Modeling Trustworthiness of Peer Advice in a Framework for Presenting Web Objects that Supports Peer Commentary

John Champaign and Robin Cohen

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

Abstract. In this paper, we present an approach aimed at enabling users to enrich their experience with web-based objects (texts or videos). In particular, we consider a social network of users offering commentary on the web objects they have experienced together with opinions on the value of this commentary being registered by peers. Within this framework, we integrate a reasoner that personalizes the presentation of these annotations to each new user, selectively limiting what is displayed to promote the commentary that will lead to the most effective knowledge gains, based on a modeling of the trustworthiness of the annotator and the similarity of peers who have found this commentary to be useful. We demonstrate the effectiveness of our approach for selective presentation of these web document annotations by constructing a simulation of knowledge gains achieved by users. Our method is shown to approach the ideal knowledge gains achieved by an optimal algorithm, far outpacing a system where a random selection of commentary is offered (as might match what users would experience if employing self-directed limiting of browsing behaviour). As a result, we offer an effective method for enhancing the experiences of users in contexts with potentially massive amounts of peer commentary.

Keywords: Trust in Social Networks, Trust-based Recommender Systems, User modeling in Social Networks, Trust of Users in Recommenders, Peer-based Intelligent Tutoring

1 Introduction

Social networks of peers allow for users to be exposed to a large variety of online material (e.g. texts and videos) suggested as valuable by others. It may be difficult at times for users to select the most appropriate content themselves, in order to enhance their learning experience, as one must cope with the overabundance of possible material from which to learn. In this research, we introduce a valuable technique known as annotations. This enables users to comment on the web objects they have experienced, in order for new users to have better information as they proceed to process the new information, towards their learning. The challenge then becomes to reason about which of these annotations should be shown to each new user.

We offer an algorithm that includes reasoning not only about the similarity of the current user and the past users who have experienced the object but also about the reputability, both of the annotator and in general of the annotation (as reflected in the collective approvals or disapprovals of that annotation, provided by others who have experienced it). Taken together, this system enables effective, streamlined presentation that can be demonstrated to offer significant knowledge gains. We demonstrate the effectiveness of our approach for personalizing the selection of the commentary on web documents by constructing simulations of the knowledge achieved by users presented with our selection of annotations, compared to what would be gained through two baselines.

In all, the user modeling that we introduce that considers trustworthiness, similarity and knowledge gains serves to provide an effective avenue to support peer-based information sharing within web-based contexts.

2 An Approach for Modeling Knowledge Gains by Users

In order to introduce our proposal for supporting commentary left on web objects together with selective presentation of that commentary to peers, we first require a framework for reasoning about the knowledge gains that users will achieve when exposed to a set of web documents.

For this, we draw from literature on intelligent tutoring systems and in particular the ecological approach espoused by McCalla [1] which promotes the construction of a history of experiences of peers with web objects which is then leveraged in the determination of what should be offered to future users in order to promote effective learning gains.

We advocate the selection of objects from a large repository based on the gains in knowledge that have been achieved by similar peers in the past. Each web object is assumed to have a target level of knowledge and then a factor known as impact that increases when the knowledge level of the user experiencing the object is closer to the target. The aim is to be delivering to each user those objects that best serve to increase their overall knowledge.

As such, we must nominally begin with a modeling of the overall knowledge level of each user which may then be compared to the target level of each web object in determining whether knowledge increases have been achieved with the presentation of each new object.[1]

3 Supporting Annotations of Web Objects

Our proposal is to allow peers to leave commentary on web objects which future peers may then be viewing, to increase their knowledge. In order to make effective decisions about which annotations to show to each new user, we require a bootstrapping phase. Here, a set of peers will be invited to leave commentary on objects, another set of peers will be exposed to all of these commentaries, leaving a thumbs-up or thumbs-down approval rating and then once this phase is complete, we will have an initial set of annotations with attached approval ratings.

[1] One might imagine, for instance, grade levels from A+ down to D- to reflect a user's current knowledge on the central topic conveyed by the web object, where we aim to increase the knowledge upon exposure to the objects we have selected for presentation. In e-learning contexts this is derived from an assessment quiz.

See Figure 1 for an example of an annotation and the approval ratings it has received.

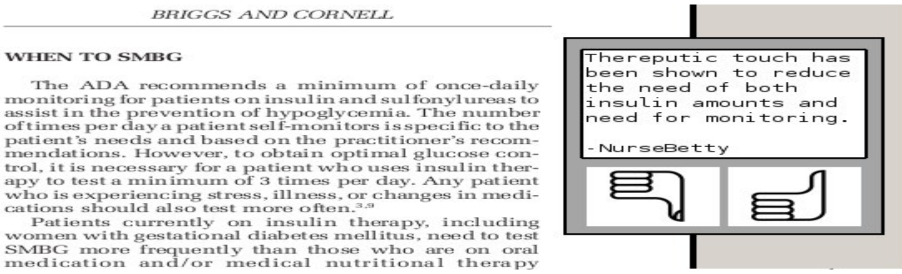

Fig. 1: Example of a low-quality annotation

All users in the system are represented in terms of their knowledge levels as well. As such, once the bootstrapping phase is over, we are now able to reason about which annotations to show a future user based on: i) the reputation of the annotation (how many thumbs up vs thumbs down) ii) the reputability of the annotator (over all annotations left by this user, what percentage received thumbs up) and iii) the similarity of the peers in terms of how they have rated the same annotations in the past. Our proposed formulae for modeling reputation and similarity and our proposed algorithm for determining which annotations to show to each new user are outlined in detail below.

3.1 Overview of the Reasoning

Our proposed model for reasoning about which annotations to show a new user u 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 user u (calculated using Algorithm 2) to value votes by similar users more highly. The reputation of each annotation for a user u is calculated by Algorithm 4 using an adjust function to scale (i) according to (ii) and those annotations with the highest reputation are shown.

3.2 Calculation of Reputability

We assume that there is currently an object in focus from the repository, for this user. For example, in the context of intelligent tutoring this object may have been selected according to its potential for enhancing the learning of this particular student, based on a modeling of that student's level of achievement and progress in the topic material to date [3]. Within a general web-based information system, the object may simply have been selected by the user for possible consideration.

Determining which of the full set of annotations left on the object should be shown to the user is inspired by the model of [2] which models trustworthiness based on a combination of private and public knowledge (with the latter determined on the basis of peers). Our process integrates i) a restriction on the maximum number of annotations

Algorithm 1: User Reputation
```
//Consider user as an annotator
calUserReputation (User u)
if number of annotations by u == 0 then
    R(u) = 0.5; //Reputation of user u
else
    R(u) = 0;
    foreach annotation a of u do
        R(u) += calcAnnRep(a);
    end
    R(u) /= num of annotations by u;
end
return R(u)∈ [0, 1];
```

Algorithm 2: User Similarity
```
similarity (User c, User 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;
    end
end
similarity = (vS − vD)/(vS + vD);
return similarity∈ [−1, 1];
```

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

Algorithm 4: Specific Annotation Reputation
```
calAnnRepSpecific (Annotation a, User u)
foreach vote on annotation do
    sim = similarity(u, voterUser);
    if vote.for then
        vF+ = 1 * sim;
    else
        vA+ = 1 * sim;
    end
end
return adjust(a.initRep, vF, vA);
```

shown per object ii) modeling the reputation of each annotation iii) using a threshold to set how valuable any annotation must be before it is shown iv) considering the similarity of the rating behaviour of users and v) showing the annotations with the highest predicted benefit.

Let A represent the unbounded set of all annotations attached to the object in focus. Let r_j^a = [-1, 1] represent the jth rating that was left on annotation a (1 for thumbs up, -1 for thumbs down and 0 when not yet rated). The matrix R has R^a representing the set of all ratings on a particular annotation, a, which also represents selecting a column from the matrix. To predict the benefit of an annotation for a user u we consider as Local information the set of ratings given by other users to the annotation. Let the similarity[2] between u and $rater$ be $S(u, rater)$. Global information contains all users' opinions about the author of the annotation. Given a set of annotations $A_q = \{a_1, a_2, ..., a_n\}$ left by an annotator (author) q we first calculate the average interest level of an annotation a_i provided by the author, given the set of ratings R^{a_i} to the a_i, as follows:

$$V^{a_i} = \frac{\sum_{j=1}^{|R^{a_i}|} r_j^{a_i}}{|R^{a_i}|} \tag{1}$$

The reputation of the annotator q is then:

$$T_q = \frac{\sum_{i=1}^{|A_q|} V^{a_i}}{|A_q|} \tag{2}$$

[2] The function that we used to determine the similarity of two users in their rating behaviour examined annotations that both users had rated and scored the similarity based on how many ratings were the same (both thumbs up or both thumbs down). The overall similarity score ranged from -1 to 1, a range which lets us scale ratings based on students being unrelated (0), highly similar (1) or highly dissimilar (-1).

which is used as the Global interest level of the annotation.

A combination of Global and Local reputation leads to the predicted benefit of that annotation for the current user. To date, we have used a Cauchy CDF[3] to integrate these two elements into a value from 0 to 1 (where higher values represent higher predicted benefit) as follows:

$$\text{pred-ben}[a, current] = \frac{1}{\pi}\arctan\left(\frac{(vF^a - vA^a) + T_q}{\gamma}\right) + \frac{1}{2} \tag{3}$$

where T_q is the initial reputation of the annotation (set to be the current reputation of the annotator q, whose reputation adjusts over time, as his annotations are liked or disliked by users); vF is the number of thumbs up ratings, vA is the number of thumbs down ratings, with each vote scaled according to the similarity of the rater with the current user, according to Eq. 4. γ is a factor which, when set higher, makes the function less responsive to the vF, vA and T_q values[4].

$$v = v + (1 * S(current, rater)) \tag{4}$$

Annotations with the highest predicted benefit (reflecting the annotation's overall reputation) are shown (up to the maximum number of annotations to show, where each must have at least the threshold value of reputation).

4 Experimentation

To begin, an algorithm needs to be run that selects the learning objects to be shown to the users. For our experimentation we assumed this to be a Collaborative Learning Algorithm for a peer-based intelligent tutoring (omitted here for brevity, see [3]). Each learning object in the system is modeled as a set of distinct knowledges k. Each annotation is attached to a learning object and is modelled with an impact factor which serves to increase or decrease the user's knowledge. For each user their knowledge level (KL) is calculated by averaging all k. On the y-axis in the below figures, we plot the average KL over all users, calling this the mean average knowledge. The x-axis represents the time of instruction. Throughout our experiments, we restrict the maximum number of annotations shown to a student to be 3. We compare our algorithms against two baselines. **Random Association** associates each student with a randomly assigned annotation; **Greedy God** chooses the best possible interaction for each student for each trial by pre-calculating the benefit of every annotation.

Our first experiment (Figure 2) is aimed at demonstrating that our approach for reasoning about annotations to show to a user is effective with a variety of possible settings for threshold that serves to cut off whether an annotator is considered to be reputable or not.

This is achieved by soliciting ratings from other students who have experienced annotations. These ratings are used to determine the reputability of both the annotations

[3] This distribution has a number of attractive properties: a larger number of votes is given a greater weight than a smaller number (that is, 70 out of 100 votes has more impact than 7 out of 10 votes) and the probability approaches but never reaches 0 and 1 (i.e. there is always a chance an annotation may be shown).

[4] In this work, a value of 0.2 was used for gamma.

themselves and the students who left them. Only annotations from reputable annotators are included in the collection of annotations that may be shown to users viewing a learning object. An annotator is labelled as reputable if his overall trustworthiness is above a set threshold. This is achieved by avoiding the case where annotations that are left are generally disliked. A trade-off exists: if the threshold is set too low, bad annotations can be shown to students, if the threshold is set too high, good annotations can be stigmatized. In order to determine an appropriate level in the context of a simulation, we examined cut-off thresholds for annotations of 0.5, 0.4, 0.3, 0.2 , 0.1 and 0.0. The maximum number of annotations shown is set at 3.

The results in Figure 2 indicate that our algorithm is still able to propose annotations that result in strong learning gains (avoiding the bad annotations that cause the random assignment to operate less favourably) with the varying cut-off thresholds. The 0.0 threshold (indicative of not having a cut-off since this is the lowest possible reputation) underperformed in the early stages of the curriculum, indicating that having a threshold is worthwhile. The other various thresholds did not display substantial differences, which suggests a low cut-off threshold, such as 0.1 should be considered to allow the greatest number of annotations to be used by the system.

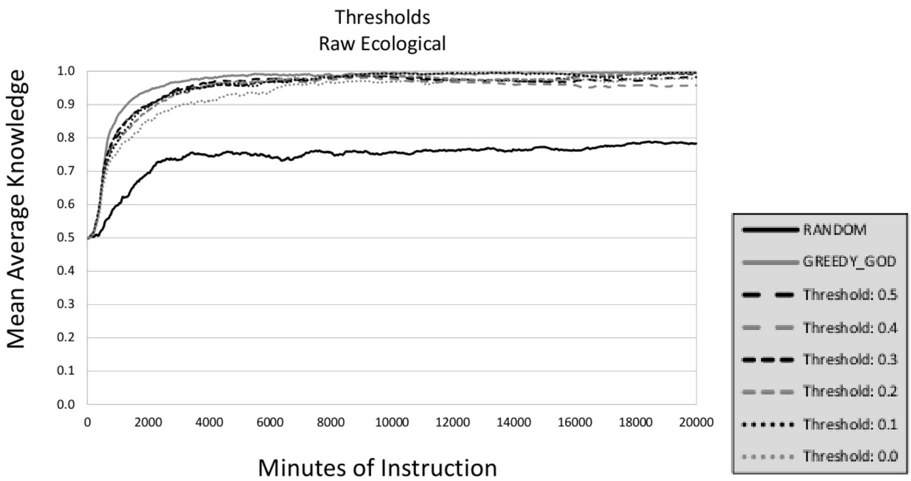

Fig. 2: Comparison of Various Thresholds for Removing Annotations

In our next set of experiments, we show that our approach for selecting annotations can cope, even when there is an abundance of poor annotators. In this experiment (Figure 3) we contrast the Cauchy technique, described above with the Tally which is simply the average interest level as calculated in Equation 1 and the Trust-Based approach, which initially recommends annotations based on the reputation of the annotator (Equation 2) but increasingly bases its recommendation on the average interest level as more votes are registered on the annotation (in this work we uniformly distributed the weight of each recommendation between completely based on the annotator reputation

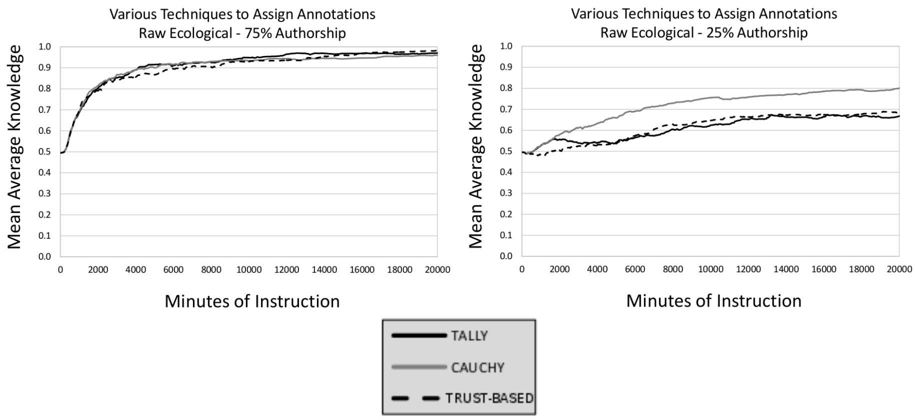

Fig. 3: Comparison of Annotation Assignment Techniques - Raw Ecological

when no votes were registered through to being completely based on the average interest level once 10 votes were registered).

The formula used for the trust-based calculation is as in Formula 5 below. The term $\frac{\sum_{j=1}^{|R^{a_i}|} r_j^{a_i}}{|R^{a_i}|}$ represents the community's rating of the reputability of the annotation. The term $\frac{\sum_{i=1}^{|A_q|} V^{a_i}}{|A_q|}$ represents the initial reputation of the annotation, which is set to be the annotator's reputation at the time the annotation was created. Over time, we want to place a greater weight on the community's view of the reputability of the annotation instead of the inherent reputations of the annotator. With each vote made on the annotation, we move the weight to a greater emphasis on the community's view of the reputability, until we reach a point where the community's perspective is the entire reputation and the reputation of the author is no longer considered. N_{min} is this point where we no longer consider the author's reputation. In our simulations we set N_{min} to be 10.

$$\text{pred-ben}[a_i, current] = min(1, \frac{|R^{a_i}|}{N_{min}})\frac{\sum_{j=1}^{|R^{a_i}|} r_j^{a_i}}{|R^{a_i}|} + max(0, (1 - \frac{|R^{a_i}|}{N_{min}}))\frac{\sum_{i=1}^{|A_q|} V^{a_i}}{|A_q|} \quad (5)$$

Authorship is a variable in the simulation which determines the percentage of time that a user will leave a good annotation. Figure 3 shows that the three techniques, Tally, Cauchy and Trust-Based each perform well in environments with many strong annotators (the 75% authorship). In an environment with many poor annotators (the 25% authorship), the Cauchy technique outperforms the other two. Regardless, we can cope even if poor annotators exist.

Even for the worst annotators, there is a chance that they will leave an occasional good comment (which should be promoted), or improve the quality of their commentary (in which case they should have a chance to be redeemed). For this experiment (Figure 4, shown for Cauchy), we considered allowing an occasional, random display of annotations to the students in order to give poorly rated annotations and annotators a second chance and to enhance the exploration element of our work. We used two baselines (random and Greedy God again) and considered 4 experimental approaches. The first used our approach as outlined above, the standard authorship of 100%, a cut-off threshold of 0.4 and a 5% chance of randomly assigning annotations. The second used an exploration value of 10%, which meant that we used our approach described above 90% of the time, and 10% of the time we randomly assigned up to 3 annotations from learning objects. We also considered conditions where annotations were randomly assigned 20% and 30% of the time. Allowing a phase of exploration to accept annotations from students who had previously been considered as poor annotators turns out to still enable effective student learning gains, in all cases. Our algorithms are able to tolerate some random selection of annotations, to allow the case where annotators who would have otherwise been cut off from consideration have their annotations shared and thus their reputation possibly increased beyond the threshold (if they offer an annotation of value), allowing future annotations from these students to also be presented[5].

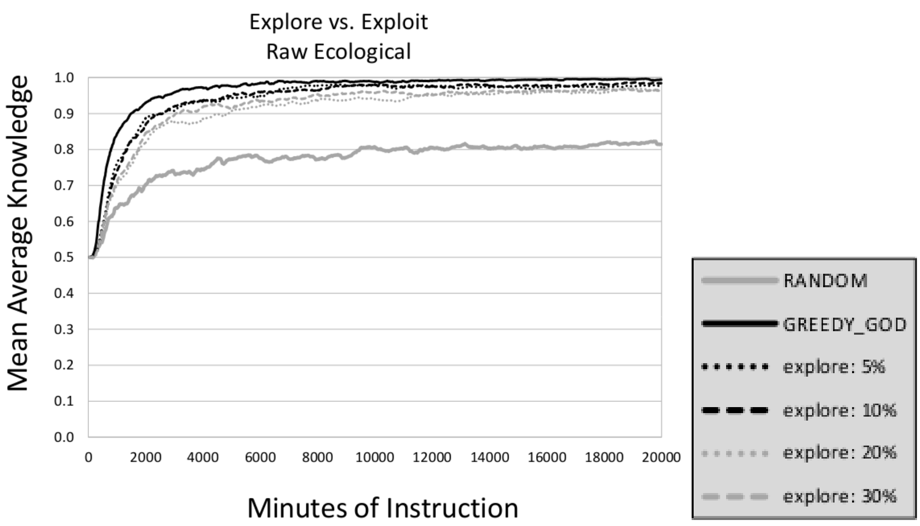

Fig. 4: Explore vs Exploit: Raw Ecological

For our next experiment we consider the following. Our trust modeling parameters combine experience-based trust and public trust together, where the balance of these two elements can shift, over time, due to the setting of their relative weights. In these experiments (Figure 5, shown for Cauchy) we chose to value considerably more (or less) the votes provided for and against (so the public opinion) to demonstrate that our approach is effective in a variety of possible scenarios.

In our simulation we used a factor of 2 for the scaling provided by student similarity and made the recommendations using the Cauchy approach. This means that a highly

[5] In this simulation the authorship ability of students remained consistent throughout the experiment. Added randomness in the assignment of annotations would be more effective in situations where the authorship ability of students varied throughout the course of a curriculum.

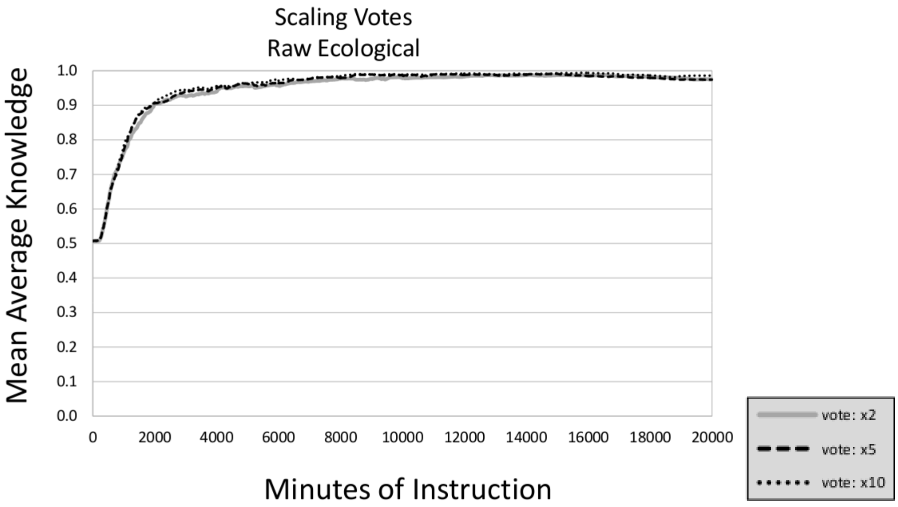

Fig. 5: Increasing Impact of Student Similarity When Scaling Ratings

similar student's (i.e. similarity of 1) rating could be doubled (and would be considered equivalent to the ratings provided by two generic students.) (This falls out of lines 5 and 7 in Algorithm 4.) A student with a similarity of 0 would have their vote be treated normally. A consistently dissimilar student would have their rating weakened. We can now explore other scaling: for example with a similarity of 1 and a scaling factor of 10, a thumbs up would be considered equivalent to 10 thumbs up ratings, with a similarity of -0.9 and and a scaling factor of 10, a thumbs down would be equivalent to 9 thumbs up ratings.

The curves appear to be very closely matched to one another, but upon closer examination it can be seen that there are small variations, although in each case all three factors provide very similar results. These results make sense. For our simulation we attached the 3 most highly rated annotations to the learning object being seen by the student. The slight differences in results show that sometimes with different vote scaling factors this will result in different annotations being assigned. It also shows that this doesn't result in a great deal of difference, since each approach is making worthwhile recommendations to the students.

Finally, in order to test the simulation software for bugs in the source code, where students rate whether or not an annotation was useful to them was replaced with a method that would randomly rate annotations. The expectation from this simulation was that each of the techniques for recommending annotations would degrade to provide results comparable with random assignment of annotations. For this simulation only a single iteration was run, which accounts for the greater volatility in the results. Figure 6 shows the techniques with the ratings replaced with random ratings. These show the expected degradation to resemble the random assignment (which may occur if students are left their own devices), confirming the value of our approach.

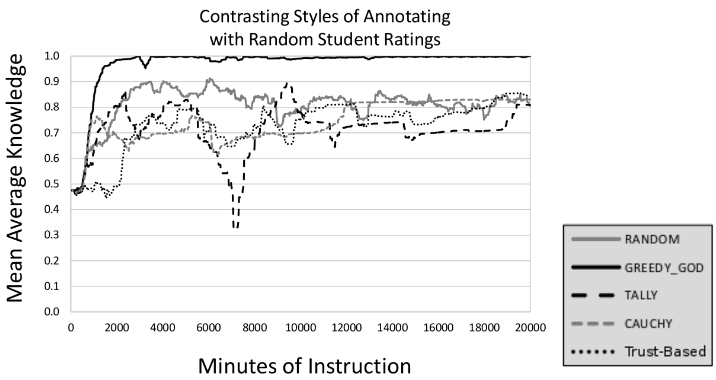

Fig. 6: Three Techniques for Recommending Annotations with Student Ratings Replaced By Random Ratings

5 Discussion

This research has proposed an approach for allowing commentary to be left on web objects and to present a streamlined, personalized selection of those annotations to users. The aim is to be supporting an increase in the knowledge gained by the users. A central element in the overall solution is the representation of the reputation of each annotator and of each annotation. The other critical component is a representation of the similarity between users. We are able to demonstrate that this framework delivers effective knowledge levels to users, compared to ones that operate with perfect knowledge. In addition, we show considerable gains beyond algorithms with random selection of objects. When users are left to make selective restrictions themselves, on the basis of their current preferences and whims, without a principled approach for ensuring increased learning of the material, then this process could indeed quite resemble what is labelled here as random. We therefore feel that we are offering an important new direction for enabling users to intelligently cope with the plethora of possible commentary being offered by their peers.

Indeed, when a social network of peers is involved, a user does eventually need to make decisions about which peers to listen to. In our framework, we allow those peers who have not been helpful in the past to be redeemed and embraced anew within the social network (through our discussion of exploit vs explore); in addition, we do not blindly accept the advice that is offered by each respected peer, separately evaluating its worth towards the learning gains to be achieved. As such, we are advocating the use of user models of peers as part of the reasoning of a user that is processing web content. Our approach is similar to researchers such as [4] which promote the modeling of peers in social networks for more effective decision making and learning by users. In contrast with other intelligent tutoring systems researchers such as [5], however, we are able to leverage past experiences rather than requiring peers to be assisting, in real-time with the learning that is achieved. As such, our users are free to browse and view the accompanying annotations on their web documents and videos, at their leisure. We are also careful to avoid possibly harmful content from peers.

Similar to collaborative filtering recommender systems, we are indeed concerned with similarity of peers as part of our selection process. But here we integrate a consideration of trust modeling as well. The particular combination of similarity and trust that we offer here is distinct from that of other researchers such as [6, 7]. Our consideration of the reputation of the annotator as well as the annotations themselves introduces a novel method for overcoming cold start issues.

Other trust modeling researchers have proposed the limiting of social networks of advisors, in order to make the trust modeling process more efficient and effective, when it is used to drive the decision making of a user [8, 9]. Our algorithms do incorporate a process of restricting the set of advice that is considered; but beyond this we also integrate the elements of similarity and modeling the tutorial value of each new object that is presented. Future work would consider integrating additional variations of Zhang's original model [10] within our overall framework. For example, we could start to flexibly adjust the weight of Local and Global reputation incorporated in the reasoning about which annotation to show to a student, using methods which learn, over time, an appropriate weighting (as in [10]) based on when sufficient Local information is available and can be valued more highly. In addition, while trust modeling would typically have each user reasoning about the reliability of every other user in providing information, we could have each student maintain a local view of every other student's skill in annotation (though this is somewhat more challenging for educational applications where a student might learn and then improve their skill over time and where students may leave good annotations at times, despite occasionally leaving poor ones as well). In general, studying the appropriate role of the Global reputation of annotations, especially in quite heterogeneous environments, presents interesting avenues for future research (since currently this value is not, in fact, personalized for different users).

Note that our work takes a different perspective on Local and Global than in Zhang and Cohen's collected work. Rather than personal experiences (Private) being contrasted with the experiences of other buyers (Public), we contrast experiences with a particular annotation (Private) to the overall experiences with the annotation's author (Public). This is a novel contribution that can be considered in the e-commerce domain. The analogous business model would be reasoning about particular products or services independently of the business. For example, if a buyer thinks Microsoft Office is excellent, Windows 7 is ok and Internet Explore is awful, she may have a neutral view of a new product from Microsoft (the company that makes each of these products), but be enthusiastic about a new release for Microsoft Office and pessimistic about a new release for Internet Explorer. This helps mitigate the cold start problem and allows a business to be evaluated in a richer manner: rather than being good or bad, reputations can provide more nuanced recommendations.

Our simulation modeled insightful students who are able to accurately rate whether a specific annotation has helped them learn. In practice, the ability of students to make this identification may not be as accurate as we have assumed in this work. In future work we are interested in exploring the impact of modeling some uncertainty in this process, where students sometimes give an incorrect rating to annotations they experience.

In addition, our aim is to offer this system to real users, in order to learn more about its usability and user satisfaction with the results that it delivers. We have already conducted a preliminary study with human users; this focused more on validating our particular intelligent tutoring algorithm that drives the selection of the web documents (against which annotations are applied) but we asked our subjects 4 questions about the annotations during their exit survey: **1)** Do you find any value in using learning objects with annotations? **2)** How likely would you be to contribute an annotation to a learning object if using a system that supported this? **3)** How often might you leave annotations? **4)** How satisfied would you be reading annotations left by previous students?

Participants were given a 11 point scale, ranging from -5 to 5 for Q1, Q2 and Q4 with the labels "less value", "unlikely" and "unsatisfied" (respectively) at -5, "neutral" at 0 and "more value", "likely" and "satisfied" (respectively) at 5. For Q3 participants were given an 11 points scale ranging from 0 to 10 with the labels "never" at 0 and "always" at 10. For the 23 participants the feedback was (question, mean, standard deviation): Q1 (2.5, 2.79); Q2 (1.63, 3.21); Q3 (4.64, 2.59); Q4 (2.32, 2.98).

Although participants were mostly neutral with respect to creating new annotations (Q2,Q3), they were positive about using a system where other students have left annotations on learning objects. This conforms to research on participatory culture (e.g. [11]) which has shown that contributors usually greatly outnumber consumers. It has been shown to be possible (e.g. [5]) to use incentives to encourage greater participation.

Acknowledgements. Financial support was received from NSERC's Strategic Research Networks project hSITE.

References

1. 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 **7** (2004) 1–23
2. Zhang, J., Cohen, R.: A comprehensive approach for sharing semantic web trust ratings. Computational Intelligence **23**(3) (2007) 302–319
3. 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)
4. Vassileva, J., Mccalla, G., Greer, J.: Multi-agent multi-user modeling in i-help. User Modeling and User-Adapted Interaction **13** (2003) 179–210
5. 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
6. Breese, J.S., Heckerman, D., Kadie, C.: Empirical analysis of predictive algorithms for collaborative filtering, Morgan Kaufmann (1998) 43–52
7. 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
8. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: Travos: Trust and reputation in the context of inaccurate information sources. Autonomous Agents and Multi-Agent Systems **12**(2) (2006) 183–198
9. Yu, B., Singh, M.P.: Detecting deception in reputation management. In: AAMAS. (2003) 73–80
10. Zhang, J., Cohen, R.: Evaluating the trustworthiness of advice about seller agents in e-marketplaces: A personalized approach. Electronic Commerce Research and Applications **7**(3) (2008) 330–340
11. Almeida, R.B., Mozafari, B., Cho, J.: On the Evolution of Wikipedia, International Conference on Weblogs and Social Media (2007)