CMSC 671, Fall 2010, Prof. Laura Zavala

Final Project

Project Descripion

For the final course project, you will work in groups of two or three students to design, implement, demonstrate, evaluate, and report on a working AI system that investigates the techniques learned in the course in more depth. Your project should identify an interesting application domain for which you can apply one or more techniques from the class, including advanced topics of your interest that were briefly discussed.
The project will have a series of deliverables (topic, draft, and final report). The intention of these deliverables is to keep you on schedule for completing the project in time, and to give you feedback on your plans.
You can choose to work on one of two types of projects:

1. Experimental study: For this type of project, the emphasis is not on programming, but rather on testing existing known technology on novel data sets, and/or comparing against other known technologies.
   For this type of project, you may use software that is available off-the-shelf (ideally, in the public domain). Of course, the less programming you do, the more work I will expect in data analysis and analysis of your experimental results.
   Note that you might still have to do some programming in order to apply the existing tools to your dataset and do your comparisons, evaluation, analysis, etc.

2. Application project: In this type of project, the emphasis will be on the implementation and application of AI technologies to some domain. If you choose this type of project, you will implement the AI techniques you chose and apply them to some specific problem.

Here are some suggested projects for both types (experimental and application). Some of these ideas are borrowed from Dr. Marie desJardins final project (http://www.csee.umbc.edu/courses/671/fall10). You can select one of these, or develop your own project. I might post more as I think of ideas for your projects.

1. A Sudoku-puzzle-solving system that uses tabu search, the min-conflicts heuristic, and/or problem-specific ordering heuristics to find solutions efficiently.
2. An HTN planner that plays a solitaire card game such as Freecell or Spider, with a reinforcement learning component that learns operator preference heuristics.
3. A constraint-based scheduler that incorporates temporal reasoning and dependency-directed backtracking.
4. A genetic algorithm for solving the Traveling Salesman Problem or for automated design of combinational logic circuits.
5. A distributed constraint satisfaction algorithm for the n-queens puzzle, map coloring problem, or a Sudoku puzzle.

6. EXPERIMENTAL project: Conduct a machine learning experiment with some existing datasets (check the UCI machine learning repository at http://archive.ics.uci.edu/ml/datasets.html). For example, demonstrate some (known) differences between different learning algorithms across different data sets. You might also want to check WEKA (www.cs.waikato.ac.nz/ml/**weka**/), a collection of machine learning algorithms (in Java).

The experimental project type seems more suitable to machine learning algorithms, since there are many public available implementations and plenty of datasets. If you want to do an experimental project on another topic, talk to me and we can discuss it.

Please note that all work submitted for this project (implementation of the AI technique for the application projects and the test cases for the experimental projects) must be your own group's work, as stated in the course academic integrity policy. If you reuse any code, it should be only to support your investigation, not to serve as the main implementation of your project. You should explicitly let me know of any code that is incorporated, as well as document it as such.
Similarly, any written material or figures that you incorporate into your report (or any other deliverables) must be explicitly quoted and cited appropriately.

Project Deliverables

(midnight time for everything)

Topic and teams	Nov. 12	List of team members, project type, area of work, and problem domain.
Design/Status Report	Nov. 19	Design of your application/experiment, current status report
Presentation of your work	Dec. 7, Dec. 9	Oral presentation of your work
Final Report and system implementation	Dec. 15	Final report for your project, implementation, distribution of work report