Personalized presentation of multimedia objects for home healthcare environments: a peer-based intelligent tutoring approach

John Champaign and Robin Cohen

University of Waterloo {jchampai, rcohen}@uwaterloo.ca

Abstract. In this paper we present an approach for reasoning about which media content from an existing repository should be presented to users. We elaborate on our technique by considering students within an e-health intelligent tutoring environment. Our approach models the benefits in socially connected learning gained by peers in order to then recommend those objects predicted to offer the best gains in knowledge for the student. This is achieved in a framework where the past learning gains of peers are modeled and recorded with the objects in the repository. We previously confirmed the value of the approach by simulating student learning. From here, we then conduct a user study comparing the learning achieved by students presented with objects selected by our algorithms, compared to a less principled approach for curriculum sequencing; this is performed for the application of home healthcare (assisting caregivers of autistic children). We provide compelling evidence for the value of our proposed vision for achieving effective peer-based tutoring: through past experiences of peers in an extensive repository.

Keywords: E-health multimedia, social recommender systems

1 Introduction

With an aging population, home healthcare solutions are becoming, by necessity, more prevalent. Caregivers and patients alike face the challenge of making medical decisions in dynamically changing environments, using whatever resources are available in the home. With copious amounts of information (e.g. text, videos, interactive systems) users benefit from methods for effectively focusing on what would be most beneficial to view.

Our research aims to provide important decision-making support in these scenarios by leveraging the learning of peers through a social networking approach. In particular, we propose that peer-based tutoring form the basis of the information imparted to homecare caregivers and patients. Distinct from other approaches to peer-based intelligent tutoring which assume an active social network of information exchange in real-time (e.g. [1]), we propose a framework that makes use of learning experienced by peers at several points in the past and allow these peers to streamline content that will be shown to future students. In essence, we seek to adopt an approach to learning that respects what McCalla has referred to as the ecological approach [2]: enabling various learning objects (texts, videos, book chapters) and adjusted version of these objects to be introduced to peers, based on the past experiences of other, similar, students with this media content (or learning objects). The core of our proposal is to enable the peers to influence the determination of learning objects which will be considered. While an initial corpus will be introduced, once a peer has experienced learning, it will be possible to suggest, for example, subdividing an existing, lengthy learning object into a smaller, cogent element, which is strongly recommended to other students.

An example scenario helps to motivate our research. Consider a diabetic patient, attempting to manage his disease. Distinct from an approach of simply posting a query to a discussion group and receiving various responses from peers (with varying degrees of reliability), one would treat this problem as one of properly teaching the patient suitable information that may be contained in a variety of online articles or instructional videos. We assume a corpus of these learning objects exists and has been experienced by other peers in the past. Pre- and post-testing of the learning achieved by these peers is conducted (for example, through an exit quiz that results in a level of understanding represented as a grade achieved, before and after the interacting with the learning object). Then, each learning object has stored with it the students who have experienced it, along with the benefit that each students obtained (an increase, or decrease, in grade level achieved).

In determining which learning object to display to a new student, we propose two distinct methods. The first focuses on presenting to new students those learning objects which produced the most benefit to like-minded peers, where the similarity between students is determined on the basis of their overall level of knowledge. This approach is motivated by collaborative filtering techniques, as performed in recommender systems [3]. For example, those learning objects which resulted in a weak understanding for other similar patients would be avoided for the new student. This system allows the object that is best suited to a particular student population to be shown to them.

Continuing with the motivating scenario of informing homecare diabetic patients, our second focus concerns the situation where there may be a particular article in a book (or some other subset of a larger learning object) on managing diabetes which is of special value. As with our algorithm for recommending learning objects, the determination of which of these smaller articles to present to a peer will be based on the learning that is experienced by others. The object would be added to the corpus and then its overall benefit to peers can be tracked. It is possible that for one population of (perhaps more advanced) students a more targeted, succinct learning object would be preferable, while for another population of students a learning object with additional explanations may be preferable. In addition, one can manage the entire corpus by eventually discarding learning objects that are not of use (garbage collection), resulting in a refined and more valuable corpus on which the learning may proceed.

In all, we believe that home healthcare can be improvedby enabling patients and caregivers to learn on the basis of the past learning of their peers, through judicious choice of material to present to the learners, which evolves over time as the learning experiences of the peer group expand.

2 Background

Intelligent tutoring systems research has increasingly incorporated student models using the machinery of user modeling[4] and has most recently progressed to encompass efforts aimed at allowing students to benefit from the learning that their peers are undergoing [1]. McCalla has proposed an ecological approach [2] for e-learning: making use of a repository of learning objects (which could be web pages, research papers, videos, simulations, etc.) with a history of interactions from previous students, in order to direct the learning of new students. This view of peer-based tutoring is distinct from the standard view where peers are assisting each other in real-time. McCalla's approach is simply a general philosophy for design; with this hand, we proceed to create specific algorithms embodying this approach, used as the basis for recommending objects to students. Using techniques inspired by collaborative filtering [3], the basis of our approach is to identify which users in a system are similar to each other, to then use past interactions with these similar students to intelligently tailor the system's interactions with the current student.

2.1 Our Approach

Algorithm 1 Pseudocode For Collaborative Learning Algorithm (CLA)

```
Input the current-student-assessment
for each learning object: do
  Initialize currentBenefit to zero
  Initialize sumOfBenefits to zero
  Input all previous interactions between students and this learning object
  for each previous interaction on learning object: do
    similarity = calculateSimilarity(current-student-assessment, interaction-initial-
    assessment)
    benefit = calculateBenefit(interaction-initial-assessment, interaction-final-
    assessment)
    sumOfBenefits = sumOfBenefits + similarity * benefit
  end for
  currentBenefit = sumOfBenefits / numberOfPreviousInteraction
  if bestObject.benefit < currentBenefit then bestObject = currentObject
end for
if bestObject.benefit < 0 then bestObject = randomObject
```

Our proposed algorithm for determining which learning objects to present to students is presented in Algorithm 1. We assume that we are tracking a set of values, *v[j,l]*, representing the benefit of the interaction for user j with learning object l. *v[j,l]* is determined by assessing the student before and after the interaction, and the difference in knowledge is the benefit. We also record for each learning object what we refer to as the *interaction history*: the previous interactions of students with that object, in terms of their initial and final assessments.[1] We assume that a student's knowledge is assessed by mapping it to 18 discrete levels: A+, A, A-, ... F+, F, F-, each representing $\frac{1}{18}$th of the range of knowledge. This large-grained assessment was used to represent the uncertainty inherent in assessing student knowledge, and only this large-grained assessment is used to reason about the students' ability in our approach.

The anticipated benefit of a specific learning object l, for the active user, a, under consideration would be [2]

$$p[a,l] = \kappa \sum_{j=1}^{n} w(a,j)v(j,l) \tag{1}$$

where *w(a,j)* reflects the similarity $\in$ (0,1] between each user j and the active user, a, and κ is a normalizing factor. $\frac{1}{|n|}$ was used as the value for κ in this work where n is the number of previous users who have interacted with learning object l. *w(a,j)* was set as $\frac{1}{1+difference}$ where difference is calculated by comparing the initial assessment of j and the current-student-assessment, and assigning an absolute value on the difference of the letter grades assigned. This is in order to obtain a similarity between 0 and 1, with 1 representing identical assessments. So the difference of A+ and B- would be 5 and the difference of D+ and C- would be 1. *v(j,l)* is also computed using a difference. Instead of a sum of the absolute differences between the initial assessments of two users, it is the sum of the difference between initial and final assessments for user j's interactions with learning object l. For example, *v(j,l)* where j is initially assessed as A+ and finally assessed at B- would be -5, while where j is initially assessed at B- and finally assessed at A+ would be 5. This is shown as the calculateBenefit function in Algorithm 1. In the absence of other criteria, a user a will be assigned the learning object l that maximizes *p[a,l]*. If the maximum *p[a,l]* is a negative anticipated benefit, a random learning object will be assigned to the user.

The CLA's value in achieving increases in knowledge to students has been confirmed by a method of simulated student learning [6, 7] achieving performance approaching that of algorithms with perfect knowledge about the students, the learning objects and the learning gains of their interactions. Below we show just one of our graphs of results where the learning of 50 students was simulated over 100 trials with 20 iterations where the mean of the average student knowledge is

[1] The algorithm would be run after an initial phase where students are learning through the use of a set of learning objects. These students' experiences would then form the basis for instructing the subsequent students.

[2] Adapted from [3, 5].

mapped. Simulated students interacted with learning objects that had varying impact depending on the student's current assessment grade where a total of 100 learning objects were included in the repository. The Raw Ecological curve embodies the CLA Algorithm. The Pilot variation allows 10% of the students to prime the system first and the Simulated Annealing variation included a cooling phase where students first had a chance of being randomly assigned a learning object; both variants are done to address cold-start problems. All three variations show very effective student learning (Figure 1). We then move on to human evaluation in order to confirm the value of our methods; necessarily we are investigating a smaller sample size (i.e. we cannot easily subject participants to thousands of learning experiences nor easily manage hundreds of participants in one study). But the learning that is accomplished is now matching the ground truth for those students (revealed through performance on assessment quizzes).

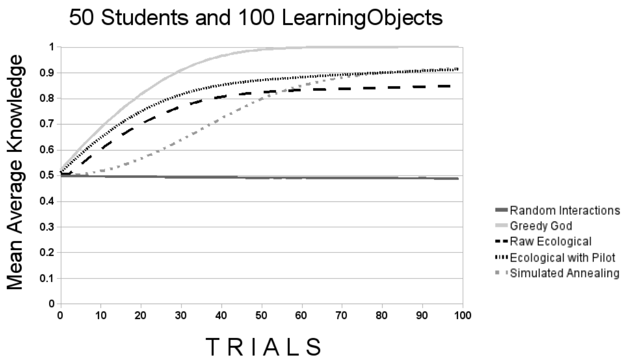

Fig. 1: Comparison of 5 Approaches for Selecting Learning Objects [6]

3 Human Evaluation

To study the effectiveness of our approach with humans we conducted a preliminary evaluation with participants at the University of Waterloo. We chose as an application domain enabling users to learn about how to care for a child with autism (which may arise as a home healthcare scenario, of interest to projects such hSITE [8], with which we are involved). Our first step was to assemble our repository of learning objects: the material that students would learn from. In collaboration with a clinical psychologist specializing in children and autism, we created 20 learning objects (16 text articles and 4 videos) that each took about 5 minutes to experience. Also in collaboration with the psychologist we created a 10-question multiple choice assessment, covering material from the learning objects. This was used to carry out the pre- and post-test assessments which serve to model learning gains in students (and form a component of our algorithm for determining which objects to present to each student).

We hypothesized that a group of students using our peer-based technique for selecting learning objects would show greater learning gains than a control group that had learning objects randomly assigned to them. The aim of our study, therefore, was to validate our proposed Collaborative Learning Algorithm for curriculum sequencing (Algorithm 1) – the centrepiece of our overall peer-based learning framework.

In order to obtain feedback from our participants about corpus division, we also explained our corpus approach to participants and then offered them the opportunity to subdivide each learning object that they were shown, as the learning proceeded.[3] We finally obtained more information during an exit survey where participants responded to questions asking them how they felt about this option of streamlining learning objects. The entire process lasted approximately 1 hour. 23 participants took part in our experiment, including graduate students, undergraduate students and staff members at the university. All were at least 18 years old, fluent in English and not an expert in autism spectrum disorders.

3.1 Procedure

Each participant experienced 5 learning objects and was assessed before and after each for a total of 6 assessments, with the first assessment before they experience any learning objects being a measure of the student's initial knowledge before seeing any learning objects. The assessments were the same 10 multiple choice questions each time.[4] The quiz was designed so that each question was covered well by different learning objects in the repository (and more than one learning object served to help a student to respond to that question). After experiencing each learning object, each participant did the assessment quiz and also answered a separate questionnaire allowing the student to propose a streamlining (division) of that learning object. At the end of the experiment each participant was given an exit survey asking them their overall feelings about streamlining and soliciting general feedback.

The first 12 participants were randomly assigned learning objects. They were used both as a control group and to provide training data for our technique. The next 11 participants experienced a curriculum sequence provided by our approach. Participants read hardcopy articles or watched videos on a provided netbook and then a "Wizard of Oz" style study was performed. For our technique, a program was written using the CLA (Algorithm 1) and the answers provided by participants in their pre-test assessments served as the current student assessment; a new recommendation for a learning object was then determined. This sequence continued until the student had experienced five different learning objects. In essence, the first 12 participants served to prime the system for

[3] These subdivisions were not used by other participants. Getting enough data, with the limited number of participants, to differentiate between original learning objects and streamlined versions would have been problematic.

[4] This was done in part to ensure that we were modeling comparable learning experiences from the participants.

the remaining participants. After this phase, each learning object in the repository had 3 experiences recorded: while the initial control group of students were shown a random set of objects, which objects would be presented to each was determined offline in a way that ensured that each object would be shown to 3 different participants. The net-benefit obtained by each subject in the control group (number of questions correct between pre and post-test) became part of that object's interaction history. For the participants in our experimental group, determining the similarity between the current student and previous peers was measured by comparing the number of questions on the assessment that were answered identically. Only the data collected from the training group was used to make recommendations to the experimental group.[5] No learning objects were shown twice to the same participant.

3.2 Results

Curriculum Sequencing We first compared the learning gains of our 11 experimental group participants, namely the post-test (their final assessment) minus the pre-test (their first assessment).

	Mean	s.d.	Mean (without P20)	s.d. (without P20)
Control	1.83	1.27		
Experimental	3.09	2.21	3.4	2.07

Table 1: Comparison of overall learning gains of users in each case

These results can be interpreted that, on average, participants in the control group got 1.83 more questions correct (out of 10) after completing the 5 learning objects and participants in the experimental group got an average of 3.09 more questions correct.

P20 was a participant who did not seem to be taking the experiment seriously, did not read learning objects fully and rushed through the experiment (finishing in about 40 minutes when most participants took about 1 hour). The data was analyzed with and without this participant's data included.

The results were statistically reliable at p=0.059 (one-sided, two samples, unequal variance t-test) which was not statistically significant. With participant 20 removed, the results were statistically reliable at p=0.027 (one-sided, two samples, unequal variance t-test) which was statistically significant.

Next, we compared the proportional learning gains of participants. This was to take into consideration the suggestion of Jackson and Graesser [9] that simple learning gains are "biased towards students with low pretest scores because

[5] Had we followed our proposed approach and continually added data for the program to make recommendations from, the final participants would have been given learning objects based on a richer repository of data and the experimental group would not have been provided with a consistent treatment.

they have more room for improvement". This is measured using [(post-test − pretest)/(10-pretest)][6].

	Mean	s.d.	Mean (without P20)	s.d. (without P20)
Control	0.530	0.452		
Experimental	0.979	1.07	1.08	1.02

Table 2: Comparison of proportional overall learning gains of users

The results were statistically reliable at p=0.10 (one-sided, two samples, unequal variance t-test) which was not statistically significant. With participant 20 removed, the results were statistically reliable at p=0.071 (one-sided, two samples, unequal variance t-test) which also was not statistically significant.

Next, we considered the per-LO learning gains of each student. Here, the change in assessment after assignment of a single learning object, were measured for each learning object experienced and the average computed. This average was then compared for the control and experimental groups.

	Mean	s.d.	Mean (without P20)	s.d. (without P20)
Control	0.367	0.253		
Experimental	0.618	0.442	0.68	0.413

Table 3: Comparison of average learning gains of users in each case

The results were statistically reliable at p=0.059 (one-sided, two samples, unequal variance t-test) which was not statistically significant. With participant 20 removed, the results were statistically reliable at p=0.027 (one-sided, two samples, unequal variance t-test) which was statistically significant.

Taken together, our results indicate that students presented with learning objects determined by our algorithm achieved greater learning gains than those who were randomly assigned objects.

Corpus Approach Each participant was invited, after the concept had been explained to them, to streamline learning objects according to the corpus approach.

In spite of being told it was up to them whether or not to streamline learning objects, only 5 out of 23 participants declined to streamline any objects. On average, participants suggested streamlined versions for 2 of the 5 learning objects they saw.[7]

Each participant was asked 3 questions about the corpus approach during their exit survey:

[6] 10 is the maximum possible score on a 10 question multiple-choice quiz.

[7] In practice, participation may be lower if there isn't a researcher sitting across the table when students are deciding whether or not to streamline; however there was clearly a willingness to engage in this activity.

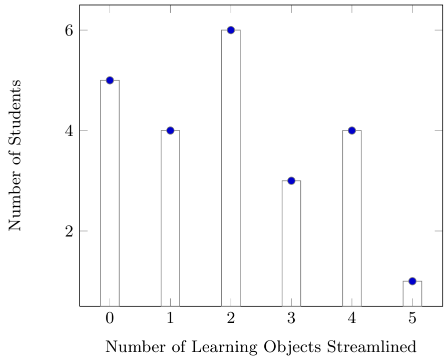

Question	Mean	s.d.
Q1	0.227	3.35
Q2	0.864	2.949
Q3	3.773	1.232

Table 4: Mean answer values to exit survey questions

1. How would you rate the difficulty of creating a new streamlined learning object?
2. How would you rate the difficulty of deciding what content to include in a streamlined version?
3. How would you rate the usefulness of a system offering a user the full version or streamlined version of content like you've seen?

Participants were given a 11 point scale, ranging from -5 to 5 with the labels "difficult" at -5, "neutral" at 0, and "easy" at 5 for Q1 and Q2 and "useless" at -5, "neutral" at 0 and "useful" at 5 for Q3.

For the 23 participants the feedback is provided in Table 4.

Although participants were mostly neutral with respect to creating streamlined versions of learning objects (Q1 and Q2), they were clearly positive about using a system where other students create streamlined learning objects for them. This conforms to research on participatory culture (e.g. [10]) which has shown that consumers usually greatly outnumber contributors. It has been shown to be possible (e.g. [11]) to use incentives to encourage greater participation.

4 Conclusion and Discussion

With an overall aim of enabling effective patient-led health management, we offer here a specific approach for peer-based tutoring that makes use of a rich interaction history to personalize delivery of content for users; this serves to assist caregivers in focusing their attention on the most valuable material and demonstrates the true potential of social recommendation for this critical application area. The human study described in this paper confirms the effectiveness of the approach in achieving knowledge gains; the exit survey also support our proposal for allowing peers to augment the repository through corpus division.

Personalization for E-Health Other work in this area, in the area of E-Health that has demonstrated the importance of personalized content delivery and of leveraging social networks as part of that learning (e.g. [12, 13]) focus on promoting healthier lifestyles by encouraging reflection and discussions within the family through the use of a collaborative platform. Our approach is aimed instead at allowing individuals to better understand their health concerns and make informed decisions. [14] proposes personalized delivery of video to users to educate about self-care of fibromyalgia. This work confirms several elements in our approach: including video objects, supporting personalized selection of objects from a corpus. Like us, their user study compared the value of their approach with one that was less personalized. One notable difference is that our tailoring is based on modeling peer-experiences.

Personalized E-learning More traditionally, intelligent tutoring systems researchers have been focused on providing opportunities for students to direct their own learning, making choices about lessons to explore [4, 15]. One significant way in which our e-learning framework differs is in its enhanced focus on determining the content to be presented to a student. In our approach, this content is determined on the basis of the experiences of other students and on the initiative of those students to introduce new subdivided learning objects into the corpus that is used as the lesson repository for the tutoring.

Our approaches to tutoring can be viewed as peer-based, to the extent that student learning is enabled by previous peer interactions. Other researchers have explored a peer-based approach to intelligent tutoring. For example, the COMTELLA project [11, 16, 1] at the University of Saskatchewan investigated recommendation of academic papers by users in a small-scale, on-line community. The system encouraged users to participate but was concerned as well with preventing information overload (too many items being posted at once). This was achieved by introducing limits on the number of posts that a student could contribute (higher or lower based on the perceived quality of the student's previous posts). In contrast, in our approach, having a wealth of possible learning objects does not necessarily detract from student learning, as our system is determining the content to be presented to each student (and our simulations verified that valuable learning objects could still be selected for students, when the corpus grew to include many objects).

Also, in contrast to efforts such as [11], in our approach each student's learning is directed by considering all experiences of previous students, thus allowing for a continuous redirection of possible content. Personalization is maintained throughout, as well. This is achieved by modeling the knowledge levels of each student and an assessment of their current overall understanding in order to perform matching to like-minded peers, for the selection of learning objects.

Previous work on collaborative learning, such as [17], has attempted to use interactions between students and the system to provide a better experience for subsequent students. The authors created a program that would capture user problem solving behaviours in the system. This data was then used to begin the development of a tutor, in what they call "bootstrapping novice data (BND)". The authors admit, however, that the task is non-trivial and reach the conclusion that that analysis must happen at multiple levels of abstraction. In contrast, our approach does not try to model specific user actions. Instead it pragmatically considers the sequence that learning material is experienced and how successful the students were.

Collaborative Recommendation Collaborative filtering recommender systems [3, 5] also make use of content selection via modeling similarity of peers.

On the surface, it might seem that recommendation techniques could be applied directly in an intelligent tutoring setting. However, whereas most recommender systems endeavour to obtain an increasingly specific understanding of a user, an intelligent tutoring system seeks both to understand a user and to enable change or growth. In addition, in contrast to positioning a user within a cluster of similar users, we would like to model a continually evolving community of peers who are operating at a similar level of knowledge.

Some of the cutting-edge areas of recommendation research are more relevant to us. The work of Herlocker et al. ([5]), which explores what not to recommend (i.e. instead of seeking highly relevant items from a set, removing irrelevant items) is perhaps relevant in our context, where peer-created learning objects which have be found to lack benefit for student learning may be worth removing from the repository.

Our work is distinct from the above approach, however, in a number of ways. Obtaining a history, and accurately categorizing a user's life events, will be a time consuming process that may be difficult to convince users to undertake. In their system user histories must be continually updated, with the ongoing issue of out-of-date user profiles. In contrast, the data used by our system should be easily gathered and will be as up-to-date as the last usage of a learning object. Finally, our modeling of students is on the basis of their knowledge levels as reflected in pre- and post-test assessments and as such reflects a more concrete representation than a cumulative code.

The Evolution of Social Relationships One valuable aspect of our approach in its management of the social network of peers is its ability to cope with a potentially large number of fellow students. This is achieved in part by first grounding the student learning in the context of a particular learning object that is most appropriate, based on the benefits in learning derived from this object by students at a similar level of knowledge.

Scaling is problematic for many approaches to real-time peer-tutoring (e.g. [1]). Our approach, like many ecological approaches, uses data from past interactions and performance improves as the size of the user base and repository of learning objects increases. A very large social network, therefore, is not a challenge at all, but instead an opportunity to provide highly personalized recommendations to students.

While computational demands do increase with a larger group of students, the time needed for such computations is small compared to the time it takes for students to complete tasks. The approach detailed in this paper could easily scale to making recommendations for large numbers of students every 1/2 hour if needed. If computation were to become a limiting factor, a straightforward adjustment would be to compute the predicted benefits for a student in between her interactions with the system instead of making recommendations on-the-fly.

References

1. Vassileva, J.: Toward social learning environments. IEEE Transactions on Learning Technologies **1**(4) (2008) 199–214
2. McCalla, G.: The ecological approach to the design of e-learning environments: Purpose-based capture and use of information about learners. Journal of Interactive Media in Education: Special Issue on the Educational Semantic Web **7** (2004) 1–23
3. Breese, J.S., Heckerman, D., Kadie, C.M.: Empirical analysis of predictive algorithms for collaborative filtering. In Cooper, G.F., Moral, S., eds.: Proceedings of the 14^{th} Conference on Uncertainty in Artificial Intelligence. (1998) 43–52
4. Kay, J.: Lifelong learner modeling for lifelong personalized pervasive learning. TLT **1**(4) (2008) 215–228
5. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. ACM Trans. Inf. Syst. **22**(1) (2004) 5–53
6. Champaign, J., Cohen, R.: A model for content sequencing in intelligent tutoring systems based on the ecological approach and its validation through simulated students. In: Proceedings of the Twenty-First International Florida Artificial Intelligence Research Society Conference, Daytona Beach, Florida (2010) 486–491
7. Champaign, J., Cohen, R.: Exploring the effects of errors in assessment and time requirements of learning objects in a peer-based intelligent tutoring system. In: 2011 FLAIRS Conference Proceedings, Palm Beach, Florida (2011)
8. Plant, D.: hSITE: healthcare support through information technology enhancements NSERC Strategic Research Network Proposal. (2008)
9. Jackson, G.T., Graesser, A.C.: Content matters: An investigation of feedback categories within an ITS. In: Proceeding of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts That Work, Amsterdam, The Netherlands, The Netherlands, IOS Press (2007) 127–134
10. Almeida, R.B., Mozafari, B., Cho, J.: On the Evolution of Wikipedia, International Conference on Weblogs and Social Media (2007)
11. Cheng, R., Vassileva, J.: Design and evaluation of an adaptive incentive mechanism for sustained educational online communities. User Model. User-Adapt. Interact. **16**(3-4) (2006) 321–348
12. Colineau, N., Paris, C.: A portal to promote healthy living within families. In: eHealth. (2010) 259–266
13. Colineau, N., Paris, C.: Motivating reflection about health within the family: the use of goal setting and tailored feedback. User Model. User-Adapt. Interact. **21**(4-5) (2011) 341–376
14. Camerini, L., Giacobazzi, M., Boneschi, M., Schulz, P.J., Rubinelli, S.: Design and implementation of a web-based tailored gymnasium to enhance self-management of fibromyalgia. User Model. User-Adapt. Interact. **21**(4-5) (2011) 485–511
15. Brusilovsky, P., Vassileva, J.: Course sequencing for large-scale web-based education. International Journal of Continuing Engineering Education and Life-long Learning **13**(1/2) (2003) 75–94
16. Conejo, R., Ruiz, A., Barros, B., Vassileva, J.: Towards a group model for learning communities. first steps with the comtella d collaborative learning community. In: ICALT. (2008) 905–909
17. Harrer, A., McLaren, B.M., Walker, E., Bollen, L., Sewall, J.: Creating cognitive tutors for collaborative learning: steps toward realization. User Modeling and User-Adapted Interaction **16** (2006) 175–209