Effective Learning Recommendations Powered by AI Engine

Xiaodong Dang
Adaptive Learning Research, Adaptemy
27 Lower Mount Street, Dublin
Ireland
xiaodong.dang@adaptemy.com

Ioana Ghergulescu
Adaptive Learning Research, Adaptemy
27 Lower Mount Street, Dublin
Ireland
ioana.ghergulescu@adaptemy.com

ABSTRACT

This paper presents the learning effectiveness evaluation of a recommender system powered by Adaptemy's AI Engine in terms of average lesson success rate and improvement per lesson. The data from over 80k lessons were used in this analysis. Three main cases are considered based on the level of teachers' guidance. The first case is when the system makes recommendation with no input from the teacher, the second case is when the system recommendations are loosely-guided by teacher input through assignment in a topic, and the third case is when the lessons are done on concepts that are specified by teachers while the system-given recommendation is ignored. In each case the results are compared between the lessons done on system-recommended concepts and the lessons done on other concepts. The results have shown that both the learning success-rate and the improvement per lesson are higher if the system-based recommendations are followed, in all the three cases.

CCS CONCEPTS

•**Applied computing** → **Education** → *Interactive learning environments;* •**Information systems** → **Information retrieval** → *Retrieval tasks and goals* → Recommender systems

KEYWORDS

Recommender systems, AI engine, learning experience, technology-enhanced learning, learning effectiveness

1 INTRODUCTION

With the amount of learning material available on the Internet, there is bigger uptake of recommender systems in Technology-Enhanced Learning [1], and a bigger need for effective recommender systems and their evaluation through real-life testing [2]. In order to support learning, recommender systems for TEL need to consider specific learning aspects which differ from recommender systems from other domains such as e-commerce [2]. In their review, Drachsler et al. [1] reviewed and classified the recommender systems in TEL in terms of relevant contributions to the field, categorising them in 7 clusters: TEL RecSys that follow collaborative filtering approaches as in other domains, TEL RecSys that propose improvements to collaborative filtering by taking into account specifics of the TEL domain, TEL RecSys that take into consideration educational constraints, TEL RecSys that explore the alternative collaborative filtering approaches, TEL RecSys that consider learning contextual information, TEL RecSys that assess the educational impact of the recommendations, and TEL RecSys that focus on recommending courses. As the authors presented in their review, the recommender system and engine could be informed by complex information from learner model, domain model and personalisation model. Furthermore, other research studies presented the need for the recommender systems to use as input complex information such as: learning ability, learner needs [3], affective and motivational state [4], ontologies about the learner and the content [5], learner context and domain [6], or to use complex engines such as neuronal networks or Bayesian networks [7]. However, there are few solutions for recommender systems in TEL that make use of complex information and that were implemented in a system used in a real- world scenario.

This paper introduces the recommender system of the Adaptemy platform that is powered by an Artificial Intelligence Engine. The Adaptemy platform performs several layers of adaptation and personalisation for students such as: personalised feedback, personalised content sequence, tailored interventions when disengagement and demotivation was detected as well as learning paths recommendations. The recommender system makes use of complex information involving the user, content, domain and context. The recommender system recommends to learners what is the next most suitable concept to study, the recommended action together with personalized guidance and evaluation. The system gives flexibility to teachers in terms of how they want to use it: system-independent recommendation with no input from them, system recommendations loosely-guided by teacher, teacher direct recommendations to the class where they overwrite the system's recommendations.

Following a layered evaluation approach as suggested by the literature [8], this paper evaluates the recommendation layer with its recommender system powered by the Adaptemy's AI engine when the system was used in real-life contexts and at large scale. The evaluation is focused on the educational impact of the recommendations to the learning effectiveness. Success-rate and improvement per lesson [1] are the two metrics that are used in the learning effectiveness evaluation depending if students are following or not following the system recommendations. The evaluation study made use of data corresponding to 4257 students from secondary schools and 80266 learning lessons finished between September 2017 and March 2018, covering 211 unique concepts of a Maths course.

The paper is structured as follows: section 2 provides an overview of the Adaptemy system, section 3 presents the evaluation methodology, section 4 analyses the results and discusses the research findings, while section 5 concludes the paper.

2 ADAPTEMY SYSTEM

2.1 Overview

Adaptemy system is an intelligent personalized learning environment that is developed based on existing research in the areas of Intelligent Tutoring Systems and Adaptive E-Learning [9]. It follows the classical architecture of an adaptive and intelligent e-learning system that makes a separation between curriculum model, content model, user model, and adaptation engine.

The curriculum model includes concepts and the relationships between them (see Fig. 1 for an example a topic map). The list of prerequisite links allows for misconception detections, complex user model and enables multiple layers of personalization and adaptation. The content model contains all the metadata about the content and up to date analytics. The rich information from the content model enables the AI engine to personalize the learning loop and to accurately update the user model. For example, each question has attached information such as: difficulty level, discriminant, probability of guessing, probability of having a slip and expected time to solve the question.

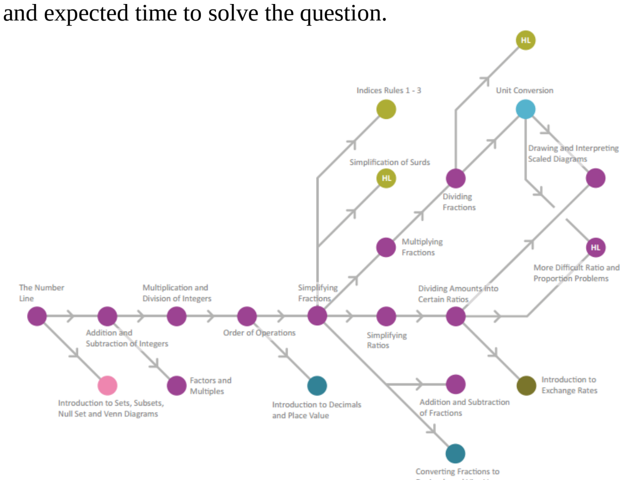

Figure 1: Example of topic map for maths curriculum.

Students learn via Adaptemy by doing lessons, each of which is on a single concept consisting of a group of questions. For each student, an ability profile on all the concepts in the curriculum is maintained in the system, which is updated by the lesson outcomes. On each concept, the ability profile is represented by a vector of 100 elements giving the probability densities of the ability level being from 1 to 100. As they finish a lesson, the ability profile of the concept that is worked on during the lesson is updated based on the direct evidence using a customised Item Response Theory (IRT) model. The profile on the other concepts is also updated based on the lesson outcome as indirect evidence through Bayesian Networks update [10]. Overnight, student forgetting is modelled, and the profile is updated. Additionally, the user model is enhanced with track information about previous work and behaviour. The AI Engine is responsible for updating the 3 models (user, content and curriculum), and for performing the adaptivity across various layers such as: content difficulty adjustment, learning loop, motivation detection, learning path recommendations.

In a previous study [9], the feasibility of integrating adaptive learning powered by the Adaptemy system in the classroom was analysed with 62 schools and 2691 students. The results showed that 97% of teachers believe that students enjoy using the Adaptemy system and want to use it at least once per week. A further study with over 10,000 students using the system for more than 6 months in over 1,700 K12 math classroom sessions was carried out to analyse Adaptemy system's learning effectiveness. The students' math ability improved by 8.3% on average per concept for an average of 5 minutes and there was a statistical significant improvement across various ability ranges [11]. Moreover, a 25% problem solving speed increase was observed for the first revision, and 38% increase for the second revision [11].

2.2 Learning Path Recommendations

In the Adaptemy system, each learner receives recommendations by default. To promote student's autonomy, the system also allows learners to reject the given recommendations and to select themselves the concept to work with. The system's recommendation has 3 parts: a specific concept to study, a specific action (i.e., learn, attempt, revise, practice), and tailored encouragements. The aim is to both present the students with specific goals and to prepare them for the given sessions.

The centre of the recommendation is the specific concept to work with. The recommendation is given based on the updated ability profile of the student at the time before a lesson is started. The engine considers which concepts have been worked on, the student's ability profile in each of the worked and unworked concepts, as well as the positions of the concepts in the knowledge graph of the course map. The engine aims at maximising the learning gain in the student's next lesson and getting the student better prepared for more advanced concepts, while not making the work too demotivational.

There are 3 types of recommendation strategies available for students and directed by teachers. The first case is when the teacher does not provide input to the recommender system and the recommendation is done by the system judgment (see Fig 2 A). The second case is when the teacher provides loose input to the recommender system through a form of assignment by providing as input the topic where students should work and the number of concepts they should work. The third case is when the teacher overwrites the system recommendations and provides the learners in the class with specific concepts though assignments.

In the first two cases, the Adaptemy system makes use of a hybrid knowledge-based recommender algorithm. The algorithm makes use of information from the learner model and curriculum model. Information from the leaner model includes learner ability and previous learning experiences with each concept, as well as learner motivation index. The algorithm makes use of the curriculum model to identify concepts that are misconceptions for students and concepts that would have a high prerequisite activation in their memory.

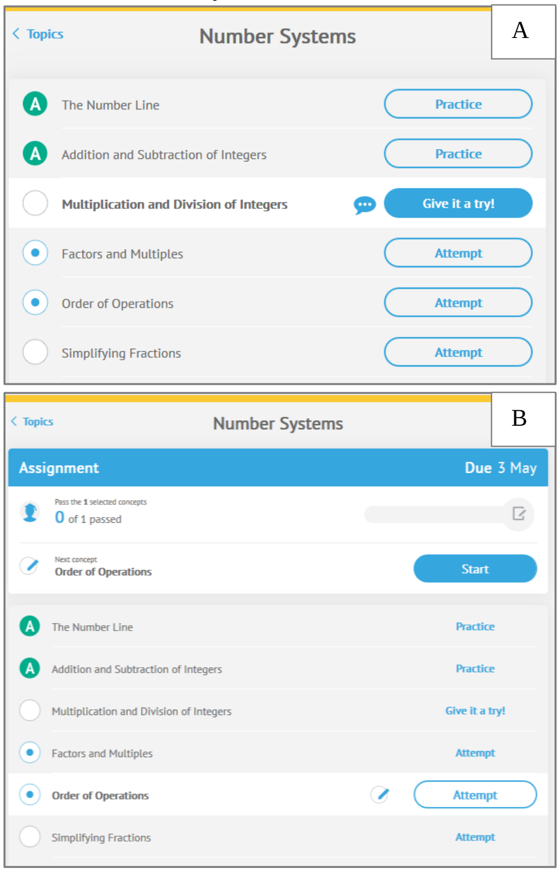

Figure 2: A) Screenshot of an example recommendation when there is no input from the teacher (case 1), and B) when the teacher assigned a specific concept and overwrote the system recommendation (case 3).

The algorithm's strategy is to reduce misconceptions, increase coverage and to increase engagement by keeping students in flow. In the first case, the system takes all concepts into consideration, while in the second case, the system will filter and use only the concepts from the recommended topic. When the teacher overwrites the system recommendations, the student will receive the concept recommended by the teacher and the system recommendation will be logged for offline analyses (see Fig. 2 B).

3 METHODOLOGY

This section details the methodology of the data processing and analysis study conducted to evaluate the recommendation component of the Adaptemy system. The data used in the study corresponded to 4257 students and 80266 learning lessons. The lessons were finished between September 2017 and March 2018, covering 211 unique concepts in the Maths course.

Two main metrics were used to evaluate learning effectiveness of the recommendations: percentage of students successfully finishing the lessons and average improvement per concept studied in a lesson. The lesson is labelled as success if the estimated ability after lesson on the worked concept is higher or equal than 60 on a 1 to 100 scale. The improvement per concept in a lesson is defined as the difference between estimated ability at the end of the lesson and estimated ability at the beginning of the lesson.

The lessons done by students correspond to one of the 3 cases depending on whether each lesson is involved in an assignment given by the teacher, and whether the assignment is made with specified concepts:

1) Lessons with no assigned concept or topic by the teacher
2) Lessons in an assigned topic by the teacher
3) Lessons on specific concepts assigned by the teacher

For all three cases, each time when a student is doing a lesson, the Adaptemy system records a log of the lesson details as well as what was the recommended concept by the AI engine right before the lesson. However, the system-recommended concepts may be of different indications in the three cases. In the first and second cases, the students have their autonomy in choosing to follow the recommendations given by the system or not. This enables to compare the effectiveness of the recommendations when students are following or not the system's recommendations.

In the third case, the students are in fact following the recommendations given by the teacher. The system-recommended concept in such a lesson means only what would have been the system recommendation at that time given the student's ability profile updated by the adaptive engine right before the lesson. This enabled to study how students are doing when the teacher recommendations matched the system's recommendation in comparison with when the teacher's recommendations did not match the system's recommendation.

Based on the course map used in the Adaptemy system and the prerequisite links, each concept is given a rank position depending on its position in the map representing its level from the very basic concepts to the most advanced concepts. The difference in this rank position between two concepts can be used to represent if one concept is more basic or more advanced (easier or harder) than the other. In this study, the rank position difference between the worked concept and system-recommended concept is used to represent whether the student is doing an easier, harder or equal-level concept than the recommended concept.

4 RESULTS AND DISCUSSION

This section presents the results analysis structured into several cases based on the type of the learning activity (i.e., study vs. assignment), and recommendation (i.e., adaptive system vs. teacher recommendations).

4.1 Case 1: Study Activity, no Teacher Recommendation

This first case corresponds to student activity with no teacher recommendations given through assignments. In this case, the system will recommend to students the concept to work with through a lesson. Students have the choice to follow the system recommendations or to choose another concept to work through the lesson. The lessons were divided in 2 categories: lessons where the students followed the system recommendation and lessons where the students did not follow the system's recommendations and they chose themselves other concepts. This case corresponds to 28416 lessons (involving 3597 students), out of which 16964 lessons (59.70%) followed the system recommendations.

As shown in Fig. 3, the success rate when following the recommendation was 78.77% and the success rate when not following was 57.41%. This difference is statistically significant based on Chi-squared = 1488.9, p < 10e-16. The overall success rate was 70.16% in student activity with no teacher recommendations given through assignment.

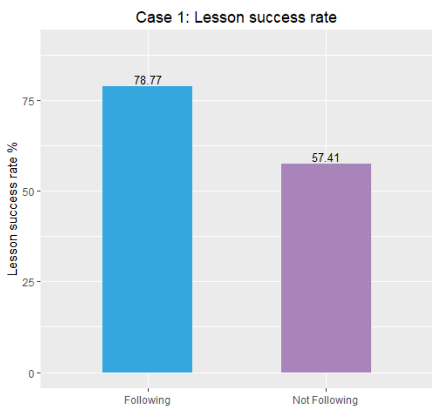

Figure 3: Percentage of students succeeding the lessons based on them following or not the system recommendations.

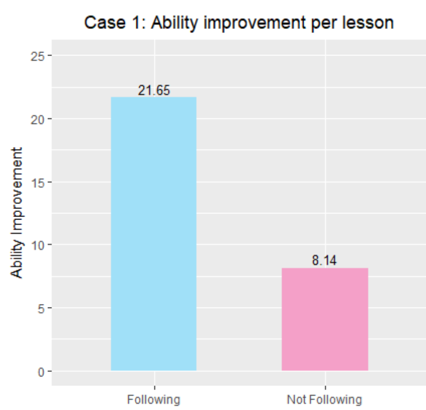

Figure 4: Improvement in student ability based on them following or not the system recommendations.

The average improvement per lesson when following the recommendation was 21.65 ability points and the average improvement per lesson when not following was 8.14 ability points (see Fig. 4). The difference was statistically significant, with t = 48.537, p < 10e-16. The overall average improvement for student activity with no teacher recommendations given to students through assignment in the system is 16.21.

4.2 Case 2: Assignment Activity, Teacher Recommends Overall Topic

The second case corresponds to student activity when the teacher gives students an assignment and recommends the overall topic of activity and the numbers of concepts to be worked. In this case, the system will recommend to students one at a time a concept to work with through a lesson until they successfully finished the number of concepts recommended by the teacher. Students have the choice to follow the system recommendations or to choose another concept to work within the topic specified for the assignment. The lessons were divided in 2 categories: lessons where the students followed the system recommendation, and lessons where the students did not follow the recommendation and they chose themselves other concepts. All concepts (recommended by the system or not) will contribute to the assignment. This case corresponds to 7756 lessons (involving 1138 students), out of which 4956 lessons (63.90%) followed the system recommendations.

The success rate when following the recommendation was 71.19% and the success rate when not following was 37.50% (see Fig. 5). Overall success rate in assignments with recommended topic was 59.03%. The difference is statistically significant with Chi-squared = 838.08, p < 10e-16.

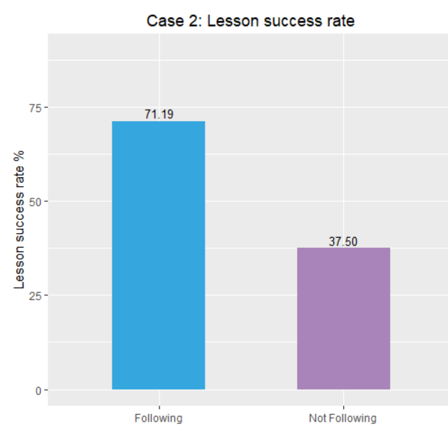

Figure 5: Percentage of students succeeding in lessons based on them following or not the system recommendations when the teacher recommended the overall topic.

The average improvement per lesson when following the recommendation was 17.96 ability points and the average improvement per lesson when not following was 2.08 ability points (see Fig. 6). The difference was statistically significant with t = 28.422, p < 10e-16. The overall average improvement in the second case was 12.23.

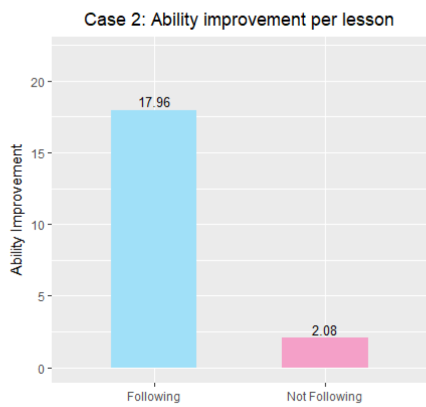

Figure 6: Improvement in student ability based on them following or not the system recommendations when the teacher recommended the overall topic.

4.3 Case 3: Assignment Activity, Teacher Recommends Specific Concepts

The third case corresponds to student activity when the teacher gives students an assignment with specified concept(s) to work on. In this case the worked concept is the one that is recommended by the teacher. The lessons are divided into two categories: first is when the concept worked is the same as the concept recommended by the adaptive learning system, and the other is when the concept worked is different from the system-recommended concept. In total 43417 lessons worked by 2649 students are included in this case, where for 16546 of the lessons (38.11%) the concept specified by the teacher matches the system recommended concept.

Figure 7: Percentage of students succeeding in lessons based on the lessons where the concept specified by the teacher matches or not the concept recommended by the system.

As shown in Fig. 7, the lesson success rate is 75.95% when the teacher assigned concept matches the system recommended concept, while the success rate is 48.11% in non-matching lessons. This difference is statistically significant with Chi-squred = 3273.4 and p < 10e-16. The overall success rate is 58.72% in this case.

As shown in Fig. 8, when the concept specified by the teacher matches the recommended concept by the system, the average improvement is 21.08 ability points, and when they do not match the average improvement is 6.13 points. The difference is statistically significant with t = 63.701 and p < 10e-16. The average improvement per lesson for this case is 11.83.

Figure 8: Improvement in student ability based on the lessons where the concept specified by the teacher matches or not the concept recommended by the system.

From the results it is shown that in every case of the lessons, the learning effectiveness is higher if the worked concept that is also the concept specified by the teacher in the 3rd case, matches the concept recommended by the system.

4.4 A closer look into all the three cases

To further investigate the effect of concept difficulty levels on the success rate of lessons, a further analysis is done by dividing the lessons not matching the system recommendations into three categories: 1) lessons done on easier concepts, 2) lessons done on same level concepts, and 3) lessons done on more difficult concepts. Fig. 9 and Fig. 10 provide a closer look into the three cases when the system-given recommendation is not followed. The dashed blue lines in the two figures represent the value when the recommendation is followed for easy visual comparison.

The results from Fig. 9 show that in all three cases the highest success rate is when the worked concept is easier or at a more basic level than recommended, and the lowest success rate is when the worked concept is harder or at a more advanced level than recommended. However, even if the worked concept is at the same level as the recommended concept, the success rate is still lower than that when the system recommendation is followed. Therefore, the higher success rate when the recommendation is followed is not only due to the difficulty levels of the concepts chosen by the recommendation.

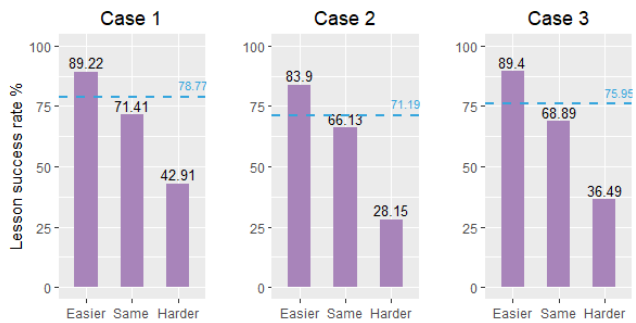

Figure 9: Lesson success rates compared across different cases when the worked concept is easier or harder than, or at the same level as the system-recommended concept.

The results from Fig. 10 show that in all three cases, no matter if the concept is of an easier, same or higher level than the recommended concept, the improvement per lesson (represented by the three bars) is still lower than that if the lesson is worked on the system-recommended concept (represented by the dashed line in blue). This is seen as a corroborative evidence that the recommendation engine does not only take into account the difficulty levels of concepts, but also the prerequisite relationships between concepts in the knowledge map. The effectiveness of the recommendation engine regarding the best learning paths is thus supported by the results here.

In all three cases, the average improvement is the lowest when the lesson is on a harder concept than recommended. In Case 1 and Case 2, the highest average improvement is seen when the lessons are done on a same-level concept as recommended. However, in Case 3, where the lessons are done on concepts assigned by the teachers, the highest average improvement is seen when the worked concept is easier than the system recommended one.

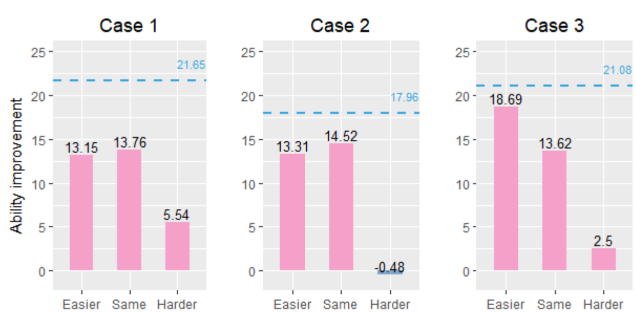

Figure 10: Average improvement per lesson compared across different cases when the worked concept is easier or harder than, or at the same level as the system-recommended concept.

Looking at Case 3 in comparison with Case 1 one can note that the improvement values are similar on the lessons where the worked concepts match the recommendation (21.65 compared to 21.08) or are on the same-level concepts as recommended (13.76 compared to 13.62). The indication is that the teachers may provide a better selection of easier concepts for the students to revise and reinforce their abilities on, than the easier concepts chosen by the students themselves. However, regarding the more advanced concepts worked in the lessons, the teachers' selection may be more over-challenging than what the students choose by themselves, and thus the improvement is even lower (2.5 compared to 5.54).

The results shown in Fig. 10 indicates that choosing the right difficulty levels of concepts to be worked on is part of the reason why working on the concepts recommended by the engine would gain higher improvement per lesson. It provides a strong evidence that the recommendation engine is giving the right difficulty levels of concepts to be worked on.

5 CONCLUSIONS

In summary, the study shows that the learning recommendations provided by the Adaptemy's AI Engine, when followed by the users, will give both a higher success rate and a higher average ability improvement than if they are not followed, which shows the effectiveness of the personalised learning path recommendation. Moreover, looking into more details regarding to concept advance levels shows that the higher effectiveness of the recommendation is not solely due to the right advance levels, corroborating that the learning path consideration takes its part of effect in the recommendation engine.

REFERENCES

[1] H. Drachsler, K. Verbert, O. C. Santos, and N. Manouselis, "Panorama of recommender systems to support learning," in *Recommender systems handbook*, Springer, 2015, pp. 421–451.

[2] M. Erdt, A. Fernández, and C. Rensing, "Evaluating Recommender Systems for Technology Enhanced Learning: A Quantitative Survey," *IEEE Trans. Learn. Technol.*, vol. 8, no. 4, pp. 326–344, Oct. 2015.

[3] I. Hulpuş, C. Hayes, and M. O. Fradinho, "A Framework for Personalised Learning-Plan Recommendations in Game-Based Learning," in *Recommender Systems for Technology Enhanced Learning*, N. Manouselis, H. Drachsler, K. Verbert, and O. C. Santos, Eds. New York, NY: Springer New York, 2014, pp. 99–122.

[4] O. C. Santos, M. Saneiro, S. Salmeron-Majadas, and J. G. Boticario, "A Methodological Approach to Eliciting Affective Educational Recommendations," 2014, pp. 529–533.

[5] J. K. Tarus, Z. Niu, and A. Yousif, "A hybrid knowledge-based recommender system for e-learning based on ontology and sequential pattern mining," *Future Gener. Comput. Syst.*, vol. 72, pp. 37–48, Jul. 2017.

[6] A. Klašnja-Milićević, B. Vesin, M. Ivanović, Z. Budimac, and L. C. Jain, "Recommender Systems in E-Learning Environments," in *E-Learning Systems*, Springer, Cham, 2017, pp. 51–75.

[7] T. K. Paradarami, N. D. Bastian, and J. L. Wightman, "A hybrid recommender system using artificial neural networks," *Expert Syst. Appl.*, vol. 83, pp. 300–313, Oct. 2017.

[8] A. Paramythis, S. Weibelzahl, and J. Masthoff, "Layered evaluation of interactive adaptive systems: framework and formative methods," *User Model. User-Adapt. Interact.*, vol. 20, no. 5, pp. 383–453, Dec. 2010.

[9] I. Ghergulescu, C. Flynn, and C. O'Sullivan, "Adaptemy – Building the Next Generation Classroom," presented at the EdMedia: World Conference on Educational Media and Technology, 2015, vol. 2015, pp. 79–88.

[10] T. Lynch and I. Ghergulescu, "Large Scale Evaluation of Learning Flow," in *2017 IEEE 17th International Conference on Advanced Learning Technologies (ICALT)*, 2017, pp. 62–64.

[11] I. Ghergulescu, C. Flynn, and C. O'Sullivan, "Learning Effectiveness of Adaptive Learning in Real World Context," presented at the EdMedia: World Conference on Educational Media and Technology, 2016, vol. 2016, pp. 1385–1390.