The Validation of an Annotations Approach to Peer Tutoring Through Simulation Incorporating the Modeling of Reputation

John CHAMPAIGN[a*], Robin COHEN[a] & Jie ZHANG[b]

[a]*David R. Cheriton School of Computer Science, University of Waterloo, Canada*
[b]*School of Computer Engineering, Nanyang Tech University, Singapore*
*jchampai@uwaterloo.ca

Abstract: In this paper, we promote a model for peer-based intelligent tutoring that leverages the past learning experiences of students with a repository of learning objects. Consistent with McCalla's ecological approach, we determine appropriate peers and appropriate learning objects to direct a new student's learning. In particular, we focus on allowing peers to provide annotations of learning objects. We revisit a procedure developed to select which annotations to present to students in order to improve their learning: one that combines a modeling of the reputation of the annotation (based on its approval or disapproval by previous students), the reputability of the annotator (based on the reputation of all annotations left by the student) and the similarity of the raters with the new student. Our focus is on developing effective validation of the procedure's benefit, using an approach of simulated student learning. This is achieved by developing algorithms in greater detail and then making particular design decisions for the simulation in order to manage the reputability of the annotators and annotations in a way that enables the best learning objects to be employed for the tutoring. We are able to demonstrate the value of our proposed approach using distinct measures of rater similarity. We conclude with a comparison to related work and a view to future directions for the research. As a result, we present an approach for interpreting data from interactions with previous students in order to influence how to interact with current and future students, to enable effective learning.

Keywords: peer-based intelligent tutoring, annotations of learning objects, trust modeling, simulating students, ecological approach to ITS

Introduction and Background

In the work of [2], a peer-based intelligent tutoring model was introduced that determined which learning objects from a given repository to present to a new student, based on which objects had provided the most advantageous benefits to learning to similar students. This Collaborative Learning Algorithm (CLA) leveraged pre- and post-tests of student learning in order to model the benefits derived from a particular learning object, and used a modeling of student knowledge in terms of letter grades, where students were similar if their grade level of achievement was sufficiently close on that letter grade scale. That work also introduced a method of validation based on simulating student learning, attaching target levels of knowledge to each learning object in the repository.

In this paper, we explore the use of a richer repository of learning objects, namely ones where previous students have elected to attach annotations. An annotation could be for example a text message including an observation, a question, references to related work, etc. Peer-based learning is then facilitated by exposing each student to some of the available annotations and allowing the student to provide a rating (thumbs up or thumbs down) to record whether the annotation provided learning benefit. The annotations and the ratings would then inform each subsequent student, as part of the learning experience.

A preliminary sketch of the algorithms that would be beneficial to employ was presented in [2], in order to determine whether previous annotations should be shown to a new student,

or not. A new element was then introduced into the framework, in an effort to select the most appropriate annotations for each student, namely a modeling of the reputation, both of the annotation itself (reflected in the totality of the ratings left behind) and of the annotator (accumulated over time, with each new annotation that student has left). This was inspired from research on the modeling of trust in multiagent systems [7] where, over time, it is possible to predict an agent's trustworthiness based on past behaviour, and where selecting trustworthy partners offers better overall decisions. In addition, the similarity of the raters with the new student is an influence on the calculation of the annotation's reputation, continuing to encompass peer-based tutoring.

In the sections that follow we describe our revised approach in detail.

1. Overview of the Model

Algorithm 1: Student Reputation
//Consider student as an annotator
calStudentReputation (Student s)
if *num of annotations by s == 0* **then**
| R(s) = 0.5; //Reputation of s
else
| R(s) = 0;
| **foreach** *annotation a of s* **do**
| |_ R(s) += calcAnnRep(a);
|_ R(s) /= num of annotations by s;
return R(s) ∈ [0,1];

Algorithm 2: Student Similarity
Similarity (Student c, Student r)
vS = 0; //num of voted same
vD = 0; //num of voted different
foreach *annotation voted by both* **do**
| **if** *current.vote == rater.vote* **then**
| | vS += 1;
| **else**
|_ |_ vD += 1;
similarity = (vS − vD) / (vS + vD);
return similarity ∈ [-1,1];

Our proposed model for reasoning about which annotations to show a new student *s* integrates: (i) the annotation's initial reputation (equal to the reputation of the annotator, as calculated in Algorithm 1 - in turn based on how much his previous annotations were liked) (ii) the current number of votes for and against the annotation, adjusted by the similarity of the rater with the student *s* (calculated using Algorithm 2) to value votes by similar students more highly. The reputation of each annotation for a student *s* is calculated by Algorithm 4 using an adjust function to scale (i) according to (ii) and those annotations with the highest reputation are shown; examples of this function will be shown in Section 2.

2. Validation of the Model

2.1 Experimental Set-up

Algorithm 3: Annotation Reputation
calAnnRep (Annotation a)
foreach *vote on annotation* **do**
| **if** *vote.for* **then**
| | vF += 1;
| **else**
|_ |_ vA += 1;
return adjust(a.initRep, vF, vA);

Algorithm 4: Specific Annotation Rep
calAnnRepSpecific (Ann a, Student s)
foreach *vote on annotation* **do**
| sim = similarity (s, voterStudent);
| **if** *vote.for* **then**
| | vF += 1 * sim;
| **else**
|_ |_ vA += 1 * sim;
return adjust (a.initRep, vF, vA);

In order to demonstrate the value of our approach for determining which annotations to show to which students, we constructed a simulation where learning objects in a repository were modeled as having a certain target level of instruction and a certain projected impact level for students, based on their level of knowledge. In a first phase, those learning objects

determined to be most appropriate for each student were selected using the CLA (so each simulated student arrived with a certain knowledge level and was offered the learning object that previous, similar peers had found most beneficial to their learning). We then reasoned about showing annotations attached to this learning object, independently deciding for each annotation whether it would be shown, based on the expected benefit it would introduce to the current student. In essence, those annotations of high overall reputation (left by annotators of high overall reputation and with endorsement from raters determined to be similar to this student carrying even greater weight) were shown to the student. The simulation continued with a rating for that annotation being left by the student, to feed back to future decisions about showing this annotation to other students. Whether the annotation was beneficial to the student was modeled by an increase in that student's average knowledge level (where an object with an appropriate target level of knowledge would result in a greater overall benefit).

For this experiment 3 curriculum sequencing approaches faithful to the CLA were run: raw ecological, ecological with pilot and simulated annealing[1]. For each of these curriculum sequencing approaches, 4 different annotations variants were used: random, greedy god, tally and Cauchy. For the random interactions, each student was randomly assigned up to 3 of the annotations attached to the learning object. For the greedy god variant, each annotation attached to the assigned learning object was used to pre-calculate the learning gains for that combination of student, learning object and annotation and the 3 determined to produce the highest average improvement in overall knowledge were assigned. For the tally approach, each student gave a thumbs up or thumbs down rating after an interaction with a learning object and annotation. These ratings were then used, with the ratings modified based on the similarity (which was calculated using the ratings in common) between the rating student and the active student, to generate a predicted benefit (value of the reputation) for that annotation using Eq.1[2]. The Cauchy variant recorded ratings for and against, similar to the tally, but instead determining the predicted benefit using a Cauchy CDF as in Eq.2 for the adjust function of Algorithm 4 (and Algorithm 3 as well). Here, x_0 is the annotation's initial reputation, x is a modifier based on vF and vA and γ is a factor which, when set higher, is less responsive to the vF and vA.

$$\frac{|\ Rating_{\ for}\ | - |\ Rating_{\ against}\ |}{|\ Rating_{\ for}\ | + |\ Rating_{\ against}\ |} \qquad (1) \qquad\qquad \frac{1}{\pi}\arctan(\frac{x - x_0}{\gamma}) + \frac{1}{2} \qquad (2)$$

Each variant was run for 20 iterations and the results averaged for each minute of simulated instruction. A total of 20,000 minutes of instruction was simulated (with the expectation that our algorithms, Tally and Cauchy, would progressively yield higher mean average knowledge over time)[3]. Random seeds were created (1 for each of the 20 iterations), and used for each of the 12 conditions so that the same students and learning objects were used in the interactions (results are from the varying conditions, not from random difference between the students or learning objects).

---

[1]**Raw Ecological** has each student matched with the learning object best predicted to benefit her knowledge; **Pilot Group** has a subset of the students (10%) assigned, as a pilot group, systematically to learning objects - these interactions are used to reason about the best sequence for the remaining 90% of the students; **Simulated Annealing** is such that during the first 1/2 of the trials there is an inverse chance, based on the progress of the trails, that each student would be randomly associated with a lesson; otherwise, the ecological approach was applied.

[2]This did not integrate the initial annotation reputation and replaced the use of Algorithm 4.

[3]Note that we also adjust the benefit provided by each learning object to be proportional to the time of instruction (i.e. taking less time provides greater benefit). The expected time is attached to each of the learning objects in the repository.

Student knowledge is updated by interaction with a learning object as:

$$\Delta UK\,(s,l,k) = \frac{l.i}{e^{(D*(s.k-l.k^2)}}$$ (3)

Here *s.k* represents the student *s*'s knowledge *k* which is randomly assigned some initial value. *l.k* represents the learning object *l*'s target level of instruction. *l.i* represents the learning object *l*'s impact which represents, for the ideal student, how much the interaction will adjust that student knowledge. The factor *D* is to influence how much the impact decays as a student's knowledge is different from the learning object's target level of instruction. Squaring the difference results in the decay being symmetrical. Each instructional system is modeled as a set of distinct knowledges *k*. For each student their knowledge level (KL) is calculated by averaging all *k*. On the y-axis in Figure 1, we plot the average KL over all students, calling this the mean average knowledge.

2.2 Results

Fig. 1. Comparison of Approaches to Sequencing Learning Objects

The 3 curriculum sequencing approaches detailed above were each run with the 4 annotation approaches (random, greedy god, tally and Cauchy). The random interaction approach to annotating was expected to give the worst performance, and it did. The greedy god approach pre-calculates every annotation and was expected to give the best performance, which was also observed.

The Cauchy and Tally approaches both balance positive and negative ratings in an attempt to preferentially recommend annotations that previous students, particularly students similar to the current student, have found useful. In each of the 3 approaches to curriculum sequencing it was seen that the Cauchy technique outperformed the Tally, suggesting that this is a worthwhile approach to consider when recommending annotations. Using the random and greedy god as baselines, the performance of annotating approaches can be evaluated by where they fall between them (any approach that underperformed the random should be immediately rejected as counter-productive, and as the greedy god is optimal it should never be outperformed). As the curves can be seen to be closer to the upper bound, it suggests that this approach is doing a reasonable job of recommending annotations. The Tally and Cauchy curves are closer to the optimal case in the Ecological with Pilot and Simulated Annealing conditions. Both of these approaches place a higher emphasis on exploration before exploitation compared to the ecological approach. This would suggest that the annotations approach is also benefiting from making exploration a priority.

The design decisions we made to generate the simulation were as follows. Each student and learning object was modeled as having 6 distinct knowledges, randomly assigned to be in the range of [0,1] and each learning object was also randomly assigned impacts in the range

of [-0.05,0.05]. The decay D was set to be 10. Since student learning was modeled on the basis of knowledge levels determined from pre- and post- tests, we included Gaussian noise with standard deviation of 0.1, to be robust in the face of assessment error. The same random number generator seed was used each time learning objects and students were created for the different experimental treatments.

Each student was randomly assigned authorship in the range [0,1] which represents the quality of annotations that student leaves. Students provide, with 100% accuracy, an appropriate rating (thumbs up or thumbs down) depending on whether an annotation helped them learn or not. Students leave a new annotation on a learning object they have just interacted with 20% of the time[4]. Negative annotations represent comments left by students that detract from the educational experience rather than enhance it and were not shown.

3. Discussion

Our peer-based tutoring is distinct in its use of the previous learning experiences of peers, rather than relying on peers to collaborate in real-time (e.g. [5]). While we embrace McCalla's ecological approach [4], we also work to specify algorithms that enable objects in learning repositories to be selected and offer a validation method based on simulated learning which contrasts with other efforts using simulated students ([3]) to predict human performance (yet our comparison of learning effects when validating uses an approach similar to [3] which confirms the value of plotting average knowledge of students). In addition, beyond using tags to infer information about students ([6]), we model the reputation of annotators and focus on the most valuable annotations. While our inspiration for the reputation modeling is the work of [7], we integrate the modeling of student similarity (related to collaborative filtering recommendation [1]) to decide which advice to follow. Possible future work includes i) experimenting with students with different skill level in ratings ii) varying thresholds for excluding poor annotations iii) demonstrating the inferiority of learning without the chance of viewing annotations iv) moving forward to user studies. In conclusion, this research offers important validation of our peer-based tutoring approach as one that creates valuable learning opportunities for students.

References

[1] Breese, J. S., Heckerman, D., and Kadie, C. 1998. Empirical analysis of predictive algorithms for collaborative filtering. Morgan Kaufmann, 43-52.
[2] Champaign, J. and Cohen, R. 2010. An annotations approach to peer tutoring. In The 3rd International Conference on Educational Data Mining, Pittsburgh, PA, USA, June 11-13, 2010. 231-240.
[3] Matsuda, N., Cohen, W. W., Sewall, J., Lacerda, G., and Koedinger, K. R. 2007. Evaluating a simulated student using real students data for training and testing. In User Modeling. 107-116.
[4] McCalla, G. 2004. 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 7, 1-23.
[5] Read, T., Barros, B., Barcena, E., & Pancorbo, J. 2006. Coalescing individual and collaborative learning to model user linguistic competences. User Modeling and User-Adapted Interaction 16, 3-4, 349-376.
[6] Vassileva, J. 2008. Toward social learning environments. IEEE Trans. Learn. Technol. 1, 4, 199-214.
[7] Zhang, J. and Cohen, R. 2008. Evaluating the trustworthiness of advice about seller agents in e-marketplaces: A personalized approach. Electronic Commerce Research and Applications 7, 3, 330-340.

---

[4]When multiple annotations are attached to an object, they may enhance or detract from one another. Due to the computational demands for reasoning about such interactions, the decision was made to reason about each annotation independently.