

Analyzing Job Postings to Improve IT Programs within Higher Education



By Adam Jon Richardson
Computer Information Technologies (CIT) Professor
Lansing Community College
July 6th, 2025



nitic.org | nitic@csc.edu

Contents

Executive Summary	2
Introduction	3
Acquiring Job Postings Data	5
Identifying Prevalent Employers	7
Analyzing Role-Based Trends	8
Monitoring Geographic Market Trends	10
Tracking Technology and Tool Trends	12
Final Encouragement	14
About the Author	15
Acknowledgements	16

Executive Summary

Information Technology (IT) programs at colleges and universities have historically relied on advisory groups and Business & Industry Leadership Team (BILT) meetings to understand current employer needs and market demands. At these meetings local industry leaders typically describe their workforce needs and vote on the importance of various forms of IT-specific knowledge and skills. While these meetings provide valuable insights and feedback, serving as the backbone for IT-focused insights, they also face some constraints due to limited time and sampling of the target population (employers).

Local job postings provide a treasure trove of data that can supplement traditional advisory/BILT processes, painting a more detailed picture and enabling a timelier understanding of local industry trends. There are many ways that faculty can obtain local job postings, and once obtained, these postings can be analyzed in numerous ways. The number of postings per organization/company can point out possible strategic alliances with the most prevalent employers. Job titles can be analyzed to identify role-based trends for entry-level hiring opportunities. Geographic locations can be used to reveal changing markets for IT opportunities. Job descriptions can be searched to track trends related to tools and technologies. Elements of a case study are described in the analysis sections to help illustrate the ideas and demonstrate utility.

“Local job postings provide a treasure trove of data that can supplement traditional advisory/BILT processes, painting a more detailed picture and enabling a timelier understanding of local industry trends.”

Introduction

Colleges and universities across the United States face the ongoing challenge of ensuring their IT programs remain current and relevant to employer needs. The rapid pace of technological change, combined with shifting industry demands and recent workforce reductions, requires educational institutions to maintain a close, constant view of the job market.

Historically, IT programs in higher education have relied on advisory meetings, which provide opportunities for administrators and faculty to network with local leaders in the IT industry once or twice a year. More recently, Business & Industry Leadership Teams (BILT) meetings have fostered a more rigorous exploration of industry trends, as votes by industry leaders on essential knowledge and skills are meant to directly impact curriculum. Additionally, BILT meetings are held two to four times a year, providing more frequent feedback. These are important forms of feedback, and they are not going away.

However, advisory groups and BILT votes/meetings still only tell part of the story. They represent a small sample of possible respondents from the greater professional community, and faculty and industry leaders are limited by time constraints. To be clear, the insights gained through BILT votes/meetings are priceless, and they inform practical ways that we can improve our programs. Even so, we must be careful about making inferences from the information shared in these meetings to the workforce demands across entire communities. Finally, the increasing speed of innovation in the IT world, powered largely by AI, can yield significant changes in the local job market in mere weeks.



These challenges are exacerbated by the time it takes to get new curriculum and courses approved. Faculty often must work to create several documents that outline any proposed changes, which typically involves significant research and collaboration with support staff. Then, they must build consensus among other faculty within their program/department. Once approved by faculty, these changes often must work their way through committees and deans and provosts, who consider accreditation and the economic viability of the proposed changes. Finally, work must be done behind the scenes to get any new courses or pathways approved for financial aid purposes. Often, because of how the curriculum approval cycles fall, course and pathway updates can take a full year to complete.

Local job postings offer a complementary data source that can address many of these limitations. Unlike advisory meetings, job postings provide access to a much larger sample of employer data, update continuously rather than quarterly, and reflect real-time market demands. Fortunately, instructors and administrators within IT programs/departments at community colleges, four-year colleges, and universities across the country have access to this resource to supplement the valuable information gained through advisory groups and BILT meetings.

“...job postings provide access to a much larger sample of employer data, update continuously rather than quarterly, and reflect real-time market demands.”

Acquiring Job Postings Data

There is a myriad of ways that faculty and administrators can gain access to job postings. Some programs already provide frequent job postings updates to students and faculty. These weekly emails or announcements are a great source of data, and they often consist of structured data that can be merged into a larger dataset with relative ease. Just remember to carefully examine any terms of service (often, using data internally is allowable, but one should always seek out legal advice).

Additionally, some colleges and universities have employers that provide the vast majority of hiring opportunities. For example, in the Greater Lansing Area, the State of Michigan provides the most hiring opportunities for new graduates with associate degrees in IT, and often these types of organizations make it relatively easy to pull in their job listings through feeds and email updates (after all, they want to find the best applicants).

One can scrape data from public websites, such as Indeed, Glassdoor, and other job postings sources. This provides the most control over the data you collect, but it also proves to be the most technically challenging. While effective, automating this type of web scraping requires a significant time commitment every week, as websites are constantly changing, and some providers actively work to make this type of scraping more difficult.

TIP: Always check the terms of service for any website from which you want scrape web content.

Finally, there are services that facilitate the retrieval of data (e.g., Bright Data, Lightcast, etc.). In fact, some higher education institutions already pay for these data sources, so many faculty and administrators already have access (it is always worth checking). These services can also be paid for directly. That said, the initial cost for “all” Indeed or Glassdoor job postings is often tens-of-thousands of dollars over the course of a year. However, careful filtering of data (e.g., job title contains “programmer”, location contains “Michigan”) can drop the associated costs to \$20 – 40 per month.

Whatever path one chooses, remember this guiding principle: a bit better is better. There is no perfect source or scheme for the collection of job postings. Anything that you do will likely improve your current understanding of market trends in the IT industry, which will empower you to improve your offerings for your students.

“In fact, some higher education institutions already pay for these data sources, so many faculty and administrators already have access (it is always worth checking).”

Case Study: Acquiring Job Postings Data

The author recently compiled a list of over 2,000 job postings for entry-level IT positions spanning from March 2022 to June 2025. These postings were collected from emailed job notifications sent to faculty and students. Python was leveraged within a Jupyter Lab notebook to merge the postings into one large spreadsheet and prepare the data for analysis.

Identifying Prevalent Employers

The importance of strong advisory groups and/or BILT groups cannot be overstated. Part of the strength of these groups relates to the representation of the most prevalent employers in the local region. After all, a common goal after completion of most IT programs within higher education is for students to be hired or to achieve some level of upskilling that helps the student achieve new career goals. This works best when community colleges know the workforce needs faced by local industry leaders.

Hiring rates among local companies and organizations are fluid, especially in the IT industry. Faculty and administrators must work diligently to keep a pulse on the employers hiring the most talent, both in terms of general types of businesses (help desk, software development), and in terms of the specific entities doing the hiring. We want to hear from these leaders.

Determining the number of job postings per employer is a relatively straight forward calculation. These counts can also be broken down by year to provide a view of trending companies and organizations. Using these insights can help you prioritize communications with the employers who will likely have the greatest impact on the success of your students.

TIP: If you only have time to do one thing related to analyzing job postings, start here! This is often the easiest analysis to accomplish, and yet it enables some of the largest gains in terms of the impact on your students.

Case Study: Identifying Prevalent Employers

The author leveraged Microsoft Excel to create a table displaying the total number of listings by employer using the aforementioned dataset. A review of the table revealed that a local university's IT department had become a leading source of hiring in the past year, and communications with this employer were initiated to promote future advisory/BILT meetings.

Analyzing Role-Based Trends

Job titles provide a straightforward form of information related to relevant job titles for entry-level IT positions. Additionally, one can break up the time periods by month, year, or custom time frames to identify role-based trends. However, although the rationale for this approach is straightforward, one typically must exercise caution while attempting to find meaningful trends. Some simple steps should always be applied, such as stripping the titles of any surrounding white space, replacing characters such as parentheses and dashes with white space, and using case insensitive searches/counts. This data cleaning and preparation helps, but we must also consider other potential issues.

There are many ways that similar positions can be titled, and this variance across companies and organizations complicates rudimentary counts of various titles. One organization might use the terms “software engineer i” to describe the same position that another company refers to as a “entry level software engineer”, and these differences in nomenclature can make similar roles seem very different if they are handled simplistically. Fortunately, there are effective approaches that can be used to deal with this level of diversity in job posting titles.

One approach involves the creation of standardized job titles, which can be manually or automatically created for each respective role, and then these standardized titles can be used instead of the varied titles found within the original job postings (e.g., “software engineer i” might be translated to “entry level software engineer”). This used to require a tremendous amount of tedious manual work, but one can now leverage generative AI (with caution, after verifying the results) to generate standardized titles for the postings. This requires a bit more technical sophistication, though, and one will have to frequently recreate the standardized titles and naming conventions shift.

TIP: Sometimes we are focused on job postings within one large organization (e.g., IBM, the State of Michigan, etc.). In this case, a simple handling of the job posting titles may work because of the controlled vocabulary.

A simpler approach in terms of dealing with the variety of job titles involves the use of n-grams, which are simply sequences of n consecutive words. While job title strings often vary significantly even when describing essentially the same role, the patterns of word sequences within the varying titles for one particular role often reveal shared patterns that can provide market insights. While less refined than the standardization approach, use of n-grams often proves simpler to implement, and it still yields helpful views of general prevalence and trends. For example, one can view counts for bigrams (sequences of two words) and trigrams (sequences of three words), and this can provide a comparison of the number of “software engineer” roles across years.

“While job title strings often vary significantly even when describing essentially the same role, the patterns of word sequences within the varying titles for one particular role often reveal shared patterns that can provide market insights.”

Case Study: Analyzing Role-Based Trends

The author leveraged Python within a Google Colab notebook to create bigrams for job titles within the aforementioned dataset. This work revealed that many job postings were focused on student positions (e.g., “student assistant”, “training information”), and subsequent analysis revealed some organizations underrepresented within our advisory/BILT group. Communications with these organizations were prioritized moving forward. Additionally, a comparison of 2022/2023 bigrams with 2024/2025 bigrams revealed strong growth in titles containing “data analyst”, which strengthened the case for the creation of new courses covering Power BI and/or Tableau, an emphasis in our recent BILT meetings.

Monitoring Geographic Market Trends

Geographic regions frequently experience broad swings in hiring patterns, both generally, and within specific markets. Students are most likely to find work when they seek out locales with the most opportunities, and faculty and administrators at colleges and universities must keep a watchful eye on the evolving hotbeds within regions. It is relatively easy to miss the impactful hiring trends in neighboring communities when job postings are unevaluated for geography. However, the good news is that monitoring geographic market trends is relatively easy to implement, and the results can have a powerful impact.

Most job postings include location information, whether in the form of specific addresses, cities, or broader regional descriptions. One can leverage this geographic data to create counts and visualizations that reveal where opportunities are concentrated and how these concentrations shift over time. Simple approaches might involve creating tables or charts that display job counts by city or county, broken down by time periods such as quarters or years (or larger timeframes if the sample of postings is small). More sophisticated approaches can incorporate mapping tools or geographic information systems (GIS) to create visual representations of hiring hotspots.

When analyzing geographic trends in light of opportunities, it is important to consider the populations of communities. Raw job counts can be misleading if they do not account for population differences. A small city with ten IT job postings might represent a more robust job market than a large metropolitan area with fifty postings, depending on the respective populations and existing workforce. One can address this by calculating jobs per capita or by focusing on percentage changes over time rather than absolute numbers, though this is a more advanced pursuit (remember, a bit better is better).

TIP: Start with simple city-by-city counts before moving to more sophisticated geographic analysis. Even basic location data can reveal surprising trends that inform advisory/BILT communications, student advising, transfer opportunities, and program marketing efforts.

Geographic trends can also inform strategic decisions about satellite campuses, online program delivery, or partnerships with other educational institutions. If analysis reveals growing IT opportunities in a neighboring county with limited educational options, this might represent an opportunity for expanded program reach.

Case Study: Monitoring Geographic Market Trends

The author leveraged Python within a Google Colab notebook to compare geographic trends between 2022/2023 and 2024/2025 (sometimes grouping years together can help when data within some categories is limited). Three cities near the community college experienced significant growth in terms of IT job postings. These trends prompted investigations into strengthening possible transfer opportunities with four-year institutions within these communities. Additionally, communications were initiated with representative employers in the region to promote attendance at upcoming advisory/BILT meetings.

TIP: Depending on the time frame chosen for comparison, some categories may contain very few instances, making it difficult to see trends. Use of larger time frames (e.g., quarters instead of months, two-year periods instead of years) can sometimes help.

Tracking Technology and Tool Trends

Perhaps no aspect of job posting analysis proves more valuable to IT curriculum development than the identification of technology and tool trends. The rapid pace of technological change in the IT industry means that the specific programming languages, software platforms, and technical skills in demand can shift significantly within just a few ~~years~~ months. Faculty and administrators need timely, data-driven insights to ensure their programs remain well aligned with employer expectations and that graduates possess the most marketable skills.

Job descriptions typically contain detailed lists of required and preferred technologies, making them an excellent source for tracking which tools and platforms are gaining or losing traction in the local market. However, extracting meaningful insights from this information often requires careful consideration of how technologies are mentioned and referenced across different postings.

The most straightforward approach involves searching job descriptions for specific keywords relating to technology and counting their frequency over time. One can create lists of relevant technologies for different IT disciplines. For example, a list of programming languages might include Python, Java, JavaScript and Go; and a list of cloud-related terms might include AWS, Azure, GCP, S3, and IaaS. One can search job descriptions for mentions of these terms quite simply (i.e., any match is considered a match). This approach works well for getting a general sense of which technologies appear most frequently in local job postings.

However, several challenges complicate this seemingly simple task. Different employers may refer to the same technology using various names, abbreviations, or versions. For example, job postings might reference "JavaScript," "JS," "ECMAScript," or specific frameworks like "React" or "Node.js" when describing similar skill requirements. Additionally, some words prove to be generally ambiguous. Remember the reference to the "Go" programming language? You may suddenly find that the Go programming language *appears* wildly popular merely due to job postings that make use of the common verb "go".

TIP: Depending on sample size, be cautious about over-interpreting short-term fluctuations in technology mentions. A single large employer posting multiple positions requiring a specific tool can skew the data significantly. Look for sustained trends over several months or quarters before making major curriculum decisions.

We've already discussed how we can standardize terminology in the job titles section, so let's focus on the issue of ambiguity. One could leverage an API to one of the frontier generative AI models, feeding in each occurrence of one of the identified terms, and then ask the AI to classify the terms as technical or other (e.g., "Go to the interview..." would not count towards the frequencies, but "One year of experience programming in Go..." would). This is a bit more technical, though, and one does pay for API use. One can also scan nearby words in the context looking for indicators of technical usage, and while not as reliable as the previous approach, it does not run the risk of added costs and proves quite simple to implement (one can leverage generative AI to generate lists of indicator words).

When analyzing technology trends, it is important to consider both absolute frequency and relative growth rates. A technology that appears in 50% of job postings represents a critical skill for students, regardless of whether its usage is growing or declining. Conversely, a technology that appears in only 10% of postings but has doubled in frequency over the past year might represent an emerging opportunity worth incorporating into elective coursework or specialized tracks.

“When analyzing technology trends, it is important to consider both absolute frequency and relative growth rates.”

Case Study: Tracking Technology and Tool Trends

The author leveraged Python within a Google Colab notebook to compare frequencies of terms associated with several sectors within IT (e.g., cybersecurity, data processing, programming, AI/ML, cloud, networking, web, etc.). Indicator words within the context of the term appearance were used to validate use considering potential ambiguity. To help visualize trends, heatmaps were created, which help other faculty and administrators scan the results. Among other findings, the results revealed a surge in Tableau and Power BI use (data processing), strong uptake of web frameworks such as Angular, React, and Vue (web), and the continued importance of teaching TCP/IP (networking).

Final Encouragement

We covered an important principle earlier in the paper: a bit better is better. We in the IT industry know this to be true, perhaps better than any other industry (our product/development cycles are built around the power of iteration and making incremental improvements). As educators it is easy to feel overwhelmed as we attempt to keep up with the changing needs of our students, of the workforce, and of the technologies that drive innovation. However, we can take solace in the knowledge that a bit better really, truly is better.

Do you have issues finding a source of job listings? Go peek around Indeed or Glassdoor and just find one new company with which you could try to open communications.

Do you lack experience using Excel, Python, Colab, or similar tools? Try using the research capabilities within generative AI tools like Perplexity to find local job posting trends.

Do you experience errors or issues as you try to implement some of these approaches, or do you just need general encouragement? Reach out to the author or other members of NITIC.

We are all on the same team!

“...we can take solace in the knowledge that a bit better really, truly is better.”

About the Author

Adam Richardson is a full-time Professor within the Computer Information Technologies (CIT) Program at Lansing Community College (LCC). Adam worked full-time in the IT industry for 16 years prior to transitioning to the academic world and possesses a Master of Information Technology degree from Virginia Tech that focused on software engineering and machine learning. Adam has earned over twenty technical certifications from CompTIA, AWS, Intel, Nvidia, and Stanford University. Adam currently teaches a broad range of IT courses at LCC, including courses focused on introductory Python, AI/machine learning, cloud computing through AWS Academy, computer science, cybersecurity, and electronics. He currently acts as the advisor for the LCC Artificial Intelligence Club, and he is an elected senator to LCC's Academic Senate who also serves on LCC's Curriculum Committee.

You can connect with Adam via LinkedIn:

<https://www.linkedin.com/in/adamjonr/>

Acknowledgements

This work was supported by the National Information Technology Innovation Center (NITIC) through work with the Data Technology Team. The author thanks NITIC for their commitment to improving IT education and workforce development.

To contact NITIC, please email:

nitic@cscs.edu



This material is based upon work supported by the U.S. National Science Foundation under Grant No. 2300188. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.