

Recommendation in Personalised Peer-Learning Environments

Hassan Khosravi
University of Queensland
h.khosravi@uq.edu.au

Recommendation in Personalised Peer Learning Environments (RiPPLE) is an adaptive, crowdsourced, web-based, student-facing, open-source platform that employs exemplary techniques from the fields of machine learning, crowd-sourcing, learning analytics and recommender systems to provide personalised content and learning support at scale. RiPPLE presents students with a repository of tagged multiple-choice questions and provides instant feedback in response to student answers. The repository of the questions is created in partnership with the students through the use of crowdsourcing. RiPPLE uses students responses to the questions to approximate their knowledge states. Based on their knowledge state and learning needs, each student is recommended a set of personalised questions. For students that are interested in providing learning support, seeking learning support or finding study partners, RiPPLE recommends peer learning sessions based on their availability, knowledge state and preferences. This paper describes the RiPPLE interface and an implementation of that interface that has been built at the University of Queensland. The RiPPLE platform and a reference implementation are released as an open-source package under the Apache 2.0 license via GitHub.

CONTENTS

1. Introduction	3
2. Related Work	4
2.1. Knowledge Tracing	4
2.2. Adaptive Learning	4
2.3. Recommender Systems in Technology Enhanced Learning	5
2.4. Reciprocal Recommendation	6
3. Formal Notation	7
4. The RiPPLE Platform	8
4.1. Creating and Answering Questions	8
4.2. Knowledge Tracing and Recommending Questions	10
4.3. Reciprocal Peer Recommendation	12
4.4. Leaderboard	13
4.5. Personal Profile	13
4.6. Instructor Page	15
5. Expected Benefits	16
6. Implementation Details of the RiPPLE Client	17
6.1. Source Code Conventions and Overview of the RiPPLE Client	17
6.2. Deployment of the RiPPLE Client	18
6.2.1. Installing Dependencies	18
6.2.2. General Configuration	19
6.2.3. Development Configuration	19
6.2.4. Production Configuration	20
7. Implementation Details of the RiPPLE Server	20
7.1. Source Code Conventions and Overview of the RiPPLE Server	20
7.2. Deployment of the RiPPLE Server	21
7.2.1. Installing Dependencies	21
7.2.2. General Configuration	22
7.2.3. Development Configuration	23
7.2.4. Production Configuration	24
8. Contributing	24

1. INTRODUCTION

Universities continue to rapidly evolve to address the needs of diverse, growing student populations, while embracing advances in pedagogy and technology. With the revolution in big data, universities are striving to utilise their rich and complex digital data on learners towards the personalisation of education.

At The University of Queensland, we have developed an adaptive, student-facing learning platform called RiPPLE (Recommendation in Personalised Peer Learning Environments) that provides personalised content and learning support at scale. RiPPLE employs exemplary techniques from the fields of machine learning, crowdsourcing, learning analytics and recommender to provide the following main functionalities:

- *Co-creation*: RiPPLE empowers learners to contribute to co-creation of learning content.
- *Knowledge tracing*: RiPPLE approximates learners' knowledge states based on their interactions with the platform. It uses a variety of visualisations to enable learners to track their knowledge states and progress.
- *Content recommendation*: RiPPLE uses the knowledge state of learners to recommend formative exercises tailored to the needs of each individual.
- *Peer Learning Support*: RiPPLE recommends peer learning sessions based on the availability, learning preferences and needs of individuals. Students are given the opportunity to provide peer learning support, seek peer learning support or find study partners.
- *Gamification*: RiPPLE uses leaderboards and a badging system to increase student motivation and performance.

The RiPPLE platform has been released as an open-source package¹ under the Apache 2.0 license and is freely available for usage in non-commercial settings. A prototype of the system is accessible through a GitHub account². RiPPLE can use the Learning Tools Interoperability (LTI) standard to be integrated into many popular learning management systems including Blackboard, Moodle and Canvas. It is designed to be scalable, allowing the system to be easily adopted by many courses across different faculties in many universities. It is also designed to be sustainable, requiring very low maintenance from the teaching team.

The rest of this paper is organised as follows: Section 2. provides an overview of the related work on knowledge tracing, adaptive learning, recommender systems in technology enhanced learning (RecSysTEL), and reciprocal recommendation. Section 3. describes the data sources, functionality, the main aims and outputs of RiPPLE using a formal notation, which can be used by educational researchers. Section 4. presents an overview of the main features and functionalities that are supported by RiPPLE including creating and answering questions, knowledge tracing and recommending questions, reciprocal peer recommendation, leaderboards, personal profiles, and a page designed for instructors. Section 5. discusses the expected benefits students, instructors, and educational researchers will receive from using RiPPLE. Section 6. provides information on the implementation of the RiPPLE Client and Section 7. provides information on the implementation of the RiPPLE Server. Finally, Section 8. provides guidelines around contribution to RiPPLE.

¹<https://github.com/hkhosrav/RiPPLE-Core/>

²<https://hkhosrav.github.io/RiPPLE-Core/#/>

2. RELATED WORK

2.1. KNOWLEDGE TRACING

Knowledge tracing is the task of modelling the knowledge state of students so that their future performance on learning activities can accurately be predicted (Corbett and Anderson, 1994). The problem of knowledge tracing was introduced, and has been heavily studied within the intelligent tutoring community (Corbett, 2001). The Bayesian Knowledge tracing (BKT) algorithm (Corbett and Anderson, 1994) is one of the most prominent methods used for knowledge tracing. BKT uses Hidden Markov Models to capture the student knowledge states as a set of binary variables representing whether or not a concept has been mastered.

BKT has received significant attention and improvement since it was first proposed. Baker et al. (2008) introduced slipping and guessing parameters; slipping refers to the situation where a student has the required skill for answering a question but mistakenly provides the wrong answer; and guessing refers to the situation where a student provides the right answer despite not having the required skill for solving the problem. Later on, Pardos and Heffernan (2011) effectively extended BKT to capture item difficulty, which led to improve prediction accuracy. More recently, Yudelson et al. (2013) further improved BKT by introducing a new set of parameters capturing prior knowledge of individual learners.

Other algorithms with comparable or superior predictive power to BKT have also been proposed for knowledge tracing. Cen et al. (2006) introduced the Learning Factors Analysis Framework, and Pavlik Jr et al. (2009) introduced the Performance Factor Analysis framework. Khajah et al. (2014) incorporated Item Response theory (IRT) into knowledge tracing, and more recently, Piech et al. (2015) and Sha and Hong (2017) used recurrent neural networks for deep knowledge tracing.

The current implementation of RiPPLE uses the knowledge tracing algorithm of Khosravi et al. (2017)

2.2. ADAPTIVE LEARNING

Adaptive learning platforms use knowledge tracing to dynamically adjust the level or type of instruction based on individual student abilities or preferences (Paramythis and Loidl-Reisinger, 2003). Use of adaptive platforms help personalise learning, and has been shown to improve or accelerate student performance (Yilmaz, 2017). At a high level of generality all adaptive learning platforms rely on four interacting models: (1) a knowledge space modelling what the learners need to know, (2) a set of knowledge states modelling what students currently know, (3) a repository of learning objects, mapped to the concepts of the knowledge space, modelling the learning activities that are available to the students, (4) a recommender system modelling the extend to which different learning activities meet the learning needs of each of the students (Essa, 2016).

There are two main types of adaptive learning platforms referred to as the publisher model and the platform model (Oxman et al., 2014). In the publisher model, the system is designed with pre-existing content, often based on textbooks from the publisher. Pearsons MyLabs (using Knewton (Jose, 2016) for its adaptive functionality), McGraw-Hills LearnSmart and ALEKS (Falmagne et al., 2006) are established examples of this model. The publisher model has been very successful in K-12, where course content often have to comply with national standards. However, the overall adoption rate of adaptive learners that use the publisher model in higher

education has been very low, and has been mostly restricted to research projects (Essa, 2016). Limitations in tailoring the course content and the high per student cost of using these systems are among the main factors that contribute to the low adoption rate.

The platform model provides a content-agnostic system infrastructure that enables the teaching team to develop and author the content of their course. Smart Sparrow (Sparrow, 2016) and many learning management systems such as Desire2Learn, Loudcloud and edX that 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 amount of new content and object tagging, which introduces a significant overhead for the teaching time. RiPPLE uses crowdsourcing to overcome this challenge, lowering the required workload associated with the adoption of adaptive learning using the platform model.

2.3. RECOMMENDER SYSTEMS IN TECHNOLOGY ENHANCED LEARNING

RecSysTEL is an active and rapidly evolving research field. For example, Drachsler et al. (2015) perform an extensive classification of 82 different RecSysTEL environments, and Erdt et al. (2015) review the various evaluation strategies that have been applied in the field. Together these articles provide recent comprehensive surveys that consider more than 200 articles spanning over 15 years. Collaborative filtering (CF) identifies similar users and provides recommendations based upon their usage patterns. CF has been extensively employed in RecSysTEL; an early LAK paper, Verbert et al. (2011) evaluated and compared the performance of different CF techniques on educational data sets, showing that the best choice of algorithm is data dependent. In a more recent study, Kopeinik et al. (2017) also concluded that the performance of the algorithms strongly depends on the properties and characteristics of the particular dataset. In combining educational data sets with social networks, Cechinel et al. (2013) used CF to predict the utility of items for users based on their interest and the interest of the network of the users around them. Similarly, Fazeli et al. (2014) proposed a graph-based approach that uses graph-walking for improving performance on educational data sets.

One important way in which RecSysTEL has been used in an educational setting is to recommend personalized learning objects. Thus, Lemire et al. (2005) used inference rules to provide context aware recommendation on learning objects, and Mangina and Kilbride (2008) recommend documents and resources within e-learning environments to expand or reinforce knowledge. Interestingly, Gomez-Albarran and Jimenez-Diaz (2009) combined content based filtering with collaborative filtering to make recommendations in a student authored repository. When recommending learning objects (e.g. questions) related to student, Cazella et al. (2010) provided a semi-automated, hybrid solution based on CF (nearest neighbor) and rule based filtering, while Thai-Nghe et al. (2011) used students' performance predictions to recommend more appropriate exercises. CF techniques (basic, biased, and tensor matrix factorization) were used to address a number of different student behaviors and to model the temporal effect of students improving over time. Recently, Imran et al. (2016) provided an automated solution to personalize a learning management systems (LMS) using advanced learners' profiles to encapsulate their expertise level, prior knowledge, and performance in the course. The approach used association rule mining to create the learning object recommendations.

Matrix factorization (MF) is one of the most established techniques used in CF; however, despite its success in RecSys, MF has rarely been used in RecSysTEL. Out of the 124 papers

referenced by Drachsler et al. (2015), only two papers directly use it (Salehi, 2013; Thai-Nghe et al., 2011). This is somewhat surprising; MF has been put to good use in EDM for generating latent profiles of student expertise and so ought to combine with RecSysTEL in a straightforward manner. Indeed, the intelligent tutoring systems that appear to be utilized in that body of work could be seen as closely related to RecSysTEL, although they tend to give students less autonomy to accept or reject the pathways chosen for them (Chen, 2008). MF is particularly powerful in modeling students' performance and knowledge, because it implicitly incorporates guess and slip factors as latent factors (Thai-Nghe et al., 2011). The current implementation of RiPPLE uses the recommender system of Khosravi et al. (2017), which employs MF.

2.4. RECIPROCAL RECOMMENDATION

Reciprocal recommender systems has a literature rich with applications in online dating (Pizzato et al., 2013), job matching (Hong et al., 2013), and social networking (Guy, 2015), each highly tailored by the best consideration of the specific needs and constraints of the immediate environment. Fundamentally they entail operationalised user preferences and seek solutions such as a list of recommendations that match those preferences to a higher degree than competing items. In reciprocal recommendation, items are usually other users whose preferences must also be fulfilled and so entail a higher level of complexity than other recommender systems (Pizzato et al., 2010). Much of the research in social recommending has been developed and evaluated in existing social networks and particularly online dating sites where the site and/or the individual users may have a substantial data history which can be used to train machine learning algorithms (Cai et al., 2011; Chen and Nayak, 2013; Kutty et al., 2014). In peer recommending the environment may be a newly formed cohort of learners (such as a first-year university course) who have no data history but can volunteer preferences on a narrow selection of relevant variables for the purposes of matching.

By their nature, recommender systems across domains will exhibit similarities to a high degree. All systems incontrovertibly share the same fundamental goal - to provide recommendations that are well-received by users according to their preferences, explicit and implicit, in an otherwise overwhelming information environment where the likelihood of users successfully finding preferred items without technological assistance is very low. However the nature of domain-specific information and definition of a successful recommendation is so heavily context- and goal-dependent that little more than the general way of thinking can be adapted or generalised from existing systems to new domains. This is particularly true of the formulation of user preference models upon which recommendations are to be based, making them necessarily bespoke.

As such, while reciprocal peer recommendation has similarities to traditional recommendations in education and reciprocal recommendations in other domains, it does have a distinct nature. Although some primary research has been done on utilising peer learning and support for improving learning and enhancing the learning experience of students, the area remains fertile for many research and development opportunities. In this paper we introduce a platform that can enable adoption of peer support systems for both large on-campus and online courses with competency-based user preference models.

The current implementation of RiPPLE uses the reciprocal peer recommendation algorithm of (Potts et al., 2018).

3. FORMAL NOTATION

This section describes the data sources, functionality, the main aims and outputs of RiPPLE using a formal notation.

USERS AND QUESTIONS: Let $U_N = \{u_1 \dots u_N\}$ denote the set of users that are enrolled in a course in RiPPLE, where u_n and $u_{n'}$ refer to arbitrary users. Let $Q_M = \{q_1 \dots q_M\}$ denote the repository of the questions that are available to users in a course in RiPPLE, where q_m and $q_{m'}$ refer to arbitrary questions. All of the events occurring in RiPPLE are logged using a set of timestamps $T_T = \{\tau_1 \dots \tau_T\}$, where τ_t refers to an arbitrary timestamp. A three-dimensional array $C_{N \times M}^T$ keeps track of question creations, where c_{nm}^t indicates that user u_n has created question q_m at timestamp τ_t . Similarly, a three-dimensional array $A_{N \times M}^T$ keeps track of question answers, where $a_{nm}^t = 1$ indicates that user u_n has answered question q_m correctly at timestamp τ_t , and $a_{nm}^t = 0$ indicates that user u_n has answered question q_m incorrectly at timestamp τ_t . In addition, a three-dimensional array $D_{N \times M}^T$ provides information on question difficulty perceptions, where d_{nm}^t is the difficulty level user u_n has expressed for question q_m at timestamp τ_t . Furthermore, a three-dimensional array $R_{N \times M}^T$ provides information on question ratings, where r_{nm}^t is the rating user u_n has expressed for question q_m at timestamp τ_t .

KNOWLEDGE TRACING AND QUESTION RECOMMENDATION: Each course consists of a set of knowledge units $\Delta_L = \{\delta_1 \dots \delta_L\}$ referred to as a knowledge space, where δ_l refers to an arbitrary knowledge unit. Questions can be tagged with one or more knowledge units; $\Omega_{M \times L}$ is a two-dimensional array, where ω_{ml} is $\frac{1}{g}$ if question q_m is tagged with g knowledge units, including δ_l and 0 otherwise. One of the initial aims of RiPPLE is to use knowledge tracing algorithms to approximate the knowledge state of students on each knowledge unit. A three-dimensional array $\Lambda_{N \times L}^T$ is used for representing students' knowledge states approximated by the system, where λ_{nl}^t represents the knowledge state of u_n on δ_l at timestamp τ_t . This information is used to produce a three dimensional-array $\Theta_{N \times M}^T$, where θ_{nm}^t shows the personalised score of question q_m for user u_n at timestamp τ_t . $\Lambda_{N \times L}^T$ can be used for recommending questions to each user.

RECIPROCAL PEER RECOMMENDATION : Let $V_J = \{v_1 \dots v_J\}$ indicate a set of time slots, which denote the weekly available time slots for scheduling a study session, where v_j refers to an arbitrary time slot. Students can hold different roles $E_3 = \{e_1, e_2, e_3\}$ for participation in peer learning sessions, where e_1 is used to present providing peer learning support, e_2 for seeking peer learning support, and e_3 for searching for study partners. e_o refers to an arbitrary role. A four-dimensional array $S_{N \times L \times 3}^T$ represents the requests of the students, where $s_{nlo}^t = 1$ indicates that user u_n has indicated interest in participating in a study session on knowledge unit δ_l with role e_o at timestamp τ_t . In addition, a three-dimensional array $Z_{U \times V}^T$ represents the availability of students, where z_{nj}^t shows the availability of user u_n for time slot v_j at timestamp τ_t . Furthermore, A three-dimensional array $P_{N \times 3}^T$ shows competency preference of students in study sessions, where P_{no}^t shows the competency preference of user u_n for roles e_o at timestamp τ_t . For example, $P_{n1}^t = \epsilon$ means that at timestamp τ_t , u_n prefers providing support to peers with a competency of around ϵ less than his own. To be able to provide meaningful recommendations, we constrain eligibility by role such that users (1) provide support to less competent learners, (2) seek support from more competent learners and (3) find study partners with relatively similar

competency to that of their own. Using all of this information, RiPPLE aims to compute a three-dimensional array $\Psi_{N \times N}^T$, where $\psi_{nn'}^t$ shows the reciprocal score between user u_n and user $u_{n'}$ at timestamp τ_t . $\Psi_{N \times N}^T$ can be used for recommending peers to one another.

Table 1 provides a summary of the notation used for describing the functionality of RiPPLE.

Users and Questions	
U_N	A set of users $\{u_1 \dots u_N\}$, where u_n and $u_{n'}$ refer to arbitrary users.
Q_M	A set of questions $\{q_1 \dots q_M\}$, where q_m and $q_{m'}$ refer to arbitrary questions.
T_T	A set of time stamps $\{\tau_1 \dots \tau_T\}$, where τ_t refers to an arbitrary timestamp.
$C_{N \times M}^T$	A three-dimensional array where c_{nm}^t indicates that user u_n has created question q_m at timestamp τ_t .
$A_{N \times M}^T$	A three-dimensional array where $a_{nm}^t = 1$ indicates that user u_n has answered question q_m correctly at timestamp τ_t and 0 if answered incorrectly.
$D_{N \times M}^T$	A three-dimensional array where d_{nm}^t is the difficulty level user u_n has expressed for question q_m at timestamp τ_t .
$R_{N \times M}^T$	A three-dimensional array, where r_{nm}^t is the rating user u_n has expressed for question q_m at timestamp τ_t .
Knowledge Tracing and Question Recommendation	
Δ_L	A set of knowledge units referred to as knowledge space $\{\delta_1 \dots \delta_L\}$ where δ_l refers to an arbitrary knowledge unit.
$\Omega_{M \times L}$	A matrix, where ω_{ml} is $\frac{1}{g}$ if question q_m is tagged with g knowledge units, including δ_l and 0 otherwise.
$\Lambda_{N \times L}^T$	A three-dimensional array, where λ_{nl}^t represents the knowledge state of u_n on δ_l at timestamp τ_t .
$\Theta_{N \times M}^T$	A three-dimensional array, where θ_{nm}^t shows the personalised score of question q_m for user u_n at timestamp τ_t .
Reciprocal Peer Recommendation	
V_J	A set of time slots $\{v_1 \dots v_J\}$, which denote the weekly available time slots for scheduling a study session, where v_j refers to an arbitrary time slot.
E_3	A set of roles $\{e_1, e_2, e_3\}$ for participation in peer learning sessions where e_1 is used to present providing peer learning support, e_2 for seeking peer learning support, and e_3 for searching for study partners. e_o refers to an arbitrary role.
$S_{N \times L \times 3}^T$	A four-dimensional array, where $S_{nlo}^t = 1$ indicates that user u_n has indicated interest in participating in a study session on knowledge unit δ_l with role e_o at timestamp τ_t .
$Z_{U \times V}^T$	A three-dimensional array in which z_{nj}^t shows the availability of user u_n for time slot v_j at timestamp τ_t .
$P_{N \times 3}^T$	A three-dimensional array in which P_{no}^t shows the competency preference of user u_n for roles e_o at timestamp τ_t .
$\Psi_{N \times N}^T$	A three-dimensional array, where $\psi_{nn'}^t$ shows the reciprocal score between user u_n and user $u_{n'}$ at timestamp τ_t .

Table 1: A summary of the notation used for describing the functionality of RiPPLE

4. THE RIPPLE PLATFORM

This section presents an overview of the main features and functionalities that are supported by RiPPLE.

4.1. CREATING AND ANSWERING QUESTIONS

There are many benefits in engaging students as partners in co-creation of educational content (Bovill et al., 2011), and in particular, in creation of multiple-choice questions (Denny et al.,

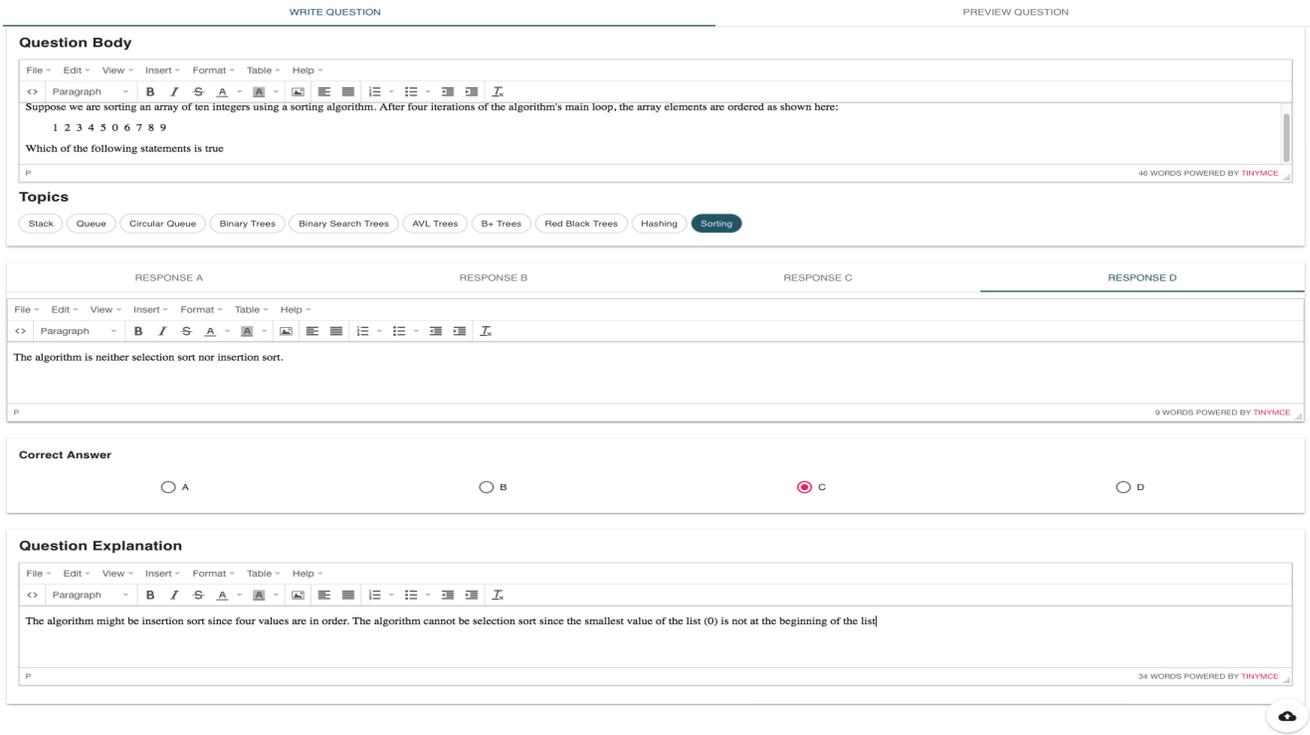


Figure 1: Overview of the question authoring page in RiPPLE.

2008). RiPPLE enables students to create questions and share them with their peers. Figure 1 shows the graphical interface used in RiPPLE for creating questions.

Creating a question includes the following steps:

1. Developing the body of the question. Text, images, videos and scientific formulas may be used in the development of the body of the question.
2. Tagging the question with one to four topics. The available topics are pre-defined by the instructor of the course.
3. Authoring the multiple-choice answers. Similar to the body of the question, Text, images, videos and scientific formulas may be used in the development of each of the multiple-choice answers.
4. Nominating the correct answer and developing a solution. An ideal solution includes rationale for the correctness of the right multiple-choice answer and the lack of correctness of the other multiple-choice answers.
5. previewing the question to make sure that it renders correctly and as expected.
6. submitting the question to be stored as part of the question repository.

Students are able to answer and rate questions that are available on the platform. Once they have answered a question, they are able to view the right answer, the distribution of how the question has been answered by their peers, and the explanation provided for the answer.

Students are then able to rate the quality and the difficulty of the question. Figure 2 shows the graphical interface used in RiPPLE for answering questions.

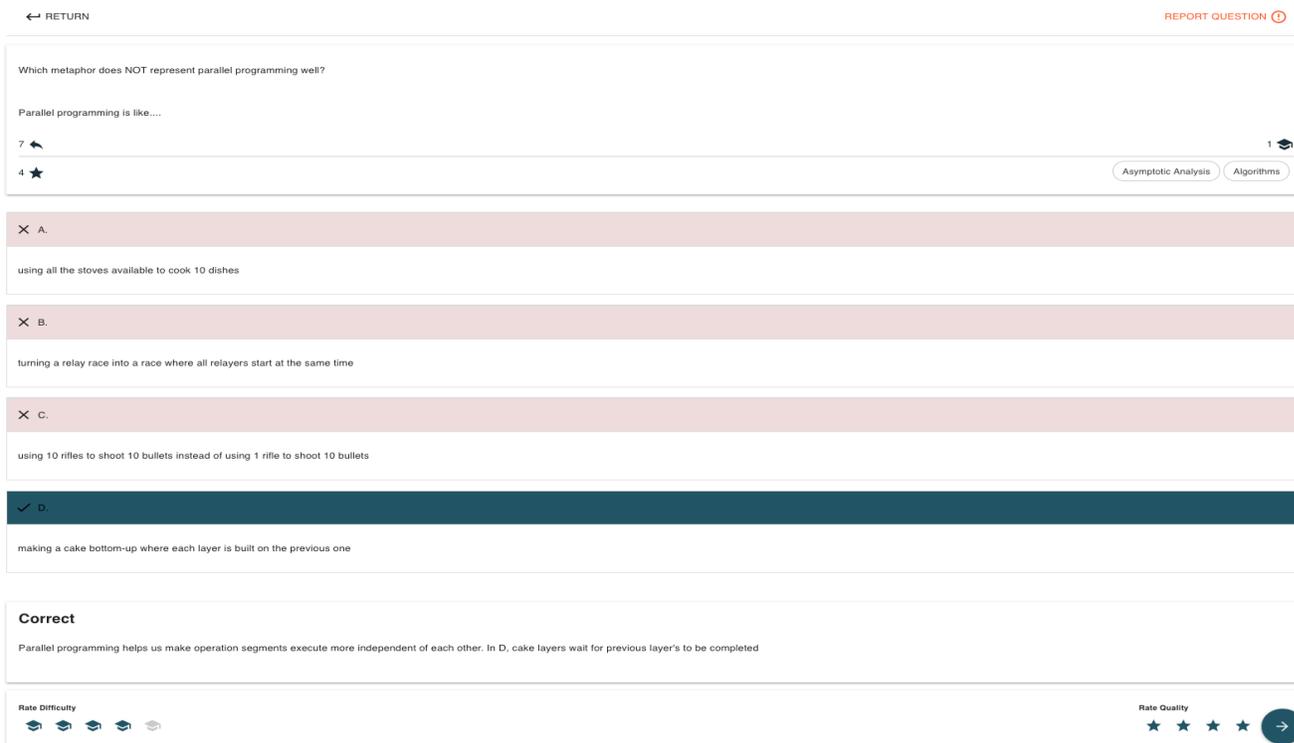


Figure 2: Overview of the question answering page in RiPPLE.

RiPPLE also relies on crowdsourcing to identify inappropriate or incorrect questions. Figure 3 shows how questions may be flagged as inappropriate in RiPPLE. Users with the “instruc-

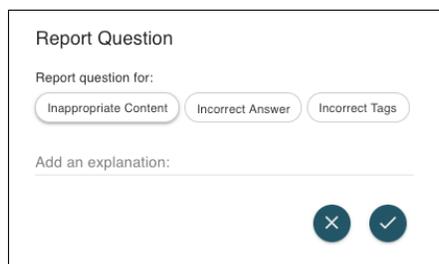


Figure 3: The interface used for flagging and reporting inappropriate questions in RiPPLE.

tor” role have the ability to view, edit and delete questions that are flagged as inappropriate.

4.2. KNOWLEDGE TRACING AND RECOMMENDING QUESTIONS

One of the main pages of RiPPLE, as shown in Figure 4, is dedicated to knowledge tracing and recommending questions. The top section of this page provides an interactive visualisation widget that enable learners to select their desired visualisation type for viewing their knowledge state.

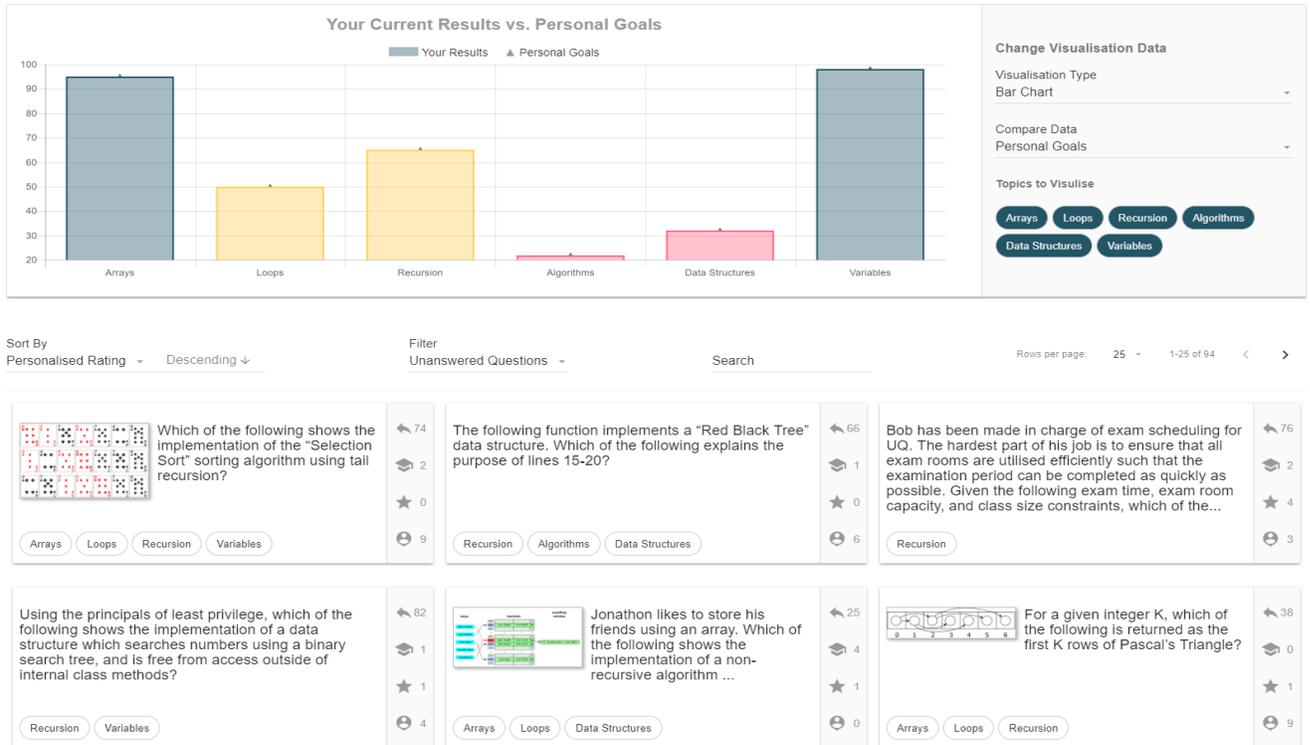


Figure 4: Overview of the knowledge tracing and question recommendation page of RiPPLE.

“Visualisation Type” allows learners to select from a range of different visualisation techniques so that they can choose a visual display that better suits their comprehension and personal preference. RiPPLE is equipped with a set of different types of visualisations: Bar charts (as displayed in Figure 4) are helpful in presenting rich information regarding users. They are simple to read and are comprehensible by a wide audience. Colour of the charts 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 topic. Radar charts are visually striking, and can add interest to what would otherwise be a dry data presentation. One of the strengths of radar charts is that they support visualisation of multiple variables consisting of measures that require different quantitative scales, which a bar chart cannot accommodate. They are very well suited for comparing the knowledge states of users students. Box plots are an effective way of displaying the distribution of data based on minimum, median, maximum, 1st quartile and 3rd quartile. Box plot displays may be preferred when displaying data of a group of learners as they determine if the data is skewed based on where the median sits within the box relative to the inter-quartile ranges. However, they are harder than both bar charts and radar charts to read and comprehend. The widget also supports the use of more recently developed visualisations that are tailored towards education. For example, it supports the use of Topic Dependency Models (Cooper and Khosravi, 2018), which use two-weighted graphs to display the knowledge state of learners not only on based on individual topics but also on combination of two or more topics.

“Compare Data” allows learners to compare their knowledge states against a range of options: “Peers” mode enables them to compare their performance with a selected distribution

(e.g. top 20%) of the peers that are currently enrolled in the course. "Previous Offerings" enables learners to compare their performance with a selected distribution of learners across all offerings of the course. "Topic to Visualise" option enables users to select topics in which their competencies are visualised.

The bottom section of this page, as shown in Figure 4, enables learners to select questions using search and recommendation functionalities. The "Sort By" option allows learners to sort questions based on their difficulty, quality, number of responses, number of comments or personalised rating. By selecting "Personalised Rating", the platform sorts the questions based on the outcome of recommender system. The "Filter" option enables users to filter the questions that are included in the results. They can request all questions (default), unanswered questions, answered questions, or wrongly answered questions to be included in the results. The "Search" option enables learners to search for questions based on specific content that may be present in the questions or multiple choice answers.

The results of the search are presented as a list of question cards, allowing users to engage with questions that best suit their needs. Figure 5 shows a sample question card. Each question card includes an overview of the question content, the topics associated with the question, and a sidebar in which the first icon shows the number of responses to the question, the second icon shows the average difficulty rating of the question, the third icon shows the average quality rating of the question, and the last icon shows the personalised rating that demonstrates the suitability of the question for each learner. Clicking on the question card will take the users to another page that would allow them to answer and rate the question.

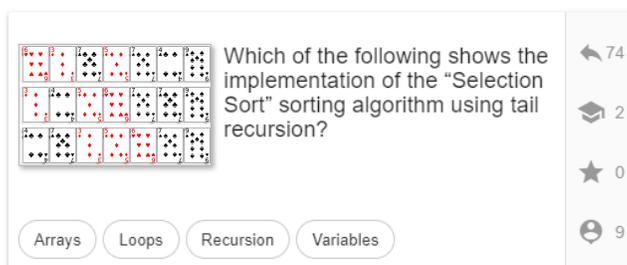


Figure 5: A sample question card in RiPPLE.

4.3. RECIPROCAL PEER RECOMMENDATION

One of the main features of RiPPLE is its ability to recommend peer learning sessions. Learners nominate their availability in hourly blocks, and their preferences for providing or seeking peer learning support and finding study partners across the range of course-relevant topics. Figure 6 shows the graphical interface used for capturing this information in RiPPLE.

The shade around each of the time slots provides an indication of the popularity of that time, where darker shades indicate a higher amount of interest in the time slot. The knowledge state of the student is provided in the form of a coloured bar chart superimposed over the list of topics. The option to provide peer support is only available for those topics in which the student meets a required competency threshold, denoted in the bar chart by the colour blue. The knowledge state of students are updated using the their cumulative performance on assessment

items progressively during the teaching period using algorithms described by in Khosravi et al. (2017).

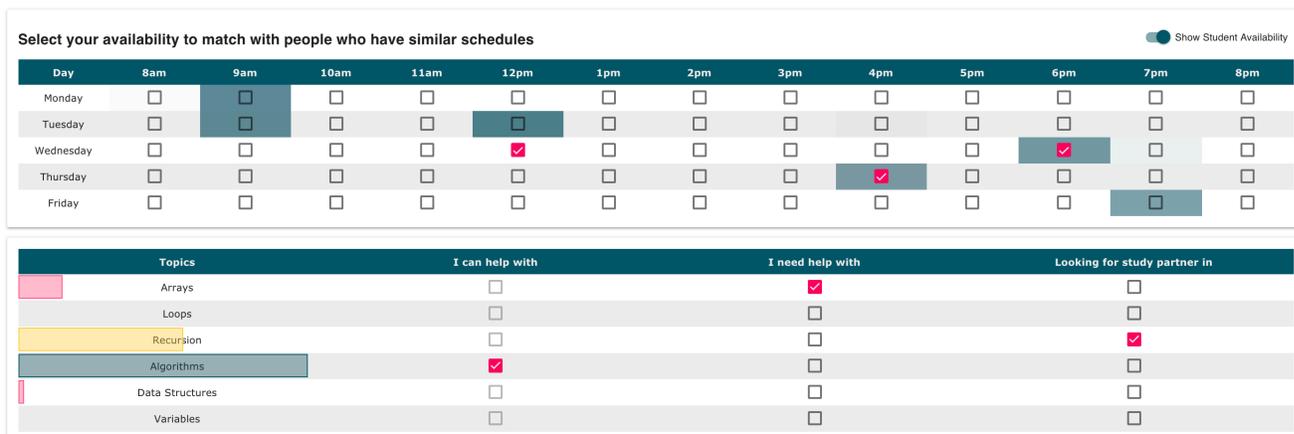


Figure 6: Overview of the peer recommendation interface in RiPPL.

Figure 7 provides an example of a peer support recommendation in the student interface, identifying the potential peer learning supporter and the two topics for which support is available. Students then have the option to ignore or request a meeting with the recommended peer at the nominated time and date.

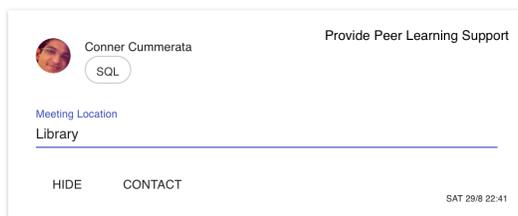


Figure 7: Example recommendation for a learner to provide peer learning support.

4.4. LEADERBOARD

A leaderboard, as its name implies, displays individuals with highest score on a give task. Use of leaderboards in education has shown to increase student motivation and engagement (Landers and Landers, 2014; Banfield and Wilkerson, 2014). The leaderboard in RiPPL 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. Figure 8 shows the leaderboard used in RiPPL.

4.5. PERSONAL PROFILE

Each student is provided with a personal profile that includes information on their engagement, achievements, notifications, and their consent on use of their data for educational research purposes.

Rank	Avatar	Name	Questions Contributed	Questions Answered	Answered Correctly	Questions Rated	Achievements Earned
1		Rebekah	0	0	0	0	0
2		Victoria	0	0	0	0	0
3		Gabrielle	0	0	0	0	0
4		Megan	0	0	0	0	0
5		Felicia	0	0	0	0	0
6		Brianna	0	0	0	0	0
7		Eric	0	0	0	0	0
8		Kelsey	0	0	0	0	0
9		Ryan	0	0	0	0	0
10		David	0	0	0	0	0

Rows per page: 10 1-10 of 20 < >

Figure 8: Overview of the leaderboard in RiPPLE

ENGAGEMENT The engagement level of students on a variety of tasks are presented using a visualisation widget. This widget enables students to compare their engagement using a visualisation type of their choice against their peers or their own targeted goals on a set of tasks. The default visualisation type uses Kiviat diagrams, which are more informally known as radar charts. Kiviat diagrams have been used extensively in visualising educational dashboard (e.g., see (May et al., 2011)) as they are able to display multivariate observations with an arbitrary number of variables (Chambers et al., 1983). The visualisation widget Figure 9 shows the chart used for presenting engagement in RiPPLE.

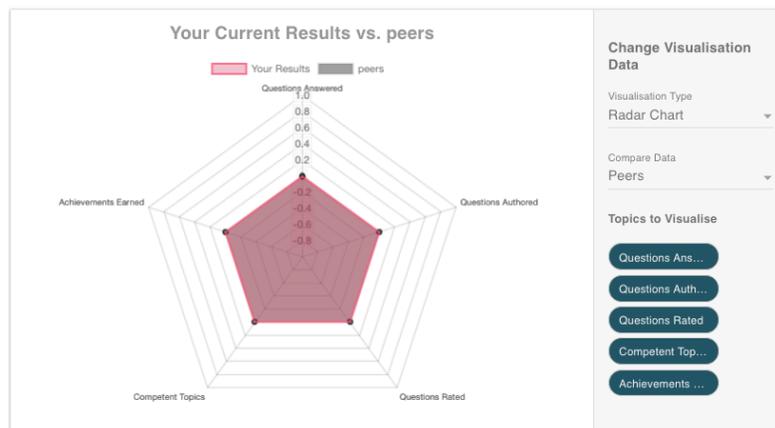


Figure 9: The visualisation widget used for showing engagement in RiPPLE

ACHIEVEMENTS RiPPLE uses Gamification and badging to increase student motivation and performance. Students are able to achieve badges in three broad categories of “Engagement Badges”, “Competency Badges” and “Peer Support Badges”. The achievement view enables

students to track their progress towards achievements. Figure 10 shows the graphical interface used for showing achievements in RiPPLE.

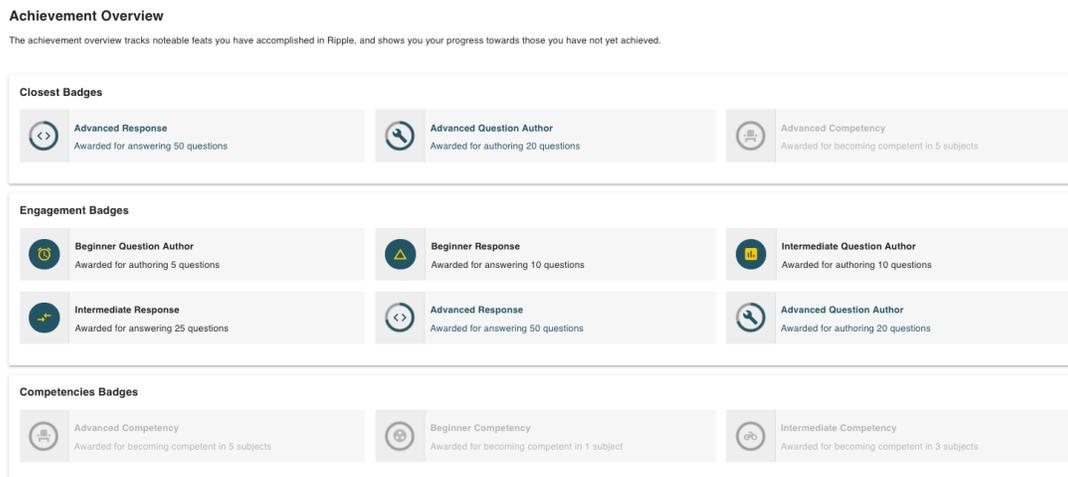


Figure 10: The graphical interface used for showing achievements in RiPPLE.

NOTIFICATIONS The notification view in RiPPLE allows students to view notifications about their achievements and up coming study sessions.

CONSENT Upon the first use of the platform, students are presented with a consent form seeking their permission to use their data to improve our understanding of the learning process and to evaluate the effectiveness of the recommended content. The consent view enables students to change their response at any time.

4.6. INSTRUCTOR PAGE

Users with “instructor” role have access to an additional page that has three main views: the course overview, consent form, and reported questions.

COURSE OVERVIEW This view enables instructors to add the set of topics that are to be used for tagging the questions. This list can be updated throughout the semester. This view also allows instructors to track the progress of each of the students that are enrolled in the course. Data related to their progress can be downloaded as a CSV or an SQL dump, which may be useful for educational data mining researchers.

CONSENT FORM This view enables instructors to develop the content of the consent form, which is to be filled by the students.

REPORTED QUESTIONS This view enables instructors to view, edit, and delete questions that have been flagged as inappropriate.

5. EXPECTED BENEFITS

This section discusses the multiple beneficial outcomes that RiPPLE provides for students, instructors, and educational data mining researchers.

RiPPLE enables students to:

- *think about ways of evaluating understanding and learning.* Designing questions requires students to think carefully about the topics of the course and focuses attention on the learning outcomes; choosing distractors requires students to consider misconceptions, ambiguity and possible interpretations of concepts; and writing explanations require students to express their understanding of a topic with as much clarity as possible, helping them develop their written communication skills and deepen their understanding (Denny et al., 2008).
- *identify their knowledge gaps.* Students often lack the requisite skills for making good decisions about what and how to study (Biggs, 1999), which can leave them undirected and time wasted. RiPPLE uses knowledge tracing algorithms to approximate knowledge states of students, enabling them to identify their knowledge gaps.
- *receive a more tailored learning experience.* Having course content that serves the needs of diverse student populations (e.g., those with differing academic ability, backgrounds, and generational expectations) is extremely challenging. RiPPLE allows students to receive a more personalised learning experience by recommending content based on their knowledge states.
- *compare themselves with their peers or against their target goals.* Students are often curious to know how they are performing compared to their peers or are keen to set personal goals and track their progress. RiPPLE uses knowledge tracing algorithms and visualisations that empower students to track their progress.
- *become more effective communicators.* RiPPLE promotes collaborative learning in a social, student-centred learning environment, where students learn to articulate their opinions, provide support for their views, and listen and relate to the views of others. Such learning communities lead to the development of “cognitive or intellectual skills or to an increase in knowledge and understanding”, or to the development of communication and professional skills (Falchikov, 2001).
- *enhance their social connectedness.* Many students, especially in their first year of university, will have difficulty navigating their new academic and social environment and adequately exploiting the available resources. RiPPLE enables students to grow their social connectedness, connecting students with peers who can adequately help each other for collaborative assessments or in reaching their academic aspirations.
- *increase their digital literacy.* Through exposure to innovative visualisation approaches, students will develop an appreciation for methods of communication and externalisation of knowledge from complex data sets.

RiPPLE enables instructors to:

- *utilize crowdsourcing to develop course content.* Implementing adaptively in a course requires a large amount of new content and learning object tagging. Through the use of crowdsourcing, RiPPLE enables instructors to develop a large data set of tagged course content.
- *provide rich and immediate feedback to students.* Instructors find it challenging to provide meaningful, rich and timely feedback at scale. RiPPLE provides immediate feedback to students on their progress and provides recommendations to help them overcome their knowledge gaps.
- *identify individual-level and course-level knowledge gaps.* Instructors often find it challenging, especially in large classes, to comprehend individual-level and course-level gaps. RiPPLE informs instructors of these gaps, so that they can update their course content accordingly.
- *identify at-risk students early in the semester.* Instructors find it challenging to identify at-risk students in large classes, while working within limited budgets. RiPPLE uses knowledge tracing algorithms to identify at-risk students.

Finally, RiPPLE enables learning analytics and educational data mining researchers to develop and validate their own knowledge tracing and recommender system algorithms. Currently, most researchers have to validate these algorithms using synthetic or historical data sets that do not provide compelling evidence that they lead to better learning. RiPPLE enables researchers to validate their algorithms in a live setting using parallel-group double-blind randomised trials or A/B Testing to determine whether recommendations lead to measurable gains.

6. IMPLEMENTATION DETAILS OF THE RIPPLE CLIENT

This section provides information on the implementation of the RiPPLE Client. Further details can be found on the Wiki³ page of the project on GitHub.

The RiPPLE client is a VueJS⁴ application augmented with Awesome Vue TS⁵ for the Typescript language. Its primary purpose is to consume the RiPPLE API supported by the RiPPLE Server (presented in Section 7.) to provide its information to the user.

6.1. SOURCE CODE CONVENTIONS AND OVERVIEW OF THE RIPPLE CLIENT

The RiPPLE client follows several conventions to ensure the scalability and managability of the software. Notably, it follows the Service and Repository, and Subscription patterns to ensure development consistency and reduce code duplication. It includes the following folders:

- The *project root* only contains files which are used for starting the build process, core application configuration, and files to dictate the behaviour of external tools interacting with the project (such as git, typescript, and npm).

³<https://github.com/hkhosrav/RiPPLE-Core/wiki>

⁴<https://vuejs.org/>

⁵<https://github.com/HerringtonDarkholme/av-ts>

- The *build directory* contains files required for the build process. This is a strange convention, but has been made popular by the Vue community and adopted by many JavaScript developers.
- The *config directory* contains files which are used to influence the build process.
- The *dist directory* is short for distributable. It contains the build files required for a production build. Files created from a development build are not placed here, and are instead stored in memory.
- The *docs directory* contains the same thing as the dist directory for each merge into master. This is to ensure that GitHub will serve the latest client build.
- The *node_modules directory* contains 3rd party project dependencies created by npm. It should not be in source control
- The *src directory* contains source files for the application.
- The *test directory* contains the specification verification of how the application operates. They are executed with *npm run test*.

Within the *src directory*, there are a few subdirectories. The *components directory* contains all of the views for the application. In this case, VueJS is acting as the view. The components directory is subdivided into multiple directories based on the purpose of the component. The *interfaces directory* contains the TypeScript interfaces for the application. These are used heavily throughout the application and typically mirror the *toJSON()* method of the RiPPLE Server models. The *repositories directory* contains entities which are responsible for retrieving the external data required by the application. It is encapsulated into a single place because it (1) provides a substitution point for testing (2) allows for a flexible system architecture for easier 'plugging-in' to changing parts (e.g., mocking data on client vs. retrieving from server), and (3) is consistent and predictable, which in turn provides for a better developer experience. Typically, the repositories are only interacted with from the services. This top-down approach is consistent with the Vue framework. The *routes directory* defines the routes of an application such that a URL will map to a Vue component. This is necessary in single-page applications which are dependent on state. The *services directory* contains a service layer which aims to encapsulate all of the business logic of the application. Given the asynchronous nature of the application, a subscription system called *Fetcher* also exists here. Finally, the *style directory* contains global CSS styles for the application.

6.2. DEPLOYMENT OF THE RIPPLE CLIENT

The RiPPLE client architecture is relatively simple, however it does need some configuration to work correctly.

6.2.1. Installing Dependencies

From your terminal, make sure you can access both nodeJS (*node -version*) and npm (*npm -version*). If not, you will be unable to continue. This project ships with a *package.json* file, which details all of the modules required to build and run the application. The simplest way to install them all is to run *npm install* from your terminal.

npm will automatically read the package.json file and install everything it requires. Generally you only need to do this once, but if new dependencies are added you will need to re-run *npm install*.

6.2.2. General Configuration

RiPPLE client side will read its configuration from what is called an *environment variable*. These variables are injected into the application at runtime by replacing all instances of *process.env.** with the name of predefined variable. These predefined variables are things such as debug mode and where to look for the API. Since environment variables can sometimes be sensitive (such as database passwords), it is best practice to keep them out of source control.

RiPPLE reads its environment variables from environment files, which are named *.env.dev*, and *.env.prod* for development and production respectively. In a fresh install, the *.env.dev* and *.env.prod* files will not exist (since they are not in source control). Instead, a *.env.example* file is provided which has all of the possible environment variable names, and example values for them.

If you make your own *.env.dev* or *.env.prod* files, you can place whatever you like in it, and your git client will automatically ignore anything inside of it. Supported environment values include the following:

- **API_LOCATION** - string: A URI which points to the RiPPLE Server. Should not have a trailing *'/'*.
- **Node_ENV** - string: A string indicating the development environment as **PRODUCTION** or **DEVELOPMENT**.

6.2.3. Development Configuration

Typically, a RiPPLE client in development mode will connect to a Development RiPPLE Server (See Section 7.). By default, this will be *http://localhost:9000*. You will need to initialize your *.env.dev* environment file to development values. Once your environment file has been created and your [dependencies installed]RiPPLE-Client—Environment) you will be able to build and test your application. The key command for active development is *npm run dev*, which builds your application and opens it in your browser.

DEVELOPMENT COMMANDS This project comes packaged with several convenience commands

- *npm run dev*: Creates a local development server, and serves your files over it. When a local file is changed your project will be rebuilt incrementally and the new code injected into your web browser.
- *npm run lint*: Runs eslint over the codebase to ensure it meets the specified JavaScript style guide.
- *npm run unit*: Runs all unit tests for the project. They are located under *./test/unit/spec/*
- *npm run e2e*: Runs all integration tests for the project. They are located under *./test/e2e/spec/*. The current webdriver used is PhantomJS (runs headlessly). It is installed locally as part of *npm install*.

- *npm run test*: Runs all unit tests and e2e tests. It is the same as running both *npm run unit* and *npm run e2e*.

HELPFUL DEVELOPMENT TOOLS The Chrome Console⁶ allows you to inspect your application at runtime. All errors will also be reported to this console which allows for easy issue identification. The console also has a debugger, which is very useful for inspecting code to make sure it behaves correctly. The Vue Devtools⁷ allows inspection of Vue components.

6.2.4. Production Configuration

Typically, a RiPPLE client in production mode will connect to a Production RiPPLE Server, which needs to be manually configured. You will need to initialize your `.env.prod` environment file to production values. An example environment file is provided in the repository, but the defaults are not suitable for a production application. Once your environment file has been created and your dependencies installed you will be able to build your application. The key commands for building is *npm run test* and then *npm run build*, which builds your application and places it into the *dist directory*. This directory is on the `.gitignore`, so any changes to it will not be committed to the project.

Once you have built your application, you will need to serve it via an actual webserver. This is most commonly done by placing it into `/var/www/htdocs/` of a machine configured to run `nginx`⁸.

This GitHub project is configured to serve the *docs directory* of the master branch on demo page⁹. You should therefore remember to copy the result of `./dist` into `./docs` before a merge into master to ensure that the latest version is being served via GitHub.

7. IMPLEMENTATION DETAILS OF THE RIPPLE SERVER

This section provides information on the implementation of the RiPPLE Sever. Further details can be found on the Wiki¹⁰ page of the project on GitHub.

The RiPPLE server is a Django¹¹ application running on Python. Its primary purpose is to provide a data gateway to user information, authentication through LTI, and personalised recommendations for the student competency

7.1. SOURCE CODE CONVENTIONS AND OVERVIEW OF THE RIPPLE SERVER

The RiPPLE server follows several conventions to ensure the scalability and managability of the software. Notably, it follows the Service and Repository, as well as Django-community favoured patterns to ensure development consistency and reduce code duplication. It includes the following folders:

⁶<https://developers.google.com/web/tools/chrome-devtools/console/>

⁷<https://github.com/vuejs/vue-devtools>

⁸<https://nginx.org/en/>

⁹<https://hkhosrav.github.io/RiPPLE-Core/#/>

¹⁰<https://github.com/hkhosrav/RiPPLE-Core/wiki>

¹¹<https://www.djangoproject.com/>

- The *project root* only contains project-specific files which are used for getting started in the project. This includes things like dependencies, licensing, and a production deployment scripts
- The *src directory* contains source files for the application. You will spend lots of time in this directory. It follows the Django Project Structure, and should contain environment files and Django bootstrap files (such as `manage.py`).
- The *ripple directory* contains the core application configuration. This directory contains files important configuration files such as `settings.py`, `urls.py` which are necessary to tie different apps together.
- The *src directory* contains many subfolders, each of which is referred to as a Django App¹². To quote from Django: "Whats the difference between a project and an app? An app is a Web application that does something e.g., a Weblog system, a database of public records or a simple poll app. A project is a collection of configuration and apps for a particular website. A project can contain multiple apps. An app can be in multiple projects." In order to promote code re-usability; the RiPPLE server makes use of projects where appropriate.
- The *services directory* inside each app contains entities which are responsible for handling the business logic of the application. They should be used exclusively by the routing controllers to ensure that tasks are always completed in the same way (e.g., using the QuestionSearch is beneficial since it will automatically ensure your search results are within the current course context). Services are encapsulated into a single place.

Each App has the following files

- *urls.py*, which contains application route definitions. Routes are not automatically added to the global application, you must also modify *ripple/urls.py* to import your app routes.
- *views.py*, which contains the application controllers.
- *tests.py*, which contains the specification verification of how the application operates. They are executed with *python manage.py test*.

7.2. DEPLOYMENT OF THE RIPPLE SERVER

The RiPPLE server architecture is relatively simple since it builds off of one of the most popular python frameworks, however it does need some configuration to work correctly.

7.2.1. Installing Dependencies

From your terminal, make sure you can access both python (`python--version`) and pip (`pip --version`). If not, you will be unable to continue.

It is highly recommended to run python through `virtualenv` to ensure you are developing in an isolated environment. Django recommends doing this

¹²<https://docs.djangoproject.com/en/1.11/intro/tutorial01/#creating-the-polls-app>

This project ships with a *requirements.txt* file, which details all of the modules required to build and run the application. The simplest way to install them all is to run *pip install -r requirements.txt* from your terminal at the project root. *pip* will automatically read the *requirements.txt* file and install everything it requires. Generally you only need to do this once, but if new dependencies are added you will need to re-run *pip install -r requirements.txt*.

7.2.2. General Configuration

The RiPPLE server will read its configuration from what is called an *environment variable*. These variables are injected into the application at runtime by replacing reading from *os.getenv()*. These predefined variables are things such as debug mode and application secrets. Since environment variables can sometimes be sensitive (such as database passwords), it is best practice to keep them out of source control.

RiPPLE reads its environment variables from environment files, which are named *.env.dev*, and *.env.prod* for development and production respectively.

In a fresh install, the *.env.dev*, and *.env.prod* files will not exist (since they are not in source control). Instead, an *.env.example* file is provided which has all of the possible environment variable names, and example values for them. If you make your own *.env.dev*, or *.env.prod* files, you can place whatever you like in it, and your git client will automatically ignore anything inside of it. Supported environment values include the following:

- **API_LOCATION** - string: A URI which points to the RiPPLE Server. Should not have a trailing *'/'*.
- **DEVELOPMENT_ENVIRONMENT** - string: A string indicating the development environment, e.g., (PRODUCTION or DEVELOPMENT)
- **DJANGO_KEY** - string: The SECRET_KEY environment variable required by django. It should be a large unique string
- **PROXY_LOCATION** - string: The subpath of the proxy_pass location if the application is being run through nginx and is not on the root path.
- **LTI_SUCCESS_REDIRECT** - string: URI to redirect to after a successful LTI validation request
- **LTI_URL** - string: URI to use in LTI validation
- **LTI_APP_KEY** - string: Application key to pass to LTI validation service
- **DATABASE_TYPE** - string: Database type to use (e.g., mysql, sqlite3)
- **DATABASE_NAME** - string: Name of database to use
- **DATABASE_HOST** - string: Host of database
- **DATABASE_USER** - string: Username to use when connecting to database
- **DATABASE_PASSWORD** - string: Password to use authenticate with on database

7.2.3. Development Configuration

Typically, a RiPPLE server in development mode will be running on the loopback address of the machine; this is `http://localhost:9000` by default.

To initialise your development environment, you have two options: (1) run `python manage.py env dev` - which will prompt to create a `.env.dev` if none exists or (2) run `cp .env.example .env.dev` - which will copy the example configuration into a development file.

Once your environment file has been created and your dependencies installed you will be able to build and test your application. The key command for active development is `python manage.py runserver`, which builds your application and enables hot-reload.

From a fresh install, the development workflow is typically:

1. `pip install -r requirements.txt` - Install project dependencies
2. `cd src` - Change to project working directory
3. `CREATE_SCHEMA` - If not using `sqlite3`, then the schema must manually be created
4. `python manage.py migrate` - Creates database schema on machine
5. `python manage.py seed` - Populates database with values
6. `python manage.py seedCourse --name courseName --course corseCode --file /path/to/JSONfile --host /host/domain` - Populates database with values from file
7. `python manage.py runserver` - Starts the server
8. edit files and make changes

RUNNING IN UNAUTHENTICATED MODE If you wish to run the application without authentication, change the `ALLOW_UNAUTHENTICATED` parameter in `src/ripple/settings.py` to `True`. This will disable token verification on login requests, and instead assign the requester a random user token when the `/users/login` endpoint is accessed. please note that the unauthenticated mode should never be enabled on a live system with real users.

DEVELOPMENT COMMANDS

- `python manage.py test`: Runs the application unit tests
- `python manage.py makemigrations`: Reads in all changes from project models and creates migration files from them
- `python manage.py migrate`: Updates the database schema to match the migration definitions
- `python manage.py seed` Seeds the database with mock data
- `python manage.py env ENV_NAME` `ENV_NAME` must be either `”prod”` or `”dev”`. It reads in the environment definition from `.env.prod` or `.env.dev` respectively, and makes it the active environment

- *python manage.py runserver*: Spawns a webserver on `http://localhost:9000` running the application with hot-reload
- *python manage.py runsslserver*: Spawns a HTTPS compatible server on `https://localhost:9000` running the application with hot-reload

7.2.4. Production Configuration

Typically, a RiPPLE server in development mode will be running on the loopback address of the machine through the use of a stable webserver (e.g., gunicorn), but with an external web server (e.g., nginx) proxying the request into the loopback address.

You will need to initialize your *.env.prod* environment file to production values, which should be kept secret. An example *.env* file is provided in the repository (*.env.example*), but *python manage.py env ENV_NAME* will also create an *.env.current* file if none exists.

Ensure that *./src/ripple/settings.py* has `ALLOW_UNAUTHENTICATED` set to `False` if you are deploying to a live system

To initialise your production environment, you have two options: (1) run *python manage.py prod* - which will prompt to create a *.production.env* if none exists and (2) run *cp .env.example .env.prod* - which will copy the example configuration into a development file

Once your environment file has been created and your dependencies installed you will be able to build and test your application. The key command for production development is *sudo -E ./deploy.sh .env.prod*, which will deploy your application and run it as a service. Running the application as a service is the best way to ensure it continues to run after an SSH disconnect

8. CONTRIBUTING

RiPPLE is an open-source system - and always will be. Collaboration, improvements, suggestions and pull requests are always welcome. In order to streamline the development process; the following guide exists to help people get started:

1. Open a ticket in the ticketing system - this will serve as a point of contact to the project maintainers.
2. Either: a. Create a branch of the format `RIPPLE-#:ticketId`. (eg. `RIPPLE-#33` for ticket #33). b. Fork the project and do your work there.
3. When the work is finished, create a pull request to indicate your changes are ready to be reviewed. In your pull request; document your changes and preferably link to your ticket.
4. Your code will be reviewed by a project contributor, and will be merged in if it aligns with the projects goals (which was hopefully discussed previously in your ticket).
5. After your code has been merged in, this documentation will be updated as appropriate.

REFERENCES

BAKER, R., CORBETT, A. T., AND ALEVEN, V. 2008. More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing. In *International Conference on Intelligent Tutoring Systems*. Springer, 406–415.

- BANFIELD, J. AND WILKERSON, B. 2014. Increasing student intrinsic motivation and self-efficacy through gamification pedagogy. *Contemporary Issues in Education Research (Online)* 7, 4, 291.
- BIGGS, J. 1999. What the student does: teaching for enhanced learning. *Higher education research & development* 18, 1, 57–75.
- BOVILL, C., COOK-SATHER, A., AND FELTEN, P. 2011. Students as co-creators of teaching approaches, course design, and curricula: implications for academic developers. *International Journal for Academic Development* 16, 2, 133–145.
- CAI, X., BAIN, M., KRZYWICKI, A., WOBCKE, W., KIM, Y. S., COMPTON, P., AND MAHIDADIA, A. 2011. Learning to make social recommendations: a model-based approach. In *International Conference on Advanced Data Mining and Applications*. Springer, 124–137.
- CAZELLA, S., REATEGUI, E., AND BEHAR, P. 2010. Recommendation of learning objects applying collaborative filtering and competencies. In *Key Competencies in the Knowledge Society*. Springer, 35–43.
- CECHINEL, C., SICILIA, M.-Á., SÁNCHEZ-ALONSO, S., AND GARCÍA-BARRIOCANAL, E. 2013. Evaluating collaborative filtering recommendations inside large learning object repositories. *Information Processing & Management* 49, 1, 34–50.
- CEN, H., KOEDINGER, K., AND JUNKER, B. 2006. Learning factors analysis—a general method for cognitive model evaluation and improvement. In *Intelligent tutoring systems*. Vol. 4053. Springer, 164–175.
- CHAMBERS, J. M., CLEVELAND, W. S., KLEINER, B., TUKEY, P. A., ET AL. 1983. *Graphical methods for data analysis*. Vol. 5. Wadsworth Belmont, CA.
- CHEN, C.-M. 2008. Intelligent web-based learning system with personalized learning path guidance. *Computers & Education* 51, 2, 787–814.
- CHEN, L. AND NAYAK, R. 2013. A reciprocal collaborative method using relevance feedback and feature importance. In *Proceedings of the 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT) - Volume 01*. WI-IAT '13. IEEE Computer Society, Washington, DC, USA, 133–138.
- COOPER, K. AND KHOSRAVI, H. 2018. Graph-based visual topic dependency models. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*.
- CORBETT, A. 2001. Cognitive computer tutors: Solving the two-sigma problem. *User Modeling 2001*, 137–147.
- CORBETT, A. T. AND ANDERSON, J. R. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* 4, 4, 253–278.
- DENNY, P., HAMER, J., LUXTON-REILLY, A., AND PURCHASE, H. 2008. Peerwise: students sharing their multiple choice questions. In *Proceedings of the fourth international workshop on computing education research*. ACM, 51–58.
- DRACHSLER, H., VERBERT, K., SANTOS, O. C., AND MANOUSELIS, N. 2015. Panorama of recommender systems to support learning. In *Recommender systems handbook*. Springer, 421–451.
- ERDT, M., FERNÁNDEZ, A., AND RENSING, C. 2015. Evaluating recommender systems for technology enhanced learning: A quantitative survey. *IEEE Transactions on Learning Technologies* 8, 4, 326–344.
- ESSA, A. 2016. A possible future for next generation adaptive learning systems. *Smart Learning Environments* 3, 1, 16.
- FALCHIKOV, N. 2001. *Learning together: Peer tutoring in higher education*. Psychology Press.

- FALMAGNE, J.-C., COSYN, E., DOIGNON, J.-P., AND THIÉRY, N. 2006. The assessment of knowledge, in theory and in practice. In *Formal concept analysis*. Springer, 61–79.
- FAZELI, S., LONI, B., DRACHSLER, H., AND SLOEP, P. 2014. Which recommender system can best fit social learning platforms? In *European Conference on Technology Enhanced Learning*. Springer, 84–97.
- GOMEZ-ALBARRAN, M. AND JIMENEZ-DIAZ, G. 2009. Recommendation and students authoring in repositories of learning objects: A case-based reasoning approach. *International Journal of Emerging Technologies in Learning (iJET)* 4, 2009, 35–40.
- GUY, I. 2015. Social recommender systems. In *Recommender Systems Handbook*. Springer, 511–543.
- HONG, W., ZHENG, S., WANG, H., AND SHI, J. 2013. A job recommender system based on user clustering. *Journal of Computers* 8, 8, 1960–1967.
- IMRAN, H., BELGHIS-ZADEH, M., CHANG, T.-W., GRAF, S., ET AL. 2016. Plors: a personalized learning object recommender system. *Vietnam Journal of Computer Science* 3, 1, 3–13.
- JOSE, F. 2016. White paper: Knewton adaptive learning building the worlds most powerful recommendation engine for education.
- KHAJAH, M. M., HUANG, Y., GONZÁLEZ-BRENES, J. P., MOZER, M. C., AND BRUSILOVSKY, P. 2014. Integrating knowledge tracing and item response theory: A tale of two frameworks. In *Proceedings of Workshop on Personalization Approaches in Learning Environments (PALE 2014) at the 22th International Conference on User Modeling, Adaptation, and Personalization*. University of Pittsburgh, 7–12.
- KHOSRAVI, H., COOPER, K., AND KITTO, K. 2017. Riple: Recommendation in peer-learning environments based on knowledge gaps and interests. *JEDM-Journal of Educational Data Mining* 9, 1, 42–67.
- KOPEINIK, S., LEX, E., SEITLINGER, P., ALBERT, D., AND LEY, T. 2017. Supporting collaborative learning with tag recommendations: a real-world study in an inquiry-based classroom project. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*. ACM, 409–418.
- KUTTY, S., NAYAK, R., AND CHEN, L. 2014. A people-to-people matching system using graph mining techniques. *World Wide Web* 17, 3, 311–349.
- LANDERS, R. N. AND LANDERS, A. K. 2014. An empirical test of the theory of gamified learning: The effect of leaderboards on time-on-task and academic performance. *Simulation & Gaming* 45, 6, 769–785.
- LEMIRE, D., BOLEY, H., MCGRATH, S., AND BALL, M. 2005. Collaborative filtering and inference rules for context-aware learning object recommendation. *Interactive Technology and Smart Education* 2, 3, 179–188.
- MANGINA, E. AND KILBRIDE, J. 2008. Evaluation of keyphrase extraction algorithm and tiling process for a document/resource recommender within e-learning environments. *Computers & Education* 50, 3, 807–820.
- MAY, M., GEORGE, S., AND 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.
- OXMAN, S., WONG, W., AND INNOVATIONS, D. 2014. White paper: Adaptive learning systems. *Integrated Education Solutions*.

- PARAMYTHIS, A. AND LOIDL-REISINGER, S. 2003. Adaptive learning environments and e-learning standards. In *Second european conference on e-learning*. Vol. 1. 369–379.
- PARDOS, Z. AND HEFFERNAN, N. 2011. Kt-idem: introducing item difficulty to the knowledge tracing model. *User Modeling, Adaption and Personalization*, 243–254.
- PAVLIK JR, P. I., CEN, H., AND KOEDINGER, K. R. 2009. Performance factors analysis—a new alternative to knowledge tracing. *Online Submission*.
- PIECH, C., BASSEN, J., HUANG, J., GANGULI, S., SAHAMI, M., GUIBAS, L. J., AND SOHL-DICKSTEIN, J. 2015. Deep knowledge tracing. In *Advances in Neural Information Processing Systems*. 505–513.
- PIZZATO, L., REJ, T., AKEHURST, J., KOPRINSKA, I., YACEF, K., AND KAY, J. 2013. Recommending people to people: the nature of reciprocal recommenders with a case study in online dating. *User Modeling and User-Adapted Interaction* 23, 5, 447–488.
- PIZZATO, L., REJ, T., CHUNG, T., YACEF, K., KOPRINSKA, I., AND KAY, J. 2010. Reciprocal recommenders. In *8th Workshop on Intelligent Techniques for Web Personalization and Recommender Systems, UMAP*.
- POTTS, B., KHOSRAVI, H., REIDSEMA, C., BAKHARIA, A., BELONOGOFF, M., AND FLEMING, M. 2018. Reciprocal peer recommendation for learning purposes. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*.
- SALEHI, M. 2013. Application of implicit and explicit attribute based collaborative filtering and BIDE for learning resource recommendation. *Data & Knowledge Engineering* 87, 130–145.
- SHA, L. AND HONG, P. 2017. Neural knowledge tracing. In *International Conference on Brain Function Assessment in Learning*. Springer, 108–117.
- SPARROW, S. 2016. Smart sparrow - adaptive elearning platform.
- THAI-NGHE, N., DRUMOND, L., HORVÁTH, T., KROHN-GRIMBERGHE, A., NANOPOULOS, A., AND SCHMIDT-THIEME, L. 2011. Factorization techniques for predicting student performance. *Educational Recommender Systems and Technologies: Practices and Challenges*, 129–153.
- VERBERT, K., DRACHSLER, H., MANOUSELIS, N., WOLPERS, M., VUORIKARI, R., AND DUVAL, E. 2011. Dataset-driven research for improving recommender systems for learning. In *Proceedings of the 1st International Conference on Learning Analytics and Knowledge*. ACM, 44–53.
- YILMAZ, B. 2017. Effects of adaptive learning technologies on math achievement: A quantitative study of aleks math software. Ph.D. thesis, University of Missouri-Kansas City.
- YUDELSON, M. V., KOEDINGER, K. R., AND GORDON, G. J. 2013. Individualized bayesian knowledge tracing models. In *International Conference on Artificial Intelligence in Education*. Springer, 171–180.