Development and Adoption of an Adaptive Learning System

Reflections and Lessons Learned

Hassan Khosravi
The University of Queensland
Brisbane, QLD, Australia
h.khosravi@uq.edu.au

Shazia Sadiq
The University of Queensland
Brisbane, QLD, Australia
shazia@itee.uq.edu.au

Dragan Gasevic
Monash University
Clayton, Vic, Australia
Dragan.Gasevic@monash.edu

ABSTRACT

Adaptive learning systems (ALSs) aim to provide an efficient, effective and customised learning experience for students by dynamically adapting learning content to suit their individual abilities or preferences. Despite consistent evidence of their effectiveness and success in improving student learning over the past three decades, the actual impact and adoption of ALSs in education remain restricted to mostly research projects. In this paper, we provide a brief overview of reflections and lessons learned from developing and piloting an ALS in a course on relational databases. While our focus has been on adaptive learning, many of the presented lessons are also applicable to development and adoption of educational tools and technologies in general. Our aim is to provide insight for other instructors that are interested in adopting ALSs or are involved in implementation of educational tools and technologies.

KEYWORDS

adaptive learning, educational technologies, crowdsourcing

ACM Reference format:
Hassan Khosravi, Shazia Sadiq, and Dragan Gasevic. 2020. Development and Adoption of an Adaptive Learning System. In *Proceedings of The 51st ACM Technical Symposium on Computer Science Education, Portland, OR, USA, March 11–14, 2020 (SIGCSE '20),* 7 pages.
https://doi.org/10.1145/3328778.3366900

1 INTRODUCTION

Educators continue to face significant challenges in providing high quality, post-secondary instruction in large online or on-campus classes [38]. A significant portion of these challenges emerges from high levels of diversity in learners' academic ability [7]. Teach-to-the-middle instruction is most commonly used as it benefits the majority of the students; however, this method does not meet the needs of learners who differ significantly from the norm.

Adaptive learning systems (ALSs) [4, 40] provide a potential solution to this problem. ALSs make use of data about students, learning processes, and learning products to provide an efficient, effective and customised learning experience for students by dynamically adapting learning content and activities to suit their individual abilities or preference. There are two main classes of ALSs. The first class, which are commonly referred to as Intelligent Tutoring Systems (ITSs) [5] use AI-based algorithms to replicate the support that is often provided by a tutor such as assisting a learner by providing personalised step by step guidance in solving a problem. Carnegie Learning's adaptive learning products such as MATHiaU and Zulama [45] are established examples of this class of ALSs. The ALSs in the second class have mostly focused on adaptively recommending learning activities to a student from a large repository of learning resources to match the current learning needs of a student. Pearson's MyLabs (using Knewton for its adaptive functionality) and McGraw-Hill's LearnSmart and ALEKS are established examples of this class of ALSs.

A consistent and growing body of knowledge has provided evidence about the effectiveness of both of these types of ALSs relative to classroom teaching or to educational systems that provide instructions and learning activities that are not adaptive [4]. For ITSs, an early empirical study in a highly cited review paper by Anderson et al. [6] found that on average using ITSs have a learning gain effect-size [12] of $d = 1.0$. A later comprehensive review by Vanlehn [52] confirmed the effectiveness of using ITSs, but reported the effect size to be around $d = 0.76$. A meta-analysis of 107 studies on ITSs involving 14,321 participants found that: ITS were associated with higher achievement relative to teacher-led large-group instruction, non-ITS computer based-instruction, and texbooks or workbooks. However, they found no learning difference between using an ITS or individualised human tutoring [34]. More recently, Xu et al. [56] provided a meta-analysis based on 19 studies on the effectiveness of ITSs on K-12 students' reading comprehension, reporting an average effect size of $d = 0.6$ which is still considered a medium to large effect size [44]. ALSs that focus on recommending learning resources have also been reported to improve learning outcomes. For example, Mojarad et al. [37] conducted a study on 3422 students from 198 offerings that have used ALEKS reporting significantly higher pass rates amongst students using ALEKS. Yilmaz [57] conducted a similar experiment with over 2000 students that have used ALEKS reporting statistically significant gains in math achievement compared to students that had not engaged with ALEKS.

ALSs are commonly designed and developed with pre-existing learning resources and activities targeted towards a particular domain. Rough estimates from the literature put the development time for earlier versions of ALSs at 200 hours of expert's time for each hour of instruction [3]. More recently, with the use of smart tools for authoring content for ALSs such as Cognitive Tutor Authoring

SIGCSE '20, March 11–14, 2020, Portland, OR, USA

ACM ISBN 978-1-4503-6793-6/20/03…$15.00
https://doi.org/10.1145/3328778.3366900

Tools (CTAT) development time has been reduced to roughly 25 hours of domain expert's time per instructional hour [3]. Nevertheless, ALSs are still very expensive to develop and challenging to scale across different domains [32].

Responding to these challenges, a new wave of content-agnostic adaptive learning systems that enables the educators to develop and author the content of their course(s) has been established. Smart Sparrow and many learning management systems such as open edX, Cerego and Desire2Learn that incorporate adaptive functionality into their course building tools follow this model. This platform model is relatively new and mostly suffers from an operational limitation rather than a technological one: To adapt to each student's unique learning needs, a large number of learning activities must be created and tagged, which introduces significant overheads for teaching staff in both time and training.

A potential solution for developing cost-effective ALSs across many domains is to adopt a crowdsourcing approach [22] and to engage students as partners [35] in creating, moderating and evaluating learning resources. This can significantly reduce the cost of development and has the potential to foster higher-order learning for students across many domains [32]. The use of crowdsourcing in education [28] and in particular in ALSs is beginning to receive significant attention. Heffernan et al. [27] have proposed employing crowdsourcing for development of adaptive learning systems; Williams et al. [55] have developed an Adaptive eXplanation Improvement System (AXIS) that uses crowdsourcing to generate, revise, and evaluate explanations as learners solve problems; and Karataev and Zadorozhny [30] have proposed a framework that combines concepts of crowdsourcing, online social networks, and adaptive systems to provide personalized learning pathways for students.

Building on the preliminary success of these research-focused systems that use crowdsourcing towards personalisation of education, we have developed a student-focused, scalable, content-agnostic adaptive learning system that relies on crowdsourcing and partnership with students for the development of learning resources that are adaptively served. Here, we provide an overview of this system, which is called RiPPLE[1] [32], and share our reflections and lessons learned from developing and piloting this system. While our focus has been on developing an adaptive learning system, many of the challenges that we have faced, which are discussed in our lessons learned, are also applicable to a broad range of educational tools and technologies. Our aim is to share insight with other instructors who would like to adopt an ALS or are involved in the implementation of educational tools and technologies.

2 THE RIPPLE SYSTEM

RiPPLE is an adaptive learning system that recommends personalised learning activities to students, based on their knowledge state, from a pool of crowdsourced learning activities that are generated and evaluated 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, and practical implementation into a platform that instructors can easily use in their courses. To date, over 3000 registered users from 15 courses have used RiPPLE to create over 7,000 learning resources and attempt or review over 250,000 learning resources. For a detailed description of RiPPLE, please see [31, 32]. Here, we only provide a brief description of the system. Figure 1 illustrates one of the main pages of RiPPLE.

Creation of a new offering. RiPPLE supports two types of roles for users: instructors and students. Instructors have the ability to create new offerings associated with a course. Once a new offering has been created, the instructors have to specify a domain model (a set of topics) for the offering. The domain model may be altered throughout the semester. 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 sharing resources with other instructors who are teaching similar courses.

Content creation and evaluation. RiPPLE enables students and instructors to create, attempt, and evaluate a wide range of learning activities that include worked examples, multiple-choice questions (MCQs) and general notes. Students and instructors are able to view/attempt and then evaluate the learning activities associated with any RiPPLE offering that they are enrolled in. For MCQs, once a user has attempted a question, they are able to view the right answer, the distribution of how others have responded, and an explanation. For worked examples, students view a question, a step-by-step solution to solving the question, and the answer. While for notes, students can view and interact with a variety of media such as images, videos, and embedded simulations.For all resources users can view the current rating, identity of the creator and user comments made about the resource. Users are able to add their own comments and rate the effectiveness of resources. RiPPLE provides instructors with multiple options to customise how resources are moderated in the platform, which are further discussed in Section 3.1.

Learner modelling. The top part of Figure 1 shows an interactive visualisation widget, which enables students to view an abstract representation of their knowledge state based on the topics that are present in the domain model. Colour of the bars, determined by the underlying algorithm modelling the learner, categorises 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. In its current state, RiPPLE makes use of the ELO rating system for approximating the knowledge state of users [1].

Content selection and recommendation. The bottom part of Figure 1 shows the page used by RiPPLE to present the available resources to students. RiPPLE enables students to sort resources based on their, recommendation, perceived effectiveness, number of responses, number of comments, or creation time. By selecting "Recommended", the platform sorts the resources based on their learning benefits to the student. It also enables students to filter the resources that are included in the results based on their tagged topics. The results of the search are presented as a list of cards, allowing students to engage with learning activities that best suit

[1] http://ripplelearning.org

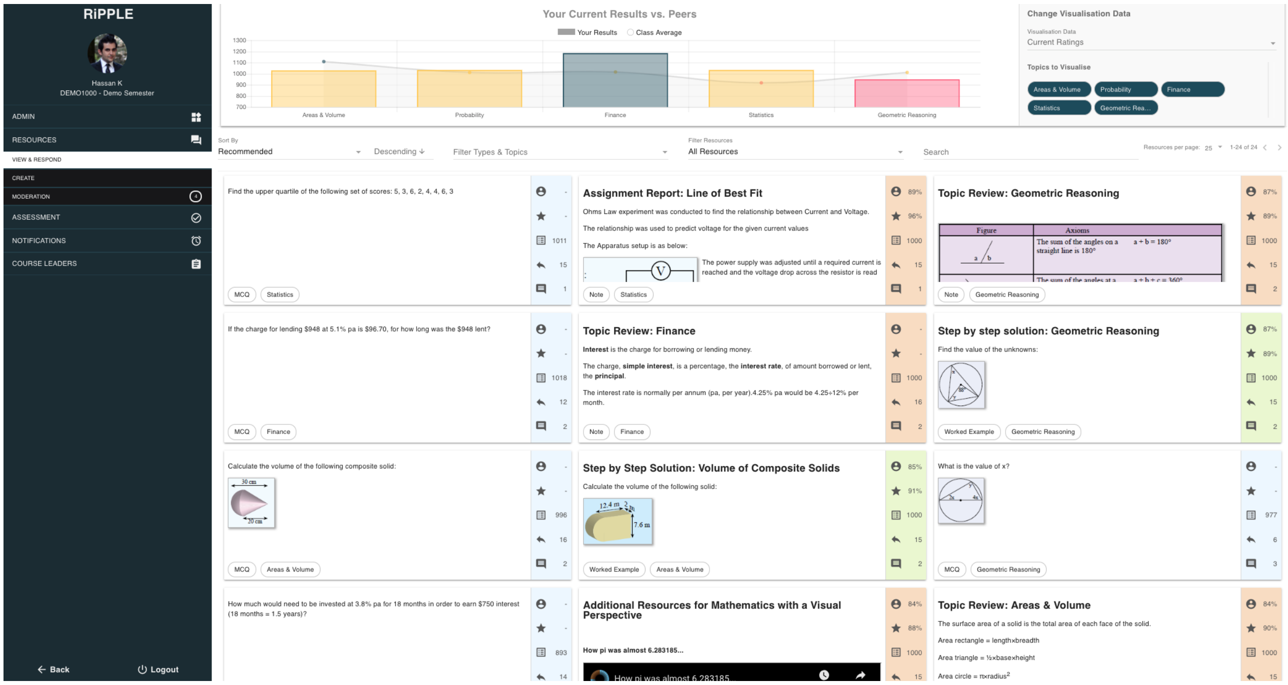

Figure 1: Overview of one of the main pages of RiPPLE.

their needs. Clicking on a resource card will take a student to a page that allows them to view, answer, rate and comment on the resources.

Badges and Leaderboard. RiPPLE makes use of badges to promote student engagement. Students are able to achieve badges in two broad categories of "Engagement Badges" and "Competency Badges". RiPPLE also uses a leaderboard which displays students with the highest score on a variety of items including the number of questions contributed, answered, correctly answered, and rated. It also displays the students with the highest number of achievements, which are presented in terms of gamified badges.

User Profiles and Course Reports. 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 who are enrolled in their RiPPLE offering, which can help them to identify the learning needs of each of their students. Additionally, RiPPLE provides the ability for instructors to download a set of course reports based on data collected by RiPPLE in their offering. These reports provide additional information on the students, questions, comments, knowledge units as well as students' attempts and views.

3 REFLECTIONS AND LESSONS LEARNED

In this section, reflecting on findings from the literature as well as synthesising our own, we present a set of lessons learned from developing and adopting RiPPLE. Our lessons learned draw insight from piloting RiPPLE in 15 courses including feedback from students and instructors that have used the system as well as data collected by the system. We, in particular, refer to survey data that was conducted in one of these courses in our discussion. This course was a graduate-level, on-campus course on database principles with 75 students at The University of Queensland. In this course, students used RiPPLE to author 632 resources and made 22440 attempts on these resources. The survey was conducted at the end of week 7 to capture students' perception of the platform. A total of 56 students completed the survey; all of the questions were based on 5-point Likert scale statements. Some of the statements and their associated results are presented within the subsections of this section.

3.1 Harness the creativity power of your students

Contemporary models of learning have emphasised the importance of higher-order learning, encouraging students to take an active role in constructing their knowledge [9, 42]. Despite, some efforts to enrich the pedagogical approaches of ALSs (e.g., [2]), by and large, the majority of the currently available ALSs still have a focus on the acquisition of conceptual knowledge and procedural skills.

One approach to engage students in higher-order learning is to invite students to adopt an active and participatory role in creating learning resources with instructors. Multiple past studies, using the popular PeerWise platform [17], have reported that students, as non-experts, indeed have the ability to create high-quality resources that meet rigorous judgemental and statistical criteria [8, 16, 24, 49, 54]. In fact, students as authors of learning resources may have an advantage over instructors: they can use the knowledge of their own previous misconceptions to create resources so there is less chance that they will suffer from an expert blind spot [39]. However, feedback from students and instructors using the platform suggested that many of the learning resources developed by students were ineffective or incorrect. As such, we introduced a set of moderation options that can be used to review the student-created content before they are publicly released. One of these options is “staff moderation", where the created resources are moderated by the instructors. This option, however, may not be feasible in large classes as it introduces a significant amount of additional workload for the instructors. Alternative options of moderation implemented in RiPPLE rely on the collective wisdom of the crowd such as “staff review on student moderation", “staff review after student appeal" or “Competent student moderation"). Engaging students with content moderation can help students develop evaluative judgement, which is arguably “one of higher education’s ultimate goals and a necessary capability for graduates“ [50].

To get a better sense of students’ view on creation and moderation, we asked them to state their agreement with the following statement in our survey: (1) Creating resources in RiPPLE is an effective way for me to develop mastery in the course content. (2) Moderating resources in RiPPLE is an effective way for me to develop mastery in the course content. Students reported a positive contributions towards their learning from creating questions (71% agreement vs. 9% disagreement) and moderating question (77% agreement vs. 11% disagreement).

Despite students’ personal beliefs and strong evidence from the learning science literature about the benefits of engaging in resource creation and moderation, based on our experience, students often require additional incentive mechanisms to engage with these activities. Mechanisms that have worked well in the courses that had piloted RiPPLE include (1) frequent reminder of the benefits, (2) inclusion of the activity as part of their assessment, (3) bonus marks, or (4) availability of gamification aspects and digital rewards. Another additional source of motivation that has worked effectively in past pilots is to encourage students to be involved in developing sample solutions to questions from past exams that are commonly not available to students.

3.2 Provide rich and transparent models of the learners

ALSs use comprehensive learner models that harness the digital traces left by learners to provide recommendations on a broad range of learning resources and activities tailored to each learner’s learning needs. However, these learner models usually operate as a “black box” where the actual learner models and the rationale for the recommendations are not shared with the learners. This in turn, can lead to trust issues when the learner does not agree with the proposed learning resources [53]. Operating as a “black box” for modelling and recommendation seem to be in particular inadequate for educational settings where educators strive to enable learners to develop their own vision, reasoning, and an appreciation for inquiry and investigation. There have been significant recent contributions in the learning analytics and educational data mining communities on the articulation of the importance of providing transparency in modelling learners [10, 29]. Open learner models (OLMs) are learner models that are externalised and made accessible to students or other stakeholders such as instructors, often through visualisation, as an important means of supporting learning [11].

To get a sense of students’ view on the open learner model in RiPPLE, we asked students to state their agreement with the following statements: (1) *Interpretation*: the visualisations used by RiPPLE for showing my mastery level across different topics is easy to understand and interpret; (2) *Rationale*: having the ability to visually see my mastery level across different topics helps me to understand the rationale behind recommendations made by the platform; (3) *Trust*: having the ability to visually see my mastery level across different topics increases my trust in the recommended questions; and (4) *Motivation*: the visualisations used by RiPPLE for showing my mastery level across different topics increase my motivation to study or further use the platform. Figure 2 presents the results, which suggest that having an open learner model was appreciated by the students. For example, one student stated the following: “I have used a similar platform in the past, however, the visualisation of my knowledge state in this platform is a great improvement on those".

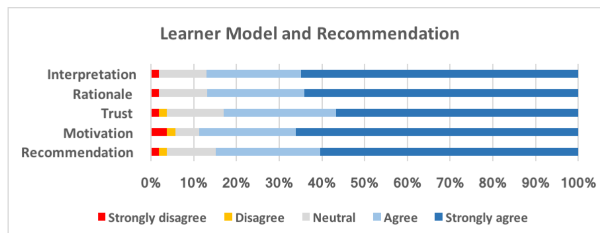

Figure 2: Results of the survey on statements related to the learner model and recommendation of learning activities.

Despite their overall popularity, we did find drawbacks in using open learner models too. Most students have reported that the availability of the OLM has increased their engagement with RiPPLE as it motivates them to improve their competency. However, feedback from some students has revealed that the availability of the OLM can also act as a source of disengagement. In particular, students have referred to engaging with more challenging practice questions as “risking a hit to your rating". Based on the request of the students, we are considering adding a practice space on RiPPLE where answering questions is decoupled from students’ OLMs. Further investigations on how we can provide open learner models without discouraging participation are underway by the authors.

3.3 Put the instructors in charge and empower them with rich analytic

It is increasingly being recognised that educational tools and technologies are not aimed at replacing instructors but rather are used as support tools to enable instructors to improve their teaching practices [13]. We have developed RiPPLE with the general principle of providing instructors authority and flexibility in terms of how the tool operates and what features and data sets are shared with students. Adopting co-creation strategies [20] where we put instructors in charge of designing the types of features they require has helped us tremendously in developing the system. Co-design and development of a set of moderation strategies that are now integrates into the system is an example of a successful feature that was co-created with the instructors.

We have found that it is also important to provide instructors with the required support for adopting educational tools and technologies. We started by providing small group workshops and technical support in terms of meeting with instructors to address their technical challenges, which was well appreciated. To provide support at a larger scale, we have now developed an online help centre[2] which includes a set of tutorial videos and a repository of tips, best practices, case studies and answers to frequently asked questions. We have also included a "Contact us" form in the tool itself which is used by instructors to contact us directly via email.

One of the greatest benefits of using ALSs and more broadly learning technologies is that they provide rich and timely actionable insight for the instructors so that they can best manage their class within the context of their own course [26]. One of our design principles has been to consider instructors as the data custodian of their RiPPLE offerings. As such, we have included a comprehensive set of reports that provides rich analytics on students and their learning process and activities. Adopting co-creation strategies [20] where we put instructors in charge of designing the types of analytics they require has been very helpful. Co-design sessions to design analytics around early detection of at-risk students are currently underway.

3.4 Provide support for ethical and low-cost educational research

While there has been a significant rise in the adoption of educational technologies in higher education, in most cases these systems are built without the aim of supporting research. Consequently, they often do not enable data harvesting or the implementation of observational or controlled experiments. To fill this gap, an important part of our mission in developing RiPPLE has been to provide support for ethical and low-cost educational research to further promote the development of evidence-based teaching practices [36].

The ethical considerations bearing on the use of student and educational data have been well studied in the field of learning analytics [21, 23]. A recent discussion paper by Corrin et al. [14] highlights the importance of careful handling of student data, providing insightful guidelines, protocols and principles. Considerable attention has been given to ensuring the compliance of RiPPLE with these principles. The following are a few examples: **Consent**: On their first use of the platform, users are presented with a consent form seeking permission to use their data to improve the academic developers' understanding of the learning process. RiPPLE allows users to change their response at any time. Regardless of their response, all users can access the platform; however, only data collected from learners who have provided and never withdrawn their consent are used for research purposes. **Transparency**: The platform provides a generic consent form to researchers and in the interests of transparency, it must be updated to describe any changes to the purpose, scope and details of planned research. **Non-maleficence**: The terms of service of using RiPPLE warn researchers against conducting research that leads to interventions which may harm a student's performance, learning experience, or simply waste their time.

[2] http://ripplelearning.org

Examples of other educational technologies that support research include PeerWise [18] and ASSISTments [27]. PeerWise has supported over 80 publications, mainly focusing on the impact of gamification and the ability of students to develop high-quality learning resources [41] and ASSISTments has enabled 27 publications, primarily looking at adaptive learning and the personalisation of feedback [25]. Their success in supporting research can be attributed to slightly different approaches. PeerWise allows instructors using the platform for teaching purposes to access data from those courses, rather than the developers retaining exclusive rights to it; the ASSISTments Ecosystem supports purposeful experimental design using Randomised Control Trials (RCTs) at low cost [27].

Inspired by the success of PeerWise and ASSISTments, RiPPLE aims to enable instructors to conduct sound, large scale educational research. RiPPLE enables instructors to collaborate with the developers of the system to conduct RCTs. It is worth noting that while there have been fiery debates about the opportunities and challenges of using RCTs in education [47, 48], they remain a gold standard test for establishing causality in many fields of educational research. To help instructors mitigate the ethical challenges of using RCTs in education, RiPPLE also supports quasi-experimental studies where students self-select whether or not to engage with an intervention. Quasi-experiments are often subject to threats to internal validity: self-selected engagement with an intervention might be influenced by specific traits or needs, meaning that students in the control group are not comparable to those in the experimental group at baseline. Propensity Score Matching (PSM) [46] may be applied to reduce baseline differences between the two groups. This method matches each student in the experiment group with a similar student from the control group, with judgements of similarity based on a set of covariates, including features of student performance (e.g., GPA), demographic (e.g., age) and behavioural engagement (e.g., learning management system logins). RiPPLE also supports observational studies by providing access to detailed analytics about student engagement (e.g., access to the platform, moderations performed, ratings provided, comments written) and performance (e.g., resources created, questions answered), through a set of interactive visualisations. Raw data can be read and downloaded as SQL and CSV files.

3.5 Provide mechanisms that motivate students to be actively engaged

There is a general consensus that motivation is regarded as one of the most important factors leading to academic success [33]. Employing mechanisms that are typically used in games in non-game contexts, commonly referred to as gamification, has been viewed as a viable option to increase participation and engagement in many different settings including education [15]. While there have been contradictory findings on the effects of interacting with gamified systems in education [51], evidence suggests that if game elements target behaviours that can improve learning then gamification can have a positive impact on student engagement and learning [19]. The use of points, levels, leaderboards and digital badges are among the most applied gamification elements in educational tools and technologies [43].

RiPPLE uses different strategies including weekly awards and badges and a leaderboard to motivate students to be actively engaged. To get a sense of students' view on gamification features of RiPPLE, we asked students to state their agreement with the following two statements: (1) The weekly *awards* motivate me to use RiPPLE and (2) The *leaderboard* motivates me to use RiPPLE. Figure 3 presents the results, which indicates that students see these strategies as effective ways for increasing their motivation. Data collected via Google Analytics complements the findings from the survey results, reporting that the leaderboard and awards pages are among the most visited pages of the platform.

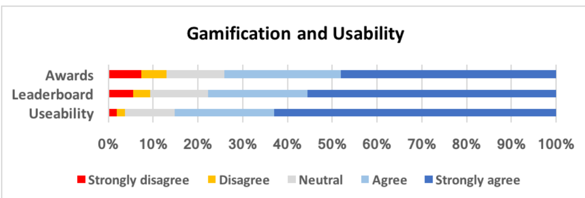

Figure 3: Results of the survey on statements related to gamification.

3.6 Do not underestimate the importance of usability, flexibility and scalability

While consideration of learning theories and pedagogical approaches are extremely important in developing ALSs and more broadly educational tools and technologies, they are not on their own sufficient for ensuring engagement from students or adoption in educational institutes. Other factors such as usability, flexibility and scalability also play an important role in the adoption of a tool.

Usability. We used a variety of methods with the aim of improving the usability of the system of which the following were the most effective: (1) observing new users interacting with the system for the first time enabled us to identify and fix some of the features that were not being used as intended; (2) holding focus groups with students and instructors helped us identify and prioritise the development of features that they found to be of most importance; (3) integration of feedback and survey forms into the system enabled us to get feedback from a larger (and often frustrated) set of users; (4) finally, adding Google Analytics to capture users' interactions and behaviour enabled us to determine the features that were underused. It also made us realise that a significant portion of the traffic of our system is coming from smartphones and tablets, which made us invest in the user experience of users on smart devices.

Flexibility. There is an increasing demand for educational tools to have the flexibility to be integrate with enterprise systems. While this was a huge development investment, we implemented two versions of the system: A stand-alone version and second version that uses the Learning Tools Interoperability (LTI) standard for integration into learning management systems. This decision has enabled us to integrate RiPPLE into Blackboard at our institute to make it more convenient for adoption while also having the opportunity to invite instructors from other universities to pilot RiPPLE without getting their universities' central IT involved.

Scalability. It is essential for any educational tool to have the ability to meet real-time demands. Depending on how broadly the system has been adopted, it may need to have the capacity to support a few hundred or millions of students using the platform simultaneously. To keep the initial cost of operation down, we started by hosting RiPPLE using the IT infrastructure supported by our institute; however, as the user-base of the tool is growing, alternatives such as using cloud computing services seems to be a more viable option for meeting unpredictable real-time demands.

4 CONCLUSION

Educational tools and technologies play an increasingly important role in higher education. In this paper, we shared our experience in developing and piloting an adaptive learning system called RiPPLE. Based on our experience, we found that (1) inviting students to create and moderate learning resources is an effective way to harness their creativity and engage them in higher-order learning; (2) providing open and transparent models of learners can help students better understand their own learning needs and improve self-regulation; (3) putting instructors in charge while empowering them with educational tools, rich analytics and the required technical support and professional development can enhance teaching and learning practices; (4) providing support for ethical and low-cost educational research can play a significant role in promoting and increasing the development of evidence-based teaching practices; (5) utilising mechanisms such as gamification that motivate students to be actively engaged can improve learning; and (6) while consideration of learning theories and pedagogical approaches are important in developing educational technologies other factors such as usability, flexibility and scalability are also critical. We hope that the lessons that we have learned while developing and piloting RiPPLE can provide insight for other instructors that are interested in adopting ALSs or are involved in the implementation of educational tools and technologies.

5 ACKNOWLEDGEMENT

The development and dissemination of RiPPLE has been supported by an innovation teaching grant by The University of Queensland.

REFERENCES

[1] Solmz Abdi, Hassan Khosravi, Shazia Sadiq, and Dragan Gasevic. 2019. A Multivariate ELO-based Learner Model for Adaptive Educational Systems. In *Proceedings of the Educational Data Mining Conference*. 462–467.

[2] Fabio N Akhras and John A Self. 2000. System intelligence in constructivist learning. *International Journal of Artificial Intelligence in Education* 11, 4 (2000), 344–376.

[3] Vincent Aleven, Bruce M. Mclaren, Jonathan Sewall, and Kenneth R. Koedinger. 2009. A New Paradigm for Intelligent Tutoring Systems: Example-Tracing Tutors. *Int. J. Artif. Intell. Ed.* 19, 2 (2009), 105–154.

[4] Vincent Aleven, Elizabeth A McLaughlin, R Amos Glenn, and Kenneth R Koedinger. 2016. Instruction based on adaptive learning technologies. *Handbook of research on learning and instruction. Routledge* (2016).

[5] John R Anderson, C Franklin Boyle, and Brian J Reiser. 1985. Intelligent tutoring systems. *Science* 228, 4698 (1985), 456–462.

[6] John R Anderson, Albert T Corbett, Kenneth R Koedinger, and Ray Pelletier. 1995. Cognitive tutors: Lessons learned. *The journal of the learning sciences* 4, 2 (1995), 167–207.

[7] James Banks, Marilyn Cochran-Smith, Luis Moll, Anna Richert, Keneth Zeichner, Pamela LePage, Linda Darling-Hammond, Helen Duffy, and Morva McDonald. 2005. Teaching diverse learners. *Preparing teachers for a changing world: What teachers should learn and be able to do* 2005 (2005), 232–274.

[8] Simon P Bates, Ross K Galloway, Jonathan Riise, and Danny Homer. 2014. Assessing the quality of a student-generated question repository. *Physical Review Special Topics-Physics Education Research* 10, 2 (2014), 020105.
[9] Benjamin S Bloom et al. 1956. Taxonomy of educational objectives. Vol. 1: Cognitive domain. *New York: McKay* (1956), 20–24.
[10] Robert Bodily, Judy Kay, Vincent Aleven, Ioana Jivet, Dan Davis, Franceska Xhakaj, and Katrien Verbert. 2018. Open learner models and learning analytics dashboards: a systematic review. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*. ACM, 41–50.
[11] Susan Bull and Judy Kay. 2010. Open learner models. In *Advances in intelligent tutoring systems*. Springer, 301–322.
[12] Jacob Cohen. 1992. A power primer. *Psychological bulletin* 112, 1 (1992), 155.
[13] Allan Collins and Richard Halverson. 2018. *Rethinking education in the age of technology: The digital revolution and schooling in America*. Teachers College Press.
[14] Linda Corrin, Gregor Kennedy, Sarah French, Simon Buckingham Shum, Kirsty Kitto, Abelardo Pardo, Deborah West, Negin Mirriahi, and Cassandra Colvin. 2019. The ethics of learning analytics in Australian higher education. (2019).
[15] Simone de Sousa Borges, Vinicius HS Durelli, Helena Macedo Reis, and Seiji Isotani. 2014. A systematic mapping on gamification applied to education. In *Proceedings of the 29th annual ACM symposium on applied computing*. ACM, 216–222.
[16] Paul Denny, John Hamer, Andrew Luxton-Reilly, et al. 2009. Students sharing and evaluating MCQs in a large first year Engineering course. In *20th Annual Conference for the Australasian Association for Engineering Education, 6-9 December 2009: Engineering the Curriculum*. Engineers Australia, 575.
[17] Paul Denny, John Hamer, Andrew Luxton-Reilly, and Helen Purchase. 2008. PeerWise: students sharing their multiple choice questions. In *Proceedings of the fourth international workshop on computing education research*. ACM, 51–58.
[18] Paul Denny, John Hamer, Andrew Luxton-Reilly, and Helen Purchase. 2008. PeerWise: students sharing their multiple choice questions. In *Proceedings of the fourth international workshop on computing education research*. ACM, 51–58.
[19] Paul Denny, Fiona McDonald, Ruth Empson, Philip Kelly, and Andrew Petersen. 2018. Empirical Support for a Causal Relationship Between Gamification and Learning Outcomes. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 311.
[20] Mollie Dollinger and Jason M Lodge. 2018. Co-creation strategies for learning analytics. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*. ACM, 97–101.
[21] Hendrik Drachsler and Wolfgang Greller. 2016. Privacy and analytics: it's a DELICATE issue a checklist for trusted learning analytics. In *Proceedings of the sixth international conference on learning analytics & knowledge*. ACM, 89–98.
[22] Enrique Estellés-Arolas and Fernando González-Ladrón-De-Guevara. 2012. Towards an integrated crowdsourcing definition. *Journal of Information science* 38, 2 (2012), 189–200.
[23] Rebecca Ferguson, Tore Hoel, Maren Scheffel, and Hendrik Drachsler. 2016. Guest editorial: Ethics and privacy in learning analytics. *Journal of learning analytics* 3, 1 (2016), 5–15.
[24] Kyle W Galloway and Simon Burns. 2015. Doing it for themselves: students creating a high quality peer-learning environment. *Chemistry Education Research and Practice* 16, 1 (2015), 82–92.
[25] Neil Heffernan. 2019. ASSISTments: As a researcher's tool. (2019). https://sites.google.com/site/assistmentsstudies/all-studies
[26] Neil T Heffernan and Cristina Lindquist Heffernan. 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 (2014), 470–497.
[27] Neil T Heffernan, Korinn S Ostrow, Kim Kelly, Douglas Selent, Eric G Van Inwegen, Xiaolu Xiong, and Joseph Jay Williams. 2016. The Future of Adaptive Learning: Does the Crowd Hold the Key? *International Journal of Artificial Intelligence in Education* 26, 2 (2016), 615–644.
[28] Yuchao Jiang, Daniel Schlagwein, and Boualem Benatallah. 2018. A Review on Crowdsourcing for Education: State of the Art of Literature and Practice. In *Pacific Asia Conference on Information Systems (PACIS)*.
[29] Ioana Jivet, Maren Scheffel, Marcus Specht, and Hendrik Drachsler. 2018. License to evaluate: Preparing learning analytics dashboards for educational practice. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*. ACM, 31–40.
[30] Evgeny Karataev and Vladimir Zadorozhny. 2017. Adaptive social learning based on crowdsourcing. *IEEE Transactions on Learning Technologies* 10, 2 (2017), 128–139.
[31] Hassan Khosravi, Kendra Cooper, and Kirsty Kitto. 2017. RiPLE: Recommendation in Peer-Learning Environments Based on Knowledge Gaps and Interests. *JEDM-Journal of Educational Data Mining* 9, 1 (2017), 42–67.
[32] Hassan Khosravi, Kirsty Kitto, and Joseph Jay Williams. 2019. RiPPLE: A Crowdsourced Adaptive Platform for Recommendation of Learning Activities. *preprint arXiv* 1910.05522 (2019).
[33] Elizabeth A Linnenbrink and Paul R Pintrich. 2002. Motivation as an enabler for academic success. *School psychology review* 31, 3 (2002), 313–328.
[34] Wenting Ma, Olusola O Adesope, John C Nesbit, and Qing Liu. 2014. Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of educational psychology* 106, 4 (2014), 901.
[35] Kelly E Matthews. 2017. Five Propositions for Genuine Students as Partners Practice. *International Journal for Students as Partners* 1, 2 (2017).
[36] James H McMillan and Sally Schumacher. 2010. Research in Education: Evidence-Based Inquiry, MyEducationLab Series. *Pearson* (2010).
[37] Shirin Mojarad, Alfred Essa, Shahin Mojarad, and Ryan S Baker. 2018. Studying Adaptive Learning Efficacy using Propensity Score Matching. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*. ACM.
[38] Catherine Mulryan-Kyne. 2010. Teaching large classes at college and university level: Challenges and opportunities. *Teaching in Higher Education* 15, 2 (2010), 175–185.
[39] Mitchell J Nathan, Kenneth R Koedinger, Martha W Alibali, et al. 2001. Expert blind spot: When content knowledge eclipses pedagogical content knowledge. In *Proceedings of the third international conference on cognitive science*, Vol. 644648.
[40] Ok-choon Park and Jung Lee. 2003. Adaptive instructional systems. *Educational Technology Research and Development* 25 (2003), 651–684.
[41] Denny Paul. 2019. PeerWise Publications. (2019). https://peerwise.cs.auckland.ac.nz/docs/publications/
[42] Michael Pressley, Eileen Wood, Vera E Woloshyn, Vicki Martin, Alison King, and Deborah Menke. 1992. Encouraging mindful use of prior knowledge: Attempting to construct explanatory answers facilitates learning. *Educational psychologist* 27, 1 (1992), 91–109.
[43] Mohd Hishamuddin Abdul Rahman, Dr Ismail, Anida Zaria Binti Mohd Noor, and Nor Syazwani Binti Mat Salleh. 2018. Gamification Elements and Their Impacts on Teaching and Learning–A Review. *The International Journal of Multimedia & Its Applications (IJMA) Vol* 10 (2018).
[44] Marnie E Rice and Grant T Harris. 2005. Comparing effect sizes in follow-up studies: ROC Area, Cohen's d, and r. *Law and human behavior* 29, 5 (2005), 615–620.
[45] Steven Ritter, Ryan Carlson, Michael Sandbothe, and Stephen E Fancsali. 2015. Carnegie learningâĂŹs adaptive learning products. *Educational Data Mining* 2015 (2015), 8th.
[46] Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1 (1983), 41–55.
[47] Ben Styles and Carole Torgerson. 2018. Randomised controlled trials (RCTs) in education research - methodological debates, questions, challenges. *Educational Research* 60, 3 (2018), 255–264. https://doi.org/10.1080/00131881.2018.1500194
[48] Gail M Sullivan. 2011. Getting off the âĂIJgold standardâĂİ: randomized controlled trials and education research. *Journal of Graduate Medical Education* 3, 3 (2011), 285–289.
[49] Sean Tackett, Mark Raymond, Rishi Desai, Steven A Haist, Amy Morales, Shiv Gaglani, and Stephen G Clyman. 2018. Crowdsourcing for assessment items to support adaptive learning. *Medical teacher* 40, 8 (2018), 838–841.
[50] Joanna Tai, Rola Ajjawi, David Boud, Phillip Dawson, and Ernesto Panadero. 2018. Developing evaluative judgement: enabling students to make decisions about the quality of work. *Higher Education* 76, 3 (2018), 467–481.
[51] Rob van Roy and Bieke Zaman. 2018. Need-supporting gamification in education: An assessment of motivational effects over time. *Computers & Education* 127 (2018), 283–297.
[52] Kurt VanLehn. 2011. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist* 46, 4 (2011), 197–221.
[53] Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of the 2013 international conference on Intelligent user interfaces*. ACM, 351–362.
[54] Jason L Walsh, Benjamin HL Harris, Paul Denny, and Phil Smith. 2018. Formative student-authored question bank: perceptions, question quality and association with summative performance. *Postgraduate medical journal* 94, 1108 (2018), 97–103.
[55] Joseph Jay Williams, Juho Kim, Anna Rafferty, Samuel Maldonado, Krzysztof Z Gajos, Walter S Lasecki, and Neil Heffernan. 2016. Axis: Generating explanations at scale with learnersourcing and machine learning. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*. ACM, 379–388.
[56] Zhihong Xu, Kausalai Kay Wijekumar, Gilbert Ramirez, Xueyan Hu, and Robin Irey. 2019. The effectiveness of intelligent tutoring systems on K-12 students' reading comprehension: A meta-analysis. *British Journal of Educational Technology* (2019).
[57] Burak Yilmaz. 2017. *Effects of Adaptive Learning Technologies on Math Achievement: A Quantitative Study of ALEKS Math Software*. Ph.D. Dissertation. University of Missouri-Kansas City.