Fostering Student Learning and Motivation: an Interactive Educational Tool for AI

Saleema Amershi, Nicole Arksey, Giuseppe Carenini, Cristina Conati, Alan Mackworth, Heather Maclaren, David Poole
Dept. of Computer Science, University of British Columbia
Vancouver, B.C., Canada
604-822-3061
{samershi, arksey, carenini, conati, mack, maclaren, poole} @cs.ubc.ca

ABSTRACT
There are inherent challenges in teaching and learning Artificial Intelligence (AI) due to the complex dynamics of the many fundamental AI concepts and algorithms. Interactive visualization tools have the potential to overcome these challenges. However, there are reservations towards adopting interactive visualizations due to mixed results on their pedagogical effectiveness. Previous work has also often failed to directly assess student preferences and motivation. CIspace is a set of nine interactive visualization tools demonstrating fundamental principles in AI. The CIspace tools are currently in use in undergraduate and graduate classrooms at the University of British Columbia and around the world. In this paper, we present two experiments aimed at assessing the effectiveness of one the tools in terms of knowledge gain and user preference. Our results provide evidence that the tool is as effective as a traditionally accepted form of learning in terms of knowledge gain, and that students significantly prefer to use the tools over traditional forms of study. These results strengthen the case for the incorporation of CIspace, and other interactive visualizations, into courses.

General Terms
Algorithms, Experimentation, Human Factors

Keywords
Algorithm visualization, pedagogy, user studies, artificial intelligence

1. INTRODUCTION
Artificial Intelligence (AI) is an important discipline within Computer Science education. However, there remains an inherent difficulty in teaching introductory AI concepts, particularly because they often involve complex dynamic algorithms [5, 7]. For this reason, AI courses stand to benefit from the incorporation of tools for interactively visualizing and animating these concepts, difficult to describe with static media alone.

CIspace [2] is a collection of Java applets and tutorials that demonstrate various concepts in AI. They can be used as tools for interactively studying material from class, for in-class demonstrations or for creating and working on homework assignments [1]. The applets cover topics ranging from various search strategies, to learning algorithms, to robot control.

CIspace is motivated by an intuition shared by many educators that visualization and animation can help students learn by demonstrating dynamic, graphical processes [5,14], and also by the fact that animations can capture student interest thus motivating them to study [4,10,11]. Yet, previous work on assessing the effectiveness of visualizations has produced mixed results [8]. These have generally come from the domain of data structures or algorithms such as sorting or searching, e.g. [3,8,10], while no studies have been done on the effectiveness of visualizations in the domain of AI. Furthermore, these existing studies have tended to focus on measuring "effectiveness in terms of knowledge acquisition" rather than by preference or motivational gains. However, the position of many advocates of interactive visualizations is that their educational effectiveness may arise not only from increased understanding of a domain, but also from increased motivation and interest in study thus encouraging learning.

The CIspace tools have been incorporated into the curriculum of undergraduate and graduate level AI courses at the University of British Columbia (UBC). We have received positive feedback from both instructors and students, which is always encouraging; however, empirical testing is necessary to provide robust evidence of effectiveness, especially as previous work has been inconclusive. For this reason, we have conducted two experiments aimed at assessing CIspace's effectiveness as a study tool. The experiments compare student study using a specific applet with a traditionally accepted form of study using written sample problems. The goals of these experiments are *i)* to ensure that the applet is at least as effective as traditional methods in terms of knowledge gain, and *ii)* to assess the applet's effectiveness in terms of user preference and motivation.

The rest of this paper is organized as follows: Section 2 presents potential benefits of the CIspace tools for teaching and learning AI. Here, previous work on algorithm visualizations is also presented in terms of teaching and learning. Section 3 introduces the Constraint Satisfaction Problem (CSP) applet. In Section 4 we discuss our experiments, including goals, procedures and results. Section 5 concludes with a summary of our main results and possible future areas of research.

2. TOOLS FOR TEACHING AND LEARNING
Any effective educational visualization tool must meet the needs of both educators and students. This involves minimizing common obstacles in adopting the tool including the time to find, integrate into a course, learn, and develop problems using the visualizations [12,13], while maximizing learning gain and

motivation. The designers of CIspace have attempted to address these needs in the following manner[1].

CIspace is composed of a collection of applets that provide coverage of fundamental topics often included in undergraduate AI course curricula. Each applet is modular and focuses on a specific concept allowing instructors the choice to use an applet or not for individual topics covered in a course. The shared look and feel of the applets is designed to reduce the time and effort required for both instructors and students to learn a new applet for each topic. CIspace also provide guidance in the form of *1)* instructions that appear on screen as the applets are used and *2)* short instructional videos demonstrating their features and use.

Each CIspace applet provides a set of pre-designed sample problems demonstrating several cases of each applet's algorithms. Providing students with sample problems can significantly increase the pedagogical value of a visualization tool [17]. The sample problems can be used as is, or modified. New problems can also be created graphically and saved for later use. This enables instructors to design examples efficiently for use in a class or for students to use in lab sessions or assignments.

Efficiency is valuable for users of any algorithm visualization tool; however, the primary objective remains its educational effectiveness. Since experiments measuring knowledge gain continue to provide mixed results, there is a need to gain further insight into this area and into the area of user preferences to lend support for visualization tools. Few previous formal studies address motivation and preference, and only do so through indirect measurements or through general observations [6,10]. While these studies are interesting, in the work presented in this paper we provide direct evidence of knowledge gain and of preference and motivation.

3. THE CSP APPLET

CSPs are generally encountered in undergraduate AI courses and are pervasive in AI. The CSP applet was an appealing choice to use in our experiments because CSPs are simple enough to be learned and tested in a short amount of time.

The problem of constraint satisfaction is, given a set of variables each with a domain (a set of values it can take on), and a set of constraints on legal assignments, find an assignment of a value to each variable that satisfies all constraints. The nature of a CSP lends to its intuitive graphical representation as a network of variable nodes and constraint arcs, making it a good candidate for visualization in an educational tool. The CSP applet (see Figure 1) models this graphical network after the static representation in the textbook Computational Intelligence [15].

The applet has two major modes, "Create" and "Solve". In "Create" mode, users can build a CSP from scratch or load a sample problem. In "Solve" mode, users can apply an algorithm for solving a CSP called the AC-3 algorithm. Both of our experiments focused the applet's ability to effectively demonstrate the AC-3 algorithm. For this reason we will only describe the relevant features of "Solve" mode.[2]

[1] For a more detailed discussion of the CIspace design see [1].

[2] For more details on the CSP applet's interface and features refer to the tutorials in [2].

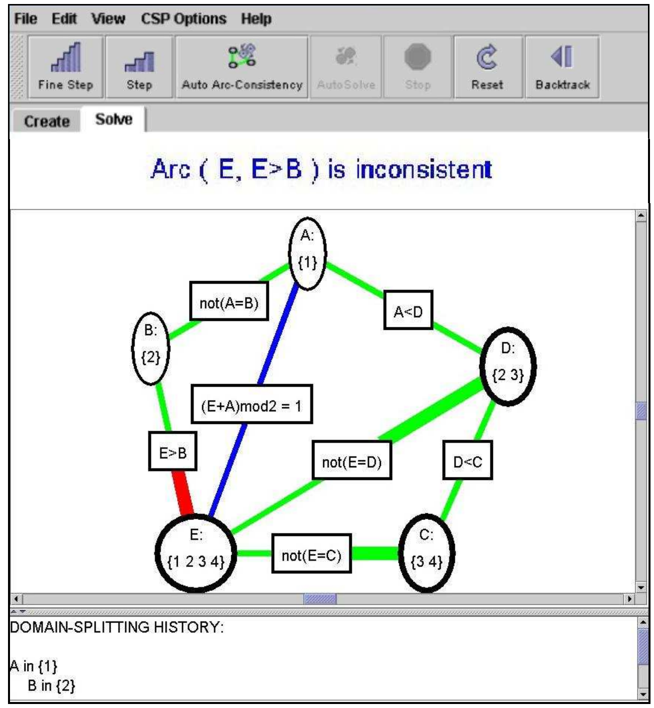

Figure 1. CSP applet with example CSP

Figure 1 shows the CSP applet's interface in "Solve" mode. The main area of the interface is used to display the graphical network. Above this, a small panel displays instructional messages about how to use the applet and about the current state of the CSP. Near the top of the applet window is a toolbar containing buttons for applying the AC-3 algorithm to the CSP. At the top, a menu bar contains menu items for basic file operations, editing and viewing of the CSP. The applet illustrates AC-3 on the network using color and animation. The AC-3 algorithm makes the entire network arc consistent by considering a set of potentially inconsistent arcs that are colored blue in the applet. Until the set is empty, an arc is removed and tested for consistency. If it is found inconsistent its color will change to red. When the arc has been made consistent, it appears in green.

Users can solve a CSP at various levels of abstraction. The "Fine Step" button allows students to apply and watch the AC-3 algorithm in detail. Fine stepping cycles through three stages: selecting an arc to test for consistency, testing for consistency, and then making it consistent. Details at each stage are displayed in the panel above the network emphasizing to the user what is happening (see Figure 1 above the network). Clicking directly on an arc in the network can also make it consistent by performing all three stages of fine stepping on the arc at once. This feature gives users control over the algorithm by allowing them to choose which arcs to make consistent rather than having the applet select arcs for them as with the fine step feature.

The fine step and direct click features enable users to move forward through the algorithm at their own pace, noted as the single most important feature of algorithm visualizations in [17].

The applet also has an Auto Arc-Consistency feature that fine steps through the CSP network to completion. The user can specify the pause duration between successive fine steps, known as the Auto Arc-Consistency speed. A faster speed is useful in giving the user an overall picture of the AC-3 algorithm; a slower speed enables users to better observe details of the algorithm.

4. EXPERIMENTS

4.1 Experiment One

One function of the CIspace tools is to help students learn AI concepts by example. Since studying by example is a conventional method of learning, our goal in this experiment is to determine effectiveness of the CSP applet in terms of knowledge gain when compared to traditional sample problems on paper.

The experiment typified a study scenario [8] in which students learned underlying theory and concepts from a textbook, and then studied related examples. The experiment was a between-subject study with the form of sample problems as the independent variable. The two conditions for the independent variable were sample problems studied using the applet and written sample problems studied on paper, referred to as the applet and non-applet group respectively. The written sample problems were modeled after the way CSP examples were illustrated in AI courses at UBC by two experienced professors prior to the introduction of the CIspace applets. These were typically demonstrated with the use of static diagrams accompanied by written steps of the AC-3 algorithm.

4.1.1 Procedure

19 students, 8 female and 11 male, were recruited for this experiment. Participants were all undergraduate students at UBC who had never taken an AI course, but had the prerequisites needed to enroll in UBC's introductory AI course including a course on basic algorithms and data structures. The experiment took three hours and participants were paid \$10/hour for their time. The time allocated for each phase of the experiment was determined through pilot studies.

All of the students were given photocopied text about CSPs from the textbook <u>Computational Intelligence</u> [15]. They were provided with two sheets of blank paper on which they could write notes if they wished. In order to guide their study, students were also given a list of topics to try to learn. They were given one hour for reading and studying the text.

The students were then given a 20 minute, closed book, pre-test containing 4 questions amounting to 19 possible marks – 10 marks of procedural type questions and 9 marks of conceptual type questions. The questions covered applications of the AC-3 algorithm to a CSP and the given study topics.

The students were then randomly divided into the two treatment groups, accounting for balanced distribution of males and females. The applet group had 10 people, 6 males and 4 females, and the non-applet group had 9 people, 5 males and 4 females.

All of the students were given 40 minutes to study three sample problems, each illustrating different aspects of the AC-3 algorithm applied to a CSP. The three sample problems were the same for both the applet and non-applet groups. The students were allowed to study the sample problems in any order and could go back and forth between them. The students were also given back their text material and notes from the earlier learning phase, which they could refer to while studying. During the applet group's study time, each student watched a three minute video describing how to use the applet and the applet's interface, but not providing extra information about CSPs. The students were told that they could watch the video as many times as they liked. After the 40 minutes, each group was given a post-test that was almost identical to the pre-test, with only a few values or arcs manipulated from the pre-test.

After the test, students were given a questionnaire specific to their treatment group in which they were asked about *i)* how confident they were in their knowledge of the topics given to them at the start of the study, *ii)* their opinions about the study materials they used, *iii)* timing during the study, and *iv)* interface specific questions (applet group only). Likert scales, ranges, and open ended questions were used.

4.1.2 Discussion of Results

The pre-test and post-test scores of the applet and non-applet groups show that both groups improved significantly[3] ($p<.015$ and $p<.005$ respectively) after having studied the sample problems, but that there was no statistically significant difference in improvement between the two groups. For the conceptual questions, the non-applet group improved 3% more than the applet group but the difference was not significant. For the procedural questions, both groups improved by 33%. The average confidence levels reported by both groups on the list of topics covered were roughly equivalent for each topic with no significant differences reported.

These results show that, in addition to being as effective a learning method as studying on paper, students learning from the applet were able to successfully transfer their gained knowledge to a traditional pencil and paper test. This is an important finding because it demonstrates that instructors can incorporate interactive visualizations in the studying portion of their courses and still test students in traditionally accepted ways.

Another interesting trend that was noticed was that females performed better when using the applet than when using the written sample problems, whereas males performed better when using the written sample problems than when using the applet. Females in the applet group improved 2.5 marks more than females in the non-applet group. And males in the non-applet group improved 2.48 marks more than males in the applet group. Neither of these trends is significant but they show a potential indication of differences in the use of visualizations by gender. Further investigation with more subjects would be required to assess if there is in fact a difference.

Table 1 shows the results of questions from the questionnaire about students' opinions on the study materials they used. Students answered the questions using a 5 point Likert scale: 5=Agree, 4=Somewhat Agree, 3=Neutral, 2=Somewhat Disagree, 1=Disagree. The only significant ($p<.04$) difference between groups was in response to the question in which students were asked about an alternate format of study from the one they used. The applet group *somewhat disagreed* that they believed having the sample problems written down on paper would have helped them learn better than with the applet, whereas the non-applet group was *neutral* when asked whether they believed watching the CSP graph change at every step would have helped them learn better than with the written problems. The non-applet group was not shown the applet.

[3] Unless otherwise stated, all tests for significance are one-tailed Student's t-tests.

Table 1. Student responses, 5 point Likert scale

Statement	Applet Group	Non-applet Group
Helped me learn the material from the book.	4.90	4.89
Difficult to follow steps of the algorithm.	2.00	2.44
Looking at examples worked out on paper would have helped me study better.	1.80	N/A
Seeing the network change at every step would have helped me study better	N/A	3.00

On average both the applet and non-applet groups reported having between *more than enough time* and *enough time* to study their sample problems. No significant difference was found. The applet group reported taking between *less than 5 minutes* to between *5-10 minutes* to learn the applet's interface. In fact, all of the students in the applet group reported that it took under 10 minutes to learn the interface and still allowed them enough time to study the sample problems within the allotted time period. This finding disproves a common reservation towards visualizations based on the idea that students may be discouraged from using visualization tools because of the apparent learning overhead that may accompany them.

From the test results of this experiment we find that the applet is just as effective, in terms of knowledge gain, as traditional methods given equal study time. However, as discussed earlier in this paper, visualizations may play another important role in that they may be better able to engage and motivate students to learn. This hypothesis leads us to our second experiment.

4.2 Experiment Two

The goal of our second experiment is to attempt to produce a clearer measurement of users' preferences between studying with the applet and studying with paper problems.

The materials used for this experiment were the same as for experiment one, except that the questionnaire was modified to produce a more in-depth gauge of user preferences and motivation. Also, a brief semi–structured interview was added at the end of the experiment to obtain richer data and possibly reveal contradictions from the questionnaire. The questionnaire was divided into two smaller questionnaires so as not to overwhelm the students. The first questionnaire focused on attitudes of students, including how they liked both forms of study and how motivated they felt using each of them. The second questionnaire included questions about applet interface issues, clarity of written sample problems and time. Likert scales, semantic differentials, ranges, and open ended questions were used in the questionnaires.

In contrast with experiment one, our second experiment was a within-subject experiment in which students were given one sample problem to study using the CSP applet and one written on paper. After using each treatment form, students were asked to explicitly choose either the applet or paper form to study an additional sample problem. We hoped the results of this choice could give us explicit quantitative preference data in addition to retrospective qualitative data obtained from the questionnaires and interview.

4.2.1 Procedure

19 students, 8 female and 11 male, were recruited for this experiment, meeting the same requirements as for experiment one. The experimental procedure was the same as in experiment one except for the phase when students studied sample problems. In this phase, students studied two sample problems for 12 minutes each, using the applet for one problem and the paper form for the other. 10 students started with the applet, while 9 started with the paper form. The 12 minutes included the time for the applet users to view the three minute video as in experiment one. For the last sample problem, each student chose a form to use on basis of preference. They were then given 16 minutes to study the third sample problem with the method of their choice. This problem was allotted more time than the previous problems because it illustrated the most complex case of the AC-3 algorithm applied to a CSP. After the experiment, in addition to answering the questionnaires, students were individually interviewed by the same experimenter, and the interview was recorded on tape.

4.2.2 Discussion of Results

13 out of 19 students chose the applet over the paper method for studying the third sample problem, although both the students that chose the applet and the students that chose the paper had significant improvements in scores from pre-test to post-test ($p<.003$ and $p<.015$ respectively). The students' opinions on the amount of time needed to study the sample problems agreed with the previous experiment. Most students felt that they had between *enough* and *more than enough* time to study each sample problem, and on average, they took between *less than 5 minutes* and *5-10 minutes* to learn how to use the applet.

Table 2. Average responses, 5 point Likert scale

	Choice/number of students		
Statement	Applet/13	Paper/6	Overall/19
I liked using the applet more than studying with the sample problems on paper.	4.67	2.80	4.11
I liked studying with the sample problems on paper more than with the applet.	2.25	4.00	2.76
Using the applet helped me more than the sample problems on paper.	4.58	2.80	4.06
The sample problems on paper helped me study better than the applet.	2.50	3.80	2.88

Table 3 shows our most interesting results the from questionnaire on student attitudes, which used the same 5 point Likert scale as described in experiment one. Overall, the students indicated that they liked using the applet more than the sample problems on paper and that they felt the applet was more helpful in studying than the sample problems on paper. Both of these statements received significantly more indications of agreement than the opposing statements ($p<.007$ and $p<.005$ respectively). Interestingly, the interviews revealed that many students who chose the paper still had positive opinions about the applet such as "the applet is for sure more motivating than the paper."

Students rated between *agree* and *somewhat agree* (4.13 on 5 point Likert scale) with the statement that they "felt more motivated to learn using the applet than studying with the sample problems on paper". This is accompanied by positive responses during the interview sessions such as "the applet makes it more interesting, it's interactive." Also, when asked whether they thought it was worth while having applets such as the CSP applet available to use in a course, students significantly ($p<.05$) answered 'Yes'.

The results of this experiment provide strong evidence that students liked and felt more motivated by the CIspace CSP applet than by traditional paper examples of CSPs. This suggests that the applet provides a positive alternative for students finding difficulty in studying with traditional methods. This also lends support to the claim that algorithm visualization tools can be beneficial for learning in that they help motivate and engage students [12].

5. CONCLUSION

CIspace has been incorporated within undergraduate and graduate AI courses at UBC, with generally positive responses from both instructors and students. Furthermore, our experiments provide evidence *i)* that the CSP applet is at least as good a study tool as traditionally accepted methods of study, and *ii)* that students prefer to study using the applet rather than with traditional methods.

Our experiments, however, only formally tested one of many learning aspects of the applet, namely the use of the CSP applet as a study tool. Some of the other CIspace applets feature quizzes and statistical analysis functions that may also prove useful for teaching and learning. Future work on these aspects and features may provide additional support for interactive educational visualization tools.

6. REFERENCES

[1] Amershi, S., Arksey, N., Carenini, G., Conati, C., Mackworth, A., Maclaren, H. and Poole, D. Designing CIspace: Pedagogy and Usability in a Learning Environment for AI. *To Appear in ITiCSE 2005.*

[2] CIspace: Tools for Learning Computational Intelligence. http://www.cs.ubc.ca/labs/lci/CIspace/.

[3] Colaso, V., Kamal, A., Saraiya, P., North, C., McCrickard, S. and Shaffer, C. Learning and Retention in Data Structures: A Comparison of Visualization, Text, and Combined Methods. In *Proc. ED-MEDIA 2002*, 2 pages.

[4] Demetriadis, S., Triatafillou, E., and Pombortsis, A. A Phenomenographic Study of Students' Attitudes Toward the Use of Multiple Media for Learning. *ITiCSE 2003,* ACM Press (2003), 183-187.

[5] Greiner, R., and Schaeffer, J. The Aixploratorium: A Vision for AI and the Web. *IJCAI 2001*, 3 pages.

[6] Hansen, S.R., Narayanan, N.H., and Hegarty, M. Designing Educationally Effective Algorithm Visualizations. *Journal of Visual Languages and Computing 13,* 3 (2002), 291-317.

[7] Hearst, M.A. Preface: Improving Instruction of Introductory Artificial Intelligence. *1994 AAAI Fall Symposium,* AAAI Technical Report (1994), 1-4.

[8] Hundhausen, C.D., Douglas, S.A., and Stasko J.T. A Meta-Study of Algorithm Visualization Effectiveness. *Journal of Visual Languages and Computing 13,* 3 (2002), 259-290.

[9] Kehoe, C. and Stasko, J.T. Using Animations to Learn about Algorithms: An Ethnographic Case Study. *Tech. Rep. GIT-GVU-96-20,* (1996), 12 pages.

[10] Kehoe, C., Stasko J., and Taylor, A. Rethinking the evaluation of algorithm animations as learning aids: an observational study. *Int. J. Human-Computer Studies 54,* 2 (2001), 265-284.

[11] Leska, C. A Study for Evaluating the Educational Impact of Algorithm Visualization in a Computer Literacy Course. http://faculty.rmc.edu/cleska/public_html/vizPlan.htm.

[12] Naps, T.L., Röβling, G. and Working Group. Exploring the Role of Visualization and Engagement in Computer Science Education. *ITiCSE 2002,* ACM Press (2003), 131-152.

[13] Naps, T.L., Röβling, G. and Working Group. Evaluating the Educational Impact of Visualization. *ITiCSE 2003,* ACM Press (2003), 124-136.

[14] Poole, D. and Mackworth, A. CIspace: tools for learning computational intelligence. *Effective Interactive Artificial Intelligence Resources Workshop IJCAI 2001*.

[15] Poole, D., Mackworth, A, and Goebel, R. Computational Intelligence: A Logical Approach. Oxford University Press, New York, 1998.

[16] Röβling, G. and Naps, T.L. A Testbed for Pedagogical Requirements in Algorithm Visualizations. *ITiCSE 2002,* 24-28.

[17] Saraiya, P., Shaffer, C.A., McCrickard D.S. and North, C. Effective Features of Algorithm Visualizations. *SIGCSE 2004,* ACM Press (2004), 382-386.