Model AI Assignments 2011

Todd W. Neller
Gettysburg College
tneller@gettysburg.edu

Marie desJardins and **Tim Oates**
University of Maryland, Baltimore County
{mariedj,oates}@cs.umbc.edu

Matthew E. Taylor
Lafayette College
taylorm@cs.lafayette.edu

Abstract

The Model AI Assignments session seeks to gather and disseminate the best assignment designs of the Artificial Intelligence (AI) Education community. Recognizing that assignments form the core of student learning experience, we here present abstracts of three AI assignments from the 2011 session that are easily adoptable, playfully engaging, and flexible for a variety of instructor needs. Assignment specifications and supporting resources may be found at http://modelai.gettysburg.edu.

Clue Deduction: an introduction to satisfiability reasoning - Todd Neller

The popular boardgame Clue (a.k.a. Cluedo) serves as a fun focus problem for this introduction to propositional knowledge representation and reasoning. After covering fundamentals of propositional logic, students first solve basic logic problems with and without the aid of a satisfiability solver (e.g. zChaff). Students then represent the basic knowledge of Clue in order to solve a Clue mystery. Several possible advanced projects are sketched if students wish to pursue the topic in more depth.

The object of this project is to give the student a deep, experiential understanding of propositional knowledge representation and reasoning through explanation, worked examples, and implementation exercises. Students are expected to:

- gain an understanding of the syntax and semantics of propositional logic, as well as general logic terminology, including "model", "(un)satisfiability", "entailment", "equivalence", "soundness", and "completeness",
- learn the process of knowledge base conversion to Conjunctive Normal Form (CNF),
- solve word problems with proof by contradiction (a.k.a. reductio ad absurdum) using resolution theorem proving,
- represent knowledge so as to complete a program implementation that performs expert reasoning for the game of Clue, and
- compare deductive learning, inductive learning, and knowledge acquisition.

Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Mastermind Course Project - Marie desJardins and Tim Oates

This assignment is a final project for an introductory AI course. Students work in small groups to design a scalable guessing algorithm for the game of Mastermind (in which players must guess a "code" of colored pegs, with limited feedback provided after each guess). Three "challenges" are presented: a fixed-size problem (4 pegs and 6 colors), a "scalability challenge" (a series of tests on problems with increasing numbers of pegs and colors), and a "learning challenge" (guessing codes generated by a "biased generator," where the students are given training data that they can use to infer the generator's bias).

The project is very engaging for students, because it is game-based and very simple to understand, yet very challenging when it comes to designing an optimal (or even practicable) solution. The potential solutions also touch on many areas of AI, so the students can be creative in applying and synthesizing what they've learned to a new problem. The three challenges give the students the opportunity to choose which aspects of the problem they most wanted to focus on. The "tournament" nature of the assignment makes students very motivated to design a good algorithm.

Reinforcement Learning in a Generalized Mario Domain - Matthew Taylor

Computer games have been used successfully in both introductory computer science courses and in general AI classes, building excitement and enthusiasm. This model assignment uses RL-Glue and the Generalized Mario domain and is suitable for a unit on reinforcement learning in either an AI or machine learning class. The Mario domain is easy to install on multiple platforms and students will be able to quickly see an example agent interacting with its environment. Students will improve their understanding of common reinforcement learning algorithms by implementing them in a complex task. Additionally, they will gain practice with the more general skills of empirically analyzing, comparing, and measuring these algorithms.

A case study analyzing this assignment also appears in this years proceedings of EAAI, titled "Teaching Reinforcement Learning with Mario: An Argument and Case Study."