

A Multiagent, Ecological Approach to Content Sequencing

John Champaign

David R. Cheriton, School of Computer Science
University of Waterloo,
Waterloo, ON, Canada
jchampai@uwaterloo.ca

Robin Cohen

David R. Cheriton, School of Computer Science
University of Waterloo,
Waterloo, ON, Canada
rcohen@uwaterloo.ca

ABSTRACT

In this paper, we present a framework designed to enable peer-based intelligent tutoring, inspired by McCalla's ecological approach[4]. In particular, we focus on the question of how to sequence the content that is presented to a student and offer an algorithm that selects appropriate learning objects (webpages, videos, research pages etc.) from a repository, for each student based on the previous experiences of similar students. We argue that in order to design an effective education environment for students, it is possible to model each student as an agent, in order to compare that student with previous students who have encountered the system (so, other agents). This perspective is emphasized in our work in particular because we choose to validate our current model using a simulation of all existing students and their learning experiences. As we have been able to gain some benefit through our simulations, we propose that multiagent systems researchers explore the approach of peer-based learning as part of their educational systems.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Algorithms, Performance, Experimentation, Human Factors

Keywords

Intelligent tutoring systems, peer-based tutoring, content sequencing, simulating students, ecological approach

1. CONTENT SEQUENCING

Two central challenges in the design of intelligent tutoring systems are compiling the material for the lessons and determining the best methods to use, for the actual teaching of those lessons. We observe in particular that it is desirable to provide a framework for determining the material to be taught that does not rely on experts hand-coding all the lessons. Indeed, that particular approach presents considerable challenges in time and effort.

Cite as: A Multiagent, Ecological Approach to Content Sequencing, John Champaign and Robin Cohen, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

We are interested in techniques for bootstrapping the system in order to initiate peer-based learning and in developing robust methods for validating the models that are presented (including the technique of employing simulated students). Once the content is in place, our efforts will be aimed at refining our model in order to enable students to benefit the most from the learning that their peers are undergoing.

We have currently developed an algorithm for reasoning about the sequencing of content for students in a peer-based intelligent tutoring system inspired by McCalla's ecological approach[4]. We record with each learning object those students who experienced the object, together with their initial and final states of knowledge, and then use these interactions to reason about the most effective lessons to show future students based on their similarity to previous students. As a result we are proposing a novel approach for peer-to-peer intelligent tutoring from repositories of learning objects.

We used simulated students to validate our content sequencing approach. Our motivation for performing this simulation was to validate that, in the experimental context, our approach leads to a higher average learning by the group of students than competing approaches. We added a modeling of the knowledge that each object is aimed at addressing (for example, an object in a first year computer science course may be aimed at addressing the knowledge of recursion). By abstracting all details from the intelligent tutoring system and the student, we defined a formula to simulate learning (Equation 1).

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

where UK is the user j 's understanding of knowledge k , I is the educational benefit (how much it increases or decreases a student's knowledge) of learning object l on knowledge k and LOK is the learning object l 's target level of instruction for knowledge k .

When running our algorithm in the simulation, each student would be presented with the learning object that was expected to bring the greatest increase in learning, determined by extracting those learning objects that had resulted in the greatest benefit for previous students considered to be at a similar level of understanding as the current students.¹

Simulated students allowed us to avoid the expense of implementing and experimenting with an ITS and human students to see the impact of our approach in contrast with al-

¹Each student in the simulation is modeled to have a current level of understanding for each possible knowledge area, a value from $[0,1]$ reflecting an overall grade from 0 to 100.

ternative approaches. In particular we contrasted our method with a baseline of randomly assigning students to learning objects and to a “look ahead” greedy approach where the learning was precalculated and used to make the best possible match. One variant we considered was a “simulated annealing” inspired approach, where greater randomness was used during the initial, exploratory phase of the algorithm, then less randomness was used once more information about learning objects had been obtained. We discovered that our approach showed a clear improvement over competing approaches and approached the ideal.

2. TESTING OUR APPROACH

For our experiment, variable numbers of simulated students and learning objects were allowed to interact over a set number of trials, arbitrarily chosen as 100. Each simulated student was randomly assigned values for each of 6 knowledges, each evenly distributed in the range $[0,1]$. A multi-dimensional structure for knowledge was used to ensure randomly generated students were distinct from one another, and to provide a rich model for simulated learning. In a real world context, this can be thought of as students who have a better understanding of one part of a course of study compared to another part of the same course. Each learning object was randomly assigned a value for the target level of instruction for each knowledge, evenly distributed in the range $[0,1]$. Impact values were assigned to learning object knowledges, randomly and evenly distributed in the range $[-0.05, 0.05]$. The values of -0.05 and 0.05 were chosen such that no single learning object could radically change a user’s knowledge level (at most it can adjust it by 5%).

Each experimental condition was repeated for 20 iterations, and the mean of the average knowledges of all students after each trial was graphed (see Results). In this context, the average knowledge might be thought of as, if a final mark for the CS 101 course was to be assigned to a student based on their current understanding of the course content, what would it be? A student’s final mark will be based on their knowledge of a number of areas, such as introductory data structures, recursion, sorting, introductory proofs, and programming in a specific language.² In order to explore the value of our approach, we graph the performance of several algorithms to select a learning object for each student, over a number of trials. After each trial, the average knowledge under each condition is compared.

Random Association: Two reference points were created to compare our approaches against. One reference was created by associating each student with a randomly assigned learning object each trial. Given that any intelligent approach to matching students with objects should outperform blind chance, this was viewed as the lower limit.

Greedy God: The other reference point, the greedy god reference point, was created by giving the algorithm full access to the fine-grained knowledge levels of the students and learning object, testing what the outcome would be for every possible interaction, then choosing the best interaction for each student for each trial. The results, based on an omniscience not typically available in real world learning environments, was viewed as a ceiling on the possible learning

²We make the simplifying assumption that all knowledges contribute equally to this progress assessment and calculate it as the average knowledge.

benefit of any approach.

The impact values and target levels of objects are used for the reasoning of the greedy god algorithm. In contrast, for the following three algorithms, our ecological approach is used to select the learning objects to be presented to users (so based on their similarity to previous students who have experienced these objects and on the benefit that these students derived). We assume that as each simulated student is assigned a learning object, that student’s interaction with the object can be used as the “previous experience” to which subsequent students are matched. Using our ecological approach, learning objects presented to students should end up being ones that have an effective combination of impact value and target level (i.e. beneficial to those previous users and at a somewhat similar level of knowledge).

Raw Ecological: For the raw ecological approach, 3 trials were run where each student was randomly assigned to a learning object. For the remaining 97 trials, each student was matched with the learning object best predicted to benefit her knowledge³. The initial three trials with random associations were used to provide rough information about the learning objects and students, which was used and refined over the course of the remaining trials.

Pilot Group: For the pilot group ecological approach, a subset of the students (10%) were assigned as a pilot group, and for their 100 trials were systematically assigned to learning objects to explore their impact. These interactions, along with the accumulation of their own interactions, were used by the remaining 90% of the students to reason about the best sequence.

Simulated Annealing: Our third ecological approach was inspired by simulated annealing, which in turn was inspired by the metallurgical approach of heating and cooling to induce change in a material. For this approach, we had a “cooling” period, which was the first 1/2 of the trials. During this period, for every student, there was an inverse chance, based on the progress of the trials, that they would be randomly associated with a lesson; otherwise, the ecological approach would be applied. For example, in the first trial, every student would be randomly associated with a learning object, but by the 25th trial, each lesson would have a 50% chance of being randomly associated. After the cooling period was over (the 50th trial), every student was repeatedly assigned to a learning object by ecological reasoning.

3. RESULTS

As seen in Figure 1, the random associations of students with learning objects is clearly and consistently shown to be an inferior approach to improving the average knowledge of a group of students, as expected. Similarly, an omniscient sequencer using perfect knowledge of students, learning objects and the outcome of a potential interaction (greedy god) can consistently produce the greatest learning benefit.

Contrasting our ecological techniques (which would each be feasible in a real educational setting) with these reference points, provides illumination on the usefulness of the ecological approach in this setting. Reasoning intelligently, in this

³Since each student’s knowledge is now multi-dimensional the difference calculated to determine benefit is now a sum of the differences of each knowledge dimension. The numeric values for knowledge are converted to one of the concrete letter grade levels before performing the computation.




Figure 1: Comparison of 5 Approaches to Sequencing Learning Objects

manner, has produced greater knowledge in a shorter number of trials for the group of students as a whole compared to a random association.

As asserted by [4], we see his predicted impact that with more learners the ecological approach’s performance improves. A correlation of improved performance with an increase in number of learning objects was also seen. This makes intuitive sense: if an intelligent tutoring system (ITS) is given a larger repository of learning object to assign, we would expect it to be able to find objects better suited to a particular student.

While Figure 1 seems to show superior performance of the ecological approach with a pilot group, it is important to remember that 10% of the class was used as a pilot group for this experimental condition. These were not included in the average assessed knowledge. Their increased knowledge, which would be roughly equivalent to the lack of increase shown by random associations, is omitted and the improved performance of the remaining students can be viewed as at the expense of the pilot group.

The “Simulated Annealing” technique was interesting as it underperformed the other two techniques during its “cooling period” but quickly gained ground after the cooling period was complete. This is due to the randomness added to the during the cooling period leading to a greater exploration of the possible interactions between learning objects and students. This improved understanding of the two groups could then be used when reasoning about which students to match with which learning objects in later trials. In the largest condition (50 students and 100 learning objects), simulated annealing matched the performance of the pilot group condition, without sacrificing 10% of the class. With the correct choice of cooling periods, this technique shows promise for delivering comparable long term performance at the expense of early progress for the entire group instead of no progress for a pilot group.

4. RELATED WORK

In our work, we propose to select learning objects to present to students, based on the previous experiences of other, similar students. As such, the tutoring of each new student 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 at the University of Saskatchewan investigated recommendation of academic papers and similar resources. In the early phases of the work [7] there was difficulty in getting

users to accurately provide metadata when entering papers in the system. Subsequent work [1] has focused on providing incentives to encourage users to interact positively with the system. In contrast, our work avoids soliciting explicit feedback, and reasons based on typical usage.

Another example of peer-based learning is that of [2], which uses collaborative filtering to match students based on their “lifelong learning”. Matching to similar users occurs based on life events, such as a specific degree at a certain university or working at a specific company. Their results are presented transparently in a user centric approach where users can investigate the “trails” of similar users. Matching was done by using string metrics where life events are encoded into a token based string which is used to reason about similarities between users. 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 contrast, our approach uses typical ITS interactions and doesn’t elicit anything specific from the user. In their system user histories must be continually updated, with the ongoing issue of out-of-date user profiles. The data used by our system can be easily gathered in real-time by usage of the system and will be as up-to-date as their last usage of a learning object.

A different style of peer-based learning, called COPPER [5], approaches the problem of ICALL. It intelligently matches students, and assigning them specific roles for their interaction, through a Bayesian approach, and allows them to help one another learn. Their approach could easily be integrated with ours, where an interaction between students is a learning object. While their approach is useful in real-time, it doesn’t allow students to independently learn from the experience of previous student interactions with the system. Our system, in contrast, reasons using the entire experiences of all previous students, not just the current, on-line students.

In our work, we have introduced experiments that simulate student learning, in order to validate our proposed model for content sequencing in intelligent tutoring systems. Other intelligent tutoring systems researchers have explored the value of simulating students.

[6] discusses the authors’ experiences with simulated students and the methods that can be used to assist in education. The authors claim that this is useful not only for providing a collaborative learning partner for a student but also for instructional developers to test systems that they develop, including early development where trials with human

students may not be feasible. The authors highlight grain-size as an important spectrum for considering simulated students. For example, an example of fine-grained knowledge in physics is knowing the existence of tension in a string, when a string is tied to a body; an example of large-grained knowledge in physics would be simply knowing the law of conservation of energy.

Our system uses a granularity outside of this range, which we would term coarse-grained. As an example, a student might be modeled as having a 0.67, which could mean, for instance, that the student has enough knowledge to receive a 67% mark in Physics 101 or that they understand enough knowledge to complete 67% of the projects.

[6] also specifically track and formally represent, for each student, their knowledge before learning, the behaviour during the learning, the instruction and the student's knowledge after learning. In contrast, we are interested in tracking behaviour with respect to learning objects, and focus on modeling the student's knowledge before and after interactions with those learning objects.

Another research group used what they call learning curve analysis to analyze how their simulated student performed [3]. They measured the accuracy of production rules, in terms of successfully matching a step in solving the problem, compared to number of training problems or frequency of learning opportunities. We follow a similar approach in the evaluation of our work, where we use the resulting learning curves to contrast educational environments.

5. REFERENCES

- [1] R. Cheng and J. Vassileva. Design and evaluation of an adaptive incentive mechanism for sustained educational online communities. *User Model. User-Adapt. Interact.*, 16(3-4):321–348, 2006.
- [2] N. Labeke, A. Poulouvassilis, and G. Magoulas. Using similarity metrics for matching lifelong learners. In *ITS '08: Proceedings of the 9th international conference on Intelligent Tutoring Systems*, pages 142–151, Berlin, Heidelberg, 2008. Springer-Verlag.
- [3] N. Matsuda, W. W. Cohen, J. Sewall, G. Lacerda, and K. R. Koedinger. Evaluating a simulated student using real students data for training and testing. In C. Conati, K. F. McCoy, and G. Paliouras, editors, *User Modeling*, volume 4511 of *Lecture Notes in Computer Science*, pages 107–116. Springer, 2007.
- [4] G. McCalla. 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*, 7:1–23, 2004.
- [5] T. Read, B. Barros, E. Bárcena, and J. Pancorbo. Coalescing individual and collaborative learning to model user linguistic competences. *User Modeling and User-Adapted Interaction*, 16(3-4):349–376, 2006.
- [6] K. VanLehn, S. Ohlsson, and R. Nason. Applications of simulated students: An exploration. *Journal of Artificial Intelligence in Education*, 5:135–175, 1996.
- [7] J. Vassileva. Toward social learning environments. *IEEE Trans. Learn. Technol.*, 1(4):199–214, 2008.