JobIQ: Recommending Study Pathways Based on Career Choices

Tomas Trescak[a], Laurence A. F. Park[b] and Mesut Kocyigit
School of Computer Data and Mathematical Sciences, Western Sydney University, NSW, Australia

Keywords: Employability, Skill Frameworks, Curriculum Development, Curriculum Benchmarking.

Abstract: Modern job markets often require an intricate combination of multi-disciplinary skills or specialist and technical knowledge, even for entry-level positions. Such requirements pose increased pressure on higher education graduates entering the job market. This paper presents our JobIQ recommendation system helping prospective students choose educational programs or electives based on their career preferences. While existing recommendation solutions focus on internal institutional data, such as previous student experiences, JobIQ considers external data, recommending educational programs that best cover the knowledge and skills required by selected job roles. To deliver such recommendations, we create and compare skill profiles from job advertisements and educational subjects, aggregating them to skill profiles of job roles and educational programs. Using skill profiles, we build formal models and algorithms for program recommendations. Finally, we suggest other recommendations and benchmarking approaches, helping curriculum developers assess the job readiness of program graduates. The video presenting the JobIQ system is available online[*].

1 INTRODUCTION

"Intelligent" technology, marking the fourth industrial revolution, is disrupting world job markets (Xu et al., 2018). Governments and organisations are trying to analyse the impact of such disruptions, analysing the *employability* profiles of 21st century workers (Daly and Lewis, 2020). A popular way of defining such profiles is by listing the in-demand *skills* for job roles. To define skill profiles, governments and organisations use an ever-increasing number of *skill frameworks*, such as SFIA[1] for science, technology and business skills or a more generic ESCO[2] framework from the European Union or the Australian Skill Framework[3] from the Australian Skills Commission.

However, research and experience show that this approach is flawed. First, the skills listed in these frameworks are not always the ones that are required for a job role, being wildly different depending on the industry, company size, or location (Holmes, 2001). Second, the skill profiles in skill frameworks are incomplete, often specifying only three to five skills for complex roles, listing outdated or missing new technologies or approaches. Last, interviews with employers show that personal values such as honesty and foundational knowledge are more critical during the employee selection process than skills (Manyika et al., 2017). Specifically, technical skills are often rated very low in importance due to their short shelf life (Collet et al., 2015b).

This paper proposes a different approach using JobIQ, our novel analytical and recommendation system. Rather than trying to invent a new generic framework, JobIQ uses "live" job market advertisements to *extract information from job advertisements*, such as the required hard skills (i.e. specialist skills), soft skills (i.e. personal skills), and knowledge (i.e. domain or technical). JobIQ processes job markets daily, obtaining a large dataset of available positions with information on employers, industries or salaries. Traditionally, related research provided no access to such datasets (Gupta et al., 2020). While due to copyright issues, we cannot publish our dataset; we present an approach to building datasets allowing for real-time analysis and projections on various aspects of the job market, such as demand for skills and knowledge. By understanding the requirements of jobs across multiple sectors, we can reach a much finer granularity of skill and knowledge requirements,

[a] https://orcid.org/0000-0002-2540-6002
[b] https://orcid.org/0000-0003-0201-4409
[*] https://www.youtube.com/watch?v=LHTW5P1tNr0
[1] https://sfia-online.org/
[2] https://esco.ec.europa.eu
[3] https://www.nationalskillscommission.gov.au

Trescak, T., Park, L. and Kocyigit, M.
JobIQ: Recommending Study Pathways Based on Career Choices.
DOI: 10.5220/0011754000003470
In *Proceedings of the 15th International Conference on Computer Supported Education (CSEDU 2023) - Volume 1*, pages 137-145
ISBN: 978-989-758-641-5; ISSN: 2184-5026

building localised, industry or employer-specific "employability" profiles. Moreover, we can detect and track emerging skills.

Complementary to "employability" profiles for job roles, JobIQ uses the same skill and knowledge extraction approach to analyse our undergraduate and postgraduate programs, extracting skills and knowledge profiles from our subjects, specialisations and programs. JobIQ compares educational profiles with employability profiles, providing recommendations for students and curriculum developers. Prospective students can explore how various degrees work towards their career goals. JobIQ also helps current students select electives and proactively helps them maintain their development to meet career goals. Last, curriculum developers can benchmark their programs concerning the job roles they support and discover opportunities to introduce new knowledge or challenges with outdated content.

Section 2 of the paper provides background information on our research, further expanding on ideas from the introduction. Section 3 of this paper explains and evaluates our approach to skill extraction from job advertisements and educational content. Section 4 presents formally defines the supporting structures and recommendation algorithms in JobIQ. Last, in Section 5 we discuss our approach and future work.

2 BACKGROUND

Modelling job roles using sets of required skills provides the opportunity to understand the retiring and emerging qualities of 21st-century job markets (Apps, 1988), (Peetz, 2019). Increasingly, we see the notion of *employability*, defining competitive employment profiles (Collet et al., 2015a), (Hinchliffe and Jolly, 2011). Consequently, world governments and organisations analyse and model employability perspectives using skill profiles to predict the growth of individual job roles and impose demands on immigration, educational institutions, or financing schemes (Peetz, 2019).

We can use several generic or specialised skill frameworks to specify job role profiles. For example, the Skills for Information Age framework (SFIA) defines technical and business roles using 147 skills, and each is further decomposed into seven areas of responsibility. The European Union, Australian Public Service, USA National Initiative for Cybersecurity Education (NICE) and other organisations and governments used SFIA to model job role skill profiles.

Similarly, the European Union created the European Skills/Competences, Qualifications and Occupations framework (ESCO), defining over 3000 occupations using 13500 skills/competencies and 11500 qualifications. More recently, in a less gargantuan effort, the Australian National Skills Commission published the second version of the Australian Skills Classification (ASC), which defines 1100 occupations mapped to ANZSCO[4] profiles using 2000 skills.

Following the governmental efforts requiring higher educations institutions to prepare "job-ready" graduates (Daly and Lewis, 2020), inspired by (Herbert et al., 2013), we intended to use skill frameworks (i.e. SFIA, ASC or ESCO) to define graduate profiles for in-demand jobs in local job markets. Subsequently, we planned to redesign and improve our undergraduate programs, matching graduate profiles to job market expectations and hoping to obtain information about the entry-level job roles' skill requirements. Finally, we aimed to validate our changes by estimating how well our redesigned program delivers the skills required by targeted job roles and subsequent careers.

Unfortunately, we were not able to fully appreciate skill frameworks for such a purpose. The SFIA framework defines job roles with very few skills (i.e. usually using three to five skills), serving only a limited understanding of the complex professional requirements. Contrary to SFIA, ESCO and ASC define numerous essential and optional skills, competencies and knowledge. Yet, we found the definitions somewhat flawed, strangely specific or lacking.

For example, the *Software Developer* occupation[5] lists "create flowchart diagram" or "perform scientific research" as essential skills, yet "use object-oriented programming" or "develop creative ideas" only as optional. We would expect this to be the other way around. Moreover, optional knowledge lists many historical or scientific languages such as Erlang, Cobol, Haskell or Smalltalk, but lacks modern in-demand languages such as GoLang, Kotlin, Rust or Julia. We understand that it is difficult to list all the currently used technologies as that list would be very long. In this case, ASC has a much better approach; instead of listing individual programming languages, it defines and maps a single technological skill, *Software development and programming languages*, albeit losing the opportunity to specify the exact required technologies.

Moreover, in stark contrast to governmental efforts that increasingly use and depend on skill frame-

[4]https://www.abs.gov.au/statistics/classificationsanzsco-australian-and-new-zealand-standard-classification-occupations/latest-release

[5]http://data.europa.eu/esco/occupation/f2b15a0e-e65a-438a-affb-29b9d50b77d1

works, (Manyika et al., 2017) surveyed numerous CEOs and members of senior and middle management, as well as reviewed multiple works dealing with employment and "employability", discovering that such a *simplistic view* is not possible. The skill requirements for the same role vary across industries or the use of underlying skill frameworks. Employers often rated technical or digital skills towards the bottom, primarily due to their "short shelf life". Similarly, (Holmes, 2001) and (Collet et al., 2015b) expose the flaws in trying to define the optimal graduate skill sets (i.e. graduate profiles) due to a large pool of diverse graduates trying to match the requirements of individual employers from varied backgrounds with unique *values*.

Additionally, (Holmes, 2001) argues that the concept of graduate employability cannot be defined by the acquisition of *measurable skills* due to the unmeasurable complexity of personal qualities. Instead of skills, (Manyika et al., 2017), (Collet et al., 2015b) or (Hinchliffe and Jolly, 2011) note the importance of *identity*, *values* and *abilities* as a driving force of graduate employment. For example, it is much more important to be trustworthy and reliable with good communication abilities than to have specific technical skills.

Such flaws of the skill-driven approach are further amplified by using context-agnostic skill frameworks designed for organisations of any size, sector or industry. Role profiles are rarely updated and become "stale" ' or obsolete. However, governments aim to help drive policy, recruitment and training by providing a data-driven classification of skills using such skill frameworks. Considering the previously mentioned research stressing the different, ever-changing needs of individual businesses and personalised preferences, such efforts have a questionable effect and informative value for employers and potential employees (Manyika et al., 2017) (Council et al., 2012).

Consequently, in this work, we present means of building datasets that allow for the automated extraction of skills and knowledge from job advertisements. Using this data, we generate highly granular skill and knowledge profiles, understanding nuances in requirements across different industries, regions or employers. Consequently, using the same extraction approach, we analyse skill and knowledge acquisition in educational activities, providing recommendations to students about subjects or courses that best cover the requirements of particular job roles.

3 SKILL EXTRACTION

Students attend university and select specific subjects to learn skills, knowledge, technologies and abilities, developing their intelligence and increasing their employment chances (for simplicity, in the rest of this paper, we write only "skills" instead of "skills, knowledge, technologies and abilities"). Employers seek out people with skills to take on specific roles. Identifying the skills required for a job or gained from taking a subject allows us to optimise the coverage of job roles available after completing a degree. Therefore, skill extraction is essential to a subject/role recommender system.

Our approach relies on understanding the skill requirements of individual jobs based on job advertisements' descriptions. Then, we can build role skill profiles for each job role, aggregating the skills from related job ads. Optimally, job advertisements would use a robust skill framework and explicitly list the required skills, knowledge and technologies. Educational institutions would use the same skill framework to specify their teaching skills and compare the demand and supply. Unfortunately, this is not the case.

As a result, we depend on skill extraction from the description of job advertisements and related metadata. Since our institution does not explicitly specify the skills covered in educational content, we also used skill extraction to create the list of skills for subjects. Please note that we only used the automated approach to prepare recommendations for subject coordinators, who further redacted and curated the list.

To automatically extract skills from text, (Kivimäki et al., 2013) used a graph-based approach, mapping the text to Wikipedia articles which map further to LinkedIn[6] skills. Unfortunately, we did not find a way how to use their method with our target ACS skill framework or any other framework. Moreover, the author's approach worked well for extracting skills from scientific articles with matching inputs in Wikipedia, less so with advertised job ads. More recently, the team at LinkedIn (Gupta et al., 2020) presented their (Bhola et al., 2020) system, which uses salience and a market-aware skill extraction system. Skills extracted by this system are not only those found in the description of the job advertisement but also those generally required by related job-role.

Moreover, the system also filters out the skills for which there is a supply on the market. While this system would possibly be a good match for our purposes, the underlying data structures used for training are not available outside LinkedIn. Unfortunately, this is the

[6]https://linkedin.com

trend of most of the related works in the area, with only a minimal number of works releasing their data (Bhola et al., 2020) (Zhang et al., 2022) and none, apart from (Zhang et al., 2022), releasing their annotation guidelines.

Other approaches (Smith, 2021) or (Zhang et al., 2022) use natural language processing to automatically build the skill database with high accuracy in detecting skills and knowledge. While this approach is interesting by having the possibility to see emerging skills, it does not fit our purpose and serves only complementary functionality to do so. The main reason is that we need to explain to our users what it means to have a specific skill and how to obtain it. Using automatically extracted emerging skills makes explanations very difficult while using established skill frameworks, we can provide them.

As a result, our requirement is for a hybrid system that detects *skills from a selected (interchangeable) skill framework*. Using simple text matching is impossible as most skill frameworks skills are defined using multiple words, which can be expressed in various ways (for example, ASC skill "Provide technical support for computer network issues"). As a result, we devised a skill-matching strategy based on matching sentence embeddings.

3.1 Skill Extraction Using Sentence Embeddings

The goal of skill extraction is to determine whether a given text contains a notion of skill from a specific skill framework. In our case, we are trying to *extract skills* from job and subject descriptions. Note that descriptions are written in natural language, often specifying activities related to applying a skill rather than providing the name of the skill. This prohibits us from quickly identifying the skills. Initial experiments using keyword extraction provided poor results. Therefore, we needed to understand the meaning of sentences, not the words themselves. The recent advancement in natural language processing using word embeddings (Jatnika et al., 2019) and deep transformer networks (Kenton and Toutanova, 2019) provided us with vector representations of sentences that capture the meaning of the sentences rather than just the words. By converting the sentences from a job description and a skill description into a set of vectors, we can identify how similar each skill is to a particular role. The vectors are created so that the similarity of two sentences is computed using the cosine of the angle between the two vectors. We use the sentence embedding approach to extract skills belonging to a specific skill framework from arbitrary descriptions (e.g. job or subject description). Formally, we define skills framework as:

Definition 1. A *skill framework* S is a set of skills or other discrete competencies, where each skill is represented in natural language.

Definition 2. The *skill-matching function* $M(t,d) \rightarrow s$, given a text t and a description d of the skill, returns a *Skill Match Strength* s, represented by a value from interval $[0,1]$ with 0 representing no match, values between 0 and 1 a partial match and 1 representing a complete match.

In our case, the skill-matching function M uses the *sentence embedding approach*. It takes the text of the job advertisement and, for a skill represented by its title and description, estimates whether that skill is mentioned (i.e. embedded) in the job description. This skill is more probably mentioned when the value of *skill match strength* is closer to 1.

Definition 3. The *skill extraction function* $E(S,M,s,t) \rightarrow S'$ given:

- S is a *Skill Framework*
- M is a *Skill-Matching Function*
- s is a minimum value of *Skill Match Strength*
- t provided text t (e.g. job-description)

extracts a subset of skills $S' \subseteq S$ from a provided text t (e.g. job-description) whose skill match strength is bigger or equal to provided s.

In our case, we take a job or subject description, and the Skill Extraction Function extracts all skills from the ASC skill framework most probably contained in the job description. Through experimental evaluation, we discovered the value 0.45 for s being the optimal skill match strength for extraction using our approach.

3.2 Evaluating Skill Extraction

The Australian Skill Classification (ASC) dataset defines over 2000 skills assigned to every role in the ANZSCO dataset. For each ANZSCO role, ASC provides a list of skills, core competencies (e.g. numeracy, literacy) and used technologies.

For example, a *Web Developer* contains *32 skill*s, such as

- Design websites or applications
- Update website content or
- Test software performance

Furthermore, the Web Designer role uses *12 technologies* such as

- Software development and programming languages,
- Graphics or photo imaging software, and
- Social media and web publishing software.

We verify the validity of our approach by analysing how well our algorithm extracts skills from job advertisements related to roles specified in the ASC dataset. We also verify whether our approach can discover new, feasible skills not mentioned in the ASC dataset.

Our dataset has 41,273 job advertisements downloaded from Australian job websites. First, we classified each job advertisement with one of the ANZSCO roles, discovering over 700 ANZSCO roles. Second, from job descriptions, we extracted the skills and used technologies. Figure 1 summarises the results of skill extraction, analysing the coverage of ASC skills. We see that we extracted 100% of the skills from ASC in almost 8% of job roles. Overall, for 73% of job roles, we extracted at least 50% of the skills defined in the ASC dataset.

Furthermore, we can see a correlation between the number of job advertisements and the accuracy of our extraction. The curve is not monotonic, declining to the group of job advertisements with 100% skill coverage. The reason is that in this group, there are advertisements for expert jobs in medicine, which provide exhaustive descriptions of activities, allowing us to extract all of the ASC skills. Job advertisements in other areas were less specific. We found no correlation with the average length of the description of the job advertisement.

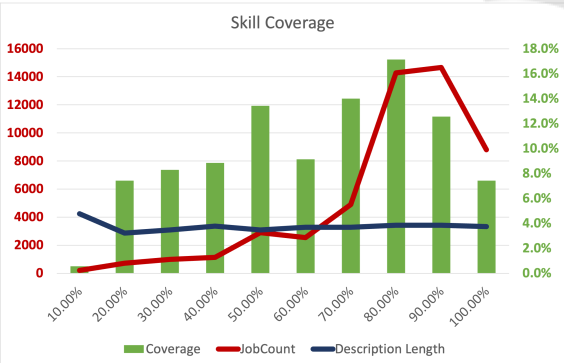

Figure 1: Completion Criteria of Mathematics Major.

The roles with high coverage of skills contained mainly medical roles such as "Resident Medical Officer" or "Registered Nurse (Aged Care)" with a detailed description of responsibilities. On the other hand, low-coverage roles, such as the "Web Developer" role, often described only the company culture and left responsibilities as assumed. For example, some of the skills from the ASC framework not detected in the "Web Developer" job ads were:

- Troubleshoot issues with computer applications or systems
- Develop diagrams or flow charts of system operation
- Prepare graphics or other visual representations of information

Such skills are usually assumed and only rarely specified in job advertisements. Consequently, in JobIQ, we use the ASC skills as a baseline skill-set for every role, extending them with detected skills. But, the JobIQ approach becomes helpful when detecting ASC skills not covered by the ASC framework role profiles. JobIQ allows for a high granularity of analysis, building personalised skill profiles by locality, employer or industry.

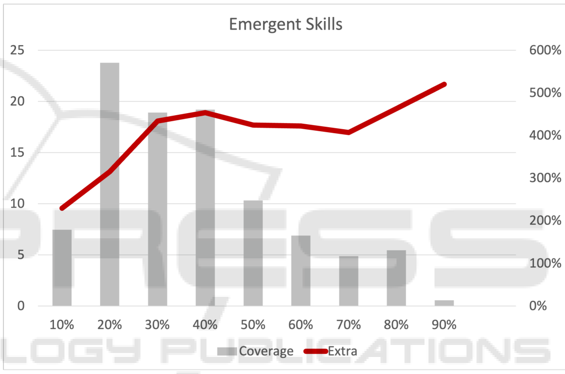

Figure 2: The emergence of skill framework roles in skill profiles.

Figure 2 depicts the number of emergent skills. This time we only considered skills that appear in at least 10% of job advertisements. We see that the ASC coverage dropped significantly, to an average of 10.8%, but detecting three times more skills than in the ASC dataset on average. For example, the system has detected 13 emerging skills for the "Software and Applications Programmers" contained in at least 10% of the job advertisements, including:

- Support individuals with diverse needs to understand, access and utilise information or services
- Deliver culturally appropriate programs, policies or services
- Maintain a working understanding of the cultural, diversity and accessibility needs of others and how this applies to the role
- Evaluate projects to determine compliance with technical specifications
- Gather information to provide services to clients

- Provide technical support for software maintenance or use

Furthermore, we used our approach to *extract ASC skills from subjects being offered at our institution*. The extraction coverage was very similar to the experimental data. We also confirmed that the skill detection failed when the subject description lacked information about related activities or outcomes. Updating the descriptions to include such information improved our algorithm's accuracy and coverage. This proved a good strategy for analysing our subject catalogue, allowing us to consult subject coordinators about possibly lacking descriptions of their subjects and providing recommendations for improvement. These improvements improve student understanding of the skills they develop during the offered subject or program.

4 RECOMMENDING STUDY PROGRAMS TO PROSPECTIVE STUDENTS

In the previous sections, we presented our approach to extracting skills from job advertisements and educational subjects. In this section, we demonstrate how we can build a recommender system to inform prospective students *who study programs deliver skills, knowledge and technology most related to their career choices*. Our approach is novel as existing methods focus on recommendations based only on student data, not considering external sources. For example, (Chaturapruek et al., 2018) provides recommendations based on historic student preferences, or (Farzan and Brusilovsky, 2006) considers student ratings. More recently, researchers employed neural networks to recommend courses and subject sequences based on enrolment data (Pardos et al., 2019) or grade prediction (Ren et al., 2019). To our knowledge, none of the existing systems considers the job market preferences.

We need a Skill Framework and a set of job roles with their skill profiles to build such a recommender system. We also need a dataset of job advertisements, where each ad specifies the description and the job role required.

Similarly to a skill profile of a job role, we can use the *skill extraction function* to extract all the skills from a job advertisement description to build a *job advertisement skill profile.*

Definition 4. Considering a *skill framework S* and a *job advertisement a*, a *job advertisement skill profile* $S' \subseteq S$ is a subset of skills from S contained within the job advertisement a.

Such a profile contains the set of skills that are found relevant only to the given advertisement. We can then filter and aggregate all the *job advertisement skill profiles* (e.g. by industry, employer) to build a highly granular *job role profile* that considers the proportion of job advertisements requiring each skill. Formally we define:

Definition 5. Considering a *skill framework S* with n roles, a *job role* and a set of *job advertisement skill profiles*, we define a *job role profile* as a vector $(c_1, c_2 \dots c_n)$ where each vector element c_k where $k \leq n$ represents the ratio of job ads containing the skills s_k and a total number of advertisements for this role.

In other words, for each skill from the skill framework, the *job role skill profile* defines how many job advertisements require this skill in proportion to all the available job advertisements for that given role. This approach allows us to model *demand* for particular skills.

For example, consider a skill framework with three skills, two job roles r_1, r_2, two job advertisements for role r_1 and three for role r_2 with job advertisement skill profiles from Table 1. Then, the *job role profiles* would be those found in Table 2.

Table 1: Job advertisement skill profiles for roles r_1 and r_2.

	Ad #	s_1	s_2	s_3
Role r_1	1	0.6	0.0	0.2
	2	0.4	0.0	0.0
Role r_2	1	0.0	0.7	0.0
	2	0.0	0.8	1.0
	3	0.0	0.2	0.4

Table 2: Job role skill profiles for roles r_1 and r_2 with *skill match strength* threshold 0.

	s_1	s_2	s_3
Role r_1	1.0	0.0	0.5
Role r_2	0.0	1.0	0.66

We can further specify that we only consider skills if their skill match strength is higher than a threshold value. Table 3 depicts how the role skill profile changes when considering a threshold value. For example, skill s_2 drops the value from 1 to 0.66 as only two advertisements match this skill with strength 0.45 or above.

With job role profiles, we can assess which skills are desired for a given job role. What we aim to do,

Table 3: Job role skill profiles for roles r_1 and r_2 with *skill match strength* threshold 0.45.

	s_1	s_2	s_3
Role r_1	0.5	0.0	0.5
Role r_2	0.0	0.66	0.33

is to recommend the educational program which will deliver those skills. To achieve this goal, similar to the *job advertisement skill profile*, we define the *subject skill profile* and a *pathway skill profile*:

Definition 6. Considering a *skill framework* with n roles, we define a *subject skill profile* as a vector $(c_1, c_2 \dots c_n)$ where each vector element c_k where $k \leq n$ represents either the skill match strength of skill s_k or zero, if that skill match strength is below the threshold value.

In other words, the *subject skill profile* provides which skills are delivered in a subject, along with a probability (i.e. strength) of their detection in a description of the subject.

Definition 7. Considering a *skill framework S* with n roles and an educational pathway with l subjects, we define *pathway skill profile* as a vector $(d_1, d_2 \dots d_n)$ where each vector element d_k where $k \leq n$ represents a *maximum* value of a c_k from all of the l subject skill profiles.

In other words, we build a *pathway skill profile* by combining all the skills delivered in all the subjects in the pathways and remembering only the maximum match strength value for each detected skill. By doing so, we understand what skills are most probably delivered in the given pathway of an educational program. For example, Table 4 depicts the pathway with only two subjects and their related subject skill profiles:

Table 4: Example pathway with two *subject skill profiles*.

	s_1	s_2	s_3
Subject_1	0.7	0.0	0.5
Subject_2	0.0	0.6	0.33

Taking the maximum value from each subject skill profile, the resulting pathway skill profile would be:

$$(0.7, 0.6, 0.5)$$

We are almost ready to recommend educational programs for a defined job role. First, we generate all possible pathways in an educational program. We can only consider representative pathways through the program if there are too many combinations. Second, we assess how well a particular pathway matches the skill requirements of a given role. We can do this by multiplying the values of a *role skill profile* with a pathway *skill profile* and then summing up all the elements.

$$\sum(c_1, c_2 \dots c_n) \times (d_1, d_2 \dots d_n)$$

Please remember that the *role skill profile* defines how important the skill s_k is for a given role, and the *pathway skill profile* defines how strongly skill s_k is covered in the pathway. We obtain a weighted value of skill importance by multiplying these two vectors. For example, if a skill is not essential, i.e. its role profile value is 0, and no matter how well this skill is covered in a pathway, it will not affect the final value.

Third, we process all the pathways and remember only the value of the *best matching pathway* as a program representative. With this approach, for a given job role, we can order all the programs based on the value of the *best matching pathway* and recommend the programs that deliver most of the desired skills for a given job.

The JobIQ recommender system pre-computes the recommendation data that is then delivered in real-time through the web interface depicted in Figure 3. We evaluated the performance of the recommendation by assessing which programs are recommended for given job roles. The recommendations for more technical programs performed as expected, recommending our IT programs for ICT roles and business programs (accounting) for business (accounting roles). More interesting recommendations appeared when assessing other roles, such as school teachers ranking our arts or criminology programs high.

4.1 Other Recommendations

Following a similar approach, we were able to deliver other recommendations that our students desired.

- The recommendation of *electives* helps current students find electives that best match their career choices. Through careful selection of often interdisciplinary electives, students can maximise their skill and knowledge uptake. The recommendation uses the same approach as program recommendations, using *subject skill profiles* instead of *pathway skill profile*.
- The recommendation of *careers* and *individual jobs* based on educational programs is an inverse recommendation strategy to program recommendation. This strategy is popular with current students who look for job opportunities towards the end of their program.
- *Benchmarking* and *comparison of programs* allow prospective students to compare the different options that programs provide concerning opportunities in current job markets.

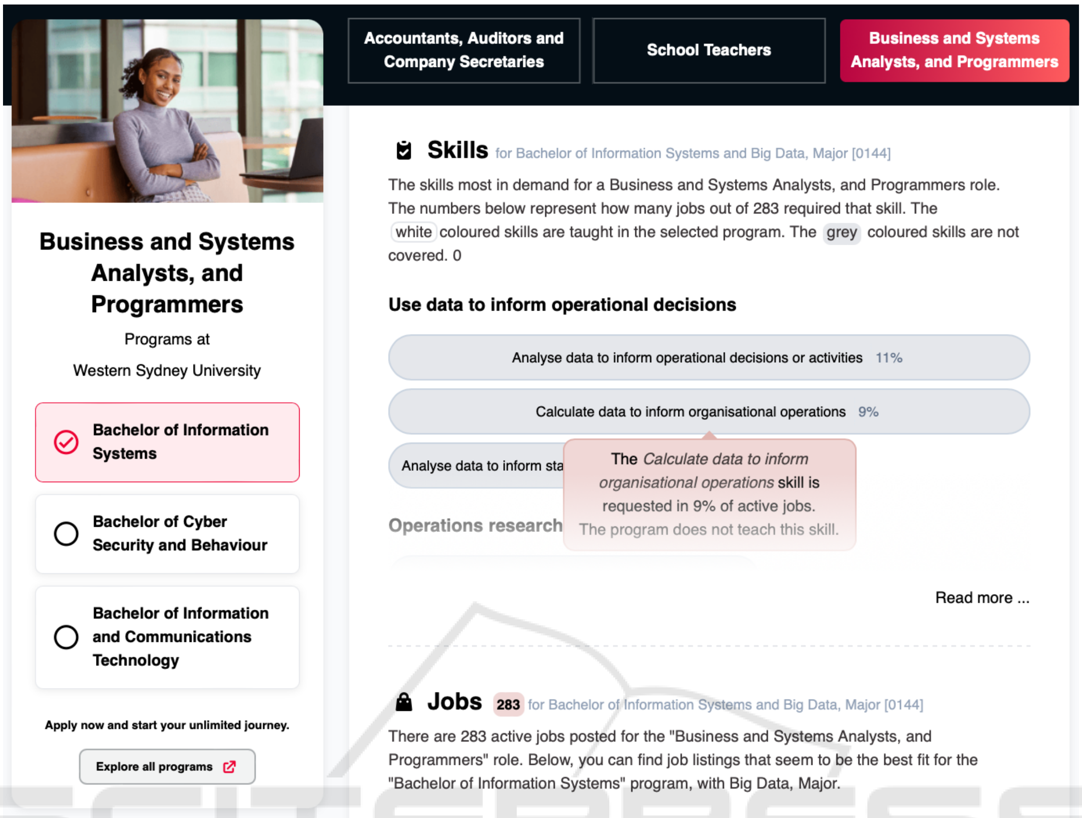

Figure 3: JobIQ system interface.

Considering that our system processes daily opportunities on job markets, we built a proactive job, and career-monitoring system that overlooks student performance, provides updates on skill development related to current opportunities and finds alternative pathways in case a student needs more practice or fails a subject.

5 CONCLUSIONS AND FUTURE WORK

Our approach delivers a novel strategy for recommending educational programs or electives based on career choices. For this purpose, we designed a skill extraction system that can compare the skills coverage based on supply (i.e. education) and demand (i.e. job markets). Overall, the accuracy of the automated extraction depends on the quality of the description. We can build highly accurate profiles only with a sufficient number of jobs in the dataset. As a side effect, our approach is helping curriculum designers to write better subject descriptions mentioning the skills and capabilities covered during the subject delivery, providing more information to students. We developed a recommender system that allows review and assigns skills to subjects, helping to order and assign correct skills to knowledge subjects. We will evaluate this approach and further optimise our skill extraction strategy as part of our future work.

We are currently working on enabling JobIQ to help with life-long learning. Since the system understands its users' capabilities, skills and knowledge, it can monitor for career opportunities and up-skilling. For example, it can provide notifications of opportunities to take a specific course to achieve the skill set necessary for a new, possibly more exciting career.

REFERENCES

Apps, J. W. (1988). *Higher Education in a Learning Society. Meeting New Demands for Education and Training.* ERIC.

Bhola, A., Halder, K., Prasad, A., and Kan, M.-Y. (2020). Retrieving skills from job descriptions: A language model based extreme multi-label classification framework. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5832–5842.

Chaturapruek, S., Dee, T. S., Johari, R., Kizilcec, R. F., and

Stevens, M. L. (2018). How a data-driven course planning tool affects college students' gpa: evidence from two field experiments. In *Proceedings of the fifth annual ACM conference on learning at scale*, pages 1–10.

Collet, C., Hine, D., and Du Plessis, K. (2015a). Employability skills: perspectives from a knowledge-intensive industry. *Education+ Training*.

Collet, C., Hine, D., and Plessis, K. d. (2015b). Employability skills: perspectives from a knowledge-intensive industry. *Education + Training*, 57(5):532–559.

Council, N. R. et al. (2012). *Education for life and work: Developing transferable knowledge and skills in the 21st century*. National Academies Press.

Daly, A. and Lewis, P. (2020). The proposed job-ready graduate package: A misguided arrow missing its target. *Australian Journal of Labour Economics*, 23(2):231–251.

Farzan, R. and Brusilovsky, P. (2006). Social navigation support in a course recommendation system. In *International conference on adaptive hypermedia and adaptive web-based systems*, pages 91–100. Springer.

Gupta, R., Liu, Y., Shah, M., Rajan, S., Tang, J., Prakash, B. A., Shi, B., Yang, J., Guo, F., and He, Q. (2020). Salience and Market-aware Skill Extraction for Job Targeting. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2871–2879.

Herbert, N., Lewis, I., and De Salas, K. (2013). Career outcomes and sfia as tools to design ict curriculum. In *ACIS 2013: Information systems: Transforming the Future: Proceedings of the 24th Australasian Conference on Information Systems*, pages 1–11. RMIT University.

Hinchliffe, G. W. and Jolly, A. (2011). Graduate identity and employability. *British Educational Research Journal*, 37(4):563–584.

Holmes, L. (2001). Reconsidering Graduate Employability: The 'graduate identity' approach. *Quality in Higher Education*, 7(2):111–119.

Jatnika, D., Bijaksana, M. A., and Suryani, A. A. (2019). Word2vec model analysis for semantic similarities in english words. *Procedia Computer Science*, 157:160–167.

Kenton, J. D. M.-W. C. and Toutanova, L. K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Kivimäki, I., Panchenko, A., Dessy, A., Verdegem, D., Francq, P., Bersini, H., and Saerens, M. (2013). A graph-based approach to skill extraction from text. In *Proceedings of TextGraphs-8 graph-based methods for natural language processing*, pages 79–87.

Manyika, J., Lund, S., Chui, M., Bughin, J., Woetzel, J., Batra, P., Ko, R., and Sanghvi, S. (2017). Jobs lost, jobs gained: Workforce transitions in a time of automation. *McKinsey Global Institute*, 150.

Pardos, Z. A., Fan, Z., and Jiang, W. (2019). Connectionist recommendation in the wild: on the utility and scrutability of neural networks for personalized course guidance. *User modeling and user-adapted interaction*, 29(2):487–525.

Peetz, D. (2019). *The realities and futures of work*. ANU Press.

Ren, Z., Ning, X., Lan, A. S., and Rangwala, H. (2019). Grade prediction based on cumulative knowledge and co-taken courses. *International Educational Data Mining Society*.

Smith, E. W. (2021). Skill Extraction for Domain-Specific Text Retrieval in a Job-Matching Platform. pages 116–128.

Xu, M., David, J. M., Kim, S. H., et al. (2018). The fourth industrial revolution: Opportunities and challenges. *International journal of financial research*, 9(2):90–95.

Zhang, M., Jensen, K., Sonniks, S., and Plank, B. (2022). SkillSpan: Hard and Soft Skill Extraction from English Job Postings. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4962–4984.