posting text without requiring pre-defined taxonomies, at a total cost of approximately \$1.00 per analysis.

To demonstrate UpSkill-Scanner's capabilities, this study applies it to entry-level positions relevant to economics undergraduate students across five U.S. regions and 51 metropolitan areas. Economics serves as an instructive case study because it is a general major (Hemelt et al., 2023) whose graduates enter diverse occupations such as financial analysis, data analysis, research, and policy, making the skill alignment question particularly complex. However, the methods, pipeline, and analytical framework presented here are immediately transferable to any discipline.

This study makes three contributions. First, it introduces a novel, freely available methodological tool that democratizes labor market skill analysis for academic departments of any size and discipline. Second, it demonstrates the pipeline through a comprehensive analysis of entry-level economics skill demand, providing an updated, regionally differentiated inventory that extends prior work (Hemelt et al., 2023; Onodipe et al., 2025). Third, it shows that high-quality skill extraction can be conducted using LLMs at minimal cost, offering a practical alternative to both proprietary platforms and GPU-intensive NLP approaches (Bothmer & Schlippe, 2023).

Literature Review

The Curriculum–Workforce Gap Across Disciplines

The misalignment between academic curricula and employer expectations is a pervasive concern across higher education, not limited to any single discipline. In economics, Hansen (2001) identified six proficiencies that graduates should achieve and found that students consistently rated themselves low on higher-order competencies. Colander and McGoldrick (2010) argued that economics programs systematically fail to consider skills required by employers in business, undergraduate education, and government. Siegfried and Walstad (2014) documented that the most conspicuous curriculum change over 30 years was the rise of econometrics as a required course, yet questions persist about whether programs develop the full range of skills graduates need (Allgood et al., 2015). The common thread is that faculty research interests and disciplinary traditions, rather than labor market signals, often drive pedagogical decisions, a structural problem that transcends any single field.

The challenge is compounded by the rising importance of digital and data literacy across all disciplines. Patterson et al. (2023) demonstrated that integrating real-world data analysis into undergraduate instruction can bridge the gap between curricular content and applied skill development. More broadly, the rapid adoption of generative AI tools by college students underscores that the skill landscape is evolving faster than traditional curriculum revision cycles can accommodate (Frydenberg et al., 2026; Patterson et al.,

2024; Rruppli et al., 2024). Departments need nimble, low-cost methods for monitoring what employers demand.

Career Outcomes and the Case for Systematic Skill Analysis

The practical stakes of curriculum–workforce alignment are substantial. In the demonstration discipline of economics, the Bureau of Labor Statistics (BLS) data shows that a master’s degree is the typical entry-level education for economists, with only approximately 17,600 such positions nationally (BLS, 2025). Given that approximately 40,000 economics bachelor’s degrees are awarded annually (Siegfried, 2024), the vast majority of graduates enter adjacent occupations such as financial analyst, market research analyst, management analyst, and data analyst positions (AEA, 2025). This pattern, where graduates primarily enter roles beyond the discipline’s traditional boundaries, is common across many fields and makes systematic skill analysis particularly valuable.

Job Posting Analysis as a Research Method

Using online job postings as revealed-preference measures of employer demand has become an established methodology in labor economics. Deming and Kahn (2018) analyzed approximately 45 million Burning Glass Technologies postings from 2010–2015, finding significant variation in skill requirements even within narrowly defined Standard Occupational Classification (SOC) codes. Hemelt et al. (2023) mapped employer skill demands to 70 college majors using the near-universe of U.S. online job ads, finding that economics is a general rather than specialized major requiring a broad mix of skills.

Most directly relevant to the demonstration, Onodipe et al. (2025) used Lightcast’s Analyst platform to aggregate job advertisements for economics-related titles from 2010 through 2022. Their key finding, that communication, management, research, and software skills dominate employer requirements, provides an important benchmark. However, their reliance on a proprietary platform limits reproducibility and accessibility. This study complements their work by providing an open-source alternative that any department or individual can utilize.

AI-Based Skill Extraction

The application of natural language processing (NLP) and LLMs to skill extraction from job postings has evolved rapidly. Bothmer and Schlippe (2023) developed SkillScanner, which used Sentence-BERT embeddings and K-means clustering to extract and compare skills. While effective, their approach required GPU-intensive computation and manual cluster interpretation. The emergence of LLMs offers an alternative to provide structured extraction via prompt engineering, where the model identifies and categorizes skills in a single pass without requiring pre-defined entity taxonomies. Recent work has applied BERT-based analysis to Federal Reserve communications (Patterson, 2024) and extended this using LLMs (Patterson, 2025), demonstrating the viability of structured text extraction across domains. These studies utilizing structured NLP and LLM-based text analysis directly informed the design of UpSkill-Scanner’s extraction pipeline.

Methodology

UpSkill-Scanner is an open-source Python pipeline designed to be used by any academic department, assessment office, individual researcher, or student seeking to understand what skills employers demand. The pipeline is discipline-agnostic as the user specifies job titles and geographic parameters as inputs, and the system produces structured skill inventories and visualizations as outputs. This section describes the pipeline’s architecture, components, and design decisions. The subsequent section demonstrates its application to entry-level economics positions.

Architecture and Design Philosophy

The pipeline consists of seven stages: (1) web scraping, (2) relevance filtering, (3) LLM-based skill extraction, (4) post-extraction reclassification, (5) LLM-based skill normalization, (6) aggregation and statistical analysis, and (7) visualization. Each stage is independently rerunnable, and expensive operations (scraping, extraction) cache intermediate results as CSV files for reproducibility. The modular architecture means that researchers can substitute components, for example, using a different LLM provider or scraping from different job boards, without rewriting the entire pipeline.

Data Collection Module

Job postings are collected from Indeed and Glassdoor using the `python-jobs` library, which programmatically queries publicly accessible job listing pages. The author followed the framework proposed by Brown et al. (2025) for conducting web scraping research in a legally and ethically responsible manner. Specifically (a) UpSkill-Scanner accesses only publicly available job posting data without creating accounts or accepting terms of service on either platform; (b) no personally identifiable information about job seekers or individual applicants was collected; (c) data were used solely for non-commercial academic research; (d) rate limiting (2-second delays between queries) was implemented to avoid overburdening servers; and (e) the analysis reports only aggregate statistical findings as no individual job postings are reproduced. This approach is consistent with Ninth Circuit precedent in *hiQ Labs, Inc. v. LinkedIn Corp.* (2022), which held that accessing publicly available web data does not violate the Computer Fraud and Abuse Act (CFAA) when no authentication barriers are circumvented.

The scraping module accepts any set of job titles and predefined regions as input. For the economics demonstration, six titles spanning three career lanes were used; a computer science department might instead search for Junior Software Engineer, Data Engineer, and DevOps Analyst, while a nursing program could search for Registered Nurse and Nurse Practitioner. The geographic scope is similarly configurable.

Data Cleaning

Raw postings are deduplicated on the tuple of job title, company name, and location, retaining the first occurrence. Postings with missing description fields are dropped. A relevance filter addresses the fuzzy matching behavior of job board search engines as it applies a curated regular expression block list, removing clearly unrelated occupations. Approximately 10% of scraped postings are removed by this filtering process.

LLM-Based Skill Extraction

Each job posting's description is submitted to Anthropic's Claude Haiku 4.5 model via the Batch API with temperature set to 0 for deterministic extraction. The model is prompted to extract skills into four categories: technical skills (programming languages, software tools, statistical methods), soft skills (interpersonal competencies), certifications (formal professional credentials requiring examination), and domain knowledge (subject matter expertise and academic degrees). The prompt includes explicit normalization rules and category definitions with examples to minimize misclassification.

To illustrate the classification logic, consider the distinction between Analytics as a soft skill and Data Analysis as a technical skill. Data Analysis refers to a specific, executable technical activity, such as applying quantitative methods to data to produce outputs, and is thus classified as a technical skill. Analytics, by contrast, most commonly appears in job postings as a descriptor of professional disposition or cognitive orientation (e.g., analytics-driven mindset, strong analytical approach) and is accordingly classified as a soft skill. This distinction is consistent with how the two terms function in employer language: one describes what a candidate can do with a tool or method, the other describes how a candidate

thinks and approaches problems. Borderline cases of this type were adjudicated by the post-extraction reclassification rules described below.

Despite explicit instructions, the LLM occasionally misclassified cloud platforms (e.g., AWS, Azure) as certifications because certification programs exist for these platforms. A rule-based post-extraction reclassification step corrects these errors by moving known platform names to the technical skills category.

Skill Normalization

Raw extraction produces synonym variants (e.g., Python, python, Python Programming). All unique skill names are collected and submitted to Claude in batches of 200 for semantic grouping, producing a canonical mapping (e.g., R Programming and R Language both map to R). This LLM-based approach is context-aware as it can distinguish R the programming language from R as an abbreviation in another context, providing an advantage over embedding-based cosine similarity methods used in prior work (Bothmer & Schlippe, 2023).

Validation

Three complementary validation approaches are employed. First, an automated grounding check verifies whether each extracted skill appears in the source job description text via exact match, substring match, abbreviation expansion, and suffix-stripped matching. Overall grounding rates are approximately 84%, with technical skills grounding at 93.8%, certifications at 94.7%, domain knowledge at 93.0%, and soft skills at 64.6%. The lower soft skill grounding rate reflects legitimate semantic inference rather than extraction error, as employers typically express soft skills through verbs and contextual descriptions rather than as explicit noun phrases.

For example, a posting stating "Collaborate with global economists to deliver forecasts" yields the extraction Teamwork; a posting describing "our collaborative culture" similarly yields Teamwork and Collaboration; and a posting instructing candidates to "Communicate clearly to both technical and non-technical audiences" yields Communication, where the noun form is absent but the competency is unambiguous. In each case, the model correctly identifies the required competency from its verbal or adjectival expression in employer language.

Second, statistical sanity checks examine the skills-per-posting distribution, category balance, and Zipf's law fit. Skill frequencies follow the expected power-law distribution, confirming well-structured extraction.

Third, a manual spot-check was conducted on a random sample of 20 job postings, in which extracted skills were evaluated against source posting text by a human reviewer. Across all four skill categories, extracted skills were judged to accurately reflect the competencies described or implied in the posting text, providing independent confirmation that the pipeline produces results consistent with human interpretation, including for soft skills where the grounding rate is lower by design.

Cost and Accessibility

A significant advantage of UpSkill-Scanner is its minimal cost. Using Claude Haiku via the Batch API (which provides a 50% discount over standard API pricing), a full analysis of approximately 1,250 postings costs approximately \$1.00 for extraction and normalization combined. The tool includes a three-layer cost protection system (i) prepaid API credits with auto-reload disabled, (ii) an environment-variable cost limit, and (iii) a real-time software-level cost tracker that aborts processing if the budget is exceeded. This cost structure makes UpSkill-Scanner accessible to individual researchers, adjunct faculty, and departments without dedicated assessment budgets, a stark contrast to proprietary platforms that may cost thousands of dollars annually.

Demonstration: Entry-Level Economics Positions

To illustrate UpSkill-Scanner’s capabilities, this study applied the pipeline to entry-level positions relevant to economics undergraduate students. This demonstration serves as a worked example that any department can replicate for their own discipline by substituting the job titles and, optionally, the geographic parameters.

Geographic Scope

UpSkill-Scanner defines five U.S. regions: Northeast, Southeast, Midwest, South Central, and West, each containing 10–11 major metropolitan areas (51 total). This geographic structure enables regional comparison of skill demand patterns across labor markets with distinct economic profiles. Table 1 presents the metropolitan areas included in each region.

Table 1. Geographic Regions and Metropolitan Areas

Region	Metropolitan Areas
Northeast	New York, Boston, Worcester, Philadelphia, Pittsburgh, Hartford, Providence, Newark, Baltimore, Albany, Buffalo
Southeast	Atlanta, Miami, Charlotte, Nashville, Jacksonville, Tampa, Orlando, Raleigh, Richmond, Charleston
Midwest	Chicago, Detroit, Minneapolis, St. Louis, Indianapolis, Columbus, Milwaukee, Kansas City, Cincinnati, Cleveland
South Central	Houston, Dallas, San Antonio, Austin, New Orleans, Memphis, Oklahoma City, Louisville, Birmingham, Little Rock
West	Los Angeles, San Francisco, Seattle, Denver, Phoenix, Portland, Las Vegas, San Diego, Salt Lake City, Sacramento

Job Title Selection

This demonstration selected six job titles based on convergent evidence from the BLS Occupational Outlook Handbook, Onodipe et al. (2025), American Economic Association career resources, and NACE first-destination survey data: Entry Level Economist, Junior Economist, Economic Analyst, Research Analyst, Data Analyst, and Policy Analyst. These titles span three primary career lanes for economics graduates: traditional economics roles, the data-driven track, and the policy track. For each city-title combination, UpSkill-Scanner requested up to 25 postings published within the preceding week. A department in another field would simply replace these titles with their own discipline-relevant searches.

Descriptive Statistics

The full pipeline collected 3,130 unique job postings across 51 metropolitan areas after deduplication and relevance filtering. These postings represent a one-week cross-sectional snapshot collected from Indeed and Glassdoor and should be interpreted accordingly. While the sample is geographically broad and title-diverse, findings reflect employer demand as expressed in online job postings during a single collection window on two major platforms. Some postings may include roles requiring graduate credentials despite entry-level title selection. Readers should treat the results as an illustrative, data-driven baseline rather than a definitive census of employer skill demand.

The extraction rate was 99.7%, meaning nearly all postings with description text yielded at least one skill. After normalization, 7,242 unique canonical skills were identified from 71,827 total skill mentions, with

an average of approximately 23 skills extracted per posting. Table 2 presents the distribution of postings by region.

Table 2. Postings by Region

Region	Postings (n)
Northeast	590
Southeast	639
Midwest	711
South Central	498
West	692
Total	3,130

National Skill Rankings

Table 3 presents the top 30 skills nationally, broken down by region, ranked by the percentage of postings in which each skill appeared at the national level. Soft skills dominate the upper ranks as communication, problem-solving, and collaboration are highly desirable by employers. Salient technical skills include Excel, data analysis, SQL, financial modeling, and data visualization. Domain knowledge areas such as financial modeling, accounting, and finance-related expertise also appeared prominently. Figure 1 visualizes the national skill distribution as a treemap, providing a summary of relative skill demand that complements the ranked frequencies in Table 3.

Table 3. Top 30 Skills by Frequency (% of Postings)

Skill	Category	National	NE	SE	MW	SC	West
Communication	Soft Skills	97.2	96.6	97.5	97.9	96.2	97.5
Problem Solving	Soft Skills	92.0	92.7	91.9	91.6	93.8	90.6
Teamwork	Soft Skills	59.0	62.5	60.9	59.4	57.0	55.5
Leadership	Soft Skills	45.4	45.1	47.3	41.9	44.8	47.8
Collaboration	Soft Skills	33.3	34.6	32.9	35.9	32.5	30.6
Excel	Technical	32.5	33.1	33.3	34.5	32.9	28.8
Data Analysis	Technical	31.4	31.7	33.0	27.3	31.5	33.7
Stakeholder Mgmt	Soft Skills	29.9	30.3	28.2	27.3	32.9	31.8
Attention to Detail	Soft Skills	28.2	28.5	27.4	29.5	26.3	28.8
Analytics	Soft Skills	25.7	27.5	26.0	27.7	23.7	23.3
Project Mgmt	Soft Skills	20.3	17.1	18.9	20.8	23.1	21.8
SQL	Technical	20.1	19.7	19.1	18.4	20.5	22.8
Financial Analysis	Domain	18.9	19.3	15.8	20.0	18.1	21.0
Financial Modeling	Technical	18.8	18.3	16.7	19.5	17.3	21.2
Accounting	Domain	16.6	17.3	15.5	18.0	14.5	17.1

Issues in Information Systems

Volume 27, Issue 2, pp. 79-94, 2026

Finance	Domain	16.6	16.8	15.5	18.7	13.3	17.9
Organizational Mgmt	Soft Skills	15.4	15.6	15.5	17.2	16.7	12.4
Presentation Skills	Soft Skills	14.7	14.7	16.1	11.1	15.7	16.5
Budgeting	Domain	14.1	15.4	14.4	14.1	11.2	14.9
PowerBI	Technical	14.1	12.2	15.0	15.9	16.7	11.0
Data Visualization	Technical	14.0	13.4	13.1	14.6	16.7	12.9
Financial Reporting	Domain	13.6	11.9	13.5	14.8	12.4	14.9
Mentoring	Soft Skills	12.6	10.8	14.1	13.8	12.2	11.6
Time Management	Soft Skills	12.2	12.5	11.7	13.5	13.7	10.0
Cross-Func. Collab.	Soft Skills	11.6	10.2	9.5	11.7	9.8	16.0
Python	Technical	11.1	10.5	9.9	10.4	12.2	12.4
Strategic Thinking	Soft Skills	11.0	11.9	10.6	8.6	10.4	13.3
Forecasting	Domain	10.9	11.0	11.0	11.5	7.6	12.4
PowerPoint	Technical	10.8	10.5	13.1	11.7	11.6	7.4
Interpersonal Skills	Soft Skills	10.7	9.8	11.9	11.8	11.8	8.4

Regional Variation

While the top-ranked skills were broadly consistent across regions, notable differences emerged. The West showed higher demand for programming languages, such as SQL and Python, relative to other regions, consistent with the concentration of technology firms in metros such as Seattle and San Francisco. The Northeast showed relatively stronger demand for teamwork, analytics, and budgeting, reflecting the collaborative environments of financial services and consulting firms concentrated in New York and Boston. The South Central region exhibited comparatively higher demand for project management, stakeholder management, PowerBI, and data visualization, potentially reflecting the prevalence of energy sector and government analytical roles in Houston, Dallas, and Austin. The Midwest led in Excel and Analytics demand, while the Southeast showed strong demand for presentation skills. These regional patterns, while modest in absolute magnitude, carry practical implications for departments seeking to tailor curricula to their local labor markets.




Figure 1. National Skill Distribution Treemap

Discussion

The economics demonstration illustrates both the power of UpSkill-Scanner as a research tool and the substantive insights it can generate for any discipline. Both dimensions are discussed below.

Methodological Implications: A Tool for Any Department

UpSkill-Scanner's primary contribution is methodological. By open-sourcing a complete pipeline, from data collection through publication-ready visualization, UpSkill-Scanner lowers the barrier to systematic curriculum–labor market alignment analysis from thousands of dollars to dollars (conditional on sample size). The pipeline's discipline-agnostic design means that the analysis presented here for economics can be replicated for any field. A physics department could search for Junior Physicist, Lab Technician, and Research Scientist; a nursing program could search for Registered Nurse and Nurse Practitioner; a computer science department could search for Software Engineer, Data Engineer, and ML Engineer. The only required changes are the input parameters, though thoughtful job-title selection is essential. Titles that are too broad, too senior, or poorly matched to a field's actual entry-level landscape will produce misleading skill profiles, and results should be validated against field-specific benchmarks before informing curricular decisions.

The use of LLM-based extraction rather than traditional NLP offers several practical advantages. Unlike embedding-based methods (Bothmer & Schlippe, 2023), the author's approach requires no GPU infrastructure, no pre-defined skill taxonomies, and no manual cluster interpretation. The LLM handles multi-word phrases, abbreviations, and synonym resolution natively. Furthermore, the context-aware

normalization step, where the LLM semantically groups synonyms, is more robust than cosine similarity methods that may conflate unrelated terms sharing surface-level similarity. These advantages make the tool practical for individual researchers and smaller departments without computational resources.

This approach builds on broader research at the intersection of emerging technologies and education for relevant skill development. For example, prior work has explored innovative pedagogical approaches, such as including technology-enhanced speaking instruction through Virtual Reality usage (Patterson et al., 2025). UpSkill-Scanner represents a natural extension of this work: rather than studying technology adoption in isolation, it provides a tool that directly connects curricular decisions to empirical labor market evidence. The high demand for communication skills suggests emerging methods, such as Virtual Reality headsets for public speaking courses, could enhance students' skills, allowing them to perform better in interviews while equipping them with the skills to succeed in their careers. Good communication will have direct effects on other soft skills, such as collaboration and teamwork, further underlying the importance of soft skills in the modern labor market.

Substantive Insights from the Economics Demonstration

The economics results align with and extend the existing literature. The dominance of Excel, SQL, and Python in employer requirements confirms what Onodipe et al. (2025) found using Lightcast data, and reinforces the argument advanced by Colander and McGoldrick (2010) that economics programs have been slow to incorporate the technical skills employers demand. The finding that Excel appears in approximately one-third of all postings is particularly notable given that prior research has documented spreadsheet tasks as being the lowest self-reported data literacy skill of incoming freshmen (Mentzer et al., 2024). Additionally, Grealis et al. (2024) report persistent gender gaps in Excel literacy among college students. UpSkill-Scanner now quantifies the employer side of that equation, allowing students from all backgrounds to understand the importance of learning software tools, such as Excel, for optimal career placement.

The prominence of SQL, a skill rarely taught in economics departments, suggests that database literacy is now a foundational competency for economics graduates. This finding supports Hansen's (2001) emphasis on the interpreting and manipulating economic data proficiency, which now increasingly requires database skills rather than just spreadsheet skills. The absence of traditional software tools, such as Stata and SAS, from the top 30, alongside Python's presence at 11.1%, suggests that the programming landscape for economists has changed considerably.

The consistent demand for soft skills such as communication, problem-solving, and collaboration across all regions reinforces the finding by Hemelt et al. (2023) that economics is a general major requiring a broad skill mix. These competencies are not easily developed through lecture-based instruction alone, suggesting that programs should incorporate more collaborative projects, presentations, and applied research experiences. The regional differences, while modest, carry practical implications for region-based curriculum design, a novel capability enabled by UpSkill-Scanner's configurable geographic parameters.

Implications for Research and Practice

For Academic Departments

UpSkill-Scanner enables any department to conduct data-driven curriculum assessment at minimal cost. It is recommended that departments run the pipeline at least annually, comparing results year-over-year to track evolving skill demand. The regional analysis capability is particularly valuable for programs whose

graduates primarily enter local labor markets, as it allows place-based curricular tailoring. For accreditation purposes, the pipeline's output as structured skill inventories with frequency data and visualizations provides the evidence-based assessment that accrediting bodies increasingly demand.

Addressing Resistance to Curriculum Change

Translating UpSkill-Scanner's output into curricular action will inevitably encounter resistance from both faculty and students. Faculty resistance commonly stems from disrupted routines and perceived threats to academic autonomy, while student resistance is often driven by role ambiguity such as when new requirements feel disconnected from professional identity (Al Kafio et al., 2025; Folkers, 2005). Both forms of resistance are best addressed through a participatory process: departments should involve faculty in running the pipeline and interpreting results collectively, rather than presenting completed recommendations for ratification. Kotter and Schlesinger (1979) identify participation and involvement as among the most effective strategies for overcoming resistance in educational organizations, and this faculty-led approach is far more durable than administrator-mandated change, with UpSkill-Scanner's empirical output serving as transparent, data-driven evidence that grounds the conversation in employer demand rather than administrative preference.

For the Economics Demonstration Specifically

Based on the economics results, the author offers discipline-specific recommendations as an example of how departments might translate UpSkill-Scanner output into curricular action. First, departments should consider requiring at least one course that teaches SQL and database management. Second, programming instruction should not be limited to a single language; exposure to Python alongside existing Stata requirements would better prepare graduates. Third, data visualization should be explicitly taught as a communication skill, potentially integrating tools like Tableau or PowerBI. Fourth, applied capstone projects that require students to collect, clean, analyze, and present real-world data would simultaneously develop technical, analytical, and communication competencies.

For the Research Community

UpSkill-Scanner opens several avenues for future research. Longitudinal studies using repeated snapshots could track how skill demand evolves over time, potentially revealing leading indicators of curricular obsolescence. Comparative studies across disciplines could identify which fields face the largest or fastest-growing curriculum-workforce gaps. The pipeline's output could also be linked to student outcome data, such as graduation rates, employment rates, and starting salaries, to assess whether curriculum-labor market alignment predicts graduate success. Finally, the tool's low cost and open-source design make it suitable for large-scale, multi-institutional studies that would be prohibitively expensive with proprietary platforms.

Limitations

Several limitations should be acknowledged. First, the data represent a temporal snapshot of approximately one week's postings from Indeed and Glassdoor. Skill demand may vary seasonally, cyclically, or in response to macroeconomic conditions, and findings should not be interpreted as capturing long-term structural demand. Replication using repeated snapshots across multiple time periods would substantially strengthen claims about persistent skill trends. Second, Indeed and Glassdoor may not fully represent all labor market segments, as government positions, academic roles, and small-firm openings may be underrepresented on these platforms. Third, the python-jobspy library depends on the current structure of job board websites; changes to site architecture could disrupt data collection. Fourth, the LLM's inference of soft skills from contextual cues (e.g., "collaborative environment" interpreted as

“Teamwork”) introduces a degree of subjectivity; while this captures important implied requirements, it yields lower grounding rates (65%) for the soft skills category. Future research could employ an LLM-as-judge framework to scale validation, enabling systematic assessment of extraction accuracy across the full posting sample. Fifth, results are sensitive to the choice of job titles, as different title selections for the same discipline may yield different skill profiles. The author mitigated this by grounding the title selection in BLS data, AEA resources, and prior research (Onodipe et al., 2025).

These limitations apply equally to any discipline’s analysis using the pipeline. However, despite targeting entry-level titles, some postings required graduate credentials (e.g., Master's degree or PhD for Economist roles), reflecting the fuzzy matching behavior of job board search engines. Users should interpret skill frequency estimates as reflecting the broader posting landscape for a given title rather than exclusively undergraduate-level demand, and are encouraged to apply seniority modifiers (e.g., "Junior," "Entry-Level") where possible to improve targeting precision.

While UpSkill-Scanner is discipline-agnostic by design, successful deployment in a new field requires careful, domain-informed job title selection and validation. Title choices directly determine which postings are collected and, therefore, which skills surface as prominent. For example, a nursing program searching Nurse will capture a fundamentally different posting landscape than one searching Registered Nurse, New Graduate. Furthermore, skill extraction quality should be spot-checked for each new discipline, as LLM classification behavior may reflect patterns from the demonstration domain. Researchers and practitioners are encouraged to ground title selection in field-specific resources such as BLS occupation data, professional association career guides, and prior literature before treating pipeline outputs as authoritative.

Conclusion

This study introduces UpSkill-Scanner, an open-source, discipline-agnostic pipeline for analyzing the skills employers demand from entry-level graduates. By combining web scraping with LLM-based skill extraction and normalization, the tool provides a transparent, reproducible, and low-cost alternative to proprietary labor market analytics platforms. This study demonstrates the pipeline through a comprehensive analysis of entry-level economics positions across five U.S. regions and 51 metropolitan areas, confirming the growing importance of programming (SQL and Python), data visualization, and communication skills, a skill portfolio that many undergraduate programs address only partially.

UpSkill-Scanner’s primary contribution, however, extends beyond any single discipline. The tool is designed so that any academic department, such as economics, computer science, nursing, business, engineering, or any other field, can conduct the same analysis by simply changing the input job titles and geographic parameters. At a cost of approximately one dollar per analysis, systematic curriculum–labor market alignment assessment is now accessible to every program, regardless of institutional resources. As the pace of skill evolution continues to accelerate, regular, data-driven monitoring of employer demand can serve as a valuable feedback mechanism for departments seeking to prepare their students for the careers that actually await them.

References

Al Kafio, M. A., Monio, R., & Banshalo, S. K. (2025). Code vs. context: STEM students' resistance to non-STEM coursework. *arXiv*. <https://arxiv.org/abs/2512.06529>

- Allgood, S., Walstad, W. B., & Siegfried, J. J. (2015). Research on teaching economics to undergraduates. *Journal of Economic Literature*, 53(2), 285–325.
- American Economic Association. (2025). Careers in economics. <https://www.aeaweb.org/resources/students/careers>
- Bothmer, K., & Schlippe, T. (2023). Skill Scanner: Connecting and supporting employers, job seekers and educational institutions with an AI-based recommendation system. In *Lecture Notes in Networks and Systems* (Vol. 581). Springer.
- Brown, M. A., Gruen, A., Maldoff, G., Messing, S., Sanderson, Z., & Zimmer, M. (2025). Web scraping for research: Legal, ethical, institutional, and scientific considerations. *Big Data & Society*, 12(4). <https://doi.org/10.1177/20539517251381686>
- Bureau of Labor Statistics. (2025). Occupational outlook handbook: Economists. U.S. Department of Labor. <https://www.bls.gov/ooh/life-physical-and-social-science/economists.htm>
- Bureau of Labor Statistics. (2025). Occupational outlook handbook: Financial analysts. U.S. Department of Labor. <https://www.bls.gov/ooh/business-and-financial/financial-analysts.htm>
- Colander, D., & McGoldrick, K. (Eds.). (2010). *Educating economists: The Teagle discussion on re-evaluating the undergraduate economics major*. Edward Elgar.
- 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.
- Folkers, D. A. (2005). Competing in the marketplace: Incorporating online education into higher education. *Information Resources Management Journal*, 18(1), 61–77.
- Frydenberg, M., Mentzer, K., & Patterson, A. (2026). The rapid rise of generative AI adoption among first-year college students. *Information Systems Education Journal*, 24(1).
- Grealis, T., Harmon, O., Patterson, A., & Tomolonis, P. (2024). Excel literacy in the classroom. *Business Education Innovation Journal*, 16(2).
- Hansen, W. L. (2001). Expected proficiencies for undergraduate economics majors. *Journal of Economic Education*, 32(3), 231–242.
- Hemelt, S. W., Hershbein, B., Martin, S., & Stange, K. M. (2023). College majors and skills: Evidence from the universe of online job ads. *Labour Economics*, 85, 102429.

- Kotter, J. P., & Schlesinger, L. A. (1979). Choosing strategies for change. *Harvard Business Review*, 57(2), 106–114.
- Mentzer, K., Frydenberg, M., & Patterson, A. (2024). Are tech savvy students tech literate? Digital and data literacy skills of first-year college students. *Information Systems Education Journal*, 22(3).
- Onodipe, G., Wilsker, A., Routon, P., Stephenson, A., & Huff, K. (2025). A skills-based approach to improving the economics curriculum: Insights from online job postings using Lightcast. *Eastern Economic Journal*, 51(2), 308–335. <https://doi.org/10.1057/s41302-025-00294-9>
- Patterson, A. (2024). Examining text consistency amongst FOMC statements and minutes subtext using deep learning. *Issues in Information Systems*, 25(3).
- Patterson, A. (2025). Assessing the alignment of FOMC statements with minutes using large language models. *Issues in Information Systems*, 26(2).
- Patterson, A., Frydenberg, M., & Basma, L. (2024). Examining generative artificial intelligence adoption in academia: A UTAUT perspective. *Issues in Information Systems*, 25(3).
- Patterson, A., Mocarsky, M., Cho, J., Harmon, O., & Calvert, C. (2023). Teaching data literacy and sports economic fundamentals using fantasy sports. *Business Education Innovation Journal*, 15(1).
- Patterson, A., Temple, C., Anderson, N., Rogalski, C., & Mentzer, K. (2025). The virtual stage: Virtual reality integration in effective speaking courses. *Information Systems Education Journal*, 23(4).
- Rruplli, E., Frydenberg, M., Patterson, A., & Mentzer, K. (2024). Examining factors of student AI adoption through the value-based adoption model. *Issues in Information Systems*, 25(3).
- Siegfried, J. J. (2024). Trends in undergraduate economics degrees, 2001–23. *Journal of Economic Education*, 55(3), 328–332.
- Siegfried, J. J., & Walstad, W. B. (2014). Undergraduate coursework in economics: A survey perspective. *Journal of Economic Education*, 45(2), 147–158.

Appendix

[Up-Skill Scanner Code Repo](#)

Figures A1 and A2 present regional skill distribution tree maps for the Northeast and West, respectively. As shown in Figure 1, the area of the rectangle is proportional to skill frequency within each region's posting sample, and color denotes the skill category. These regional tree maps should be read

comparatively: differences in tile size across regions for the same skill (e.g., SQL appearing notably larger in the West than in the Northeast) reflect meaningful variation in employer demand across labor markets. Readers are encouraged to compare each regional figure against Figure 1 to identify where a given region diverges from the national baseline.




Figure A1. Northeast Region Treemap

Relative to the national average, the Northeast shows stronger demand for Teamwork, Analytics, and Budgeting, consistent with the concentration of financial services and consulting roles in New York and Boston.




Figure A2. West Region Treemap

SQL and Python appear with greater relative frequency in the West than in any other region, consistent with the concentration of technology firms in Seattle and San Francisco