# CS 886 - AI for Healthcare Applications

## Peer-Based Educational Tutoring Systems: A Corpus-Oriented Approach

**John Champaign**
**Robin Cohen**

**Oct 1st, 2010**

***David R. Cheriton School of Computer Science University of Waterloo***

# Presentation Overview

- Journalism: Inverted Pyramid
  - ITS Overview
  - Relevance to Healthcare Applications
  - Details of my work
  - Related Work
  - Future Work

# ITS Overview

- Active area of AI research since 1960's
- Wide variety of systems built on diverse techniques
  - Teaching physics using a natural language tutor
  - Teaching soldiers in Iraq intercultural skills in a VR environment
  - Iteratively working with students to transform a high-level, natural language query to SQL
- What haven't we seen more widespread deployment of ITS?

# ITS Overview (cont'd)

- Persistent (high) cost of development
  - 200 – 1000 hours of development time / hour of instruction (Murray, 1999)
- Possibility of decreasing costs by leveraging peer interactions
- Which peers should form basis of tutoring?
- What content should be presented to a student?

# ITS Architecture

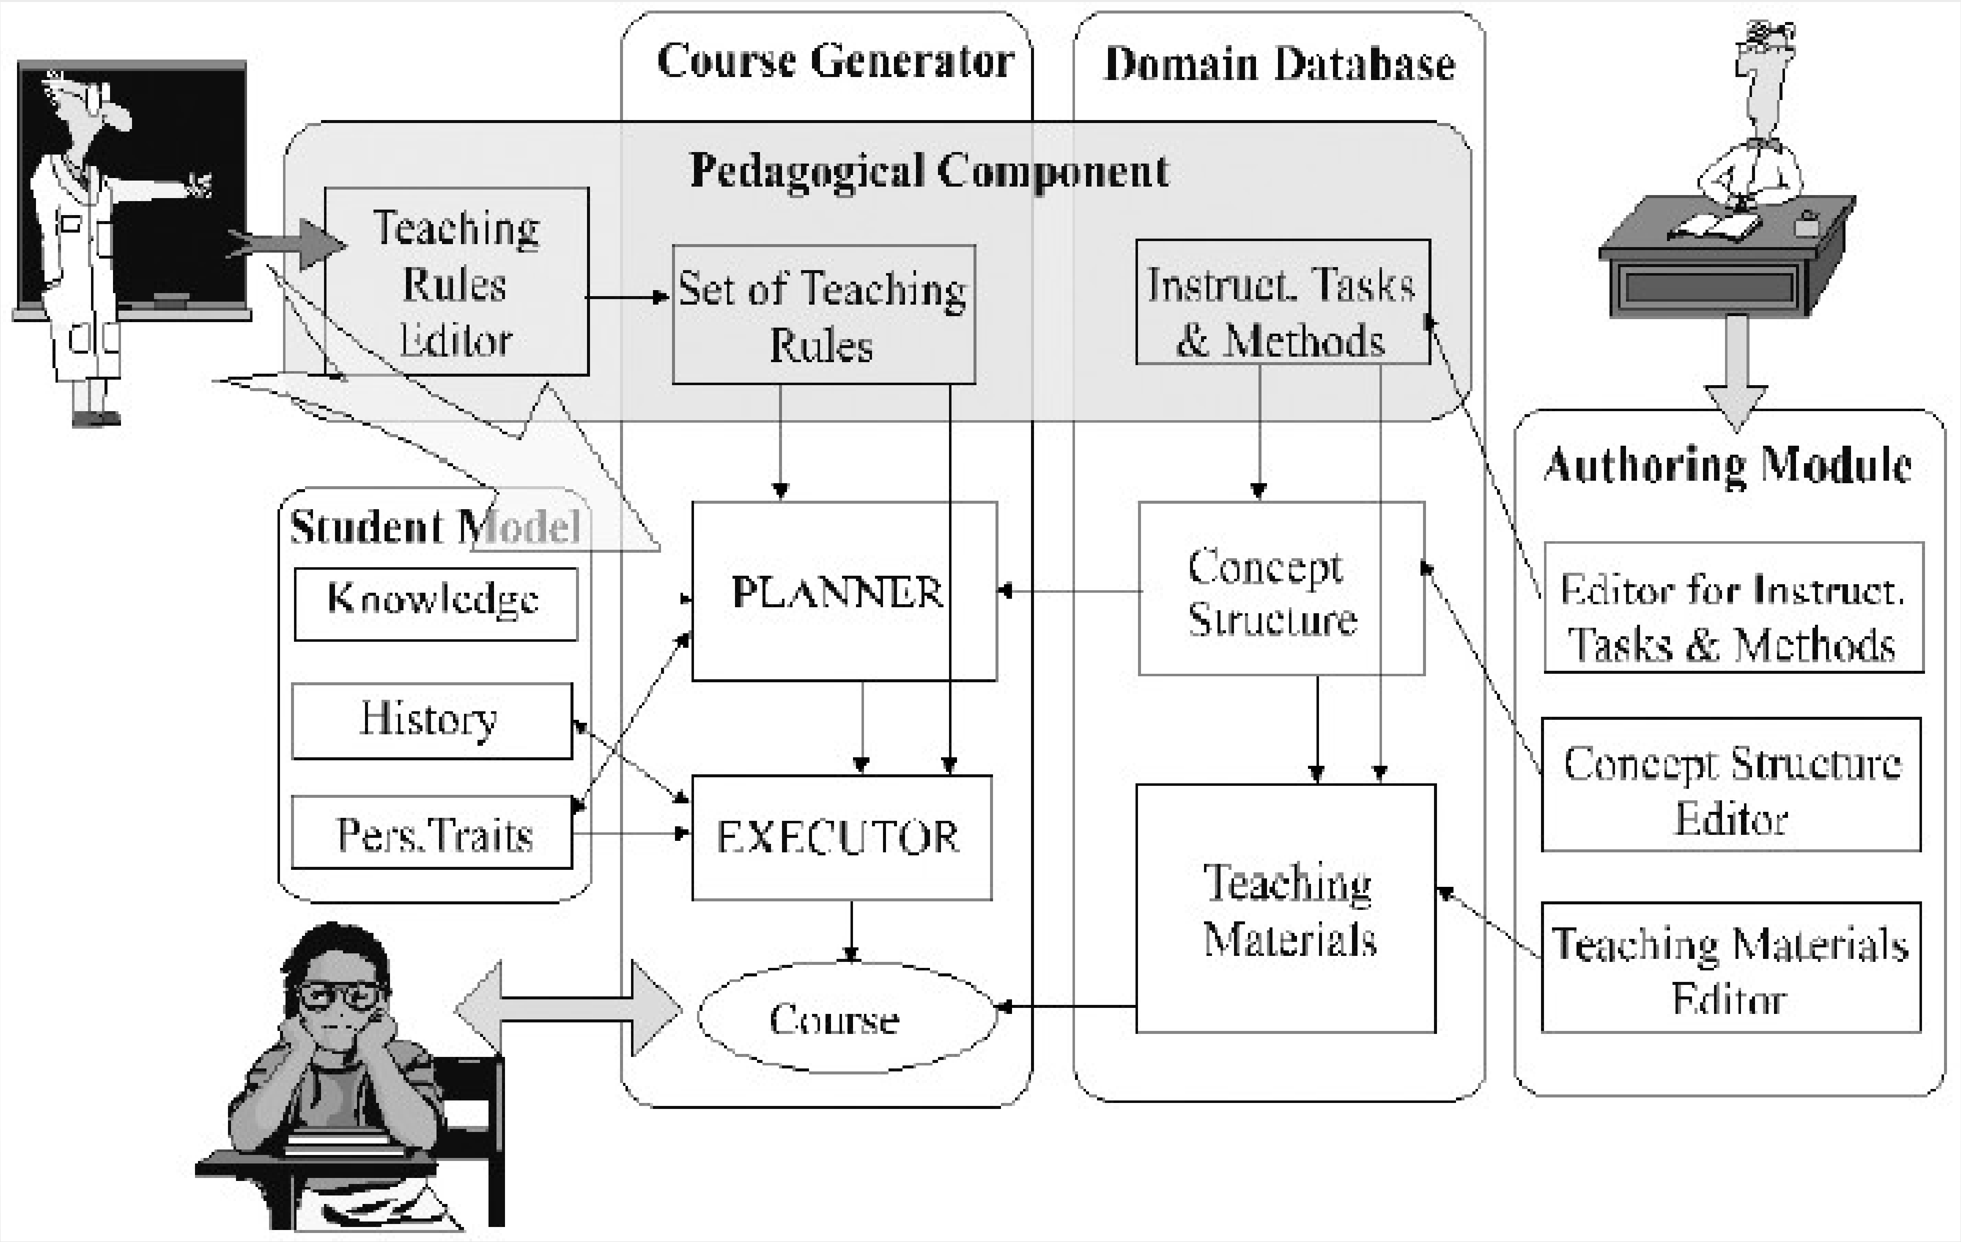

# Heath Care

- hSITE
  - Healthcare Support through Information Technology Enhancements
  - Diane Doran
- *"Deliver the right information to the right people at the right time"*
- Projecting our work into this domain
  - Home Healthcare Nurse
  - Recently Diagnosed Patient

# Motivating Examples - Patient

- A patient goes to the doctors office with symptoms
- Is diagnosed with diabetes
- Handed a pamphlet and sent on his way

# Motivating Examples - Patient

- Instead, with our approach he's given access to an on-line system
- It models his understanding, and by matching with similar, previous students tailors a curriculum for him
- Intelligently shows useful annotations previous students have left (and allows him to leave comments for subsequent students)
- After usage, he can suggest refinements to the lessons he uses, which are then evaluated and may be shown to future students

# Motivating Examples – Home Healthcare Worker

- Nurse gets reassigned from one speciality
  - Oncology -> Gerontology
- Under current system, a “dispatcher” is available by phone, can try to get needed information
- Very overworked
- Often very limited equipment (not even cell phones)

# Motivating Examples – Home Healthcare Worker

- Under our system, nurse can retrain on his own schedule
  - Anywhere there is a PC
- Given the opportunity to interact with other nurses through annotations on lessons, leave and answer questions, share knowledge
- Some element of “expert system”, where the system can be used to provide decision making knowledge on site
  - Perhaps with a client's PC

# Motivating Examples - Doctor

- Doctors INCREDIBLY busy
  - Unable to devote much time to secondary projects, such as patient educations or healthcare worker training
- System automates many of the things that would historically have to be hand coded by an expect (such as curriculum sequence)
- Allows population to make refinements, which can be over-ruled by instructors / experts

# McCalla's Ecological Approach

*“Attaching models of learners to the learning objects they interact with, and mining these models for patterns that are useful for various purposes.”*

- Philosophical
- Example learning objects:
  - Books, web pages, research articles, videos
- Real-time process
  - Interaction history is interpreted as needed

# My Work

- Leverage interactions with students to simplify ITS development
- Ecological, Peer-Based Approach to
  - Curriculum Sequencing
  - Annotations
  - Corpus Approach

# ITS Architecture We're Interested In

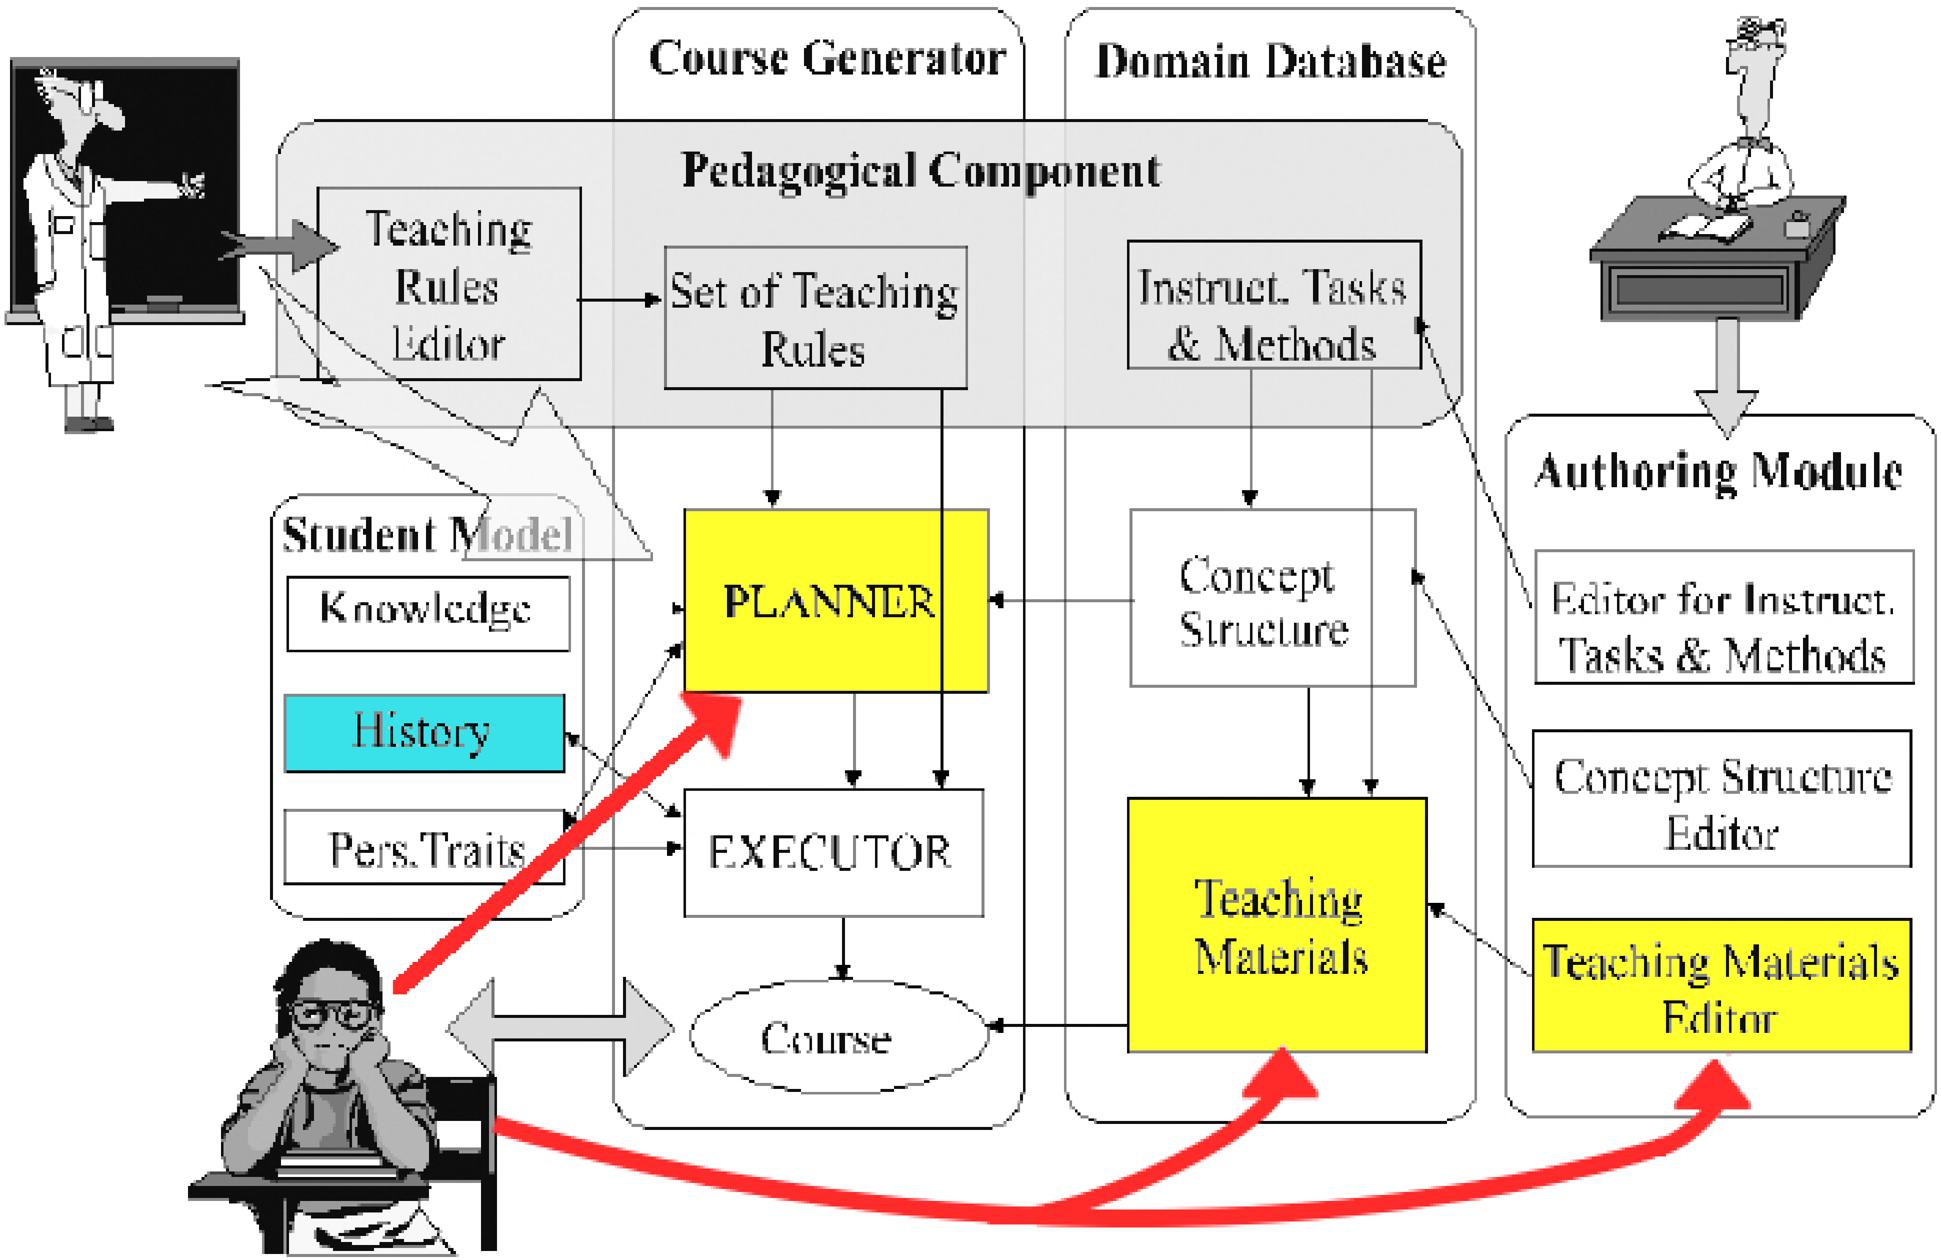

# Curriculum Sequencing

- Given a set of learning objects and a group of students, over multiple iterations, which object should be assigned to each student?
- Collaborative filtering inspired approach, where learning objects that were useful to a similar student in the past are assigned to each student
- Pre-tests and post-tests applied after each interaction (and these are used to reason about similarity between users)

# Curriculum Sequencing

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

- p[a,l] : anticipated benefit to active user, *a*, from interacting with a given learning object *l*
- ∑ : consider interactions of all previous students with the learning object *l*
- w(a,j) : how similar the student *j* was to the active user
- v(a,j) : value of the interaction to student *j*
- K : a normalizing factor

# Example

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

- Similarity = $\frac{1}{1+difference}$
- LO[1; S1(B,C), S3(B,A+)]; LO[2; S2(B-,A)]
- S4, current assessment B+
- Difference of B & B+ is 1, difference of B+ & B- is 2
- Similarity of S1 (or S3) & S4: $\frac{1}{1+1}=\frac{1}{2}$
- Similarity of S2 & S4: $\frac{1}{1+2}=\frac{1}{3}$

# Example (cont'd)

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

- LO[1; S1(B,C), S3(B,A+)]; LO[2; S2(B-,A)]
- S4, current assessment B+

- p[S4,LO1] = $\frac{1}{2} * (\frac{1}{2} * -3 + \frac{1}{2} * 4) = \frac{1}{4}$
- p[S4,LO2] = $(\frac{1}{3} * 4) = \frac{4}{3}$
- System recommends LO2 for S4

# Curriculum Sequencing

- Random Baseline
- Greedy God
- Raw Ecological
- Ecological with Pilot
- Simulated Annealing

- Students / LO conditions : 10/10, 50/50, 50/100

# Curriculum Sequencing

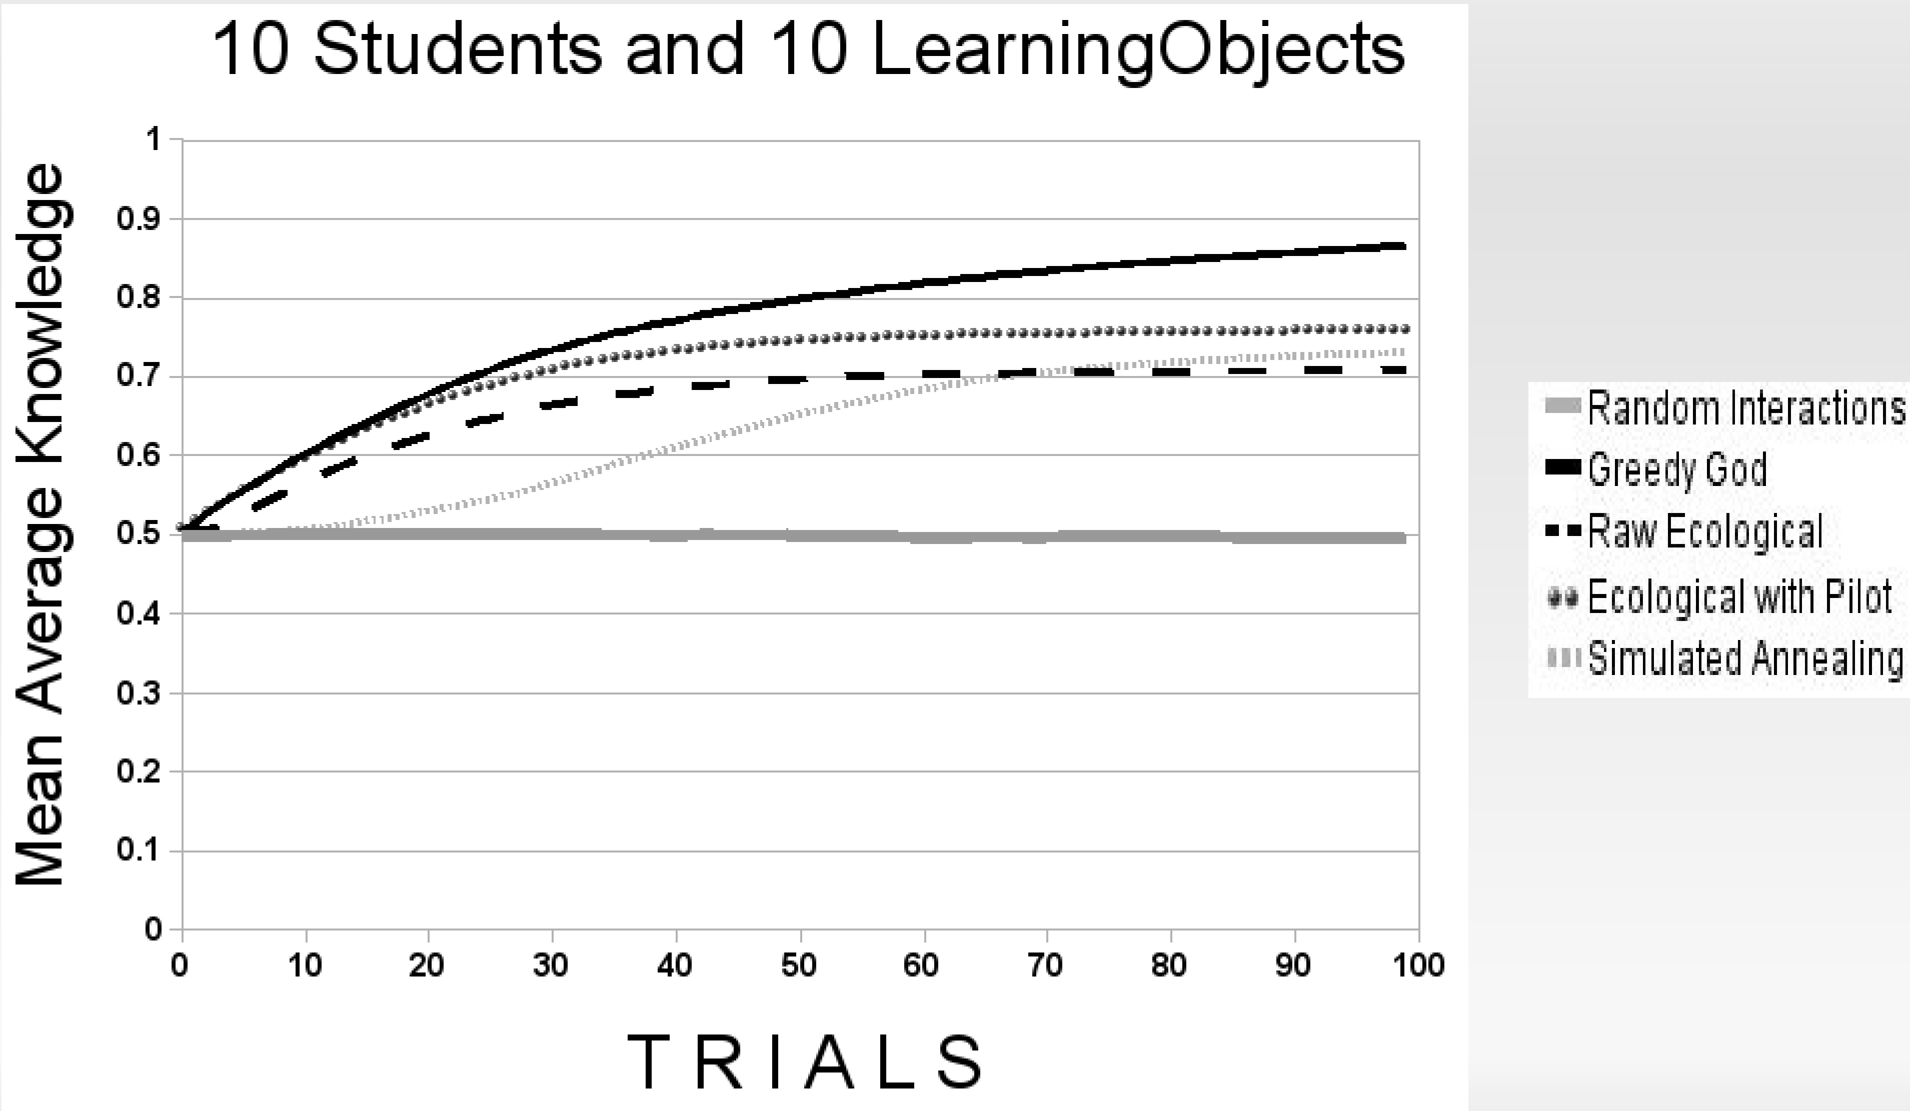

# Curriculum Sequencing

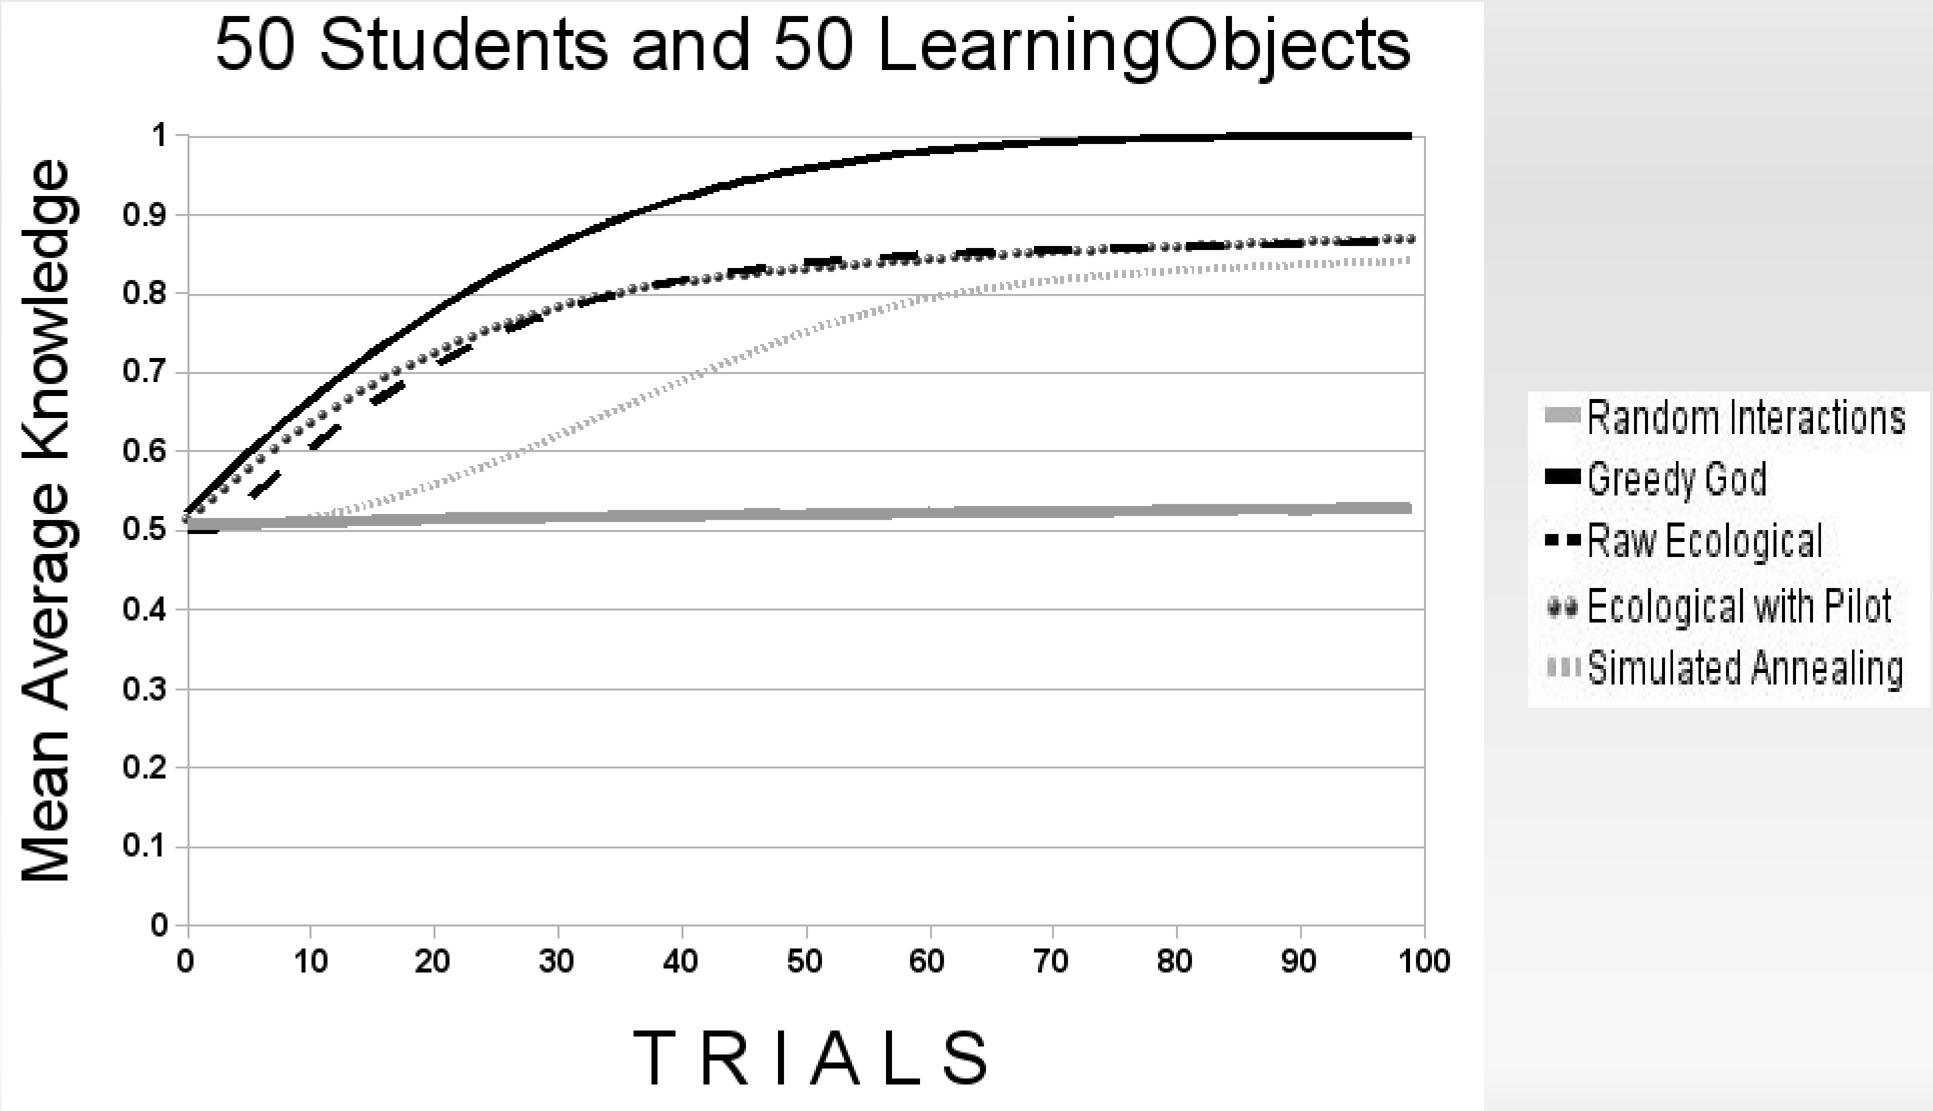

# Curriculum Sequencing

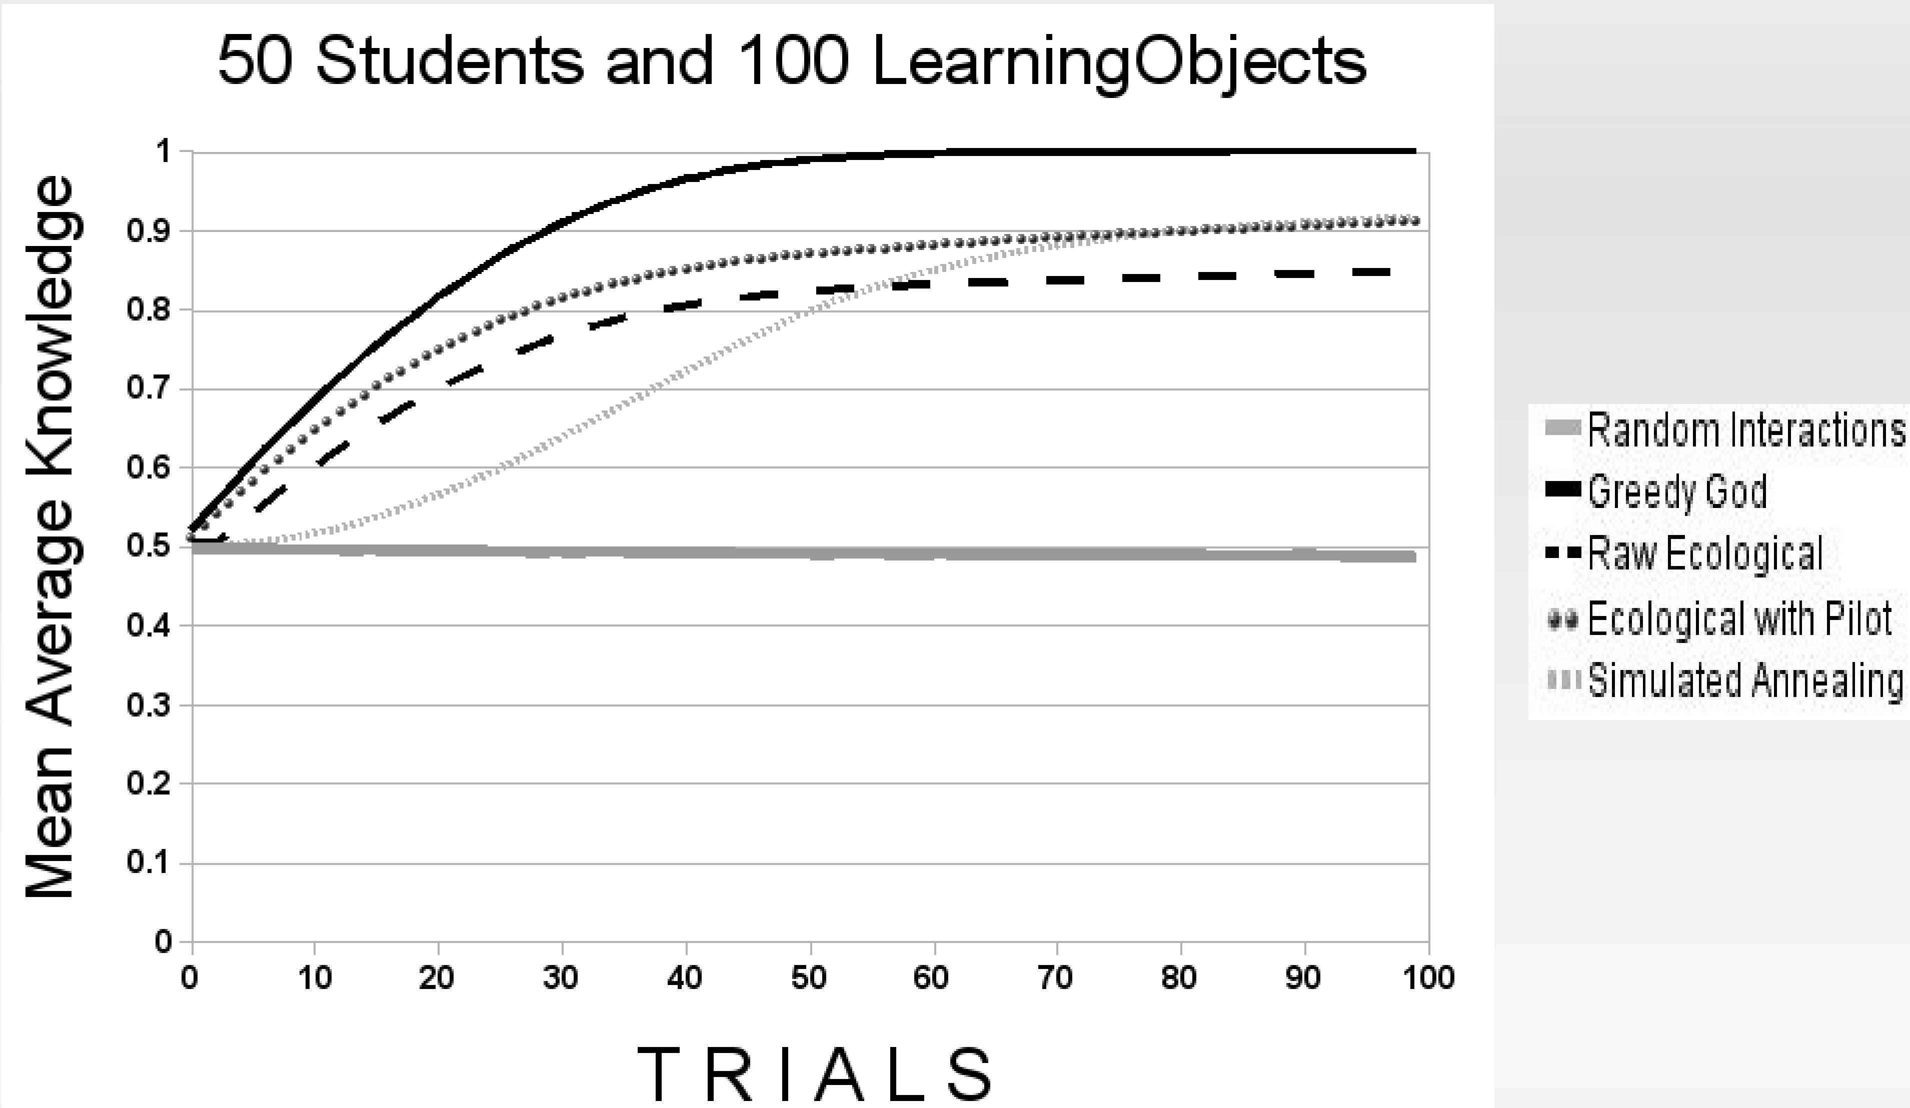

# Simulation (of learning)

$$\Delta UK[j,k] = \frac{I[l,k]}{1 + (UK[j,k] - LOK[l,k])^2}$$

- ΔUK[j,k] : The change in a user, *j*'s, knowledge *k*
- I[l,k] : Impact of the lesson (ranging from – to +)
- UK[j,k] : User *j*'s current knowledge of *k*
- LOK[l,k] : Target level of instruction of learning object *l* for knowledge *k*

# Simulated Learning Example

$$\Delta UK[j,k] = \frac{I[l,k]}{1 + (UK[j,k] - LOK[l,k])^2}$$

- Given a student with a knowledge of 0.65 who is using a learning object with a target level of instruction of 0.74 and an impact of 0.03
- After the interaction, the students knowledge will change by:
- $\frac{0.03}{1+(0.65-0.75)^2}=0.0297$
- Student's knowledge after interaction will then be 0.6797

# Simulation

- During the simulated “course of instruction” simulated students matched with variety of LO
- Learning occurs, as detail on previous slide
- The results of these interactions are used to reason about which students to match with which learning objects in the future

# Simulation vs. Real Students

- Results acknowledged to be relevant only within the context of simulated students
- Extrapolations to human students need to be made VERY cautiously
- Low cost way to run very large experiment
- Results can guide later explorations
- Verification with human students in multiple domains at a later date worthwhile

# Related Work – Simulated Students

- Confirms value of simulation
- VanLehn
  - Applications of simulated students: An exploration (1996) Journal of Artificial Intelligence in Education
  - Simulated student that would test ITS and collaborative with human students
  - Lower the cost of development
- Matsuda
  - Based on ITS student usage logs
  - Predict student actions
  - Learning curves

# Annotations

- Allow students to leave messages for subsequent students which could be questions, insights, or reflections on the learning object
- Intelligently show annotations to students who would benefit from them
- Distinct from tagging, folksonomies or ontologies
  - Attempting to allow peers to clarify information in learning objects for one another

# Annotations

- Model covers:
  - Tracking public, ongoing reputation of a specific author based on annotations left in the past
  - Tracking the reputation of a particular annotation, initially based on its author, then tracked independently
  - Similarities between students, based on ratings of annotations
  - Whether an annotation is shown or not, based on past ratings, the relationship between the current student and the rating students and the context of the annotation

# Annotations Approach

- Reputation of a student initially set as 0.5
- Afterwards based on the extent to which the student's previous annotations have been found useful by other students
- The authoring student's reputation is used as the initial reputation of an annotation, which is afterwards independently tracked and modeled with its own reputation

# Annotations Approach (cont'd)

- Students' ratings of the value of the annotation serve to adjust the overall reputation of the annotation (and indirectly, the annotation author's reputation)
- Once a student has provided a set of ratings on annotations, it is then possible to reason about their similarity to other students, based on mutually rated annotations
- Allows the probability that an annotation be shown to be determined by both the overall quality of that annotation, as rated by all students, adjusted by the similarity between the current student and the students who have previously rated the annotation

# Annotations Example 1

- Bob has left his first annotation (a question asking for the definition of a term on the slide) on a CS 115 slide and no one has rated it yet:  his reputation is 0.5
- Carol has left three annotations, two on slides and one on a video of the CS 115 instructor explaining Scheme structures
- After receiving ratings from students presented with the learning object these annotations have reputations of 0.476, 0.551 and 0.704
- Carol has a reputation of 0.577 (the average of these)

# Annotations Example 2

- Allison and Carol each leave an annotation on a new set of example exercises that have been posted
  - Allison asks if the solutions have been posted, and Carol replies that they will be posted on the March 18th after the assignment due date
- Allison's new annotation starts with a reputation of 0.5 (no one has rated her first annotation yet – Carol didn't bother) and Carol's has a reputation of 0.557 (there haven't been any additional ratings to her annotations since her reputation was calculated)
- If Carol’s comment is shown, Allison’s will be as well to maintain the thread of the conversation

# Annotations Example 3

- Bob has given a “Thumbs up” (indicating approval) to annotations *A, B, F & G* and a “Thumbs down” to annotations *D, M & Y*. Carol has given a “Thumbs up” to annotations *F, M, N*, and *Z* and a “Thumbs down” to *C, G*, and *Y*
- Bob and Carol's similarity is:

$$\frac{2(\textit{they agreed on their assessment of } F, Y) - 1(\textit{they disagreed on their assessment of } G)}{2 + 1(\textit{the number of annotations both have rated})}$$

$$= 1/3$$

# Annotations Example 4

- Bob is viewing the learning object annotation Z is attached to
- It has a reputation of 0.670 and Carol is the only other student he has a similarity rating with who has previously rated it
- She gave it a thumbs up, which leads to the annotation being given a higher reputation (say 0.687) because of her relationship to Bob (her vote for or against is strengthened)
- This means there is a 68.7% chance that this annotation will be shown to him when he uses the learning object

# Annotations Example 5

- While Bob uses the learning object, he gives the annotation Z a “Thumbs Down” (he doesn't find it useful). This lowers the annotation's reputation from 0.670 to 0.638.
- His similarity to Carol is also lowered (now that they disagree on 2 annotations) and they will now have a similarity of 0 (their preferences will no longer influence one another).
- If either gives another rating that conflicts with the other, they will begin to negatively recommend annotations for one another (what one likes will be less likely to be shown to the other).

# **iHelp** (Bretzke, Vassileva - 2003)

- Extension to COMTELLA
  - Recommender, incentives, tagging and ontologies, behaviours
- Creates “chat rooms” attached to learning objects, students automatically added
- “White board” style collaboration
- Temporal difference
  - Their work is synchronous, ours is asynchronous

# **Vicarious Learner** (Lee, Dineen, and McKendree - 1998)

- Automatic identification of worthwhile dialogs
- Critical thinking ratio (positive vs. negative)
- Populating a database, not tied to learning objects
- Similarities between students rather than ratio
  - We consider student modeling and a personalized recommendations, while their approach provides an absolute ranking
- Future work:  Integrate their critical thinking ratio into our algorithm

# Next Steps

- Looking to specifically incorporate trust and reputation work into our models (such as public / private reputation)
- Simulation for Validation
- Healthcare:  Replication with human students

# Incentive to Participate

- Possible criticism that students won’t leave annotations to help other students
- 4 possible approaches:
  - Social capital perspective, intrinsic reward
  - COMTELLA style incentive mechanisms
  - Leaderboard
  - World of Warcraft style achievements

# Work in AI Fields of Interest

- Seth’s et al. work on trust and credibility
- Recommender Systems (Breese, 1998; Herlocker, 2004; Resnick, 1997)
- Zhang et al. work on trust and reputation

# Corpus

- Idea is to provide tools to allow students to do some authoring, based on learning objects
- E.g. Given a book as a learning object, a student could highlight chapters which were more useful
- System reasons about the segmented learning objects and contrasts with the original, whole object

# Peer-Based Tutoring (and how we're different)

- COMTELLA
  - Recommending research papers
  - Tagging, folksonomies, ontologies
  - Incentives for co-operative behaviours in a learning community
- COPPER:
  - Intelligent Computer Assisted Language Learning
  - Reasoning about which active students should be matched with other students using Bayesian networks, multidimensional stereotypes and group modeling

# Student Modeling

- Closely related to user modeling
- Evolving knowledge:  constantly changing
- Scrutability of model has pedagogical impact
- Stereotypes and Communities

# Work in ITS of Interest

- Graesser's (2007) emotion work
  - Impact of boredom, engagement, frustration, confusion, delight, surprise and neutral emotions on learning
- VanLehn (1996) work with simulated students to provide collaborative learning partners and to test ITS
- Matsuda's (2007) Learning Curves

# Work in AI Fields of Interest

- Model Evaluation (Jastrzembski, 2009)
- Preference Elicitation (Qin, 2008; Berry 2007)
- Tailored Text Summaries (Paris 2009)

# Next Steps

- Currently exploring a richer learning model for the annotations work
- Some of the algorithms being consider for reasoning about annotations and corpus are preliminary – intend to get feedback from various conferences about approach
- Looking to specifically incorporate trust and reputation work into our models (such as a public / private reputation)

# Next Steps (cont'd)

- Time constraint:  incorporate length of time interacting with learning object
- Consider how content sequencing problem scales, and the effect of various alternatives to introducing students and learning objects on different schedules
- Introduce error to the assessment, and evaluate how robust our approach is to faulty data

# Next Steps (cont'd)

- Investigate more sophisticated similarity measurements with annotations work
- Investigate alternative learning simulations with annotations (e.g. non-sequential concepts)
- Investigate “bundling” of learning objects, where instead of breaking a learning object apart, students can recommend they be joined
- Simulation for Validation

# Conclusion

- Investigation of peer-based approaches
- Exploration of ecological approach and new concrete techniques
- Validation through simulation (value)
- Self-adjusting system that adapts based on the ongoing experiences of peers
- Projecting work into healthcare domain
- Exploration of AI subfields, such as trust modeling and credibility

# Questions?

- Comments?
- Concerns?