Towards Skills-based Curriculum Analytics: Can we automate the recognition of prior learning?

Kirsty Kitto
Connected Intelligence Centre,
University of Technology Sydney
Sydney, NSW
kirsty.kitto@uts.edu.au

Nikhil Sarathy
Connected Intelligence Centre,
University of Technology Sydney
Sydney, NSW
nikhil.sarathy@student.uts.edu.au

Aleksandr Gromov
Connected Intelligence Centre,
University of Technology Sydney
Sydney, NSW
aleksandr.gromov@uts.edu.au

Ming Liu
Connected Intelligence Centre,
University of Technology Sydney
Sydney, NSW
ming.liu@uts.edu.au

Katarzyna Musial
Advanced Analytics Institute,
University of Technology Sydney
Sydney, NSW
katarzyna.musial-gabrys@uts.edu.au

Simon Buckingham Shum
Connected Intelligence Centre,
University of Technology Sydney
Sydney, NSW
simon.buckinghamshum@uts.edu.au

ABSTRACT

In an era that will increasingly depend upon lifelong learning, the LA community will need to facilitate the movement and sharing of data and information across institutional and geographic boundaries. This will help us to recognise prior learning (RPL) and to personalise the learner experience. Here, we explore the utility of skills-based curriculum analytics and how it might facilitate the process of awarding RPL between two institutions. We explore the potential utility of combining natural language processing and skills taxonomies to map between subject descriptions for these two different institutions, presenting two algorithms we have developed to facilitate RPL and evaluating their performance. We draw attention to some of the issues that arise, listing areas that we consider ripe for future work in a surprisingly underexplored area.

CCS CONCEPTS

• **Information systems → Document collection models**; **Ontologies**; *Data exchange*; *Decision support systems.*

KEYWORDS

curriculum analytics, lifelong learning, recognition of prior learning, semantic spaces, skills ontologies

ACM Reference Format:
Kirsty Kitto, Nikhil Sarathy, Aleksandr Gromov, Ming Liu, Katarzyna Musial, and Simon Buckingham Shum. 2020. Towards Skills-based Curriculum Analytics: Can we automate the recognition of prior learning?. In *Proceedings of the 10th International Conference on Learning Analytics and Knowledge (LAK '20), March 23–27, 2020, Frankfurt, Germany*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3375462.3375526

ACM ISBN 978-1-4503-7712-6/20/03…$15.00
https://doi.org/10.1145/3375462.3375526

1 LIFELONG LEARNING

The modern conceptualisation of employment is rapidly shifting. While our parents planned to work for the same company for life, our children can expect to change career many times [1]. Rather than completing a set amount of schooling early in life to achieve a qualification, the 4th industrial revolution [2] means that people will increasingly need to return to the education sector: (i) for further training to enhance a skill base (ii) to gain new skills as their field is disrupted by technology, or (iii) to reskill so that they might transition into new careers. In this context, helping students to achieve success becomes a project in lifelong learning [21]. However, the problem of data transferability raised by Ryan Baker in his LAK'19 keynote[1] quickly emerges within this context.

How might the field of Learning Analytics (LA) work to support a more seamless movement of learning data across a lifetime of learning? Numerous barriers exist when we start to consider this scenario, such as, data interoperability [15], the ethics of data privacy vs access [24], institutional silos [23], and political and standards based discrepancies across different nations [16]. Here, we will focus upon transfer credit and the recognition of prior learning in higher education, considering ways in which LA might help us work towards at least semi-automating a process that is known to be a resource intensive bottleneck for many universities [16]. We will start with a consideration of the problem.

1.1 The recognition of prior learning

What happens when a student attempts to change from one study pathway to another? They might be moving from a degree offered at one institution to another university, or perhaps they are applying for admission into a Masters degree. How can the new institution be sure that the student has indeed satisfied any requirements associated with this transfer? Some institutions will offer prior credit for a selection of courses, enabling students to complete their studies faster, or to study more advanced courses due to a waiving of pre-requisites. However, determining whether a student has indeed satisfied the required level of expertise is a non trivial exercise. Considerable effort is often devoted to testing students

[1] https://www.youtube.com/watch?v=DDhPa6lVogY

upon entry, or to employing a workforce to manually map competencies between various, constantly changing, degree programs. Sometimes this mapping enables automatic recognition from prespecified pathways, but it often involves chasing down previous completion requirements for new applicants on a case by case basis, an extremely resource intensive task. This problem is further exacerbated in a lifelong learning context. Consider the increasing numbers of people who are now starting to return to higher education, often after years of work experience and substantial professional development — how are their skills and competencies to be verified? People re-entering education from the workforce often have a strong portfolio of work that can be used in the recognition of prior learning (RPL) [16], but this creates even more logistical challenges than the problem of granting credit transfer. Portfolios of work are rarely shared across institutions, and it is rare to see curriculum details made publicly available, let alone assessment details. Similarly, companies seldom release details of staff training programs. This siloing of data makes the job of those working to validate claims for recognition of learning particularly difficult. And even when curriculum information is available, it is difficult to compare the training received in one environment with that obtained in another. Micro-credentials and badges are often presented as a possible future solution to this problem [17] but this solution still requires a mapping of those badges into various institutional pathways.[2] Similarly, blockchain has been mooted as a mechanism for granting credit [14] but a similar problem of mapping any credential between institutions remains.

In what follows we will refer to all of these problems under the broad banner of RPL, while recognising that this is a slight abuse of a term that is usually applied to the problem of providing credit to people who are seeking to re-enter formal study from the workforce. While the problem of RPL has long been acknowledged, we consider it likely that lifelong learning will lead to a growing expectation that institutions can efficiently award RPL. However, as is well documented in critical infrastructure studies [6], the challenge of achieving common ground to describe curriculum should not be under-estimated. What solutions have been proposed?

1.2 Are graduate attributes a solution?

A wide range of education sectors around the world have used qualification frameworks to make graduate attributes portable beyond the institution in which they were recognised. Competency frameworks consisting of lists of graduate attributes have been released by a number of national bodies, such as the Australian Qualifications Framework (AQF), Scottish Credit and Qualifications Framework (SCQF), and the German Qualifications Framework (Deutsche Qualifikationsrahmen, DQR) at the national level, and the European Qualifications Framework (EQF) at a regional level [22]. At the institutional level, we often see various mappings proposed between: assessment of learning outcomes in different domains [37]; curriculum and content [4, 18]; and graduate attributes [38]. However, it is costly to both manually create and maintain these mappings, meaning they can rapidly lose their relevance as the content of a course changes over time. Even more problematic, there is no universal qualifications framework, which means that different jurisdictions describe graduate capabilities in different ways. Consequently, while graduate attributes (GA) are an important part of describing student capabilities at an institutional level, they fail to solve the challenge of sector wide RPL.

[2] Although we note that *Badgr pathways* holds promise for detailing these mappings — see https://www.concentricsky.com/articles/detail/introducing-badgr-pathways.

1.3 Contribution

Here, we investigate the potential of combining Natural language processing (NLP) with a defined skills ontology to work towards at least partial automation of the RPL process. The contribution of this paper is twofold: (i) First, we provide two approaches that can be used to map between publicly available subject descriptions. (ii) Second, we provide two methods for evaluating the resulting mappings. This will help to further compare methods as our approach matures.

Some work has been attempted in this direction previously, and will be discussed in Section 2, where we will see that it has tended to be restricted to a particular sub-domain or field of knowledge. Section 3 will present a series of Proof of Concept (PoC) experiments investigating possibilities for mapping between curricula obtained from two different higher educational institutions. This section will also discuss the evaluation framework we have used to judge the success of the resultant mappings. We will finish with an consideration of promising avenues for future work in this new emerging sub-field of curriculum analytics (Section 4).

2 AUTOMATIC CURRICULUM MAPPING

Some early efforts to map curriculum into ontologies was completed in the program of work that sought to represent shareable and reusable learning objects (LOs) using ontologies (by e.g. Verbert et al. [40], and Gašević et al. [10]). However, this approach relied upon human effort to perform annotation of the LOs, and has largely halted. It is worthwhile considering *why* such a potentially useful practice never obtained widespread adoption. Firstly, it appears that too much extra effort was required to complete these mappings by hand, secondly, Verborgh and Vander Sande [41] suggests that the messiness of the open world wide web has meant that terms consistently map to contradictory meanings, and different terms are used for the same object, making it difficult to maintain a consistent mapping as the knowledge domain increases in size. While some fields have well defined knowledge structures that can be mapped into largely uncontroversial taxonomies (e.g. mathematics and physics) this is not the norm. And yet the bulk of the work that has been completed in mapping curriculum assumes precisely this type of formalised epistemological structure — something that is not likely to be achievable in the formalised mappings required by the early days of the semantic web [6].

More flexible mappings are required, and attempts to address this problem have emerged over the years. Most of these approaches make use of (semi)automated methods to map curriculum into some sort of taxonomy. For example, Gibson et al. [12] gave a simple demonstration of the fact that it is possible to map curriculum documentation into widely understood educational constructs (e.g. Bloom's taxonomy) using simple stemming and text analysis. Similarly, Xun et al. [44] made use of TF-IDF to map IT curriculum documentation into two ontologies: the Skills Framework for the

Information Age (SIFA) ontology, and Bloom's taxonomy. Somewhat related, Pardos et al. [34] perform a detailed study of the pathways students follow through Berkeley curriculum, building a recommendation system that aims to help them choose a path that completes course requirements faster. However, none of these methods have been applied to the problem of RPL.

Thus NLP potentially has the necessary flexibility for addressing the RPL problem, but has not been applied to the problem of mapping curriculum between institutions. Those methods that *have* been applied are manual, and so resource intensive, leading to their ongoing status as a research curiosity. How might we start to make use of the best characteristics of both methods? Perhaps the best developed field using NLP to model curriculum has arisen in mapping computer science curriculum into defined taxonomies, a topic which we shall now explore.

2.1 Mapping computer science curriculum

The vast majority of work completed in applying techniques from NLP to curriculum mapping has tended to concentrate upon Computer Science (CS) curriculum. This is not just due to the necessary expertise occurring predominantly among those who teach CS. Importantly, the ACM and IEEE have provided a standard curriculum series (i.e. a taxonomy of what should be taught in CS) for over 40 years, with a view towards helping faculty to design a curriculum for their universities that fits into this broader professional context. The Computer Science Curricula 2013 (CS2013) [3] and the Computer Engineering Curricular Guideline (CE2016) [9] are examples of these frameworks. Each represents a Body of Knowledge (BoK) consisting of a set of knowledge areas (KAs), both of which contain about ten more specialised knowledge units (KUs) represented by a short document. A number of studies have explored this curriculum structure, analysing its evolution in time [13, 28].

Closer to the problems associated with RPL, a body of work has attempted to map the curricula from various education institutions into the KAs defined by a range of CS frameworks. For example Kawintiranon et al. [20] matched full curriculum documentation (i.e. including course materials) for Computer Networks, Operating Systems, and Systems Architecture subjects taught at a university into the CE2016 framework. A dictionary was constructed by extracting KUs from various KAs in CE2016, and used to formulate queries for various search APIs and so retrieve external documents. These were processed using TF-IDF to obtain keyword matrices describing courses, and thence to compute association scores between CE2016 and the university curricula. The same methods were later used to demonstrate the relationships between the curricula for 5 institutions, showing a fair amount of similarity. An ingenious approach was provided by Sekiya et al. [36], and further developed in Matsuda et al. [29]. This work makes use of simplified, supervised Latent Dirichlet Allocation (ssLDA) to estimate the relative weights of the Knowledge Areas (KAs) of CS2013 in various CS curricula. This method projects a subject syllabus to a point in the KA space defined by CS2013, which works to relate the strength of the connection between the syllabus and the corresponding KA. This method was used to model a full curriculum pathway by finding the centre of all points corresponding to a syllabus in a degree program. The authors were able to demonstrate a range of geographic clusters in the curricula of 50 leading computer science departments from around the world. Similarly, Dai et al. [7] used Labelled Latent Dirichlet Allocation (L-LDA) [35] to map curriculum into CS2013 at the level of KUs with the aim of enabling personalised learning object recommendation in MOOCs. Despite their promise, the dependency of these methods upon an ontology that is defined only for computer science and engineering subjects restricts its utility for the broader RPL problem.

In summary, NLP holds promise for mapping specific subsets of curricula into identified taxonomies. However, the approaches adopted to date are CS-centric. To our knowledge such approaches are yet to be tested on other disciplinary curricula, and far less in the open-ended use case where RPL potentially spans a university's entire curriculum. How can we scale solutions like this up? Doing so requires a more general curriculum structure which can apply across the entire range of fields taught in the education sector.

2.2 Skills taxonomies as a portable curriculum representation

Scaling curriculum models such that they can effectively support RPL across an institution, and eventually across a lifetime of learning, requires a generalisable representation of the skills and knowledge that students gain during their education i.e. some sort of universal ontology. And yet these have already failed to achieve widespread adoption in many sectors where the Semantic Web had aspirational use cases. Specifically, in the EdTech community the workload required to map educational materials into a standardised format makes ontologies prohibitively expensive to maintain. Manual approaches seem to have consistently failed to achieve widespread adoption.

Learning from the progress in CS curriculum mapping introduced above, we have been investigating publicly available taxonomies which have been created to link skills, competences, qualifications and occupations together into formal ontologies and taxonomies. For example, the European Skills, Competences, Qualifications and Occupations (ESCO) ontology[3] was created through manual mappings in an attempt to make qualifications portable within the EU. A similar mapping, O*Net[4] was created in the United States, and ANZSCO[5] has been generated for the Australia and New Zealand context. These taxonomies each provide a basis for tagging the skills taught in various educational domains. Critically, as recruitment processes have increasingly shifted to web platforms such as LinkedIn[6], SEEK[7] and Monster[8] it has become possible to collect data about occupations and the skills that are considered essential for them at scale. This data has been used to populate employability related ontologies by vendors such as Burning Glass technologies[9] who use a proprietary ontology for describing skills and their relationship to various job roles. Models are starting to emerge that seek to make use of these datasets to construct job

[3]https://ec.europa.eu/esco/portal/skill
[4]https://services.onetcenter.org/reference/
[5]http://www.abs.gov.au/ANZSCO
[6]https://www.linkedin.com/jobs/
[7]https://www.seek.com.au/
[8]https://www.monster.com/
[9]https://www.burning-glass.com

search engines [33] and order ePortfolios over a lifetime of learning [5]. However, we are not aware of any studies that have made use of these frameworks to understand the curriculum of an educational institution at scale, a gap that this work seeks to address. An initial attempt that we made to automate RPL using the L-LDA approach of Matsuda et al. [29] but using ESCO as a broader taxonomy than the CS2013 taxonomy failed to converge. The taxonomy is large, making that approach computationally expensive. Implementing this method also requires substantial NLP expertise, which makes it unlikely to be feasible across most institutions (which often do not possess the necessary expertise in central units). In what follows we discuss a series of experiments we have conducted to find a more computationally tractable and easy to use approach.

3 PROOF OF CONCEPT EXPERIMENTS

In this section we discuss a series of experiments that we have performed to explore the practicality of using NLP for mapping between subject descriptions from two different institutions UTS and UNSW, chosen for their openly available curriculum documentation. Acknowledging that there is no common international vocabulary for referring to curriculum, we define the following conventions:

- **Course:** A degree into which a student enrols (e.g. a Bachelor of Science, or a Master of Information Technology).
- **Subject:** A specific unit of study that a student undertakes during a course (e.g. COMP2110, LANS and Networking).

We have built on a product released by Burning Glass (BG) to explore the potential of "off the shelf" tools to assist with RPL. Of particular relevance to our challenge, BG provides a tool that can be used to tag curriculum structures (i.e. textual descriptions of courses and subjects) using a proprietary skills ontology that they curate. This ontology consists of 17,420 skills, organised into 663 coarser grained skill clusters, and mapped into 28 high level fields of expertise. Thus, this ontology covers a broad range of skills, making it a possible replacement of CS2013, capable of representing the entire curriculum of a university. In what follows we will make use of the BG curriculum API which when called on a subject's Title/Description data returns a JSON response with a list of `careers`, `potentialSkills`, and `skills` that are considered likely to be taught in that subject (as judged by a set of similarity scores which are calculated using a proprietary algorithm — see Figure 2 for an example of this returned object).

How promising is this tool for tackling the RPL problem? Our work was guided by the following two research questions (RQ):

- **RQ1:** How can an off the shelf tool that tags curriculum using a skills ontology be used to provide RPL mappings between the curriculum structure of two universities?
- **RQ2:** Which methods show the most promise for automating RPL between two institutions?

We decided to restrict our analysis to two subsets of each institution's offerings in the Information Technology (IT) and Medicine domains, as two examples of substantially different curriculum structures. This decision was both pragmatic, and considered the business process by which RPL is usually awarded, where an applicant's record is manually examined in the context of an application for entry to a specific degree program or study area (i.e. not an entire institution's curriculum). We consider this step justified, as it is likely that the bulk of RPL automation could be managed by matching subject offerings across defined fields, however, we acknowledge that there are situations for which mapping to fields of study is likely to cause problems. For example, new fields, or transdisciplinary degree subjects that aim to build skill sets not well represented by traditional fields are likely to be poorly served by this approach. This is a problem that we leave for future work.

Table 1: The breakdown of Courses and Subjects offered in IT and Medicine at the two different institutions whose curriculum is examined here. Courses are further specified with Undergraduate/Postgraduate (U/P) numbers.

Institution	Field	Courses	U	P	Subjects
UTS	IT	12	7	5	132
	Medicine	31	8	23	295
UNSW	IT	7	4	3	118
	Medicine	24	11	13	185

Table 1 summarises the structure of the datasets examined from the two institutions chosen for our study. Curriculum details were selected from the detailed information provided in each institution's online handbook.[10]

3.1 Experiment 1: Using the Burning Glass taxonomy in a common semantic space

Our first experiment explored the utility of using the BG ontology as a basis state in a high dimensional *semantic space* [27], where each skill listed in the BG ontology was taken as a basis vector for an abstract "BG space". Thus, the description for each subject was sent to the BG content tagger, and the returned (sparse) skills list was used to represent that subject as a vector in "BG space" (see Figure 1). Modelling subjects from both UTS and UNSW using the same representation enabled us to generate ranked lists of possible matching subjects at the second institution for every subject offered the first institution (and vice versa).

3.1.1 Method. A depiction of the process that we followed is given in Figure 2 for a UTS subject: *Software Architecture*.[11]

- STEP 1: First, we scraped the data for a given subject from the institution's handbook.
- STEP 2: We sent the resulting subject name and its description (which was usually about a paragraph long) to BG, which returned a JSON response of `skills` weighted with a "score" (ranging from 0–1).
- STEP 3: We then mapped each subject into a "BG vector" by taking the `skills` as an ordered list, and using score associated with each `skill` as its projection in that basis direction.
- STEP 4: The cosine similarity was then calculated between each BG vector describing a subject offering at UTS and each offered at UNSW. This was used to generate a ranked list of closest matches for each subject from the other institution.

[10] For examples of the data used see http://www.handbook.uts.edu.au/ and http://legacy.handbook.unsw.edu.au/general/2018/SSAPO/previousEditions.html.
[11] See http://handbook.uts.edu.au/subjects/48433.html for the handbook entry.

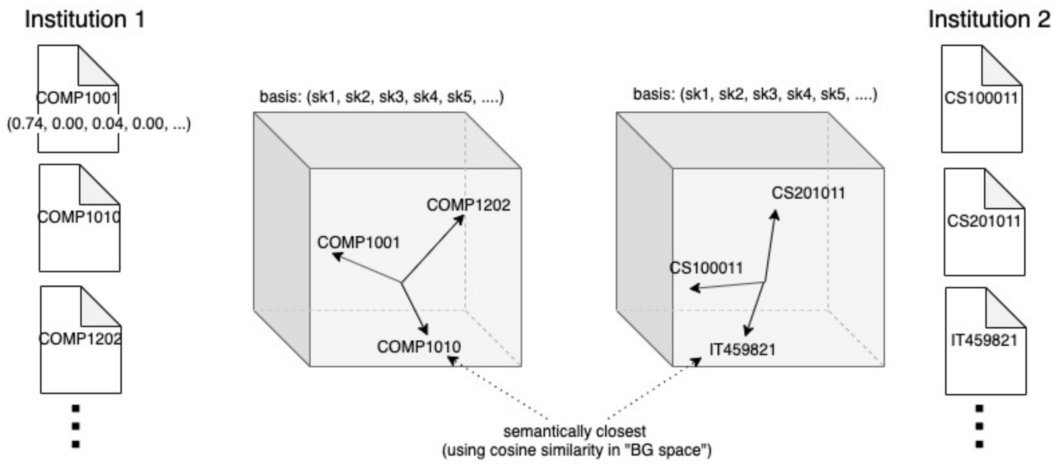

Figure 1: A schematic representation of the method followed for experiment 1. Descriptions of subjects are mapped into a BG space for each institution, each of which have the BG skills vectors as a shared basis. RPL candidates are found by ranking subjects in the new institution using a cosine similarity score.

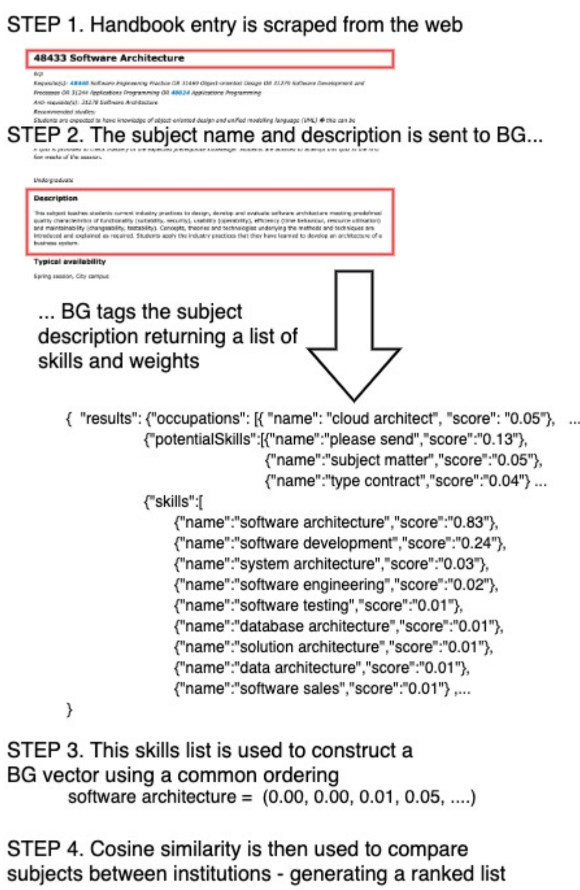

Figure 2: Mapping a subject (in this case Software Architecture, offered by UTS) into a "BG vector".

3.1.2 Results. Some example results are listed in Table 2, which gives the mappings obtained for 4 subjects over the two specialties (IT and Medicine) chosen for our study. The ranked lists below the first subject at the top of Table 2 list the subjects at the second institution which are evaluated by BG as most similar to that subject. Any person with sufficient expertise in a field (in this case IT or Medicine) can inspect the list provided and use the subject titles in the list to determine which subjects are most likely to be well aligned. For example, people with an IT background would generally agree that the "Computer Systems Fundamentals" subject offered by UNSW is likely to have more similarities with the "Software Architecture" and "Programming Fundamentals" subjects that are given as top ranked matches from UTS than the lower listed subjects. (Although see the Evaluation section below for a more rigorous exploration of this assertion.) Thus, this approach shows some promise for developing a decision support tool that would help a person tasked with awarding RPL to find subjects in their institution which are likely matches for the subjects an incoming student had studied at another one. This is a very encouraging outcome given the simplicity of the approach adopted (which we consider implementable at many institutions).

However, a number of weaknesses with the current approach can also be observed in Table 2. For example, many of the subjects in the ranked list are unlikely to be good matches for RPL purposes, even if some overlap is evident in the content that they are likely to teach. It is also apparent that many of these less related subjects are ranked higher than subjects which are likely to be a far closer match. Consider for example the high ranking of the "Algorithmic Verification" subject in the second column. It is unlikely that this subject is a closer match than e.g. "Software System Design and Implementation" (ranked third) or "Computer Systems Fundamentals" (ranked seventh). This problem emerges right at the point where BG returns a list of `skills` along with their similarity weights. Examining Table 3 provides us with an understanding of the underlying cause of this issue. We can see that there are cases where skills

Table 2: Indicative subject mappings from UTS to ranked list of candidate subjects at UNSW (and vice versa). Rankings are provided using the BG similarity scores.

Information Technology (IT)				Medicine			
Computer Systems Fundamentals (UNSW)		Software Architecture (UTS)		Advanced Epidemiology (UNSW)		Foundations in Public Health (UTS)	
• Software Architecture	0.03	• Algorithmic Verification	0.61	• Foundations in Public Health	0.97	• Biostatistics	0.94
•Programming Fundamentals	0.01	• User Interface Design and Construction	0.5	• Introduction to Public Health	0.97	• Public Health Capstone	0.93
• LANS and Routing	0.0	• Software as a Service Project	0.5	• Indigenous Public Health	0.96	• Ethics in Public Health	0.87
• Business Process and IT Strategy	0.0	• Software System Design and Implementation	0.48	• Fundamentals of Epidemiology and Population Health	0.96	• Physical Activity and Nutrition in International Contexts	0.86
• Internet Science	0.0	• Robotic Software Architecture	0.48	• Social Perspectives of Public Health	0.96	• Public Health Optometry	0.85
• Security Fundamentals	0.0	• Software Engineering Fundamentals	0.47	• Health Project and Program Management and Evaluation	0.95	• Epidemiology	0.84
• Software Defined Networks	0.0	• Computer Systems Fundamentals	0.44	• Advanced Epidemiology	0.76	• Health Promotion	0.75
• Game Design Studio 2	0.0	• Design Project A'	0.44	• Delivering Best Palliative Care Any Place Every Time	0.67	• Public Health Policy and Programs	0.72
• Advanced Routing Principles	0.0	• Programming Challenges	0.41	• Epidemiology and Population Health	0.65	• Introduction to Community Eye Health	0.69
• Deep Learning and Convolutional Neural Network	0.0	• Advanced Operating Systems	0.4	• Advanced Assessment and Diagnosis	0.19	• Principles of Prevention and Public Health Surveillance	0.66
• Multilayer Switched Networks	0.0	• Data Services Engineering	0.39	• Evidence for Nursing	0.04	• Social Business for Public Health	0.66
• Network Servers	0.0	• Software Construction: Techniques and Tools	0.38	• Core Practice for Physiotherapists	0.03	• Influencing Health Beliefs and Health Behaviours	0.61
• Cyber Security for Mobile Platforms	0.0	• Object-Oriented Design & Programming	0.38	• Sub-acute Rehabilitation	0.03	• Optometry, Medicine & Patient Management	0.55
• IoT Security	0.0	• Computer Science Project	0.38	• Foundations of the Australian Healthcare System	0.03	• Health Leadership	0.54
• Web Systems	0.0	• Management and Ethics	0.38	• Leading Change in Health Services and Practice	0.03	• Community Eye Health Project	0.3

core to the content taught in a subject are ranked too low, with skills that should not be considered closely related ranked above them. To improve these results we need a method for giving higher prominence to these core skills, a problem that we will return to in Experiment 2.

Perhaps most problematic for our current results, note that in Table 2, if we take the top listed UTS suggestion for "Computer Systems Fundamentals" at UNSW (which is "Software Architecture"), and perform the reciprocal mapping back to UNSW then this brings up a ranked list which does not include the original "Computer Systems Fundamentals" subject. The matches offered are not symmetrical, a problem that could lead to quite different RPL pathways if the current model was fully automated. This problem forms the basis of one evaluation method we have utilised below (see section 3.2.1).

One final observation concerns the low magnitude of the similarity scores that have been obtained for subjects in some of the lists. The subject descriptions at UNSW are often very short, which means that few details are provided to inform students about what is actually taught in that subject. Accordingly, our algorithm is also provided with far fewer details that it can use to extract BG skills, leading to poor matches into the curriculum offerings at UTS. Thus we see some promise for the current approach as a quick way of evaluating the quality of information provided about a subject in various university handbooks, and encouraging institutions to document their curriculum more completely.

3.2 Evaluation

Two methods have been used to evaluate our results.

3.2.1 Evaluation I — Recall for a reciprocal mapping. One evaluation metric that we adopted considered the number of times that a subject was returned in the top-5 list in a reciprocal mapping. This is a useful metric because it demonstrates how likely a RPL mapping between two institutions is likely to maintain integrity (and so tests for the occurrence of the problem raised above in section 3.1.2).

This evaluation proceeded in the following manner: (i) For any one subject in a given institution, we recorded the top-5 subjects returned as matches in the second institution. (ii) For each of those matching subjects, a list of the top-5 subjects from the original institutions were extracted. (iii) Any subject from the first institution that was matched in the reciprocal matching was recorded as a true positive (tp). (iv) Any subject from the first institution that was not matched in the reciprocal matching was recorded as a false negative (fn). This procedure gave the reciprocal recall results (defined here as $r - recall = \frac{tp}{tp+fn}$) recorded in Table 4. A number of points are worth noting. Firstly, we note that the recall scores are well above chance, but are not particularly good. The approach is not yet accurate enough for automating RPL. Also, the mappings from UTS to UNSW result in a substantially lower recall. This result supports our observation above that the generally less detailed subject descriptions provided by UNSWare a problem for automating RPL, as information from UTS is consistently being lost in the UTS $\rightarrow$ UNSW mapping.

3.2.2 Evaluation II — Human rankings. A second evaluation considered a subset of subjects mapped between institutions, asking humans (n=6, the authors of this paper) to rank the likelihood that

Table 3: Often we see responses returned from the BG curriculum API with a very low similarity score for skills taught in a subject. This means the skills vectors constructed using these scores are not appropriately weighted to account for that skill.

Subject	Skills and Scores
LANS and Routing	{cisco: 0.99}, {wiring: 0.76}, {cisco switching: 0.17}, {routers: 0.16}, {hardware experience: 0.08}...
IoT Security	{simulation: 0.97}, {middleware: 0.25}, {information systems: 0.04}, {transportation systems: 0.01}...
Internet Science	{research: 0.95}, {teamwork/collaboration: 0.72}, {experiments: 0.03}, {creativity: 0.00}, {online research: 0.00}...
Evidence for Nursing	{research: 0.92}, { clinical research: 0.03}, {patient care : 0.08}, {teaching: 0.01}, {public health and safety: 0.01}...
Advanced Epidemiology	{public health and safety: 1.00}, {research: 0.96}, {epidemiology: 0.41}, {surveillance: 0.13}, {clinical research: 0.06}...

Table 4: Results of mapping from one institution to the other and then back to the first institution. The reappearance of the original subject in the top-5 results of the final mapping was used to evaluate the reciprocal recall ($r - recall$) of the method explored in Experiment 1.

UTS→ UNSW	n	# mapped	$r - recall$
IT	132	52	.394
Medicine	295	60	.203

UNSW→ UTS	n	# mapped	$r - recall$
IT	118	53	.449
Medicine	185	73	.395

a matched IT subject at the second institution was a good match for potentially corresponding subjects at the first.

The evaluation proceeded as follows: (i) A survey chart for the human raters was prepared listing a set of subjects selected from the handbooks of both UTS (5 subjects) and UNSW(3 subjects). (ii) For each subject at a given institution, a list of the top-5 subjects returned from our experiments was generated. This was then augmented with up to 5 extra subjects to reach a baseline count of at least 10. (iii) Raters evaluated each of the paired subjects for their perceived similarity on a scale of 1–5 from 'least relevant' to 'highly relevant'. An inter-rater reliability score $\alpha = 0.465$ was calculated using Krippendorff's measure of inter-rater reliability [25]. This demonstrates some agreement between the raters, but was not substantial (further illustrating the difficulty of the RPL problem). We defined a score of 3 or higher as a judgement of a subject as relevant (with anything less declared not relevant). (iv) A ranked list of subjects was then generated by averaging the raters responses. (v) The two ranked lists were then compared, and the percentage agreement was calculated over lists of size 1, 3, and 5. The results of this evaluation for Experiment 1 are listed in Table 5.

3.3 Experiment 2: Augmenting BG with a concept based similarity model

This experiment was driven by the observation that BG tends to consistently return skills that are highly representative of a specific curriculum offering with a relatively low score (see Table 3 for some examples). We sought to improve the relevance of the BG skills returned for each subject.

3.3.1 WordNet enhanced skill based curriculum mapping (WESCM). Exploring the results of Experiment 1, we observed a number of limitations with the BG skills returned. Some skills were weighted too highly and appear out of context — they are noisy. For example, words such as `wiring` are highly weighted in subjects such as "LANS and Routing", and so override skills that are more closely related to the subject (e.g. `cisco switching`). Secondly, there are clusters of skills that tend to represent a particular job-market competency but not necessarily all of them are representative of the subject learning outcome (SLO). For example, BG skills like `software development` and `software maintenance` appear to be related in a generic software engineering context but are less well related to a more specific subject like "Software Architecture". We believe that these problems occur because the proprietary Burning Glass model tends to give higher weights to the skills frequently appearing in the training set, regardless of the curriculum content.[12] That is, BG appears to return skills that appear to be "corpus-specific" rather than "content-specific". Interestingly, in their study Mihalcea et al. [30] suggest that for short texts knowledge based semantic approach have a lower error rate when compared to corpus-based approaches that use a vector based similarity model.

We developed the WESCM algorithm to augment the BG results by improving the focus upon the *content* of subjects. We hypothesised that key concepts outlined in the subject description should also provide an indication of the skills that a student learns in that subject, and that relevant BG skills should be weighted higher if they have a strong semantic relation to *key phrases* in a subject description. For example, in "LANS and Routing", phrases like: "local area network (LAN) hardware and physical layer standards" should lead to a prioritization of the LAN term, as it is a key concept occurring multiple times throughout the subject description. In contrast, terms like "hardware" and "physical" should be deprioritized as they are less related to the key concepts in a LAN based subject, despite being common in the BG corpus (which is developed by scraping worldwide job advertisements which will have a higher frequency of these terms. Keyphrase extraction techniques [11, 42, 43] provide a method for identifying the key concepts in a subject description. We used WordNet [32] to filter out these "noisy" skills and assign new semantic weights to returned skills based on the semantic distance between key concepts and skills.

We implemented the WESCM algorithm as follows:

STEP 1: Key sentence extraction. This step focuses on extracting key sentences from a subject description using text summarization. Inspired by the work of Mihalcea and Tarau [31], we used the TextRank algorithm to extract the summary sentences. This method of extractive summarization uses

[12] Although we cannot be sure as the BG algorithm is a proprietary black box.

Table 5: Human ranking evaluation results for Experiments 1 and 2. The 1_{BG}, 1_{WESCM}, 3_{BG}, 3_{WESCM}, 5_{BG}, 5_{WESCM} metrics give percentage agreement between the BG and WESCM predictions and the human raters for lists of size, 1, 3, and 5.

UTS → UNSW	1_{BG}	1_{WESCM}	3_{BG}	3_{WESCM}	5_{BG}	5_{WESCM}
48024 Applications Programming	.000	.000	.333	.666	.200	.666
32524 LANS and Routing	1.000	1.000	.666	1.000	.600	.666
32146 Data Visualisation and Visual Analytics	.000	.000	.000	.000	.000	.000
42028 Deep Learning and Convolutional Neural Networks	.000	.000	.000	.000	.000	.222
48433 Software Architecture	.000	.000	.000	.333	.200	.666
UNSW → UTS						
COMP9021 Principles of Programming	.000	1.000	.000	.666	.000	.400
COMP9311 Database Systems	1.000	1.000	.333	.666	.200	.400
COMP2511 Object-Oriented Design & Programming	.000	1.000	.000	.333	.000	.400

a graph-based ranking algorithm which represents text as a graph, where each sentence is a node. The sentence similarity is determined by interconnecting the edges based on the word token overlap between the sentences and finally ranking the sentences based on overall weights [39, 45]. A threshold value of four was required as the minimum number of sentences for a viable subject description.

STEP 2: Key-phrase extraction. The key sentences extracted from the previous step are stripped of stop words and then tokenized. Bi-grams and tri-grams are generated from the tokenized terms which are then further filtered to only select noun-phrases. (Noun phrases are assumed to be a good way of extracting key concepts from a subject description, as they tend to emphasise things that a subject actually aims to teach.)

STEP 3: WordNet based skill-weight computation. This step focuses on augmenting the BG skills returned for a subject with a new semantic weight using a WordNet based similarity measure. To achieve this we use the Wu-Palmer (wp)[13] and the Path distance (pd)[14] metrics to compute a similarity score between each BG skill and the keyphrases returned from STEP 2. The words from each of the BG skill and keyphrase were mapped to the WordNet ontology and the closest synonyms were obtained. Using the resulting synonym set an average of the wp and pd scores was computed for each BG skill against each of the keyphrases extracted. Finally, the max of the weight scores (BG or WESCM based score) was computed as the overall semantic weight score. This enabled us to keep highly weighted BG scores, but to reorder the lower ranked BG skills, with ones that the WESCM algorithm rated highly given a higher ranking.

STEP 4: Skill based curriculum mapping. In this step we followed the procedure used in Section 3.1 to construct new "WESCM vectors", using them to re-map the curriculum of both institutions. We again computed the cosine similarity each subject's WESCM vector to those of the other institutions' WESCM vectors, which resulted in new rankings for subjects that could be considered as RPL candidates.

[13]The Wu-Palmer (wp) similarity metrics work on three different factors considering the two given word concepts, $C1$ and $C2$. The first two are the depth of each of the concept which is denoted as $depth(C1)/depth(C2)$, while the third is the depth of least common subsume (LCS) of the two word concepts [30]

$$Sim_{wu} = \frac{2 \cdot depth(LCS)}{depth(C1) + depth(C2)}. \quad (1)$$

[14]The path distance is the measure of distance between two nodes in a hypernym hierarchy (i.e. the number of nodes between two words [19]).

Table 6: Reciprocal recall ($r - recall$) obtained for the new method explored in Experiment 2. Results show substantial improvement over Experiment 1 (see Table 4.)

UTS → UNSW	n	# mapped	$r - recall$
IT	132	65	.492
Medicine	295	147	.498

UNSW → UTS	n	# mapped	$r - recall$
IT	118	67	.568
Medicine	185	131	.708

3.3.2 Evaluation. The same evaluation regimen was followed as previously used for Experiment 1 (see Section 3.2). In Table 6, we see that reciprocal recall is substantially improved in both directions. The human rankings evaluation is given in Table 5, and again shows a marked improvement in results.

We see that the enhanced mappings obtained using the WESCM algorithm appear to rank more relevant subjects more highly. This can be attributed to the concept-based similarity mapping of the BG skills to the content of the subject. However, some subjects failed to show a substantial improvement (e.g. Computer Architecture was still mapped to irrelevant subjects like "Interaction Design"). This is because WordNet does not always necessarily identify the semantic relationship in all fields of knowledge that might be taught at a university. This is hardly surprising, The synsets provided by WordNet are inherited from natural language. It is quite likely that subject descriptions have a more unique structure, and so will have some semantic relations that are still being lost. This brings us to a consideration of future work.

4 DISCUSSION AND FUTURE WORK

While the results of our analysis in the previous section are promising, some important points must be noted for future work.

Firstly, much more work on the evaluation of this approach remains to be completed. The two methods so far proposed (reciprocal mappings and human evaluation) have only been attempted for two fields (IT and Medicine) and only for a mapping between two institutions. Furthermore, the human evaluation has only been completed for a small section of subjects, and requires a rigorous implementation with independent evaluators to remove any chance of potential bias (although the protocol's use of filler subjects was designed to minimise this problem). The human evaluation has concentrated upon IT topics because the team lacked the requisite medical expertise to provide judgements of the resulting mappings. This further evaluation also remains for future work. Furthermore, the IRR value obtained between the human raters was relatively low, demonstrating the difficulty of the RPL problem. More work is required to establish a ground truth dataset that could be used to evaluate our approach, a task that we reserve for the future. Nonetheless, the contribution of a preliminary evaluation framework provides the LAK community with an initial set of baseline performance figures that can be used to judge improvements in future work as it emerges. We also acknowledge that we are yet to demonstrate that our approach works for fields beyond IT and Medicine. Space limitations made it infeasible to present results from more fields of study, but our initial exploratory analysis of other curriculum mappings is promising.

We believe that our approach can be further improved by using ESCO [8, 26] to extract skill-weight computations. Future work will investigate the potential utility of enhancing our approach through the use of a space constructed using this corpus. It seems likely that adding this further complexity will integrate two descriptions of skills that have different characteristics with the concept based approach that was implemented using WordNet. It should also be possible to extend this approach with richer information about subjects. The descriptions used in the experiments reported here are very short, but university handbooks often include assessment details, learning outcomes and other rich information. Making use of this richer curriculum information would likely improve the robustness of our results, although there is likely to be a computational cost that arises with this extra processing, as well as a social transformation that is required to get more institutions publicly releasing their curriculum information.

In any event, the approach demonstrated here is highly promising, and we consider it quite likely that a decision support tool implementing this approach could assist those working to award RPL at large institutions, as long as curriculum information is made publicly available and kept up to date.

It is also possible that this form of skills based curriculum analytics could help to rectify other curriculum problems. For example, the same approach could potentially be used *within* one institution to identify redundant subjects or overlapping curriculum components. This would enable institutions to find places where students might be dissatisfied with repeating content. Automatically flagging such potential overlaps for further investigation would enable more informed decisions to be made about whether subjects should be redesigned. Similarly, this approach could find concept drift in subjects, thus helping whole of course designers to discover places where assumed content is not being taught etc. Our approach also shows promise for evaluating the quality of subject descriptions, identifying ones that are not sufficiently detailed and so potentially unhelpful to students. For example, the results of Experiment 1 has suggested that the UNSW subject descriptions were not as detailed as they might be — a point that has been confirmed by further investigation.

Finally, it is worth commenting upon our choice of a proprietary set of black box tools offered by Burning Glass (BG) for this investigation. We made the decision to use the content tagger offered by BG due to its simplicity and speed. This tool returns skills lists quickly, making it feasible to rapidly label the curriculum offerings of an entire institution. We believe it is possible to construct similar semantic spaces from first principles using NLP and an open ontology (e.g. ESCO), but that approach would likely be slower, and require far more expertise than most central teams possess. We thus consider our off the shelf approach more practical and scalable. Ideally services will start to emerge that perform the analysis we have prototyped here, making this approach even easier to use.

5 CONCLUSIONS

As a pain point for many institutions, it is somewhat surprising that more work on automating the mapping of curriculum into a common format has not been completed by the Learning Analytics and Knowledge (LAK) community. This paper has provided an initial step towards automating a process (RPL) that is well known to be onerous, resource intensive, and time consuming, and yet is going to become increasingly necessary in an era of lifelong learning.

Returning to the Research Questions that motivated our study, we find that the discussion of Section 3 has provided us with a number of insights. Firstly, in responding to RQ1, we have found a simple way to make use of an off the shelf tool to construct a semantic space that shows promise for supporting institutions with the RPL process. Few institutions have central teams with expertise in NLP and Information Retrieval, so finding a fairly simple way in which this problem might be semi-automated is a highly promising result. Implementing such an approach at scale would require widely accessible curriculum documentation, but is in principle possible. Better still, we could envisage a scenario where institutions might choose to publicly share their curriculum vectors in order to facilitate RPL. As we have explored RQ2, we have started to learn more about which approaches show most promise for supporting the automation of RPL. It seems most likely that a hybrid approach will yield the best results, where a collection of skills based ontologies are augmented with semantic information about concepts and content taught in a subject.

In summary, this paper has provided the LA community with a new avenue for supporting institutions as they seek to improve their curriculum structures. We think that *skills based curriculum analytics* shows considerable promise for delivering LA that provides rich insights about how we might improve the curriculum, and hence the learning environments, of our students.

ACKNOWLEDGMENTS

We acknowledge the support of Burning Glass in provisioning the API tools that were used in this study.

REFERENCES

[1] 2015. *Australia's future workforce?* Technical Report. Committee for Economic Development of Australia (CEDA). Available at: http://www.ceda.com.au/research-and-policy/policy-priorities/workforce.

[2] 2016. *The future of jobs: Employment, skills and workforce strategy for the fourth industrial revolution.* Technical Report. World Economic Forum (WEF), Geneva, Switzerland. Available at: http://reports.weforum.org/future-of-jobs-2016/.

[3] ACM/IEEE-CS Joint Task Force on Computing Curricula. 2013. *Computer Science Curricula 2013.* Technical Report. ACM Press and IEEE Computer Society Press. https://doi.org/10.1145/2534860

[4] LM Aroyo, G-J Dolog, M Kravcik, A Naeve, M Nilsson, F Wild, et al. 2006. Interoperability in personalized adaptive learning. *Educational Technology & Society* 9, 2 (2006), 4–18.

[5] Alan Mark Berg, Jiri Branka, and Gábor Kismihók. 2018. Combining Learning Analytics with Job Market Intelligence to Support Learning at the Workplace. In *Digital Workplace Learning*. Springer, 129–148.

[6] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences.* MIT press.

[7] Yiling Dai, Yasuhito Asano, and Masatoshi Yoshikawa. 2016. Course Content Analysis: An Initiative Step toward Learning Object Recommendation Systems for MOOC Learners.. In *Proceedings of the 9th International Conference on Educational Data Mining (EDM2016)*, Min Chi Tiffany Barnes and Mingyu Feng (Eds.). 347–352.

[8] Johan De Smedt, Martin le Vrang, and Agis Papantoniou. 2015. ESCO: Towards a Semantic Web for the European Labor Market.. In *LDOW@ WWW*.

[9] Eric Durant, John Impagliazzo, Susan Conry, Robert Reese, Herman Lam, Victor Nelson, Joseph Hughes, Weidong Liu, Junlin Lu, and Andrew McGettrick. 2015. CE2016: Updated computer engineering curriculum guidelines. In *2015 IEEE Frontiers in Education Conference (FIE)*. IEEE, 1–2.

[10] Dragan Gašević, Jelena Jovanović, and Vladan Devedžić. 2007. Ontology-based annotation of learning object content. *Interactive Learning Environments* 15, 1 (2007), 1–26.

[11] Luit Gazendam, Christian Wartena, and Rogier Brussee. 2010. Thesaurus based term ranking for keyword extraction. In *2010 Workshops on Database and Expert Systems Applications*. IEEE, 49–53.

[12] Andrew Gibson, Kirsty Kitto, and Jill Willis. 2014. A cognitive processing framework for learning analytics. In *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge*. ACM, 212–216.

[13] Richard Gluga, Judy Kay, and Raymond Lister. 2012. PROGOSS: Mastering the curriculum. In *Proceedings of The Australian Conference on Science and Mathematics Education (formerly UniServe Science Conference)*.

[14] Alexander Grech and Anthony F. Camilleri. 2017. *Blockchain in Education.* Technical Report. Joint Research Centre (JRC). https://doi.org/doi:10.2760/60649 Available at: http://publications.jrc.ec.europa.eu/repository/bitstream/JRC108255/jrc108255_blockchain_in_education(1).pdf.

[15] Dai Griffiths, Tore Hoel, and Adam Cooper. 2016. *D7.4 Learning Analytics Interoperability: Requirements, Specifications and Adoption.* Technical Report. LACE.

[16] Judy Harris, Joy Van Kleef, and Christine Wihak. 2014. *Handbook of the recognition of prior learning: Research into Practice.* National Institute of Adult Continuing Education. https://www.learningandwork.org.uk/resource/handbook-of-the-recognition-of-prior-learning/

[17] Daniel T Hickey and James E Willis. 2017. *Where Open Badges Appear to Work Better: Findings from the Design Principles Documentation Project.* Technical Report. http://www.badgenumerique.com/wp-content/uploads/2017/08/DPD-Project-Final-Report-Dan-Hickey-Willis-May-2017.pdf

[18] Jelena Jovanovic, Dragan Gasevic, Christopher Brooks, Vladan Devedzic, Marek Hatala, Timmy Eap, and Griff Richards. 2007. Using semantic web technologies to analyze learning content. *IEEE Internet Computing* 11, 5 (2007).

[19] Jaap Kamps, Maarten Marx, Robert J Mokken, Maarten De Rijke, et al. 2004. Using WordNet to measure semantic orientations of adjectives.. In *LREC*, Vol. 4. Citeseer, 1115–1118.

[20] Kornraphop Kawintiranon, Peerapon Vateekul, Atiwong Suchato, and Proadpran Punyabukkana. 2016. Understanding knowledge areas in curriculum through text mining from course materials. In *2016 IEEE International Conference on Teaching, Assessment, and Learning for Engineering (TALE)*. IEEE, 161–168.

[21] Judy Kay. 2008. Lifelong learner modeling for lifelong personalized pervasive learning. *IEEE Transactions on Learning Technologies* 1, 4 (2008), 215–228.

[22] James Keevy and B Chakroun. 2015. *Level-setting and recognition of learning outcomes: The use of level descriptors in the twenty-first century.* Technical Report. http://unesdoc.unesco.org/images/0024/002428/242887e.pdf

[23] Kirsty Kitto, Julieanne Cutrupi, Mark Philips, Gabrielle Gardiner, Moein Ghodrati, and Simon Buckingham Shum. 2019. The connected university: connectedness learning across a lifetime. In *Higher Education and the Future of Graduate Employability: A Connectedness Learning Approach*, R. Bridgstock and N. Tippett (Eds.). Edward Elgar Publishing, Cheltenham, 230–249.

[24] Kirsty Kitto and Simon Knight. 2019. Practical ethics for building learning analytics. *British Journal of Educational Technology* (2019). https://onlinelibrary.wiley.com/doi/abs/10.1111/bjet.12868

[25] Klaus Krippendorff. 2004. *The content analysis reader* (2 ed.). Sage.

[26] Martin le Vrang, Agis Papantoniou, Erika Pauwels, Pieter Fannes, Dominique Vandensteen, and Johan De Smedt. 2014. Esco: Boosting job matching in europe with semantic interoperability. *Computer* 47, 10 (2014), 57–64.

[27] Kevin Lund and Curt Burgess. 1996. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior research methods, instruments, & computers* 28, 2 (1996), 203–208.

[28] Linda Marshall. 2012. A Comparison of the Core Aspects of the ACM/IEEE Computer Science Curriculum 2013 Strawman Report with the Specified Core of CC2001 and CS2008 Review. In *Proceedings of Second Computer Science Education Research Conference (CSERC '12)*. ACM, New York, NY, USA, 29–34. https://doi.org/10.1145/2421277.2421281

[29] Yoshitatsu Matsuda, Takayuki Sekiya, and Kazunori Yamaguchi. 2018. Curriculum Analysis of Computer Science Departments by Simplified, Supervised LDA. *Journal of Information Processing* 26 (2018), 497–508.

[30] Rada Mihalcea, Courtney Corley, Carlo Strapparava, et al. 2006. Corpus-based and knowledge-based measures of text semantic similarity. In *Aaai*, Vol. 6. 775–780.

[31] Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*. 404–411.

[32] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.

[33] Rohit Muthyala, Sam Wood, Yi Jin, Yixing Qin, Hua Gao, and Amit Rai. 2017. Data-driven Job Search Engine Using Skills and Company Attribute Filters. *arXiv preprint arXiv:1709.05050* (2017).

[34] Zachary A Pardos, Zihao Fan, and Weijie Jiang. 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 (2019), 487–525.

[35] Daniel Ramage, David Hall, Ramesh Nallapati, and Christopher D Manning. 2009. Labeled LDA: A supervised topic model for credit attribution in multi-labeled corpora. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*. Association for Computational Linguistics, 248–256.

[36] Takayuki Sekiya, Yoshitatsu Matsuda, and Kazunori Yamaguchi. 2015. Curriculum analysis of CS departments based on CS2013 by simplified, supervised LDA. In *Proceedings of the Fifth International Conference on Learning Analytics And Knowledge*. ACM, 330–339.

[37] Sam Sellar, Greg Thompson, and David Rutkowski. 2017. *The Global Education Race: Taking the Measure of PISA and International Testing.* Brush Education.

[38] David Spencer, Matthew Riddle, and Bernadette Knewstubb. 2012. Curriculum mapping to embed graduate capabilities. *Higher Education Research & Development* 31, 2 (2012), 217–231. https://doi.org/10.1080/07294360.2011.554387 arXiv:https://doi.org/10.1080/07294360.2011.554387

[39] George Tsatsaronis, Iraklis Varlamis, and Kjetil Nørvåg. 2010. SemanticRank: ranking keywords and sentences using semantic graphs. In *Proceedings of the 23rd International Conference on Computational Linguistics*. Association for Computational Linguistics, 1074–1082.

[40] Katrien Verbert, Joris Klerkx, Michael Meire, Jehad Najjar, and Erik Duval. 2004. Towards a global component architecture for learning objects: An ontology based approach. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*. Springer, 713–722.

[41] Ruben Verborgh and Miel Vander Sande. 2020. The Semantic Web identity crisis: in search of the trivialities that never were. *Semantic Web Journal* (2020). https://ruben.verborgh.org/articles/the-semantic-web-identity-crisis/

[42] Jinghua Wang, Jianyi Liu, and Cong Wang. 2007. Keyword extraction based on pagerank. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 857–864.

[43] Ian H Witten and Olena Medelyan. 2006. Thesaurus based automatic keyphrase indexing. In *Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL'06)*. IEEE, 296–297.

[44] Law Sheng Xun, Swapna Gottipati, and Venky Shankararaman. 2015. Text-mining approach for verifying alignment of information systems curriculum with industry skills. In *Information Technology Based Higher Education and Training (ITHET), 2015 International Conference on*. IEEE, 1–6.

[45] Jen-Yuan Yeh, Hao-Ren Ke, and Wei-Pang Yang. 2008. iSpreadRank: Ranking sentences for extraction-based summarization using feature weight propagation in the sentence similarity network. *Expert Systems with Applications* 35, 3 (2008), 1451–1462.