

RiPPLE: A Crowdsourced Adaptive Platform for Recommendation of Learning Activities

Hassan Khosravi¹, Kirsty Kitto², Joseph Jay Williams³

Abstract

This paper presents a platform called RiPPLE (Recommendation in Personalised Peer-Learning Environments) that recommends personalized learning activities to students based on their knowledge state from a pool of crowdsourced learning activities that are generated by educators and the students themselves. RiPPLE integrates insights from crowdsourcing, learning sciences, and adaptive learning, aiming to narrow the gap between these large bodies of research while providing a practical platform-based implementation that instructors can easily use in their courses. This paper provides a design overview of RiPPLE, which can be employed as a standalone tool or embedded into any learning management system (LMS) or online platform that supports the Learning Tools Interoperability (LTI) standard. The platform has been evaluated based on a pilot in an introductory course with 453 students at The University of Queensland. Initial results suggest that the use of the RiPPLE platform led to measurable learning gains and that students perceived the platform as beneficially supporting their learning.

Notes for Practice

- Adaptive learning systems (ALSs) dynamically adjust the level or type of instruction based on individual student abilities or preferences to provide a customized learning experience.
- However, ALSs require access to a large repository of learning resources, which are commonly created by domain experts. This makes them expensive to develop and challenging to scale across different domains.
- We introduce an ALS called RiPPLE that uses a crowdsourcing approach to engage students in creating, moderating, and evaluating learning resources. This both reduces the cost of content generation and has the potential to foster higher-order learning across many domains.
- Initial results suggest that using RiPPLE may lead to measurable learning gains and that students perceived the platform as beneficially supporting their learning.

Keywords

Adaptive learning, crowdsourcing learning content, recommender systems

Submitted: 28.11.18 — **Accepted:** 03.09.2019 — **Published:** 13.12.2019

Corresponding author ¹ Email: h.khosravi@uq.edu.au Address: Institute for Teaching and Learning Innovation, The University of Queensland, St Lucia QLD 4072, Australia. ORCID: 0000-0001-8664-6117

² Address: Connected Intelligence Centre, University of Technology Sydney, 15 Broadway, Ultimo NSW 200, Australia. ORCID: 0000-0001-7642-7121

³ Address: Department of Computer Science, University of Toronto, 40 St. George Street, Room 7224, Toronto, ON M5S 2E4, Canada

1. Introduction

Educators continue to face significant challenges in providing high-quality postsecondary instruction in large online or on-campus classes (Mulryan-Kyne, 2010). A significant portion of these challenges emerge from high levels of diversity in learners' academic ability (Banks et al., 2005). Teach-to-the-middle instruction is most commonly used; however, this method does not meet the needs of learners who differ significantly from the norm. Adaptive learning systems (ALSs) (Park & Lee, 2004) provide a potential solution to this problem. ALSs make use of data about students, learning processes, and learning products to dynamically adapt learning content and activities to suit students' individual abilities or preferences. A growing body of knowledge provides evidence about the effectiveness of ALSs in supporting student learning (Anderson, Corbett, Koedinger, & Pelletier, 1995; VanLehn, 2011; Ma, Adesope, Nesbit, & Liu, 2014).

The most successful adaptive platforms tend to focus on a specific domain (see Section 2) and also require a tremendous investment of time by experts during the development of curriculum content, which makes it difficult to scale these approaches across many domains. As a viable alternative, we see great promise in trying to leverage ideas from crowdsourcing and learning science to create adaptive systems by having students themselves generate content that can be adaptively served. The benefits of engaging students in content creation are twofold. First, students transform from passive recipients of content to active creators of course material. Previous studies have reported that placing the responsibility of content creation in the hands of students reinforces and deepens their understanding of the course content (Barak & Rafaeli, 2004; Draper, 2009), highlights the significance of representing their work in a clear and logical fashion (Denny, Hamer, Luxton-Reilly, & Purchase, 2008), encourages reflection on the course objectives (Purchase, Hamer, Denny, & Luxton-Reilly, 2010), and enhances their conceptual understanding (Bates, Galloway, & McBride, 2012). Second, the creativity of students themselves is used to develop large repositories of learning resources. Previous studies have demonstrated that students can create high-quality learning resources that meet rigorous judgmental and statistical criteria (Walsh, Harris, Denny, & Smith, 2018; Tackett et al., 2018; Denny, Hamer, & Luxton-Reilly, 2009; Galloway & Burns, 2015; Bates, Galloway, Riise, & Homer, 2014).

This paper discusses the design and implementation of a platform called RiPPLE (Recommendation in Personalised Peer Learning Environments)¹ that adopts a students-as-partners approach (Matthews, 2017) to provide adaptive learning functionalities while enabling learners to be more cognitively involved in their learning. The platform provides easily accessible data, which will help the learning analytics (LA) community research a number of topics of interest, including the type of content generated by students for their peers, the way students interact with peer-generated questions, and student use of recommendation systems and open learner models to direct their learning.

The platform was initially evaluated based on a pilot in an introductory course with 453 students at The University of Queensland. Initial results suggest that the use of the RiPPLE platform may lead to significant and measurable learning gains; the mid-semester score of the experimental group using the platform was approximately 8 percentage points (total effect size: $d = 0.54$ (Cohen, 1992), $p < .001$) higher than the control group that was not using the platform. Additionally, a student survey ($N = 55$) shows that a large fraction of the respondents believe that RiPPLE helped them to study more efficiently and effectively.

In what follows, we introduce RiPPLE to the LA community. Section 2 provides a background for the use of adaptive learning and crowdsourcing in education. Section 3 introduces the RiPPLE platform, and Section 4 evaluates the platform in an authentic learning context. Finally, Section 5 presents concluding remarks. We start with a consideration of ALSs and how they can help students navigate their way through large amounts of learning content.

2. Background: Adaptive Learning and Crowdsourcing in Education

ALSs (Park & Lee, 2004) dynamically adjust the level or type of instruction based on individual student abilities or preferences to provide an efficient, effective, and customized learning experience. At a high level of generality, there are two main classes of ALSs. The first class, commonly referred to as intelligent tutoring systems (ITSs) (Anderson, Boyle, & Reiser, 1985), uses AI-based algorithms to replicate the support that is often provided by a tutor by providing personalized step-by-step guidance in solving a problem. Carnegie Learning's MATHiaU (Ritter, Carlson, Sandbothe, & Fancsali, 2015) is an established example of this class of ALSs. The second class of ALSs focuses on adaptively recommending learning activities from a large repository of learning resources to a student to match their current learning ability. Pearson's MyLabs (using Knewton (Jose, 2016) for its adaptive functionality) and McGraw-Hill's LearnSmart and ALEKS (Falmagne, Cosyn, Doignon, & Thiéry, 2006) are established examples of this class of ALSs. RiPPLE is also representative of this second class.

A consistent and growing body of knowledge provides evidence about the effectiveness of both classes of ALSs (Anderson et al., 1995; VanLehn, 2011; Ma et al., 2014). For example, a comprehensive meta-review by VanLehn (2011) reported that on average, students using ITSs have a learning gain effect size of $d = 0.76$ relative to classroom teaching without tutors. For ALSs that focus on recommending learning resources, Yilmaz (2017) and Mojarad, Essa, Mojarad, and Baker (2018) have reported improvement on student performance while using ALEKS or the popular ASSISTments Ecosystem (Heffernan & Heffernan, 2014).

At a high level of generality, many ALSs rely on the following interacting components (Essa, 2016):

1. *Domain model*: A knowledge space modelling what the students need to know. The domain model is commonly presented as a set of knowledge units that are “elementary fragments of knowledge for the given domain” (Brusilovsky, 2012).
2. *Learner model*: An abstract representation of students, often “overlying” their knowledge state on the knowledge space defined in the domain model (Brusilovsky & Millán, 2007). The learner model may estimate a student's ability level on

¹<http://ripplelearning.org>

different knowledge units based on their performance and interactions with the system. Importantly, open learner models (Bull & Kay, 2010), which are learner models that are externalized and made accessible to students or other stakeholders, can be particularly effective in helping students learn (Bodily et al., 2018).

3. *Content repository*: A repository of learning resources that may include assessment-based and learning-based items designed to help the learner acquire the knowledge represented by the domain model. Each learning resource is tagged with knowledge units defined in the domain model.
4. *Recommender engine*: An engine that utilizes information from the learner model and the content repository to select learning activities for each student that will be most likely to advance their learning of the domain knowledge.

Commonly, ALSs are developed using the publisher model (Oxman, Wong, DV X Innovations, & DeVry Education Group, 2014). In this model, the platform is designed with pre-existing learning activities, often based on textbooks from a publisher. Pearson's MyLabs (using Knewton (Jose, 2016) for its adaptive functionality) and McGraw-Hill's LearnSmart and ALEKS (Falmagne et al., 2006) are established examples of this model. The publisher model has been successful in K-12, where course content follows a simpler structure and often has to comply with national standards. However, higher education has been slow to embrace these systems, with adoption mostly restricted to research projects (Essa, 2016). The focus on specific restricted domains, limited flexibility for educators to tailor the learning activities to their context, and the high costs associated with the use of these platforms have all contributed to their low adoption in higher education.

Responding to these limitations, an alternative has been established, referred to as the platform model (Oxman et al., 2014). The platform model provides a content-agnostic system infrastructure that enables educators to develop and author the content of their course. Smart Sparrow (Smart Sparrow, 2016) and many learning management systems (LMSs), such as Desire2Learn, Loudcloud, and edX, which incorporate adaptive functionality into their course-building tools, follow this model. The platform model is relatively new and mostly suffers from an operational limitation rather than a technological one; implementing adaptivity in a course requires a large number of new learning activities and object tagging, introducing significant overhead for teaching staff in both time and training. To overcome the limitations of both of these models, RiPPLE leverages ideas from crowdsourcing in education (Solemon, Ariffin, Din, & Anwar, 2013) by having students themselves generate and evaluate content that can then be adaptively served.

The use of crowdsourcing in education alongside insights from the students-as-partners approach (Matthews, 2017) makes way for respectful, mutually beneficial learning partnerships where students and staff work together. Successful examples of such partnerships have led to co-creation of curricula (Bovill, 2013); marking criteria (Meer & Chapman, 2014); and assessment items via the popular PeerWise platform (Denny et al., 2008), which has inspired RiPPLE. The use of crowdsourcing in ALSs is also beginning to receive attention. For example, Heffernan et al. (2016) propose employing crowdsourcing within the popular ASSISTments platform; Williams et al. (2016) present an Adaptive eXplanation Improvement System (AXIS) that uses crowdsourcing to generate, revise, and evaluate explanations as learners solve problems; and Karataev and Zadorozhny (2017) propose a framework that combines concepts of crowdsourcing, online social networks, and adaptive systems to provide personalized learning pathways for students. However, this preliminary work is yet to realize the full potential offered by crowdsourcing in ALSs or more broadly in education. An important motivator in developing RiPPLE is to provide support for ethical and low-cost educational research on the use of crowdsourcing in education and ALSs.

3. The RiPPLE Platform

In this section, we provide an overview of the main functionalities of RiPPLE. Section 3.1 provides an overview of how a new RiPPLE offering can be created. Section 3.2 presents how learning activities are created, attempted, and evaluated in RiPPLE. Section 3.3 then introduces the open learner model used by RiPPLE, which enables students to view an abstract representation of their knowledge state. Section 3.4 demonstrates the available features for resource selection and recommendation, and Section 3.5 describes the information available under personal profiles and course reports. Prototypes of the platform, demonstrating the student view and the instructor view, and further information are available on the web page of the platform².

3.1 Creation of a RiPPLE Offering

RiPPLE can be integrated into many popular LMSs (including Blackboard, Moodle, and Canvas) using the Learning Tools Interoperability (LTI) standard. Once integrated into an LMS, RiPPLE can be added as an LTI tool (via a link) to any course within that system. RiPPLE supports two types of roles for users: instructors and students. The role of a user is determined using the LTI standard based on their role in the LMS course that is hosting RiPPLE. For example, assuming that RiPPLE is hosted in a course on Blackboard, users with the Instructor, Teaching Assistant, or Grader role in that course would be given the

²<http://ripplelearning.org>

instructor role on RiPPLE. Similarly, users with the Student, Guest, or Observer role in that course on Blackboard would be given the student role on RiPPLE. Once a tool link has been added to an LMS course, a user with the instructor role can click on the link to create a RiPPLE offering for the course.

The university name, course code, course name, course semester, and teaching start date are all automatically completed based on the information of the course captured in the LMS. The instructor is required to enter a set of topics, representing the knowledge units, which define the knowledge space of a course. They can do this by importing topics from other RiPPLE offerings, creating new ones, or modifying an existing list of topics (e.g., renaming, deleting, or changing the order of the knowledge units). The list of topics may be altered throughout the semester.

RiPPLE can also be used as a standalone system without integration into an LMS. In this case, the instructor would (1) create a personal account in order to create a RiPPLE offering; (2) manually enter the university name, course code, course name, course semester, teaching start date, and course topics; and (3) add other users (e.g., students and instructors) to the created RiPPLE offering. Two methods are supported for adding other users to an offering: by invitation, where users receive an email to join the offering, or by self-enrolment, where users receive an access code that can be used to manually sign up for the course.

Once a RiPPLE offering has been created, an instructor can import resources from other RiPPLE offerings. This enables instructors to import resources from their past offerings as well as share resources with other instructors who are teaching similar courses. The resource import page allows instructors to search for resources from other offerings based on many options, including university name, course name, offering ID, topics, resource rating, resource type, and keywords.

3.2 Content Creation and Evaluation

RiPPLE integrates insight from classical and contemporary models of learning to engage students in activities across many of the higher-level objectives of the cognitive domain of Bloom's taxonomy (Bloom, Englehart, Furst, Hill, & Krathwohl, 1956). For example, RiPPLE enables students to author learning activities that help them think carefully about the concepts and learning outcomes of the course, formulate distractors for multiple-choice questions (MCQs) (which requires them to analyze misconceptions their peers might have), and explain their understanding of a concept to other students who chose the wrong answer. More recent models of learning also provide evidence that engaging students in elaborating on the learning content (Pressley et al., 1992; King, 1992), such as developing evaluative judgment (Tai, Ajjawi, Boud, Dawson, & Panadero, 2018), can increase learning. In RiPPLE, these elaborations can take many forms, including clarifying an idea, constructing an original explanation of a concept, judging the quality of an activity, or comparing and contrasting an activity with other activities.

RiPPLE enables students and instructors to create, attempt, and evaluate a wide range of learning resources that are tagged with one or more topics pre-assigned by users who have instructor status in the course. The resources that can be created currently include MCQs, worked examples, and general notes. Creating an MCQ includes (1) developing the body of the question, (2) tagging it with one or more topics, (3) developing the multiple-choice answers, (4) nominating the correct answer, and (5) writing an explanation of the solution. Creating a worked example includes (1) developing the body of the question, (2) tagging it with one or more topics, (3) developing a series of steps that work through answering the question, and (4) developing a final solution for the question. Creating a general note is more flexible and can be used to develop a broad range of learning resources, including mini-lessons (teaching a concept or applying a theoretical concept to real-world examples), summary notes, and cheat sheets. In all cases, text, tables, images, videos, and scientific formulas can be used to create learning resources. Authors can also view their learning resources before submitting and edit their submitted resources at a later time using RiPPLE.

An important assumption made by RiPPLE is that students, as non-experts, can create high-quality resources. While we were not able to find longitudinal or meta-reviews that support this claim, there seems to be adequate evidence suggesting that students can create high-quality learning resources that meet rigorous judgmental and statistical criteria (Walsh et al., 2018; Tackett et al., 2018; Denny et al., 2009; Galloway & Burns, 2015; Bates et al., 2014). In fact, students as authors of learning resources may have an advantage over instructors: they can use their knowledge of their own previous misconceptions to create resources so there is less chance that they will suffer from an expert blind spot. However, it is very likely that some of the learning resources developed by students may be ineffective, inappropriate, or incorrect (Bates et al., 2014). As such, to effectively utilize resources developed by students, a selection and moderation process is needed to identify the quality of each resource. RiPPLE provides multiple options for moderation of the created resources. One of these options is "staff moderation," where the created resources are moderated by the instructors before they are publicly released and added to the repository of the learning activities for the offering. However, this may not be feasible in large classes. Other options for moderation in RiPPLE rely on the collective wisdom of the crowd and methods that are generally used to review academic articles (e.g., the "competent student moderation" option) or content moderation on social networks (e.g., "flagging inappropriate content"). This raises an interesting research question: can non-experts accurately evaluate the quality of a learning resource? A paper by

Whitehill, Aguerrebere, and Hylak (2019) provides evidence that non-expert subjective opinions may be able to accurately determine the quality of a learning resource, and that machine-learning algorithms may be used to infer the reliability of an individual's opinion, which further increases the accuracy of the results. Further investigations about this topic are underway by the authors. Our initial results are aligned with the findings of Whitehill et al. (2019).

Once a resource has been added to a course repository, students and instructors from that offering can view, attempt, and then evaluate the resources. For MCQs, once a user has attempted the question, they can view the right answer, the distribution of how the MCQ has been answered by their peers, and the explanation provided for the question. For all resources, users can identify the author and view the current rating and comments made about the resource. Users can add their own comments and rate the effectiveness of the resource.

3.3 Learner Modelling

In its current state, RiPPLE makes use of the Elo rating system, which was developed initially to rate chess players. In the educational context, the Elo rating system is employed to conduct a paired comparison among students and questions as competitors. If the question is answered correctly, the student's rating increases and the rating of the question decreases. If the question is answered incorrectly, the student's rating decreases and the rating of the question increases. The update to the ratings is proportional to the difference between the ratings of the student and the question. If a student who is highly rated correctly answers a low-rated question, then only a few rating points will be transferred from the low-rated question to the student. However, if a student with a low rating correctly answers a high-rated question, then many rating points will be transferred. One of the main benefits of using the Elo rating system in educational settings is its simplicity and self-correcting behaviour (Pelánek, 2016). In the standard Elo-based learner model, both students and items are considered as identical rivals and are both modelled by one parameter. In this model, the student parameter indicates the global proficiency level of the student on the entire domain. We make use of a multivariate Elo-based model (Abdi, Khosravi, Sadiq, & Gašević, 2019) that uses independent parameters to model a student's proficiency level on each individual concept in the domain. This is commonly referred to as using an "overlying" learner model, which is a widely accepted approach in modelling learners in ALSs (Brusilovsky & Millán, 2007; Essa, 2016).

Open learner models are learner models that are externalized and made accessible to students or other stakeholders, such as instructors, often through visualization, as an important means of supporting learning (Bull & Kay, 2010). They have been integrated into a variety of learning technologies, such as learning dashboards (Bodily et al., 2018), ITSs (Ritter, Anderson, Koedinger, & Corbett, 2007), and ALSs (Morrison & DiSalvo, 2014), to help students and instructors monitor, reflect on, and regulate learning (Bull & Kay, 2010).

Figure 1 shows one of the main pages of RiPPLE, and in particular its open learner model in the top panel. This interactive visualization widget enables students to view an abstract representation of their knowledge state based on the knowledge units present in the domain model. The proposed OLM was designed based on the following two principles: (1) The OLM and the recommendation results need to be placed close to one another on the interface. This was to help students better understand the rationale behind the recommendations of the system. (2) The utilized visualization must be easy to understand by the majority of the users. A range of visualizations — including bar charts, line charts that demonstrate progress over time, zoomable Treemap, and Topic Dependency Models (Khosravi & Cooper, 2018; Cooper & Khosravi, 2018) — have been incorporated and tested by earlier versions of RiPPLE. Based on the results of multiple rounds of usability tests, bar charts and line charts have been adopted in the latest version of the system. Accordingly, the "Visualisation Data" drop-down enables students to select between two models visualizing their knowledge state: viewing their current knowledge state and tracking changes to their knowledge state over time. The figure demonstrates how a student can view their current knowledge state. Each bar represents the competency of the student in one of the knowledge units of the course. The colour of the bars, determined by the underlying algorithm modelling the learner, categorizes competencies into three levels: red demonstrates inadequate competency in a topic, yellow demonstrates adequate competency with room for improvement, and blue demonstrates mastery in a knowledge unit. The model also shows the average competency of the entire cohort over each knowledge unit using a line graph. Abdi et al. (2019) show that opening the model in RiPPLE increases students' motivation to use the platform and increases their trust in the recommendations provided by the platform. RiPPLE's open learner model also provides additional insight for instructors on individual student-level or class-level gaps and competencies that can be used to improve item and course design.

3.4 Resource Selection and Recommendation

The bottom part of Figure 1 presents the current graphical interface used for resource selection and recommendation in RiPPLE. A set of filters is available to help students search the resource repository. The "Filter Types & Topics" option enables students to filter the resources included in the results based on their type (i.e., MCQs, worked examples, or notes) and tagged topics. This may help students target their learning toward particular topics or types of resources. The "Filter Resources" option enables

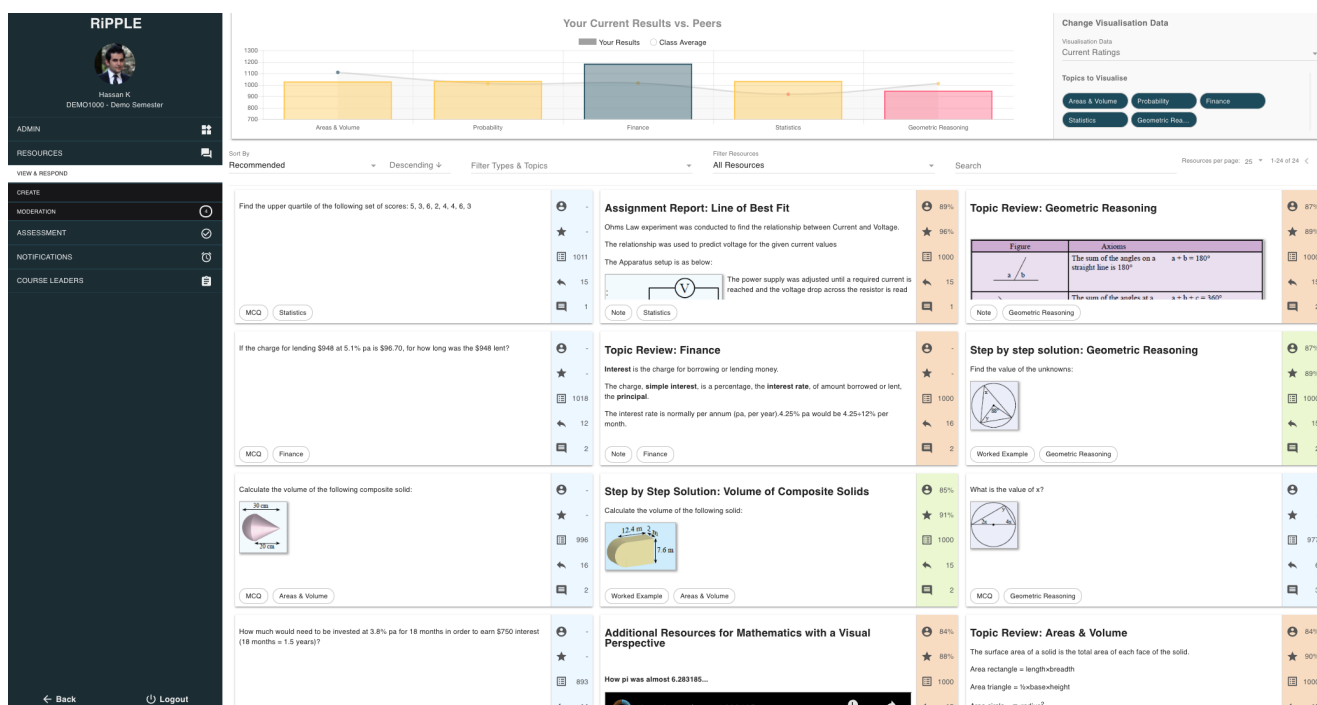


Figure 1. Overview of open learner model and the resource selection and recommendation features in RiPPLE.

students to search for previously attempted (for revision), not attempted (for engaging with new content), and incorrectly answered (for reviewing content they previously found challenging) resources. This filter also enables students to view deleted resources they created themselves. This includes resources that they have personally deleted as well as those that have been deleted through moderation. Being able to access their deleted resources, students can revise and resubmit these resources. The “Search” option enables students to search for resources based on specific content that may be present in the resource. The results of the search are presented as a list of learning resource cards that satisfy the specified filters.

The “Sort By” option allows students to sort the returned resources based on their difficulty, quality, number of responses, or personal fit (“Recommended”). If “Recommended” is selected, the platform sorts the resources based on their learning benefits to the student. This helps students make a good decision about what to study next (Biggs, 2012), while keeping them in the decision-making loop, thus maintaining autonomy. In its current state, RiPPLE employs the collaborative filtering algorithm proposed by Khosravi, Cooper, and Kitto (2017) to create personalized recommendations for individual students that address their interests and their current knowledge gaps. The results reported in Khosravi et al. (2017) suggest that this algorithm outperforms baseline recommender systems in recommending learning resources.

Each resource card includes an overview of the resource content, the topics associated with the resource, and a sidebar in which the first icon shows the personal fit of the resource for the user (approximated by the recommendation engine). The second icon shows the general effectiveness of the resource (based on users’ ratings of the resource). The third icon represents the difficulty level of the resource (approximated by the multivariate Elo-based model (Abdi et al., 2019)). The fourth icon shows the number of times the resource has been attempted, and the fifth icon presents the number of comments that have been written about the resource. Clicking on the resource card will take the user to another page that allows them to attempt and rate the resource.

3.5 Student Profiles and Course Reports

It is increasingly being recognized that educational tools and technologies should not aim to replace instructors but rather provide support tools to enable instructors to improve their teaching practices (Collins & Halverson, 2018). As such, it is important for ALSs and learning technologies more broadly to provide rich and timely actionable insights to instructors so that they can best manage their class within the context of their own course (Heffernan & Heffernan, 2014). It is equally important to provide rich feedback for students so that they can take ownership of their learning (Kitto, Lupton, Davis, & Waters, 2017). In RiPPLE, student profiles and course reports are designed to provide rich analytics for instructors and students.

Student Profiles Each student is provided with a personal profile that includes information on their achievements, engagement, and knowledge state. Instructors have access to the profiles of all of the students enrolled in their RiPPLE offering, which can help them identify the learning needs of each student. RiPPLE uses badges to help students track their progress. Students can achieve badges in two broad categories: “Engagement Badges” and “Competency Badges.”

The engagement levels of students on a variety of tasks are presented using a visualization widget that enables them to use Kiviat diagrams, more informally known as radar charts, to compare their engagement against their peers. Kiviat diagrams have been used extensively in visualizing educational dashboards (e.g., see May, George, and Prévôt (2011)) to display multivariate observations with an arbitrary number of variables (Chambers, Cleveland, Kleiner, & Tukey, 1983). Currently, engagement is computed based on the number of resources authored, resources answered, resources rated, and achievements earned.

Course Reports RiPPLE allows instructors to download a set of course reports based on data collected by RiPPLE in their offering. These reports provide additional information about students, resources, comments, knowledge units, and students’ attempts at resources. The data sets downloadable with these reports may be used to efficiently conduct educational research on many topics, including crowdsourcing, learner modelling, and recommender systems, with the consent of the users, at a low cost. Upon the first use of the platform, users are presented with a consent form that seeks their permission to use their data to improve our understanding of the learning process and to evaluate the effectiveness of the recommended content. RiPPLE enables users to change their response at any time. All users, regardless of their response, can use RiPPLE, and only data collected from users that have provided consent and have never changed their response will be used for research purposes.

4. Evaluation

An evaluation of the platform has been conducted using the following research questions:

- What measurable learning gains can be found when students use RiPPLE?
- What evidence can be found that RiPPLE supports students in identifying resources that they think are effective?
- Do students perceive RiPPLE as beneficial to their learning?

To answer these questions, RiPPLE was piloted in an on-campus course on relational databases with 453 students at The University of Queensland³. The course covers many topics that are generally included in an introductory course on relational databases, such as relational models, Structured Query Language (SQL), and database security. Data from the use of RiPPLE from the first five weeks of this course are used in this evaluation. To split students into the RiPPLE and non-RiPPLE groups, a threshold engagement of attempting at least 30 questions on RiPPLE was considered. Using this constraint, 215 (or 47%) of the students were assigned to the RiPPLE group, and the remaining 238 students (or 53%) were assigned to the non-RiPPLE group. Students in the RiPPLE group authored 351 MCQs (an average of 1.6 questions per student) and made 20,540 attempts (an average of 96 attempts per student).

4.1 Measuring Learning Gains

In this section, we investigate whether using RiPPLE leads to measurable learning gains. We considered randomly assigning students to use RiPPLE versus a control, but this raised practical and ethical challenges (Sullivan, 2011). In particular, since all students were in the same course, there was a high likelihood that some would realize there is a separate system, and it might conceivably be considered unfair to give one group the experience of RiPPLE over another. We therefore applied a quasi-experimental approach to analyzing students in the course who self-selected to engage (experimental group) with RiPPLE versus those who did not engage (control group).

More precisely, we hypothesized that the choice of students to engage or not with RiPPLE might be influenced by a set of their personal characteristics (covariates). As a result, baseline covariates of students in the control group may be substantially different from those in the experimental group, which may bias the results of the investigation. We used propensity score matching (PSM) (Rosenbaum & Rubin, 1983) to account for these potential differences in baseline covariates of the students. This method matches each student in the experimental group with a student from the control group such that the two students are similar based on their baseline covariates. For this experiment, current GPA, age, residency status (domestic or international), and program level (bachelor’s level or master’s level) of the students were selected as the covariates of the model. The PSM method using the potential outcomes framework, as described by Austin (2011), was used to match students. A nearest-neighbour search was used to conduct a one-to-one matching and to find the smallest distance for the match within a predefined threshold ($\epsilon = 0.05$). Samples from the experimental group that did not have a match within the specified threshold in the control group were removed.

³Approval from our Human Research Ethics Committee (#2018000125) was received for conducting this evaluation on RiPPLE.

To cater for students’ self-selection, the course used two rubrics to compute final grades. In one of the rubrics, the final exam and RiPPLE, respectively, had 40% and 10% contributions to the final grade. In the other rubric, the final exam and RiPPLE, respectively, had 50% and 0% contributions to the final grade. The final grade of a student was computed as the maximum of the values computed using the two assessment rubrics. The grade associated with RiPPLE included the following two criteria:

- *Answering and creating questions.* Students participated in a total of four rounds of authoring and answering questions. They received a total of two marks for each round of participation, once they correctly answered 10 questions (one mark) and created at least one correct and effective question (one mark). The effectiveness of questions was evaluated by the teaching staff and student peers.
- *Overall rating.* The overall rating of a student had a 2% contribution to their grade. This was computed as an average rating calculated across all of the knowledge units set by the instructor. The mark associated with their overall rating was computed by $Min\left(Max\left(0, \frac{rating-1000}{100}\right), 2\right)$. The Min and Max functions are used to ensure a mark between 0 and 2.

Grades from a midterm examination taken during week 7 were used to compare the learning gains of the control and experimental groups. The midterm contributed 15% toward the final grade of the course. As also reported by Mojarad et al. (2018), comparisons of two possible breakdowns of RiPPLE students (experimental group) and non-RiPPLE students (control group) are presented:

1. all RiPPLE students compared to all non-RiPPLE students, and
2. RiPPLE students compared to matched non-RiPPLE students.

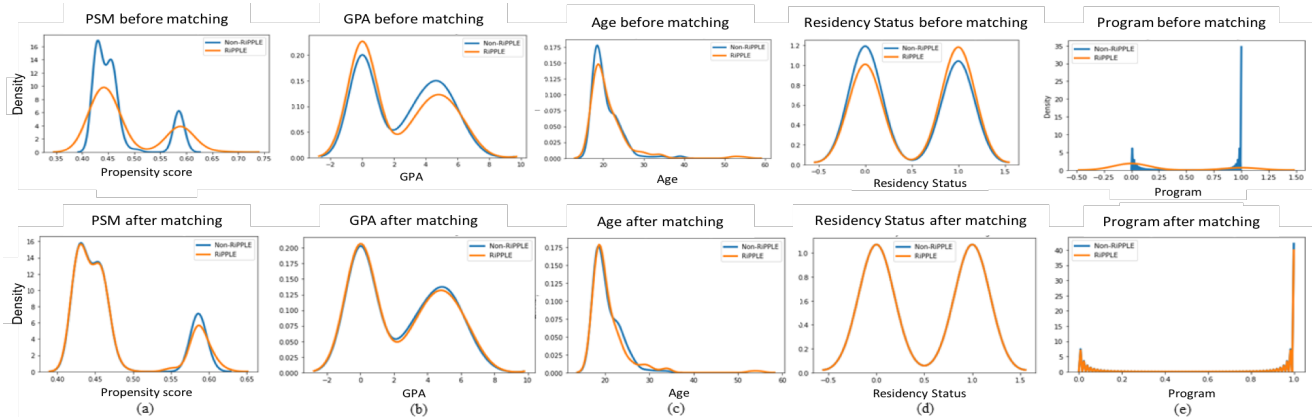


Figure 2. Distribution of propensity score and covariates before and after matching. Table 1 reports the number of RiPPLE and non-RiPPLE members before and after matching. The figure illustrates that after matching, the two groups are similar across the four nominated covariates.

Figure 2(a)–(d) represents the distribution of each of the defined covariates (GPA, age, residency status, and program level) independently before and after matching, which clearly illustrates the effectiveness of PSM in diminishing the differences in the distribution of covariates between the experimental group and the control group. The overall propensity score of participants in both groups before and after matching is presented in Figure 2(e), which again illustrates that the participants in the experimental group were well matched to their control group counterparts.

Table 1 summarizes the statistics related to midterm exam scores for each group before and after matching, where μ and σ represent the average and standard deviation of the exam scores, respectively. A *t*-test was used to evaluate the statistical significance of improvement in the learning gains. The average exam scores before matching for the students in the experimental group ($\mu = 73\% \pm 16\%$) were significantly higher than those of the students in the control group ($\mu = 65\% \pm 20\%$, $d = 0.44$, $p < .001$). After matching, a total of 198 students were selected to participate in each group; 7 participants were removed from the experimental group because no match was found for them. Students’ exam scores after matching revealed the same pattern, where the students in the experimental group had a statistically significantly better performance ($\mu = 74\% \pm 15\%$) than the students in the control group ($\mu = 64\% \pm 21\%$, $d = 0.54$, $p < .001$). The effect sizes of the two experiments support the claim that using RiPPLE does indeed provide measurable learning gains; however, the effect is considered to be at a medium level (Ma et al., 2014) and less effective than many of the studies discussed in Section 2.

Table 1. Comparing the Exam Scores of RiPPLE and Non-RiPPLE Students

Comparison	RiPPLE			Non-RiPPLE			Effect size
	N	μ	σ	N	μ	σ	<i>d</i>
Non-matched	215	73%	16%	238	65%	20%	0.44
Matched	198	74%	15%	198	64%	21%	0.54

4.2 Identifying Effective Resources

This section investigates whether RiPPLE supports students in identifying resources that they think are effective. To do so, we employed the effectiveness ratings provided by the students themselves. Figure 3 shows that roughly 60% of the time students provided a four- or five-star rating for a resource that they had attempted. Note in particular that five-star ratings were given 10 percentage points more than any other rating. In contrast, only around 17% of ratings were one or two star. This provides some evidence that the recommendation engine of RiPPLE was successful in enabling students to find resources that they thought were effective.




Figure 3. Students’ ratings of resources suggest that RiPPLE was successful in helping them identify resources that they think are effective.

4.3 Perceived Benefits for Students

A survey, advertised via an announcement on Blackboard, was conducted at the end of week 7 to capture students’ perception of the platform. A total of 55 students completed the survey. The survey asked students to indicate their agreement with a set of statements using a five-point Likert scale. These statements were related to the following three broad categories: creation and evaluation of learning activities, visualization and recommendation, and the effectiveness of the platform itself. The actual statements and the corresponding label used here for each statement are presented in Table 2, which is available in Appendix A.

Creation and Evaluation Figure 4 shows survey results concerning student perceptions about creation, completion, and evaluation of learning activities. Students reported that there were positive contributions to their learning from creating questions (63% agreement vs. 13% disagreement), writing solutions (63% agreement vs. 13% disagreement), answering questions (67% agreement vs. 12% disagreement), and rating questions (56% agreement vs. 16% disagreement). In their written comments, a few students mentioned that a moderation process on the questions before they become available to the students would enhance the platform (“the moderation setting was set to no moderation for this course”). One student mentioned that a “question originality/creativity” rating may help reduce the number of isomorphic questions.




Figure 4. Results of the survey on questions related to creation and evaluation of learning activities.

Learner Model and Recommendation Figure 5 shows survey results related to learner modelling and recommendation. Students reported that the open learner model increased their motivation to study or further use the platform (64% agreement vs. 15% disagreement) and increased their trust in the system (62% agreement vs. 14% disagreement). In their written comments, a few students mentioned that they would appreciate more transparency on how their knowledge state is computed, which is in line with recommendations from previous publications on open learner models (Bull, Ginon, Boscolo, & Johnson, 2016). The majority of students indicated that the “Recommend Question” feature helped them find effective questions (62% agreement vs. 11% disagreement). A few students identified a potential limitation that RiPPLE shares with many similar systems: because students cannot explicitly specify learning activities that they do not want to practise, they continued to have activities re-recommended despite purposely ignoring them in the past.




Figure 5. Results of the survey on questions related to the learner model and recommendation.

Platform Effectiveness Figure 6 shows the survey results relating to student-perceived effectiveness of the platform. Most students reported that they found the platform easy to use and navigate (73% agreement vs. 8% disagreement) and that it made them more aware of their current knowledge state (69% agreement vs. 13% disagreement). Most students also reported that RiPPLE stimulated them to study more effectively (62% agreement vs. 15% disagreement) and that they would like to use RiPPLE in other courses (52% agreement vs. 17% disagreement). Two limitations related to the general effectiveness of the platform were reported a number of times by students throughout the semester, which may have contributed to the lower level of interest in using the platform in other courses. The first limitation was related to the competency of students being decreased because of incorrectly answering a poorly worded or incorrect question. This limitation has now been resolved — students will now regain their lost competencies from engaging with a poorly designed question once the question has been deleted from the platform. Another alternative for resolving this issue is to use the moderation feature. The second limitation relates to the lack of availability of an environment in RiPPLE where students can freely attempt questions without being concerned about their competencies being decreased. We are currently investigating ways that students’ concerns about their competencies being decreased can be addressed.




Figure 6. Results of the survey on questions related to the effectiveness of the platform.

Most students perceived RiPPLE as beneficial to their learning; however, a significant number of students did not agree with the statements given in the survey. Focus groups and usability studies are underway to determine how we may further improve RiPPLE.

5. Conclusion

This paper integrates insights from past research on crowdsourcing in education, learning sciences, and adaptive learning to present a platform called RiPPLE that can be used to collect data for LA research. Pedagogically, this tool encourages students to move from the passive acquisition of knowledge and skills to a more active learning mode, where they create learning activities for their peers. Thus, RiPPLE supports students in learning how to learn by working to improve their knowledge state

through the recommendation of learning activities that best suit their learning needs, helping them visualize their progress along the way using open learner models. An initial evaluation of the platform in a class of 453 students has provided preliminary evidence that using RiPPLE (i) may lead to measurable learning gains and (ii) may support students in finding resources that they think are effective. Also, there is evidence that students perceived that using the platform benefited their learning.

There are, however, several limitations in the evaluation of the tool that restrict the generalizability of the results: (1) The quasi-experimental approach, even with the addition of propensity score matching, is prone to self-selection bias. Future studies aim to receive ethics approval to conduct a true randomized control trial experiment, seeking to provide more evidence in support of the claim that this tool results in demonstrable learning gains. (2) The current study explores the benefits of adopting RiPPLE using only one specific cohort: students with a background in computer science that have formal training in algorithmic literacy. Future work aims to investigate the benefits of adopting RiPPLE across more widely varying disciplines. This, in turn, may lead to the scaffolding of the platform or development of training material (e.g., short videos) so that students across all disciplines can use RiPPLE more effectively. (3) The current study only explores the benefits of using RiPPLE from the students' perspective. Future work aims to also explore the views of instructors that have used RiPPLE.

As a tool that is designed for easy integration (via LTI) with most standard LMSs, RiPPLE provides opportunities for both researchers and instructors to investigate the utility of applying crowdsourcing and adaptive learning methods in live teaching scenarios without requiring the ongoing and intensive generation of learning activities and sophisticated maps of how they interrelate. It is our hope that this will facilitate the wider use of an important technique for helping students actively participate in personalizing their learning experience while still acquiring skills and knowledge that are considered essential by instructors. Thus, RiPPLE shows promise for advancing LA research in important areas, such as adaptive and personalized learning, student creativity and critical thinking, and learner modelling.

Declaration of Conflicting Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors declared no financial support for the research, authorship, and/or publication of this article.

References

- Abdi, S., Khosravi, H., Sadiq, S., & Gašević, D. (2019). A multivariate Elo-based learner model for adaptive educational systems. In C. F. Lynch, A. Merceron, M. Desmarais, & R. Nkambou (Eds.), *Educational Data Mining 2019: 12th International Conference on Educational Data Mining (EDM2019)*, 2–5 July 2019, Montreal, Canada (pp. 228–233). International Educational Data Mining Society. Retrieved from <http://educationaldatamining.org/edm2019/proceedings/>
- Anderson, J. R., Boyle, C. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science*, 228(4698), 456–462. <https://dx.doi.org/10.1126/science.228.4698.456>
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *Journal of the Learning Sciences*, 4(2), 167–207. https://dx.doi.org/10.1207/s15327809jls0402_2
- Austin, P. C. (2011). An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 46(3), 399–424. <https://dx.doi.org/10.1080/00273171.2011.568786>
- Banks, J., Cochran-Smith, M., Moll, L., Richert, A., Zeichner, K., LePage, P., ... McDonald, M. (2005). Teaching diverse learners. In L. Darling-Hammond & J. Bransford (Eds.), *Preparing Teachers for a Changing World: What Teachers Should Learn and Be Able to Do* (pp. 232–274). San Francisco, California, USA: Jossey-Bass.
- Barak, M., & Rafaeli, S. (2004). On-line question-posing and peer-assessment as means for web-based knowledge sharing in learning. *International Journal of Human-Computer Studies*, 61(1), 84–103. <https://dx.doi.org/10.1016/j.ijhcs.2003.12.005>
- Bates, S. P., Galloway, R. K., & McBride, K. L. (2012). Student-generated content: Using PeerWise to enhance engagement and outcomes in introductory physics courses. *AIP Conference Proceedings*, 1413(1), 123–126. <https://dx.doi.org/10.1063/1.3680009>
- Bates, S. P., Galloway, R. K., Riise, J., & Homer, D. (2014). Assessing the quality of a student-generated question repository. *Physical Review Special Topics — Physics Education Research*, 10(2), 020105. <https://dx.doi.org/10.1103/physrevstper.10.020105>
- Biggs, J. (2012). What the student does: Teaching for enhanced learning. *Higher Education Research & Development*, 31(1), 39–55. <https://dx.doi.org/10.1080/07294360.2012.642839>

- Bloom, B. S., Englehart, M., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). *Taxonomy of Educational Objectives: Handbook I, Cognitive Domain*. New York, USA: David McKay.
- Bodily, R., Kay, J., Aleven, V., Jivet, I., Davis, D., Xhakaj, F., & Verbert, K. (2018). Open learner models and learning analytics dashboards: A systematic review. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge (LAK '18)*, 5–9 March, 2018, Sydney, Australia (pp. 41–50). New York, USA: ACM. <https://dx.doi.org/10.1145/3170358.3170409>
- Bovill, C. (2013). Students and staff co-creating curricula: A new trend or an old idea we never got around to implementing? In C. Rust (Ed.), *Improving Student Learning through Research and Scholarship: 20 Years of ISL* (pp. 96–108). Oxford, UK: Oxford Centre for Staff and Learning Development. Retrieved from <http://eprints.gla.ac.uk/82348>
- Brusilovsky, P. (2012). Adaptive hypermedia for education and training. In P. J. Durlach & A. M. Lesgold (Eds.), *Adaptive Technologies for Training and Education* (p. 46–66). Cambridge University Press. <https://dx.doi.org/10.1017/CBO9781139049580.006>
- Brusilovsky, P., & Millán, E. (2007). User models for adaptive hypermedia and adaptive educational systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The Adaptive Web: Methods and Strategies of Web Personalization. Lecture Notes in Computer Science* (Vol. 4321, pp. 3–53). Springer Berlin Heidelberg. https://dx.doi.org/10.1007/978-3-540-72079-9_1
- Bull, S., Ginon, B., Boscolo, C., & Johnson, M. (2016). Introduction of learning visualisations and metacognitive support in a persuadable open learner model. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge (LAK '16)*, 25–29 April, 2016, Edinburgh, Scotland (pp. 30–39). New York, USA: ACM. <https://dx.doi.org/10.1145/2883851.2883853>
- Bull, S., & Kay, J. (2010). Open learner models. In R. Nkambou, J. Bourdeau, & R. Mizoguchi (Eds.), *Advances in Intelligent Tutoring Systems* (Vol. 308, pp. 301–322). Springer Berlin Heidelberg. https://dx.doi.org/10.1007/978-3-642-14363-2_5
- Chambers, J. M., Cleveland, W. S., Kleiner, B., & Tukey, P. A. (1983). *Graphical Methods for Data Analysis*. Boca Raton, Florida, USA: CRC Press.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112(1), 155–159. <https://dx.doi.org/10.1037/0033-2909.112.1.155>
- Collins, A., & Halverson, R. (2018). *Rethinking Education in the Age of Technology: The Digital Revolution and Schooling in America*. New York, USA: Teachers College Press.
- Cooper, K., & Khosravi, H. (2018). Graph-based visual topic dependency models: Supporting assessment design and delivery at scale. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge (LAK '18)*, 5–9 March, 2018, Sydney, Australia (pp. 11–15). New York, USA: ACM. <https://dx.doi.org/10.1145/3170358.3170418>
- Denny, P., Hamer, J., & Luxton-Reilly, A. (2009). Students sharing and evaluating MCQs in a large first year engineering course. In *20th Annual Conference for the Australasian Association for Engineering Education (AAEE '09)*, 6–9 December 2009, The University of Adelaide, Adelaide, Australia (pp. 575–580). Barton, Australia: Engineers Australia.
- Denny, P., Hamer, J., Luxton-Reilly, A., & Purchase, H. (2008). PeerWise: Students sharing their multiple choice questions. In *Proceedings of the Fourth International Workshop on Computing Education Research (ICER '08)*, 6–7 September 2008, Sydney, Australia (pp. 51–58). New York, USA: ACM. <https://dx.doi.org/10.1145/1404520.1404526>
- Draper, S. W. (2009). Catalytic assessment: Understanding how MCQs and EVS can foster deep learning. *British Journal of Educational Technology*, 40(2), 285–293. <https://dx.doi.org/10.1111/j.1467-8535.2008.00920.x>
- Essa, A. (2016). A possible future for next generation adaptive learning systems. *Smart Learning Environments*, 3(1), 16. <https://dx.doi.org/10.1186/s40561-016-0038-y>
- Falmagne, J.-C., Cosyn, E., Doignon, J.-P., & Thiéry, N. (2006). The assessment of knowledge, in theory and in practice. In R. Missaoui & J. Schmidt (Eds.), *Formal Concept Analysis. Lecture Notes in Computer Science* (Vol. 3874, pp. 61–79). Springer Berlin Heidelberg.
- Galloway, K. W., & Burns, S. (2015). Doing it for themselves: Students creating a high quality peer-learning environment. *Chemistry Education Research and Practice*, 16(1), 82–92. <https://dx.doi.org/10.1039/c4rp00209a>
- Heffernan, N. T., & Heffernan, C. L. (2014). The ASSISTments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education*, 24(4), 470–497. <https://dx.doi.org/10.1007/s40593-014-0024-x>
- Heffernan, N. T., Ostrow, K. S., Kelly, K., Selent, D., Van Inwegen, E. G., Xiong, X., & Williams, J. J. (2016). The future of adaptive learning: Does the crowd hold the key? *International Journal of Artificial Intelligence in Education*, 26(2), 615–644. <https://dx.doi.org/10.1007/s40593-016-0094-z>
- Jose, F. (2016). *White Paper: Knewton Adaptive Learning: Building the World's Most Powerful Recommendation Engine for Education*. Retrieved from <https://cdn.tc-library.org/Edlab/Knewton-adaptive-learning-white-paper-1.pdf>

- Karataev, E., & Zadorozhny, V. (2017). Adaptive social learning based on crowdsourcing. *IEEE Transactions on Learning Technologies*, 10(2), 128–139. <https://dx.doi.org/10.1109/TLT.2016.2515097>
- Khosravi, H., & Cooper, K. (2018). Topic dependency models: Graph-based visual analytics for communicating assessment data. *Journal of Learning Analytics*, 5(3), 136–153. <https://dx.doi.org/10.18608/jla.2018.53.9>
- Khosravi, H., Cooper, K., & Kitto, K. (2017). RiPLE: Recommendation in peer-learning environments based on knowledge gaps and interests. *Journal of Educational Data Mining*, 9(1), 42–67. Retrieved from <https://jedm.educationaldatamining.org/index.php/JEDM/article/view/239>
- King, A. (1992). Facilitating elaborative learning through guided student-generated questioning. *Educational Psychologist*, 27(1), 111–126. <https://dx.doi.org/10.1207/s15326985ep2701g>
- Kitto, K., Lupton, M., Davis, K., & Waters, Z. (2017). Designing for student-facing learning analytics. *Australasian Journal of Educational Technology*, 33(5), 152–168. <https://dx.doi.org/10.14742/ajet.3607>
- Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, 106(4), 901–918. <https://dx.doi.org/10.1037/a0037123>
- Matthews, K. E. (2017). Five propositions for genuine students as partners practice. *International Journal for Students as Partners*, 1(2). <https://dx.doi.org/10.15173/ijsap.v1i2.3315>
- May, M., George, S., & Prévôt, P. (2011). TrAVIS to enhance online tutoring and learning activities: Real time visualization of students tracking data. *Interactive Technology and Smart Education*, 8(1), 52–69. <https://dx.doi.org/10.1108/17415651111125513>
- Meer, N., & Chapman, A. (2014). Co-creation of marking criteria: Students as partners in the assessment process. *Business and Management Education in HE*, 1–15. <https://dx.doi.org/10.11120/bmhe.2014.00008>
- Mojarad, S., Essa, A., Mojarad, S., & Baker, R. S. (2018). Studying adaptive learning efficacy using propensity score matching. In A. Pardo et al. (Eds.), *Companion Proceedings of the 8th International Conference on Learning Analytics and Knowledge (LAK '18)*, 5–9 March 2018, Sydney, Australia. Society for Learning Analytics Research.
- Morrison, B. B., & DiSalvo, B. (2014). Khan Academy gamifies computer science. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education (SIGCSE '14)*, 5–8 March 2014, Atlanta, Georgia, USA (pp. 39–44). New York, USA: ACM. <https://dx.doi.org/10.1145/2538862.2538946>
- Mulryan-Kyne, C. (2010). Teaching large classes at college and university level: Challenges and opportunities. *Teaching in Higher Education*, 15(2), 175–185. <https://dx.doi.org/10.1080/13562511003620001>
- Oxman, S., Wong, W., DV X Innovations, & DeVry Education Group. (2014). *White Paper: Adaptive Learning Systems*. Integrated Education Solutions. Retrieved from <http://kenanaonline.com/files/0100/100321/DVxAdaptiveLearningWhitepaper.pdf>
- Park, O. c., & Lee, J. (2004). Adaptive instructional systems. In D. Jonassen (Ed.), *Handbook of Research on Educational Communications and Technology* (2nd ed., pp. 651–684). Mahwah, New Jersey, USA: Lawrence Erlbaum Associates Publishers.
- Pelánek, R. (2016). Applications of the Elo rating system in adaptive educational systems. *Computers & Education*, 98(C), 169–179. <https://dx.doi.org/10.1016/j.compedu.2016.03.017>
- Pressley, M., Wood, E., Woloshyn, V. E., Martin, V., King, A., & Menke, D. (1992). Encouraging mindful use of prior knowledge: Attempting to construct explanatory answers facilitates learning. *Educational Psychologist*, 27(1), 91–109. <https://dx.doi.org/10.1207/s15326985ep27017>
- Purchase, H., Hamer, J., Denny, P., & Luxton-Reilly, A. (2010). The quality of a PeerWise MCQ repository. In *Proceedings of the Twelfth Australasian Conference on Computing Education (ACE '10)*, 1 January 2010, Brisbane, Australia (Vol. 103, pp. 137–146). Darlinghurst, Australia: Australian Computer Society, Inc.
- Ritter, S., Anderson, J. R., Koedinger, K. R., & Corbett, A. (2007). Cognitive Tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review*, 14(2), 249–255. <https://dx.doi.org/10.3758/BF03194060>
- Ritter, S., Carlson, R., Sandbothe, M., & Fancsali, S. E. (2015). Carnegie Learning's adaptive learning products. In O. Santos et al. (Eds.), *Educational Data Mining 2015: 8th International Conference on Educational Data Mining (EDM2015)*, 26–29 June 2015, Madrid, Spain. Retrieved from http://www.educationaldatamining.org/EDM2015/proceedings/edm2015_proceedings.pdf
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55. <https://dx.doi.org/10.1093/biomet/70.1.41>
- Smart Sparrow. (2016). *Smart Sparrow — Adaptive eLearning Platform*. Retrieved from <https://www.smartsparrow.com/platform>

- Solemon, B., Ariffin, I., Din, M. M., & Anwar, R. M. (2013). A review of the uses of crowdsourcing in higher education. *International Journal of Asian Social Science*, 3(9), 2066–2073.
- Sullivan, G. M. (2011). Getting off the “gold standard”: Randomized controlled trials and education research. *Journal of Graduate Medical Education*, 3(3), 285–289. <https://dx.doi.org/10.4300/JGME-D-11-00147.1>
- Tackett, S., Raymond, M., Desai, R., Haist, S. A., Morales, A., Gaglani, S., & Clyman, S. G. (2018). Crowdsourcing for assessment items to support adaptive learning. *Medical Teacher*, 40(8), 838–841. <https://dx.doi.org/10.1080/0142159X.2018.1490704>
- Tai, J., Ajjawi, R., Boud, D., Dawson, P., & Panadero, E. (2018). Developing evaluative judgement: Enabling students to make decisions about the quality of work. *Higher Education*, 76(3), 467–481. <https://dx.doi.org/10.1007/s10734-017-0220-3>
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221. <https://dx.doi.org/10.1080/00461520.2011.611369>
- Walsh, J. L., Harris, B. H., Denny, P., & Smith, P. (2018). Formative student-authored question bank: Perceptions, question quality and association with summative performance. *Postgraduate Medical Journal*, 94(1108), 97–103. <https://dx.doi.org/10.1136/postgradmedj-2017-135018>
- Whitehill, J., Aguerrebere, C., & Hylak, B. (2019). Do learners know what’s good for them? Crowdsourcing subjective ratings of OERs to predict learning gains. In C. F. Lynch, A. Merceron, M. Desmarais, & R. Nkambou (Eds.), *Educational Data Mining 2019: 12th International Conference on Educational Data Mining (EDM2019)*, 2–5 June 2019, Montreal, Canada (pp. 462–467). Retrieved from <http://educationaldatamining.org/edm2019/proceedings/>
- Williams, J. J., Kim, J., Rafferty, A., Maldonado, S., Gajos, K. Z., Lasecki, W. S., & Heffernan, N. (2016). Axis: Generating explanations at scale with learnersourcing and machine learning. In *Proceedings of the Third (2016) ACM Conference on Learning @ Scale (L@S 2016)*, 25–26 April 2016, Edinburgh, Scotland (pp. 379–388). New York, USA: ACM. <https://dx.doi.org/10.1145/2876034.2876042>
- Yilmaz, B. (2017). *Effects of Adaptive Learning Technologies on Math Achievement: A Quantitative Study of ALEKS Math Software* (Unpublished doctoral dissertation). University of Missouri-Kansas City.

6. Appendix A

A survey was conducted at the end of week 7 to capture students’ perception of the platform, with a total of 55 students responding. The survey asked students to use a five-point Likert scale to indicate their agreement with the statement given in Table 2.

Table 2. Statements Used in the Course Survey Described in Section 4.3

Creation and Evaluation	
Creating Questions	Creating questions on RiPPLE helps me reflect on the learning objectives of the course.
Writing Solutions	Writing solutions for my authored questions helps me develop a deeper understanding of the concept(s) targeted by the question.
Answering Questions	Answering questions on RiPPLE helps me develop a deeper understanding of the concept(s) targeted by the question.
Rating Questions	Rating my peers’ questions enables me to incorporate higher-order cognitive skills (e.g., appraise, judge, critique, support) in my own learning.
Writing Comments	Writing comments on my peers’ questions enables me to incorporate higher-order cognitive skills (e.g., appraise, judge, critique, support) in my own learning.
Learner Model and Recommendation	
Motivation	The visualizations used by RiPPLE for showing my knowledge state increase my motivation to study or further use the platform.
Trust	Having the ability to visually see my current knowledge state increases my trust in the recommended questions.
Recommendation	The recommend question feature helped me find effective questions.
Platform Effectiveness	
Usability	RiPPLE is easy to use and navigate.
Awareness	RiPPLE makes me aware of my current learning situation.
Stimulation	RiPPLE stimulates me to study more effectively.
Future use	I would like to use RiPPLE in my other courses.