S A M Y R A D 2026

SMART ALGORITHMS & MODELS: YOUR REACH TO ADVANCED DIALOGUES · 5–6 OCTOBER 2026 · SEVILLE, SPAIN
SAMYRAD OPEN-ACCESS PROCEEDINGS | PUBLISHER: INTELLIGENT NUDGE

PAPER • OPEN ACCESS

Skill (Mis)Alignments: A case study of a recent educational program against industry needs

Juan A. Talamás-Carvajal (e-mail: juan.talamas@tec.mx), Gustavo De Los Rios (e-mail: A01410922@tec.mx), and Patricia Caratozzolo (e-mail: pcaratozzolo@tec.mx)

[1] Institute for the Future of Education, Tecnologico de Monterrey, Mexico

[2] School of Engineering and Sciences, Tecnologico de Monterrey, Mexico

Article DOI: 10.68286/samyrad2026.356

ABSTRACT

The persistent misalignment between Higher Education curricula and rapidly evolving industry requirements challenges institutions seeking to prepare graduates for dynamic labor markets. This study presents a data-driven approach for assessing curriculum–industry alignment through a case study of a recently developed Information Technology engineering program and occupational demand in the INFOCOMM (Informatics and Communication) industry. We extracted Knowledge, Skills, and Abilities (KSAs) from educational program content and job postings from Mexico and the United States using Natural Language Processing and Large Language Models, then classified them using a domain-specific Infocomm taxonomy. We then used weighted Jaccard similarity to quantify alignment between the educational program and selected occupations over seven months. Results reveal substantial variation across occupations and KSA dimensions. Skills showed the highest overall alignment, followed by Abilities, whereas Knowledge exhibited the greatest misalignment. Occupations closely associated with software engineering generally showed stronger correspondence with the curriculum, while more specialized or adjacent occupations showed substantially lower alignment. The findings show that curriculum–industry alignment is multidimensional and occupation-dependent, not adequately captured by a single aggregate skill-gap measure. The proposed approach provides universities with an evidence-based mechanism for identifying curricular strengths, gaps, and potential areas for targeted intervention as industry requirements evolve.

KEYWORDS Higher Education, Informatics and Communication, Skills, and Abilities, Natural Language Processing, Large Language Models

How to cite: Talamas-Carvajal, J. A., De Los Rios, G., & Caratozzolo, P. (2026). Skill (mis)alignments: A case study of a recent educational program against industry needs. SAMYRAD 2026 Open-Access Proceedings. Intelligent Nudge. https://doi.org/10.68286/samyrad2026.356

Introduction

One of the longest-standing gaps that educational institutions face is the disconnect between what Higher Education Institutions (HEIs) teach and what the targeted industries need. Recent graduates are often told they will “learn the job when they get the job” and that much of what they were taught is either obsolete or not appropriate for certain roles. This misalignment stems from 2 main factors: industries that constantly push for innovation and optimization, and the slow response time of Educational Institutions to changes outside them [1]. Paradoxically, while many state-of-the-art discoveries occur within these institutions, the adoption of novel technologies and strategies is much slower [2]. Recent studies have further highlighted that digital transformation, automation, and Industry 4.0 paradigms continue to accelerate changes in workforce requirements, intensifying the need for agile educational adaptation [3]. While this gap has generated tension among industry, universities, and students, it also creates an opportunity for educational

innovation. Although eliminating the gap may not be feasible because of the need for technological maturation, proprietary industrial knowledge, accreditation constraints, and institutional inertia, it can certainly be reduced. One frequently proposed strategy is to develop stronger collaborations between HEIs and industry stakeholders, enabling curricula to evolve in response to emerging professional demands [3]. Such collaborations can also improve graduate employability, facilitate workforce transition into emerging sectors, and support regional economic competitiveness. A promising approach to reducing this gap is to systematically analyze competencies and workforce skills. The European Commission Joint Research Center defines competencies as the "*demonstrated ability to use knowledge, know-how, experience and – job-related, personal, social or methodological – skills, in work or learning situations and in professional and personal development*" [4]. Competencies are commonly operationalized through classifications such as Knowledge, Skills, and Abilities (KSAs), which provide a structured framework for analyzing workforce preparedness and occupational requirements [5].

In this study, we use the term "Skills" as a general reference to the components that comprise competencies within the KSA framework. The term "Knowledge" is self-explanatory, but we differentiate Skills and Abilities as follows: a Skill is a capacity necessary to complete a particular task, and an Ability relates to having the sensory, physical, psychomotor, or cognitive means to perform an action. Any competent professional carries a specific set of KSAs necessary to perform effectively in their occupational domain. Consequently, individuals possessing those same KSAs are theoretically more likely to succeed in equivalent professional roles. Educational institutions can therefore mitigate curriculum obsolescence by continuously monitoring evolving industry-relevant KSAs and incorporating them into academic programs through targeted learning experiences, competency-based curricula, and updated course content. This competency-centered perspective has become increasingly important as labor markets evolve more rapidly and interdisciplinary technical profiles become more common. One way this alignment has strengthened in recent years is through automated extraction and analysis of skills from labor market data. Advances in Natural Language Processing (NLP) and Machine Learning (ML) have enabled large-scale analysis of job postings, allowing researchers and policymakers to identify emerging competencies and changing occupational demands with greater accuracy and speed. Taxonomies such as ESCO in Europe and SINCO in Mexico, along with recent industrial skill studies, have demonstrated the usefulness of computational approaches for mapping workforce requirements and supporting educational planning [6], [7], [8], [9], [10]. Most existing studies focus on labor market analysis or curriculum design independently, rather than systematically comparing the two ecosystems. Despite these advances, little research directly compares educational program content with large-scale industry skill taxonomies to identify concrete curricular gaps and intervention opportunities. Methodologies are still needed to systematically evaluate how well modern educational programs align with rapidly evolving industry ecosystems, such as the Infocomm (Informatics and Communication) industry, which encompasses information technologies, software services, digital communication systems, telecommunications, and networking.

In this paper, we present a case study demonstrating how a taxonomy developed for the Infocomm industry can be used to evaluate educational programs and identify gaps, obsolete topics, and core KSAs relative to industry needs. We also show how the results can inform interventions that improve curriculum alignment and strengthen graduate preparedness. To guide this study, we propose the following research questions:

RQ1. *To what amount do modern educational programs match the needs of their target industry in terms of developing the necessary skills for jobs in that industry?*

RQ2. *Can the differences found be leveraged towards improving educational programs, student preparedness, and/or collaborations between industry and HEIs?*

Background

The Skill Gap

As mentioned before, a persistent gap exists between what students learn in school and what they need to succeed in their occupations after graduation. This gap reflects differences in KSAs between what students graduate with and what industry requires. In short, it is a difference in competence. This phenomenon has been widely discussed in educational policy, workforce development, and labor economics literature, particularly in the context of digital transformation and rapidly evolving technological ecosystems [4]. This gap goes beyond what might traditionally be called "corporate culture," meaning the rules, attitudes, and processes specific to individual workplaces. Instead, it reflects structural differences between educational preparation and occupational expectations. The European Union defines this phenomenon as the "Skill Gap" or "Skill Needs," which remains the subject of extensive research and policy development initiatives [4]. Similar concerns have also been identified in studies focused on Industry 4.0, digital economies, and emerging technical professions, where workforce demands evolve faster than institutional curricula [3].

This issue is problematic because no single solution exists, and many factors widen the gap. The time, investment, and effort required to develop the necessary skills are among the biggest barriers to entry for newly graduated students entering the workforce. Furthermore, the increasingly interdisciplinary nature of technical occupations means graduates are expected to have not only technical knowledge but also communication, adaptability, teamwork, and problem-solving skills. Therefore, both industry and HEIs have an interest in ensuring students are as prepared as possible before entering the job market. A seemingly straightforward way to address this issue is to continuously develop specific KSAs as needed. For example, continuous professional training could keep the workforce up to date on technological breakthroughs, practical courses could strengthen operational competencies, and holistic educational approaches could support the development of broader abilities. However, the challenge lies not only in developing KSAs but also in identifying which KSAs are most relevant, emerging, or strategically important for a given occupational ecosystem. This challenge has motivated the development of large-scale labor market analysis methodologies that systematically

identify and monitor workforce competencies using computational approaches and data-driven models.

Automated Skill Extraction and Labor Market Intelligence

One current solution is to automate skill extraction from job markets. Traditional analyses of occupational skills were often manual, time-intensive, and hard to scale because job postings are unstructured. While basic information such as job titles and company names tends to follow relatively standardized formats, the most relevant information for workforce analysis, namely the competencies being requested, is commonly embedded within *free-form* job descriptions. Recent advances in Natural Language Processing (NLP), Machine Learning (ML), and Large Language Models (LLMs) have enabled the automated extraction, classification, and analysis of competencies from large collections of online job postings. These approaches allow researchers and policymakers to process massive volumes of labor market information from platforms such as *LinkedIn*, *Indeed*, and *JORA* far more efficiently and consistently than manual methods. Several initiatives have demonstrated the value of computational skill extraction methodologies. Examples include the ESCO taxonomy developed by the European Commission, the SINCO occupational classification system developed by INEGI in Mexico, and recent studies focused on the Automotive and Infocomm industries [6-10]. These initiatives support workforce planning, educational policy development, occupational mobility analysis, and strategic upskilling efforts. One key advantage of automated skill extraction approaches is their ability to capture much larger and more representative samples of industry needs than traditional manual analyses. As a result, the insights generated from computational approaches can more accurately reflect the dynamic evolution of labor market demands. Additionally, continuous monitoring of online job postings enables near-real-time identification of emerging technologies, evolving professional profiles, and evolving competency requirements. Despite these advances, extracting skills from labor market data alone does not fully address the educational alignment problem. Most existing approaches focus primarily on industry demand, while comparatively few studies compare extracted competencies with educational program content and graduate profiles. Consequently, methodologies that bridge labor market intelligence, curriculum analysis, and educational evaluation remain needed.

Competency Taxonomies and Educational Alignment

Competency taxonomies play a critical role in connecting labor market analysis with educational program evaluation. By organizing competencies into structured categories such as KSAs, these frameworks provide a common language for systematically comparing workforce requirements and educational outcomes. Taxonomies such as ESCO and SINCO have demonstrated that standardized competency representations can support multiple objectives, including occupational mobility, reskilling, upskilling, workforce forecasting, and educational planning [9], [10]. In educational contexts, competency frameworks also facilitate curriculum mapping, learning outcome assessment, and identifying gaps between intended graduate profiles and actual industry requirements.

The growing availability of computational methods for competency extraction has enabled more sophisticated alignment methodologies between industry and academia. In particular, similarity-based approaches and semantic matching techniques enable researchers to compare educational content with labor market requirements at scale, revealing both strengths and weaknesses in academic programs. However, research still has limited focus on applying these methodologies directly to detailed educational program structures, including course objectives, content, and graduate profiles. Most existing studies focus on labor market analysis or curriculum design independently, rather than on systematic comparisons between the two ecosystems. This limitation becomes especially relevant in rapidly evolving sectors such as the Infocomm industry, where educational programs must continuously adapt to technological transformation and shifting occupational demands.

This study addresses this gap by using an Infocomm competency taxonomy to compare workforce requirements with the content and graduate profiles of a modern engineering program. Through this approach, we aim to identify curricular gaps, core competencies, and opportunities for improving alignment between HEIs and industry needs.

Materials and Methods

The present study comprises three main stages: skill extraction from job postings obtained from online sources relevant to the targeted industry, Skill extraction from the educational program(s) being evaluated, and, finally, a comparison/matching methodology to measure the amount of shared and missing Skills between them. The Infocomm taxonomy has already been developed, and its results are shown in [6]. The process described below has been validated in [11,12] through multiple review rounds and cross-validation across annotators.

Stage 1: Labeling and processing pipeline for job postings

The process of building the job postings dataset is composed of several steps: data sweeping or scanning (scraping) for relevant job postings on several job search websites (LinkedIn, JORA, Indeed), atomic idea extraction related to the KSA-O framework from the retrieved job postings, and classification of such ideas into a Skill, Knowledge, Ability, or Other. In the case of “Other”, this refers to requirements for the job that do not fall into the regular KSA terminology, but are nevertheless required. For example, specific credentials, availability, or certifications. This final step also groups variants of a similar Skill into one using a representative term.

We scrape the data through the APIFY platform and retrieve it as JSON or CSV files containing the information summarized in Table 1. The extracted features provide the necessary information for Skill extraction and classification and record the industry's evolution over time.

After obtaining enough data, we used a local version of ChatGPT called TecGPT, based on GPT-4o, to process the information, extract the essential KSA-Os, and place them into a structured Python list. We used a set of specific prompts developed after several iterations. These prompts did the following: create a structured list for readability with

requirements, certifications, essential skills, and other specific knowledge or expertise as separate items. The following is the base prompt used for skill extraction from job descriptions using GPT-4:

System: You are an information-extraction system specialized in identifying skills in job descriptions. Your task is to recognize both technical skills and interpersonal skills explicitly mentioned in the text.

User: *The following is the job description. Extract all explicitly mentioned skills and return them as a Python list.*

The model then returns the results via a .txt file, which can then be added to the dataset. The next step in the pipeline is classifying the recovered atomic ideas into specific KSAs under the taxonomy. This is done by running the output from the local GPT system through a fine-tuned (one-shot prompted) GPT-4o model via the OpenAI API. In this process, we use the Infocomm taxonomy as a reference through a process called Retrieval-Augmented Generation [13], producing a set of Skills, Knowledge, and Abilities aligned with what we found while developing the original taxonomy. To avoid instances where a Skill is expressed in similar terms ("teamwork" and "works in teams", for example), we created hypernym- and hyponym-based groups [14], which were then used in the RAG process. This generalization aimed to convert industry Skills into a form that could be more easily compared with Higher Education Skills.

Table 1. Extracted features from job posting websites

Feature name	Description
salary	Yearly salary range being offered
postedAt	Time since the job was posted
externalApplyLink	Link to direct application page
positionName	Job position being offered
jobType	Full-time or part-time
company	Company offering the job
location	Address of the job being offered
rating	Rating of the job posting
reviewsCount	Number of reviews of the job posting
urlInput	URL used to search for the job
url	URL of the job position
id	Unique ID of the job posting
scrapedAt	Date of scraping
postingDateParsed	Date of publishing
description	Overall description of the job posting
descriptionHTML	HTML version of the description
searchInput	Keyword specified to the scraper if not using a URL
isExpired	Indicates whether the job is still available at the time of scraping
q_extracted	Text query used to scan for a particular position
q_extracted_norm	Normalized version of the search query

To bridge the language differences between HEIs and industry, we took two additional generalization steps: we obtained canonical (more general) hypernyms [15], [16] for the obtained KSAs, and we separated them into parent categories. This way, we can properly identify specific KSAs that might initially register as "not matching" or "missing" due to granularity. One common example is specific programming languages or structures. Cases like SQL, DHCP, and Python can be properly categorized into general groups, and then into parent domains like "Databases", "Networks", and "Programming Language", respectively. This allows for a much clearer and fairer comparison between the two sets of KSAs. Using the previous process, we scraped seven months' worth of job postings.

This resulted in a total of 4032 job postings analyzed over the seven months, as shown in Table 2. The overall distribution of KSAs in the dataset was 66% Knowledge, 25% Skills, and 9% Abilities.

Table 2. Job postings and KSAs per Month

Month	Job Postings	K	S	A
Jul-25	540	11420	8969	5174
Aug-25	519	11209	9061	4971
Sep-25	560	12062	9550	5427
Oct-25	624	14569	10616	5715
Nov-25	547	12036	9235	5213
Dec-25	520	11523	8663	4959
Jan-26	722	15604	11694	6811

Stage 2: Skill extraction from educational program

Building the educational program dataset followed the same process as the job postings dataset. Because the educational models came from an official source, we did not need the initial data-scraping step; the remaining steps matched those described above: atomic idea extraction and classification of those atomic ideas into KSAs via RAG. As with the job postings dataset, a local version of ChatGPT performed atomic idea extraction to process the information, extract the essential KSA-Os, and place them into a structured Python list for later classification using the same taxonomy used for the job posting Skills.

One important detail is that the educational programs used are the "detailed" versions. This means we have not only the names and order of all the courses that comprise the majors, but also the general objective of those courses, the course content, and the specific topics covered in each course. We also recovered the graduate profile for each major, which provides a general description of the capabilities of students who graduate from the major, as well as a summary of the most relevant KSAs they will obtain through the program. The program selected for this comparison is Engineering in Information Technology (ITC), which prepares students for jobs in the Infocomm industry.

Stage 3: Matching methodology

Having obtained both the industry and Educational Institution KSAs, we proceeded to the comparison/matching stage. One aspect of the comparison is that a full comparison of the KSA portfolios of both industry and HEIs is not valuable, as the complete set of KSAs in the industry is much larger than what any institution can teach. Instead, we compare the top Skills, Abilities, and Knowledge needed by industry occupations against the top ones developed by the educational programs. We use the Weighted Jaccard Similarity for this comparison. Jaccard Similarity (or Naïve Jaccard) is a statistical measure that quantifies how similar two sets are. It equals the ratio of the intersection size to the union size. Jaccard similarity ranges from zero to one: zero indicates the sets share no elements, and one indicates they are identical. We use the similarity index to calculate document similarity. In this case, a document is the set of KSAs from both industry and HEIs [17]. Jaccard Similarity is defined as:

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|}$$

Weighted Jaccard extends the Jaccard index to numerical vectors as well, if these vectors are non-negative. Weighted Jaccard measures similarity by "calculating the sum of minimum weights divided by the sum of maximum weights for each dimension, resulting in a value between 0 and 1 that represents the overlap magnitude". It is commonly used to compare documents in data mining. Weighted Jaccard is defined as follows:

$$J_w(x,y) = \frac{\sum_i min(x_i, y_i)}{\sum_i max(x_i, y_i)}$$

We obtained different similarity scores for the program against selected occupations for each monthly extraction. In cases where weight scores were obtained, the weights were defined as *0.5 Knowledge + 0.3 Skills + 0.2 Abilities*, based on the observed KSA frequency, as well as the impact of the different categories in determining job matches. It is important to note that the "Others" category was discarded because it doesn't represent skills, but traits like experience, certifications, or place of residence, for example.

Results

Similarity matching between the ITC program curriculum and job postings yielded the following: we observed high match scores (~50% knowledge) for Software and Artificial Intelligence Engineer occupations, and low scores (~5 or less) for ICT security and marketing manager occupations. These results cover 7 months of job postings, during which the match percentage changes each month. Fig. 1 shows these results as a heat map summarizing the results by occupation and month. One important detail is that some months had no job postings for one or more occupations; in those cases, the space was left blank to indicate the absence of a match score. Additionally, most occupations averaged 20% to 30% similarity, while only the first-ranked ones averaged above 40%.

Occupations are listed by the average rank they achieved across all months; for example, Software Engineer averaged 42% similarity throughout the experiment. Some cases show similarity varying from a high score of 43% one month to a low score of 12% another, as in the Data Scientist occupation. The occupations with the fewest registered appearances were Computer System Analyst and Data Architect, with each case missing 5 and 4 months, respectively.

To avoid cases where missing months might disproportionately affect rankings, we excluded missing similarity scores when computing average scores. Additionally, because job postings follow industry needs rather than a structured calendar, we assess whether occupations released less often have their knowledge needs met.

For Knowledge, we observed that the occupations with the best overall matches were those directly related to software (software engineer, software architect, software developer, etc.) or those with "engineering" in their title. These two groups consistently ranked in the top 50% of all occupations, while occupations further from classical engineering had the lowest similarity scores. Marketing manager, product designer, and database administrator, for example, are occupations adjacent to classic engineering jobs but still require related knowledge.

On Skills, we observed higher overall matches, with scores around 55% and a higher bottom of 10%. Aside from a few isolated cases, almost all occupations had their base Skill needs covered to an average of around 40%. Fig. 2 makes it clear that Skills are more commonly matched between industry needs and the educational program being compared. Aside from a few cases, the general tendency seen with Knowledge matches is maintained: engineering occupations were more commonly seen toward the top occupations, while more adjacent cases were seen toward the bottom.

Finally, for Abilities (Fig. 3), we observed a higher ceiling than for the other two components of the KSA framework, with some cases achieving a 60% similarity rate, and a similar floor to the Knowledge case, with some cases hovering around the 5% mark. Overall, Abilities fell between the previous two components, with better similarity than the Knowledge case but lower than Skills. In this case, there appears to be no difference made towards specific occupations that mention "engineering" as part of their description, unlike the previous two cases.

Overall matching, considering all 3 aspects of the KSA framework, showed similarity scores averaging 30% to 35%, mostly due to the influence of the Knowledge aspect. This is partly because, in the blended case, we prioritized Knowledge due to its importance.

Discussion

As mentioned before, the very nature of the Skill Gap makes it a complex and persistent problem. This is reflected in our matching algorithm results, as the highest match we observed was around 62% for the Ability comparison. The score alone might raise concerns; the issue is exacerbated by the fact that this specific program was scheduled for a 2026 implementation date. It is important to note that the program was accredited by CACEI (Council of Accreditation of Engineering Teaching for its initials in Spanish), ABET (formerly Accreditation Board for Engineering and Technology, now working under the acronym exclusively), FIMPES (Federation of Mexican Private Higher Education Institutions), along with experienced teachers and members of the industry, and the national educational organism (SEP). This speaks to the complexity of building an educational program that meets requirements from different government agencies and accreditation institutions, provides the foundational knowledge needed to train future professionals to understand the underlying principles of their work rather than being purely technically trained, and delivers the state-of-the-art skills industries require.

Another possible explanation is the fast-changing nature of the INFOCOMM industry. New programming languages, hardware advances, the rapid evolution of AI tools, and global industry interest make it highly volatile. The speed at which this industry changes has far outpaced the traditional structure of a university degree, further widening the Skill Gap between newly graduated workers and industry needs. While the case of INFOCOMM-related programs is clearer and the surrounding circumstances make it more pronounced, this phenomenon has been happening across almost all University programs, especially engineering-based ones [19].

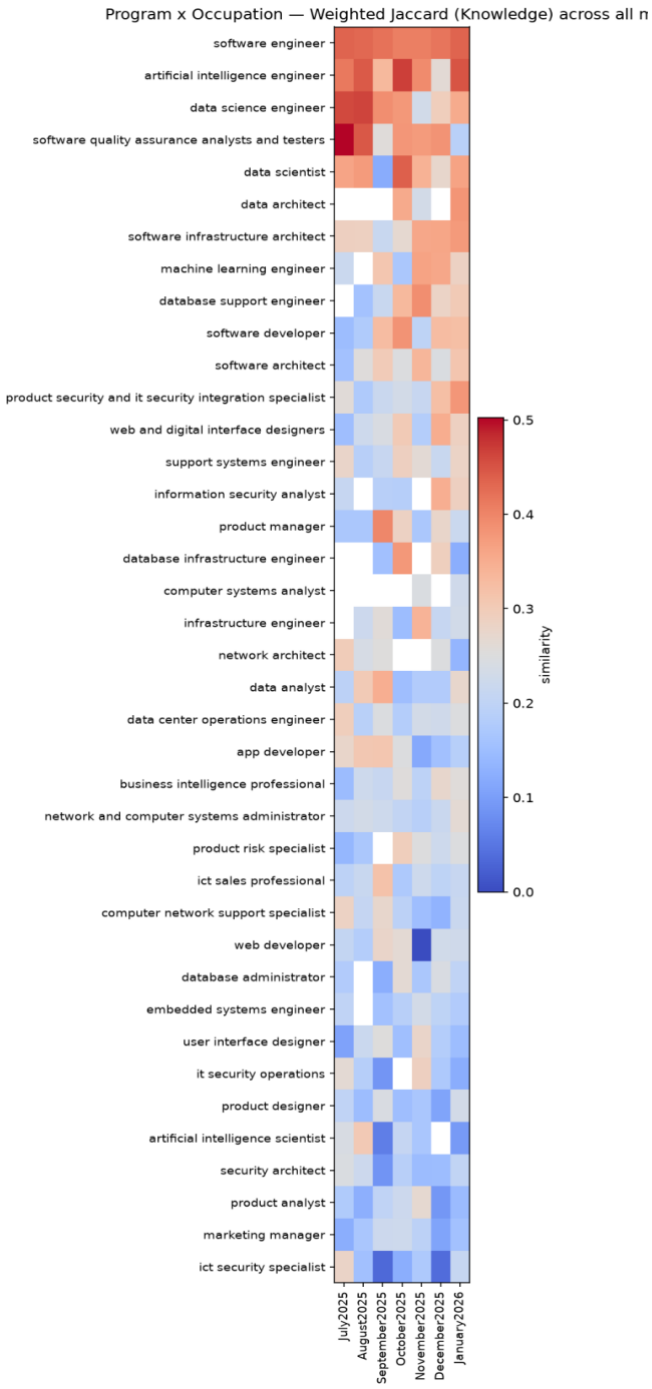


Fig. 1. Weighted Jaccard similarity scores between Job Postings and the ITC program per month for Knowledge.

Similarly, but in the opposite direction, the skill gap (specifically in the INFOCOMM sector) might be reversed: academia is ahead of industry (at least in the INFOCOMM sector). This is the case specifically for legacy systems, systems that have aged past their time and are now considered outdated, but still perform mission-critical processes: "*large software systems that we don't know how to cope with but that are vital to our organization*" [20]. In this case, job postings show skills that are "outdated" relative to the industry's current state. Examples include the need for platform-specific software like .NET, hardware-specific needs like a PLC running exclusively on Windows XP or requiring Visual Basic, or network architectures that don't support modern industry practices. This generates an interesting dilemma both for universities and the industry: the industry benefits from following the state-of-the-art and best practices, as it assures security, reliability, and compatibility with most modern needs, but a large percentage of their systems are outdated and would require extensive work and investment to modernize, leading to legacy systems that require maintenance and personnel. Meanwhile, universities aspire to prepare students for their future careers and often look to current best practices and developments for guidance, leaving legacy systems behind, as their name implies.

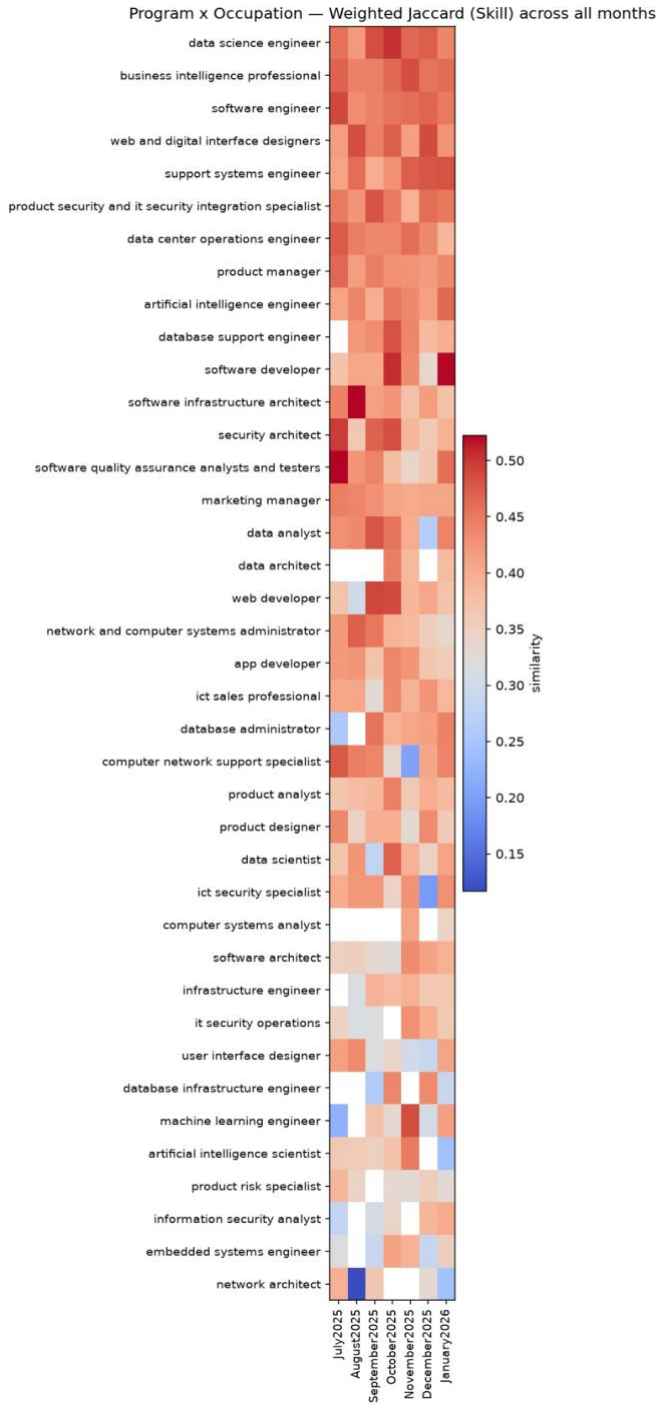


Fig. 2. Weighted Jaccard similarity scores between Job Postings and the ITC program per month for Skills.

Skills and Abilities represent a different situation altogether, as they reflect broader aspects of professional life that do not change as easily. Teamwork, Time management, attention to detail, and many other Skills and Abilities are, to some degree, evergreen and, as such, not a problem for the Skill Gap. Low skill-matching scores can occur because skills

are not always stated explicitly in a course's description or objectives. This omission does not mean the skill is not developed; it is developed indirectly, as there is no such thing as a “teamwork learning course” in most University courses.

It is also important to consider the nature of each component of the KSA framework and industry needs: while Knowledge is desired, to what degree is it necessary? In some cases, Knowledge is required only to the level of a technician, and not an engineer. Similarly, Skills might be more favored, as they represent the specific capabilities needed. Finally, Abilities are broader capabilities that are generally desired and commonly developed in many educational programs. Our similarity scores show a clear pattern: Skills are most similar between educational programs and job postings, while Knowledge is the least similar.

One final aspect worth discussing is the presence of occupations with very diverse similarity scores across months. For example, the Data Scientist occupation showed similarity scores ranging from ~40% to ~10% in consecutive months for Knowledge. We believe this month-to-month variability reflects the industry's diverse needs, even within the same occupation, and the industry's constantly changing nature.

Regarding RQ1: We found that, even in the best-case scenario, no cases existed where the program covered every single industry requirement. This could be due to the specialization of different businesses versus the more general knowledge base most universities develop in their students, the fast-changing nature of the industry (where the state-of-the-art shifts month to month), or the considerable need for legacy system support in the industry.

In terms of RQ2: The reality of the Skill Gap between recent graduates and Industry needs is clear, and it is unlikely to be solved by following the traditional structure of an educational program. Instead, it may be worth moving toward a specialization period in the later stages of degrees: cover the basics of the Skills required for employment in the initial and middle stages, then prepare students through elective courses, preferably developed in collaboration with industry. In this regard, it might benefit students and their future employers if these courses also functioned as micro-credentials, so they can demonstrate their experience and Skills.

Regarding limitations: the job postings were obtained from websites mainly serving the North American and Mexican markets and, as such, are tailored to those regions. Other markets may prioritize different Skills, so the results might differ. The educational program, while accredited by several organizations, does not represent all Information Technology majors. However, it is the program offered by one of the top universities in its region.

Conclusion

The results of this study suggest that curriculum–industry alignment should not be understood as a binary condition in which an educational program either does or does not satisfy industry needs. Instead, alignment varies substantially across occupations and across the Knowledge, Skills, and Abilities that constitute professional competence.

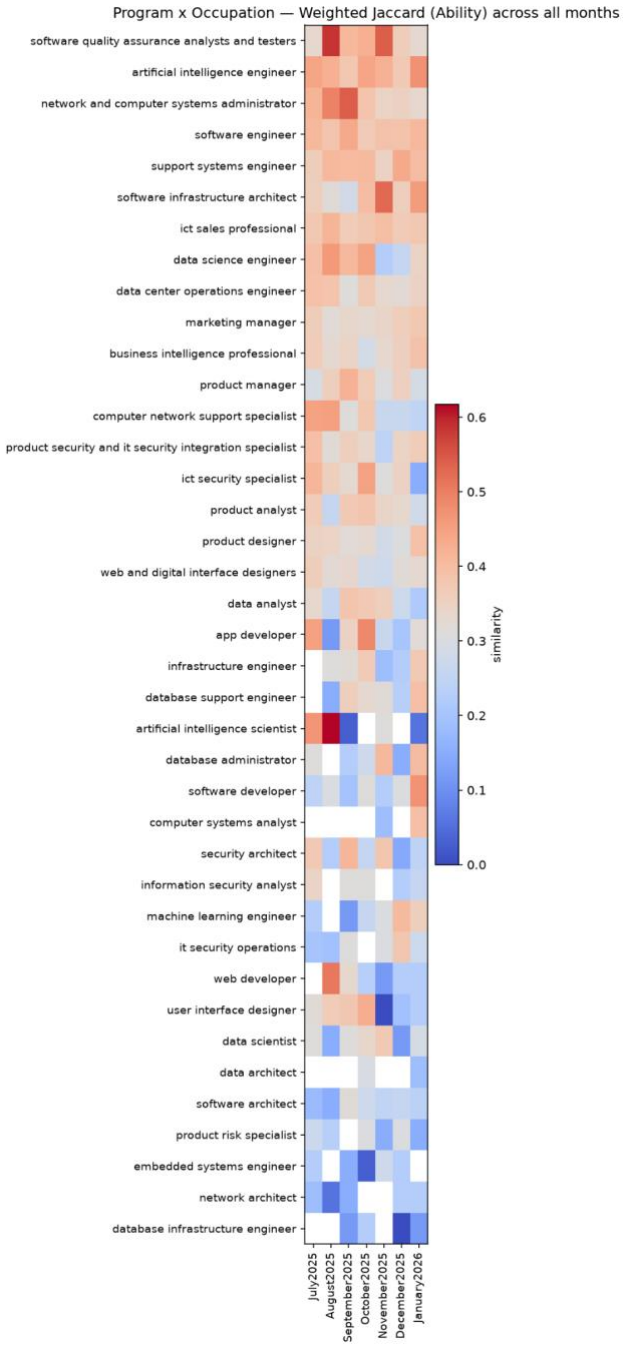


Fig. 3. Weighted Jaccard similarity scores between Job Postings and the ITC program per month for Abilities.

This distinction matters in rapidly evolving sectors such as INFOCOMM, where universities must balance foundational education with increasingly specialized and dynamic occupational requirements. From an educational perspective, the objective should therefore not necessarily be to maximize curriculum similarity with current job postings. Incorporating every industry requirement into a degree program would be neither feasible nor desirable, particularly when some requirements are highly specialized, rapidly changing, or tied to legacy technologies; systematic evidence of alignment and misalignment can help institutions distinguish competencies that belong in the core curriculum from those better addressed through electives, industry collaborations, professionalizing courses, certifications,

micro-credentials, or workplace learning. Beyond the specific Information Technology program examined here, the methodology provides a mechanism to periodically assess how educational offerings align with evolving labor-market demand. Applying the approach longitudinally and across programs, institutions, industries, and geographic contexts could support more evidence-informed curriculum review and reveal whether gaps are persistent, emerging, or temporary. Future work could also incorporate educational evidence beyond formal course descriptions, including learning activities and assessment outcomes, to better capture Skills and Abilities that may develop implicitly. Ultimately, the value of curriculum–industry alignment analysis lies not in eliminating the skill gap, but in making it visible and actionable. Combining labor market intelligence with curriculum evidence can provide universities with a more dynamic basis for deciding what should remain foundational, what should evolve, and where complementary learning pathways are needed to prepare graduates for changing professional environments.

Acknowledgment

This research was funded by the Challenge-Based Research Funding Program, Grant No. FP_IFE_EDU_S_2, of the Institute for the Future of Education, Tecnologico de Monterrey, Mexico.

References

[1] UNESCO, *Transforming higher education: global collaboration on visioning and action*. UNESCO, 2026. doi: 10.54675/SNJW1822.

[2] M. Schmidt, J. K. McDonald, and S. Moore, “The research we don’t need will persist until we dismantle the systems that sustain it,” *J Comput High Educ*, vol. 37, no. 2, pp. 507–542, Jun. 2025, doi: 10.1007/s12528-025-09446-4.

[3] D. Horváth and R. Zs. Szabó, “Driving forces and barriers of Industry 4.0: Do multinational and small and medium-sized companies have equal opportunities?,” *Technological Forecasting and Social Change*, vol. 146, pp. 119–132, Sep. 2019, doi: 10.1016/j.techfore.2019.05.021.

[4] Cedefop, *Terminology of European education and training policy*. Publications Office of the European Union, 2024. Accessed: May 27, 2026. [Online]. Available: https://data.europa.eu/doi/10.2801/991753

[5] P. O. Caratozzolo Martelliti, Y. V. Rueda Castro, M. Gutierrez Aguilar, J. D. Azofeifa Ugalde, L. J. González Gómez, "Discover tomorrow, shape today : shaping skills report", Reporte/report, ITESM, 2024. [Online]. Available: https://hdl.handle.net/11285/703615

[6] Azofeifa, Jose Daniel, Valentina Rueda-Castro, Luis Jose Gonzalez-Gomez, Sonia M. Gómez-Puente, Julieta Noguez, and Patricia Caratozzolo. "Unlocking the future of infocomm workforce: a visual KSA matrix taxonomy approach to education and occupational profiles." In *2024 IEEE Global Engineering Education Conference (EDUCON)*, pp. 01-09. IEEE, 2024. https://doi.org/10.1109/EDUCON60312.2024.10578632

[7] H. G. Ceballos Cancino, S. Butt, G. d. l. Ríos Alatorre, D. P. Madera Espíndola, "Skills report on automotive industry in Mexico : shaping skills framework", Reporte/report, ITESM, 2025. [Online]. Available: https://hdl.handle.net/11285/703348

[8] S. Butt, H. G. Ceballos Cancino, M. Fung, "Shaping skills for the Mexican workforce : Data-based insights into the Mexican automotive industry.", Reporte/report, ITESM, 2025. [Online]. Available: https://hdl.handle.net/11285/704782

[9] European Commission, “European Skills, Competences, Qualifications and Occupations (ESCO),” 2025. Accessed: May 27, 2026. [Online]. Available: https://esco.ec.europa.eu/en/classification/skill_main

[10] INEGI (Mexico), “National Occupational Classification System, SINCO,” 2020. Accessed: May 27, 2026. [Online]. Available: https://en.www.inegi.org.mx/app/biblioteca/ficha.html?upc=702825198411

[11] L. J. Gonzalez-Gomez et al., “Dynamic taxonomy generation for future skills identification using a named entity recognition and relation extraction pipeline,” Frontiers in Artificial Intelligence, vol. 8, July 2025, doi: 10.3389/frai.2025.1579998.

[12] Butt, S., Alatorre, G. D. L. R., González Gómez, L. J., & Ceballos, H. G. (2025). Semi-supervised taxonomy expansion and completion in dynamic taxonomies. Applied Sciences, 15(12), 6517.

[13] M. Klesel and H. F. Wittmann, “Retrieval-Augmented Generation (RAG),” *Bus Inf Syst Eng*, vol. 67, no. 4, pp. 551–561, Aug. 2025, doi: 10.1007/s12599-025-00945-3.

[14] C. S. G. Khoo and J.-C. Na, “Semantic relations in information science,” in *Annual Review of Information Science and Technology*, Information Today, Inc., 2005, pp. 157–228.

[15] P. Langlais and T. L. Gao, “RaTE: a Reproducible automatic Taxonomy Evaluation by Filling the Gap,” *Proceedings of the 15th International Conference on Computational Semantics*, pp. 173–182, 2023.

[16] S. Meisenbacher and P. Norlander, “Building a Custom Taxonomy of AI Skills and Tasks from the Ground Up with Job Postings,” *Proceedings of the Second Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, pp. 117–130, 2026, doi: 10.18653/v1/2026.customnlp4u-1.11.

[17] J. Noble, “What is Jaccard Similarity? | IBM.” Accessed: Aug. 19, 2026. [Online]. Available: https://www.ibm.com/think/topics/jaccard-similarity

[18] V. Garousi, G. Giray, E. Tuzun, C. Catal, and M. Felderer, “Closing the Gap Between Software Engineering Education and Industrial Needs,” *IEEE Software*, vol. 37, no. 2, pp. 68–77, Mar. 2020, doi: 10.1109/MS.2018.2880823.

[19] P. Mhaske, B. Bhattacharjee, N. Haldar, P. Upadhyay, and A. Mandal, “Bridging digital skill gaps in the global workforce: A synthesis and conceptual framework building,” *Research in Globalization*, vol. 11, p. 100311, Dec. 2025, doi: 10.1016/j.resglo.2025.100311.

[20] S. M. Hussain, S. N. Bhatti, and M. F. Ur Rasool, “Legacy system and ways of its evolution,” in *2017 International Conference on Communication Technologies (ComTech)*, Apr. 2017, pp. 56–59. doi: 10.1109/COMTECH.2017.8065750.