

COMPELLING EDUCATIONAL OFFERINGS: A STUDY ON THE EFFICACY OF SKILLS IDENTIFICATION PLATFORMS WITH COURSE DESCRIPTIONS

J.W. Lai, L. Zhang, Y.S. Chan, C.C. Sze, F.S. Lim

Nanyang Technological University (SINGAPORE)

Abstract

Course descriptions are often the first point of contact between students and the courses they are considering. These descriptions are crucial in communicating the benefits of the course and offer insights into how a particular course can contribute to a student's academic and career aspirations. By detailing what skills and knowledge the course will impart, universities can align their offerings with the student's future goals, making their courses more appealing. In this evaluative study, we present our analysis of the skills extracted by three skills tagging platforms on fifty course descriptions across four disciplines from a university in Singapore. The analysis would focus on (1) casting a spotlight on available skills tagging platforms, (2) outlining a human-in-the-loop form of skills validation, (3) reporting the efficacy of various platforms in performing skills tagging, and (4) forecasting avenues for academic and learning analytics research with skills representation in the course syllabus. Our evaluation revealed a synergistic effect with the three platforms in enhancing the overall skill tagging efficacy and in achieving a balance between precision and recall.

Keywords: Higher education, skills extraction, learning analytics, natural language processing application

1 INTRODUCTION

Course descriptions are often the first point of contact between students and the courses they are considering. Research has shown that traditional course descriptions prioritize informative purposes over promotional aspects, while online descriptions have more promotional elements while still delivering information to prospective students [1]. These descriptions are crucial in communicating the benefits of the courses and offer insights into how a particular course can contribute to a student's academic and career aspirations [2]. By detailing what skills and knowledge the course will impart, universities can align their offerings with the student's future goals, making their courses more appealing and improving enrolment rates [3]. They are thus essential touchpoints in this decision-making process, an encapsulated narrative bridging the gap between educational offerings and students' career aspirations.

Further, we are at a stage of development in artificial intelligence (AI), where it is possible to harness its capabilities to close the gap at scale. The critical question is: How can universities strategically leverage AI to analyse and optimize their course content, ensuring it is in-sync with the dynamic demands of the industry? This involves tagging courses with relevant skills and identifying emerging trends and gaps in their current curriculum. By using AI to perform skills tagging and augmenting this with the workforce data, universities can gain valuable insights into how well their offerings align with the evolving needs of the job market. This predictive analysis can inform universities about the skills in high demand, allowing them to tailor their courses to equip students with future-ready competencies proactively. Simultaneously, it can support students in understanding and more effectively market their skill sets to future employers.

The task of crafting enticing course descriptions while clearly representing the skills to be acquired is essential for educational institutions worldwide today. As industries rapidly change, courses and course descriptions must be updated to reflect the skills in demand and provide up-to-date and market-relevant education [4], [5]. Skill representation in the form of a set of tags or keywords (as additions to the course description) is a salient way to communicate the relevance of courses for the evolving demands of industries and the job market. Good skills representation in course descriptions helps bridge the gap between the skills employers seek and the skills that job seekers possess. When courses accurately convey the skills students would acquire, each student can make informed decisions to personalise their learning paths with the university [6], [7]. In turn, this could lead to better employability outcomes as there is a better match between the demand and supply of skills. When the skills for various job roles

and levels are accurately captured in course descriptions, it could also further support individuals in taking relevant courses for career advancement, upskilling, and lifelong learning. In the grand scheme, skills representation in course descriptions could enable a workforce that adapts to evolving job roles and support job mobility [8].

As educational landscapes and the industry demand in specialized skills continues to grow, the communicative role of course descriptions and their associated skill representation would only become more important. In the preceding section, we introduced the pivotal role that course descriptions and effective skill representation can play in attracting potential students, informing their decisions, and aligning educational offerings with industry needs. In Section 2, we delve deeper into the existing body of research that underscores the critical interplay between these elements. By examining studies in course marketing strategies, skill representation techniques, and the application of machine learning in education, we aim to unravel the complex dynamics that shape the effectiveness of course descriptions and set the stage for our exploration into the efficacy of three existing machine learning algorithms in course skills representation with course descriptions. Sections 3 and 4 will cover our approach, evaluation, and perspective to answering the research question. By answering the research questions, we aim to (1) cast a spotlight on the purpose of skills extraction, (2) outline a human-in-the-loop form of validating the extracted skills, (3) provide an analysis of the efficacy of various platforms in performing skills tagging, and (4) forecast avenues for learning analytics research with skills representation in the course syllabus. Finally, we conclude in Section 5.

2 LITERATURE REVIEW

Course descriptions are critical information tools for educational institutions and potential students. Research has shown that well-crafted course descriptions can significantly impact student enrolment decisions by providing clear and informative insights into the course content, learning outcomes, and practical skills students will acquire. Mourey et al. [3] investigated the direct impact of course descriptions on student enrolment decisions. They found that not only does course description affect enrolment intention, but it also affects subjective interpretations of course interest, especially for lower-level courses. However, one should note that this effect is less prevalent for students taking higher-level courses, where the effects of learning outcomes and workload play a greater factor. In the same study, Mourey et al. found that the course title has few meaningful effects. Therefore, it suggests that course descriptions have more marketable value for prospective students than course titles, motivating our intent to study how course descriptions are crafted.

The graduate employment rate is often used to assess the quality of university provision, suggesting an intricate link and expectation that institutes of higher learning should prepare their students through skills-based or career-ready education [9], [10], [11], [12]. Such a model of education aims to equip students with the tools they need to succeed in their chosen careers and make meaningful contributions to the job market [13], [14]. While educational institutions aim to impart skills, herein lies the question: do students know what skills they have learned? Consequently, this affects students' knowledge of their skill gaps. There is a gap in the literature regarding answers to this question. Therefore, to answer this question, we first need to perform skills tagging of courses to perform analyses.

As skills tagging and its associated research in learning analytics are developing [15], [16], we are at a crossroads where the technology is ready and would require evaluation. Furthermore, most institutions have a data pool of course descriptions but not necessarily the associated data on skills tagging. A viable solution to ensuring that all courses are tagged with their associated skills is to request that the course instructors submit a list of skills for the respective courses they are teaching. However, this exercise requires substantial man-hours and grows with the number of courses the institution provides. The alternative, which can reduce the man-hours required to perform manual skills tagging of existing course descriptions, is machine learning. Modern advances in machine learning, particularly tagging algorithms [17], [18], [19], [20], provide the ability to extract and suggest skills based on a body of text input, such as the course description. Institutional and commercial options for skills tagging are available to varying degrees of usage and context. In this research, we will utilize three skills extraction platforms: JobTech Talent Skills Ontology [21], Lightcast Open Skills Taxonomy [22], and SkillsFuture Singapore Skills Extraction Algorithm [23]. These three platforms are chosen because they represent private limited access, private open access, and public open access options respectively. Due to the proprietary nature of the algorithm driving the skills tagging platforms, the names of the extraction platforms are stated here for transparency of the research and subsequently referred to as Platforms A, B, and C, but not in the order mentioned.

The key research question is the efficacy of these machine learning algorithm in identifying the skills related to the courses. Specifically, we would like to find out 1) how similar are these machine learning approaches in identifying skills, and 2) how well do they align with the subject matter experts' perception of the skills that are developed in the courses?

3 METHODOLOGY

This section outlines the methodological approach undertaken to analyze the dataset of course descriptions and extract skills for further analysis. The research methodology encompasses data collection, cleaning, category selection, random sampling, and skills extraction. A dataset comprising 3468 course descriptions was compiled from a dataset made available to the research team by Nanyang Technological University, Singapore. The dataset represents a comprehensive range of academic disciplines and undergraduate course offerings. Initial data cleaning procedures were executed to eliminate incomplete, duplicate, or legacy entries. We also removed complete entries that were not in the English language, resulting in a refined dataset of $N = 2762$ course descriptions from four categories corresponding to disciplines: Business ($N_B = 262$ course descriptions), Engineering ($N_E = 823$), Humanities, Arts and Social Sciences ($N_H = 1396$), and Science ($N_S = 281$).

It was a monumental challenge to include all subject matter experts. Instead, we opted for a pilot study. A random sampling strategy was employed to ensure a representative analysis of skill extraction across four disciplines. A sample of 25 courses was randomly selected from each category, resulting in 100 course descriptions forming the sample for this study. The 100 randomly selected course descriptions underwent a skills extraction process subjected to the three distinct skills extractor algorithms. These algorithms employ natural language processing techniques to identify and extract skills from the text of the course descriptions.

Following the skills extraction process, an analysis was conducted to assess the efficacy of the skills extractor algorithms in identifying skills within the course descriptions. We performed a two-step data-cleaning process to reduce the output size. First, for each of the 100 randomly selected course descriptions, an exhaustive list of skills generated by all three skills extractor algorithms was compiled. This list collated the skills identified by each algorithm within the context of the respective course description. Second, natural language processing was performed to reduce degeneracy. An example of how a course's tagged skills are compiled is illustrated in Fig. 1.

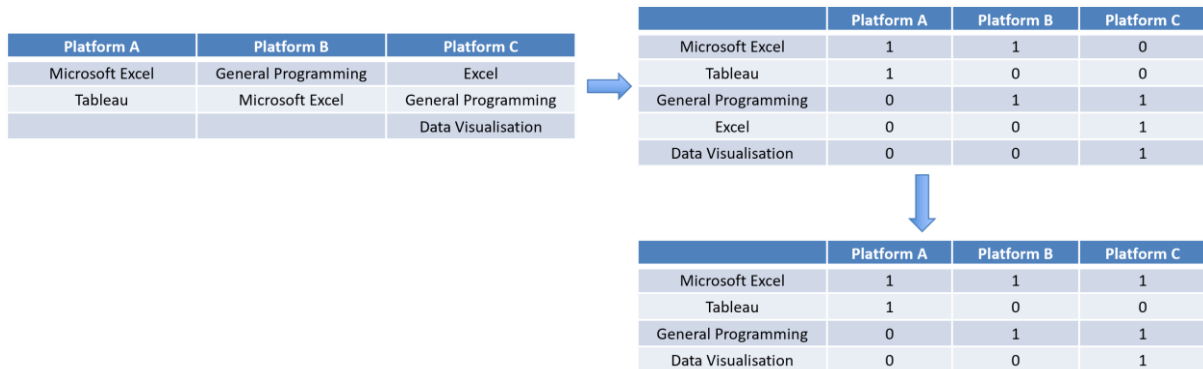


Figure 1. An example of how the skills tagged by three platforms is compiled. The first step is to translate the skills into a table. Skills tagged by the platform are indicated with "1". The second step is to reduce degeneracies. In this example, "Microsoft Excel" and "Excel" refer to the same skill and, therefore, are combined into a single entry.

Next, a human-in-the-loop validation process was initiated to ensure the relevance of the extracted skills. Collaboration was sought from subject matter experts such as the course and program coordinators of the sample courses to validate the efficacy of the extracted skills. These domain subject matter experts are persons who had crafted and approved the course descriptions. Each coordinator was provided with the compiled list of extracted skills associated with their courses (the first column of the final output of Fig. 1). The evaluation of the list was conducted with reference to the actual course description, assessing how the skills corresponded with the curriculum. For this validation process, an in-house platform was developed (refer to Fig. 2). Within this platform, coordinators had access to the data

involved in tagging skills, including the course aims, the learning outcomes, and the content. The evaluation process comprised the following stages:

1. Inclusions: Coordinators verified the presence of each skill in the course curriculum. They checked the skills that are covered by the course.
2. Exclusions: Coordinators also excluded skills falsely identified by the skills extractors and irrelevant to the course, leaving the skill unchecked.
3. Additions: In addition to verifying the extracted skills, coordinators were instructed to add skills not identified by the skills extractors but relevant to the course content.

The screenshot shows a web application titled "Skills Tagging Validation Exercise". It has a dark blue header with a home icon and the title. Below the header, there's a section for "Course Code: Course Title" with a "See my course details" button. The main content area is divided into two steps:

- Step 1: Review list and select all skills taught in your course.** This section contains a list of skills with checkboxes:
 - ☒ Algorithm Design
 - ☐ Blockchain
 - ☐ Blockchain Technology
 - ☐ Business Finance
 - ☒ Cloud Technology
 - ☒ Computational Complexity
 - ☒ Computational Thinking
 - ☐ Computer Engineering
 - ☐ Computer Networks
- Step 2: Key in additional skills. (Please separate each skill with a semicolon ;)** This section has a text input field containing "Decomposition; abstraction;" and a "Save" button.

Figure 2. An in-house web app that displays skills tagged to each course. Course coordinators used this web app to validate tagged skills and propose skills not extracted by the algorithm.

At the end of the validation exercise, conducted over two months, 50 courses completed the exercise. Statistical analysis of binary classification was performed on these courses to evaluate the efficacy of the three skills extraction platforms. As we are examining each skills extraction platform independently, only three binary classifications were meaningful:

- True positive (TP): Skill tagged by the algorithm and validated 'True' by the course coordinator.
- False positive (FP): Skill tagged by the algorithm but indicated as 'False'.
- False negative (FN): Skill not tagged by the algorithm but indicated as additions during the validation process.

Obtaining true negatives (TN) is not applicable in this study, as a skill not tagged by the algorithm cannot be validated as false. Five statistical metrics were thus identified to be relevant in this study:

- Precision: High precision means that when the extraction algorithm tags a skill, it will likely be correct. A precision value closer to 1 indicates the ability of the extraction algorithm to avoid over-tagging with irrelevant skills. The precision is given by

$$\text{Precision} = \frac{TP}{TP + FP}$$

- False Discovery Rate (FDR): Complements precision by measuring the proportion of incorrect positive predictions. The FDR is important for understanding the proportion of skills incorrectly tagged by the extraction algorithm. This is crucial in maintaining the integrity and reliability of the tagging system. The FDR is given by

$$\text{FDR} = \frac{FP}{TP + FP} = 1 - \text{Precision}$$

- Recall: High recall is crucial to ensure that the extraction algorithm does not miss important skills, particularly in job matching or educational needs assessment. The recall is given by

$$\text{Recall} = \frac{TP}{TP + FN}$$

- False Negative Rate (FNR): Complements recall and is crucial for understanding how often the extraction algorithm did not tag a relevant skill. In applications like personalized learning or career development, a high false negative rate could lead to missed opportunities for skill development recommendations. The FNR is given by

$$\text{FNR} = \frac{\text{FN}}{\text{TP} + \text{FN}} = 1 - \text{Recall}$$

- F₁ Score: Harmonic mean of precision and recall. An F₁ Score closer to 1 is better. The F₁ Score is given by

$$\text{F}_1 \text{ Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Since the FDR and FNR are complementary to the precision and recall, respectively, in subsequent sections, conclusions made for FDR and FNR will follow from mentions of precision and recall, respectively, without explicitly calculating those values.

4 RESULTS & DISCUSSION

In this section, we present the results of our analysis of the fifty courses that underwent skill validation. The distribution of skills tagged to each course is captured in Table 1 and Fig. 3. Platform A extracted 163 skills, Platform B 515 skills, and Platform C 393 skills from the fifty courses. Separately, the course coordinators added a total of eighty skills during the validation. We further looked at various combinations of the skills tagged by the three platforms and performed an equivalent analysis. The results from the validation exercise are collated in Tables 2 and 3.

Table 1. Compiled statistics showing the distribution of skills tagged by each platform.

Platform	A	B	C
# Skills Tagged	163	515	393
Mean # Skills Tagged	3.26	10.3	7.86
Standard Deviation	3.049	6.538	6.068




Figure 3. A stacked bar plot showing the frequency of the number of skills tagged by each platform. Correspondingly, a box plot showing the distribution of the number of skills tagged.

Table 2. Compilation of precision, recall, and F_1 Score for all three skills extraction platforms. The scores in bold denote the platform that performs best for that metric.

Platform	A	B	C
Precision	0.6564	0.7068	0.7074
Recall	0.5722	0.8198	0.7765
F_1 Score	0.6114	0.7591	0.7403

Referring to Table 2, Platform C demonstrates outstanding performance in precision (consequently FDR), indicating that a significant proportion of skills it extracted as relevant are indeed correct. At the same time, a low FDR implies that the skills it identifies are reliably relevant. Platform B excels in recall (consequently FNR) and returns the highest F_1 Score. Its high recall indicates its effectiveness in identifying a large proportion of relevant skills, and the complementary low FNR signifies its effectiveness in minimizing missed relevant skills. The specific requirements of skill tagging should guide the choice of platform. If minimizing false positives is crucial, Platform C is preferable because of its high precision. This might be important when students need to be accredited based on the skills they have learned, where the institution's reputation hinges on producing students with the necessary skill set for the job market. Conversely, if capturing as many relevant skills as possible is the priority, even at the risk of including some irrelevant ones, then Platform B, with its high recall, is more appropriate. This might be important to help students gain insights into a broad spectrum of their skills, some of which they might not have actively considered. This can help them recognize transferable and auxiliary skills that are valuable in various career paths. Thus, an integrated approach that leverages the strengths of each platform could offer a more nuanced solution. Such a strategy could combine Platform B's ability to identify a broad range of relevant skills with Platform C's strength in minimizing false positives.

In our analysis of the results, we keep in mind factors that could affect the validity of the results. Firstly, the disparity in the number of skills tagged influences the performances in precision and recall. For example, extracting more skills (as seen with Platform B) might enhance recall due to the converse of choice overload [24], [25]. When presented with more skill options, the tendency is to pick according to perceived benefit or desirability bias. A course coordinator might choose most of the skills shown to them as it might put the course in a better light. In contrast, extracting lesser skills (as seen with Platform A) could significantly lower recall. Tasked to identify skills without the visible presence of options, people might struggle to identify them because one has to recall them from memory, which is cognitively more demanding, thus leaving the option to provide additional skills empty. In our study of the 50 courses, 23 returned with at least one additional skill, among which nine validated courses had more than three additional skills and only three with more than 5. One needs to be aware that these external behaviors will be present in a human-in-the-loop form of validation.

Table 3. Compilation of precision, recall, and F_1 Score for all three skills extraction platforms. The scores in bold denote the platform that performs best for that metric.

Platform	$A \cup B$	$A \cup C$	$B \cup C$	$A \cup B \cup C$	$A \cap B$	$A \cap C$	$B \cap C$	$A \cap B \cap C$
# Skills	611	496	732	808	67	60	176	40
Precision	0.6890	0.6875	0.6940	0.6795	0.7463	0.7333	0.7614	0.7000
Recall	0.8403	0.8100	0.8639	0.8728	0.3846	0.3548	0.6262	0.2593
F_1 Score	0.7572	0.7437	0.7697	0.7641	0.5076	0.4783	0.6872	0.3784

Referring to Table 3, the columns denoted by the unions refer to a hypothetical platform with an algorithm that extracts the skills extracted by the union combinations of the three platforms. The columns denoted by the intersections outline the scenario where one extraction platform may be used to validate the outputs of the other. It is clear that by using combinations of the extraction platforms, we may achieve better results. In our analysis, $B \cap C$ gives the best precision, $A \cup B \cup C$ gives the best recall, and $B \cup C$ with the highest F_1 Score. In particular, taking the intersection of the platforms always produces significantly better precision, and taking the union of the platforms gives significantly better recall than performing skills extraction in isolation. While the recommendation here is to employ both Platforms B and C to achieve the best harmonic mean of precision and recall, the use of each validated data set should correspond to the context that applies to the users. For example, while it can be argued that Platform B alone achieves very similar F_1 Score to that of the union of both B and C, there may be

reasons to focus on the higher recall performance especially if we wish to reduce the burden of course coordinators from the cognitively demanding task of identifying missing skills. A comprehensive study on user preference might shed light on the approach institutions should take.

4.1 Implications and Future Work

In the rapidly evolving landscape of higher education, aligning academic curricula with industry needs is paramount. Advances in AI offer a novel approach to enhancing this alignment, particularly by tagging skills in course descriptions. Having outlined and demonstrated the efficacy of AI-assisted skill tagging and subsequent validation by course coordinators, we now focus on its potential implications and opportunities for future related work.

At the institutional level, such an extraction-validation exercise can facilitate and measure the impact of course-skill alignment on critical student outcomes such as employability, skill acquisition, and career progression. This would involve linking skills extraction with real-time job market data by analysing job descriptions and required skills in the job market and comparing them with those extracted from course descriptions. Such integration can help institutions assess how well academic curricula align with industry requirements, ensuring that their students are learning skills in high demand. By understanding this alignment, skills tagging can provide valuable feedback to educational institutions, enabling them to fine-tune their programs to better prepare students for the workforce. Longitudinal studies can be implemented to track how skills evolve in academic curricula. This continuous monitoring can provide insights into changing skills and industry demands, helping the institution update its courses to be ahead of the curve and prepare students to meet current and future job market needs.

For individual courses, course design support can be provided to educators in designing and updating academic curricula to ensure each course is relevant, comprehensive, and aligned with educational goals and industry standards. This involves a multi-faceted approach depending on the course's emphasis on the personal, social, workplace, and academic relevance. Data-driven insights through data and learning analytics to provide insights into current trends in industry demands and skill requirements. A gap analysis between the skills currently taught in a course and those emerging in the relevant field can inform curriculum and course designers on adapting course content. Efforts can begin with defining clear, skill-based learning objectives for each course.

Lastly, the exercise can lead to the development of a recommendation system for personalized learning pathways to fulfil each student's unique needs and career goals [5]. Institutions can first map out a comprehensive landscape of skills taught across different courses and disciplines using the skills extraction exercise data. Students can then undergo an evaluation to identify their interests and skills gaps. This could include aptitude tests, career interest surveys, and past academic performance analysis. Based on the evaluation, a recommender system can propose a personalized set of courses and electives that align with the student's career aspirations and learning needs. For example, a student interested in a career in digital marketing might be guided to take courses focusing on essential skills such as social media analytics, content strategy, market research, emergent skills such as generative AI and data ethics as well as fundamental skills such as quantitative and qualitative research methods. Aligning the learning pathways with projected industry trends ensures that students acquire skills that will be valuable in the job market in the future.

Before concluding, we would like to highlight the limitations of the current study which lay the ground for future research. Firstly, the sample size is small. A follow-up study involving more courses could potentially reveal more nuanced implications for each of the four disciplines. Secondly, course and programme coordinators may not always have the same understanding as the industry. It could be worth the while to involve industry partners as subject matter experts in the validation process and determine if comparable results are obtained. Finally, and in a similar vein, it could be important to understand if students perceive they have gained the identified and validated skills after completing the courses as part of the validation process.

5 CONCLUSIONS

In conclusion, we have employed a systematic approach to analyse a dataset of course descriptions and extract relevant skills for further examination. The research methodology encompassed stages of data collection, cleaning, category selection, random sampling, and skills extraction. The data was meticulously cleaned to ensure quality and relevance, and a stratified random sampling strategy was employed to ensure representation across four academic disciplines. The core of this study involved the

application of three skills extractor algorithms, each employing natural language processing techniques. We then incorporated a human-in-the-loop validation process involving course coordinators as domain subject matter experts to validate the relevance of the extracted skills. This collaborative process included verifying, excluding, and adding skills, ensuring the validity of the final dataset. Our analysis focused on evaluating the efficacy of the three skills extraction platforms through statistical analysis of binary classification. Platforms B and C have comparable precision, indicating their effectiveness in correctly identifying relevant skills. At the same time, Platform B excelled in recall and the F₁ Score, demonstrating its ability to identify a broad range of relevant skills. The combined analysis of all three platforms revealed a synergistic effect, enhancing the overall skill tagging efficacy and achieving a balance between precision and recall. Lastly, we outlined potential uses of such a process for strengthening institutional outcomes at the macro level, adapting course design at the meso level, and recommending personalized learning pathways at the micro level.

CONFLICT OF INTEREST

The authors declare that the course description data used in this study is the intellectual property of NTU. As employees, we were granted access to this data for research purposes. The authors have ensured that the data has been used strictly within the ethical guidelines and permissions granted by NTU.

ACKNOWLEDGEMENTS

This project is supported by the EdeX Grant from the Centre for Teaching, Learning and Pedagogy, NTU. The authors would like to thank the faculty of NTU who contributed their time and expertise to validate the skills.

REFERENCES

- [1] H. T. Thanh and N. T. M. Tam, "Generic Structure Potential of Course Description," *VNU Journal of Foreign Studies*, vol. 37, no. 1, pp. 83-98, 2021.
- [2] P. Gómez, M. J. González, F. Gil, J. L. Lupiáñez, M. F. Moreno, L. Rico and I. Romero, "Assessing the relevance of higher education courses," *Evaluation and Program Planning*, vol. 30, no. 2, pp. 149-160, 2007.
- [3] J. Mourey, M. M. Markley and S. K. Koernig, "Dazzling Descriptions and Tantalizing Titles: How Simple Versus Complex Course Information Influences Course Selection," *Journal of Marketing Education*, vol. 44, no. 1, pp. 100-112, 2021.
- [4] E. W. Mainardes, M. J. Silva and M. J. C. de Souza Domingues, "The development of new higher education courses," *Service Business*, vol. 4, no. 3, pp. 271-288, 2010.
- [5] K. Kitto, "How can EdTech support graduate employability?," in *2022: ASCILITE 2022 Conference Proceedings: Reconnecting relationships through technology*, Sydney, 2022.
- [6] D. Miliband, "Choice and voice in personalised learning," in *Schooling for tomorrow: Personalising education*, Paris, OECD Publishing, 2006, pp. 21-30.
- [7] C. Bruce, H. Hughes and M. M. Somerville, "Supporting Informed Learners in the Twenty-first Century," *Library Trends*, vol. 60, no. 3, pp. 522-545, 2012.
- [8] L. Hendricks, "Cross-country variation in educational attainment: structural change or within-industry skill upgrading?," *Journal of Economic Growth*, vol. 15, no. 3, pp. 205-233, 2010.
- [9] A. Peck, "Mapping Career-Ready Skills Through Student Leadership Programs," *New Directions for Student Leadership*, vol. 2018, no. 157, pp. 71-83, 2018.

- [10] H. M. Oraison, L. Konjarski and S. T. Howe, "Does university prepare students for employment? Alignment between graduate attributes, accreditation requirements and industry employability criteria," *Journal of Teaching and Learning for Graduate Employability*, vol. 10, no. 1, pp. 173-194, 2019.
- [11] N. R. Aljohani, A. Aslam, A. O. Khadidos and S.-U. Hassan, "Bridging the skill gap between the acquired university curriculum and the requirements of the job market: A data-driven analysis of scientific literature," *Journal of Innovation & Knowledge*, vol. 7, no. 3, p. 100190, 2022.
- [12] C. M. Mason, H. Chen, D. Evans and G. Walker, "Illustrating the application of a skills taxonomy, machine learning and online data to inform career and training decisions," *The International Journal of Information and Learning Technology*, vol. 40, no. 4, pp. 353-371, 2023.
- [13] K. R. Kelley and S. K. Buchanan, "College to career ready: Innovative practices that lead to integrated employment," *Journal of Vocational Rehabilitation*, vol. 46, no. 3, pp. 327-332, 2017.
- [14] M. Cheng, O. Adekola, J. Albia and S. Cai, "Employability in higher education: a review of key stakeholders' perspectives," *Higher Education Evaluation and Development*, vol. 16, no. 1, pp. 16-31, 2021.
- [15] S. Moore, H. A. Nguyen and J. Stamper, "Leveraging students to generate skill tags that inform learning analytics," in *Proceedings of the 16th International Conference of the Learning Sciences-ICLS 2022*, Hiroshima, 2022.
- [16] S. Moore, H. A. Nguyen and J. Stamper, "Leveraging students to generate skill tags that inform learning analytics," in *Proceedings of the 16th International Conference of the Learning Sciences-ICLS 2022*, Hiroshima, 2022.
- [17] D. M. Blei, A. Y. Ng and M. I. Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research*, vol. 3, pp. 993-1022, 2003.
- [18] D. Lin and X. Wu, "Phrase Clustering for Discriminative Learning," in *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, Singapore, 2009.
- [19] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu and M. Zhu, "A Practical Algorithm for Topic Modeling with Provable Guarantees," *arXiv prepr.*, vol. arXiv:1212.4777 [cs.LG], pp. 1-26, 2012.
- [20] A. Hassani, A. Iranmanesh and N. Mansouri, "Text Mining using Nonnegative Matrix Factorization and Latent Semantic Analysis," *arXiv prepr.*, vol. arXiv:1911.04705 [cs.LG], pp. 1-26, 2020.
- [21] JobTech, JobTech - We help people stay employable, Accessed 2 September 2023. Retrieved from: <https://jobtech.co/>.
- [22] LightCast, Lightcast Open Skills Taxonomy, Accessed 2 September 2023. Retrieved from: <https://lightcast.io/open-skills>.
- [23] SkillsFuture Singapore, Skills Extraction Algorithm (SEA), Accessed 2 September 2023. Retrieved from: <https://www.tpgateway.gov.sg/plan-courses/skills-extraction-algorithm>.
- [24] M. Zhao, F. Javed, F. Jacob and M. McNair, "SKILL: A System for Skill Identification and Normalization," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 2, pp. 4012-2017, 2015.

- [25] L. M. Desimone, A. Stornaiuolo, N. Flores, K. Pak, A. Edgerton, T. P. Nichols, E. C. Plummer and A. Porter, "Successes and Challenges of the "New" College- and Career-Ready Standards: Seven Implementation Trends," *Educational Researcher*, vol. 48, no. 3, pp. 167-178, 2019.