

CISpace: tools for learning computational intelligence

David Poole and Alan Macworth

Department of Computer Science,
University of British Columbia,
2366 Main Mall,

Vancouver, B.C., Canada V6T 1Z4,

poole@cs.ubc.ca, mack@cs.ubc.ca,

<http://www.cs.ubc.ca/spider/poole>

<http://www.cs.ubc.ca/spider/mack>

April 18, 2001

Abstract

This paper describes *CISpace: tools for learning computational intelligence*, a collection of interactive tools for learning AI, and how it fits in with other such tools. CISpace can be found at <http://www.cs.ubc.ca/labs/lci/CISpace/>.

1 Introduction

CISpace is a collection of tools designed to help learning about AI concepts. It is intended that there are exploratory tools as well as interactive online tutorials. It began as an offshoot of our textbook *Computational Intelligence: a logical approach*.

2 Current Applets

We currently have the following leaning tools:

- **Graph Searching** this lets the user interact with a variety of uninformed and heuristic search strategies using their own graphs (drawn on a canvas) on using a number of predefined graphs. Currently we have six search strategies ranging from depth-first search to A* search, each where we can have loop detection, multiple-path pruning, or no pruning.

- **Constraint Satisfaction via Arc Consistency** for solving constraint satisfaction problems. The user can interactively or automatically prune domains via arc consistency or can split domains.
- **Hill Climbing** for solving constraint satisfaction problems using local search techniques. This uses the same graph representation as the arc consistency applet. It lets us have random restarts, choose heuristics to decide which variable to select next and which value to select for it. We can plot how the number of unsatisfied edges changes through a single run of the hill climbing algorithm, as well as plot the run time distribution of a number of runs.
- **Neural Network** for training a neural network. This lets us design the network and train it on some data. We can plot how the total error evolves for each step of the gradient descent. Once trained we can view the parameters and try it on new examples.
- **Belief network** that lets the user build and query a belief network (Bayesian network).
- **Robot Control** that lets us design and/or modify a hierarchical controller for a simple robot. It uses a logic programming language to specify the controller. You can watch how the controller works in a 2-D environment that can be changed dynamically.
- **CILog** is a simple logic programming system. Unlike the other learning tools, this is written in Prolog and does not run as an applet. It can be seen as pure Prolog, with declarative debugging and expert-system facilities (why, how, whynot, ask-the-user, etc). It also allows for negation as failure, consistency-based diagnosis, and abduction (although the current form does not work with both negation as failure and assumables).

We are also expanding these to be part of a collection of interactive tutorials. As a first prototype see: <http://www.cs.ubc.ca/spider/poole/ci/test/csp/csp.html>. The aim is that we have interactive tutorials for a number of topics. One of the reasons for the interactive tutorials is that different students have different learning styles; some student like free-form interactive environments, while others get lost in them. Some like to be explained everything from the bottom up, but others are too impatient. We are planning on developing various ways to teach the material so that we can cater to various learning styles (and to eventually do research on different learning styles and how to present material to each person in a way that is appropriate for them).

This summer (2001) we are planning on improving all of the applets, improving the functionality and making improvements based on our experience in using them for teaching AI courses. The biggest changes will be to the belief network applet (expanding it to include multi-stage decision problems), and the robot control applet (perhaps expanding it to include multiple agents).

3 Relationship to an interactive AI resource

We don't believe it is possible or desirable to have a set of official AI teaching resources where we keep the best of the resources on each topic. There are a number of reasons for this:

- Different people have different views of what is a better tool. What is a better tool depends on the learning style of the student and what it is you actually want to teach them.
- There are reasons why we have chosen not to include some applets. Not all ideas are equally important. Adding more features does not necessarily make a tool better. AI is an ongoing research field. One of the main roles of research is to prune ideas that don't work. Unfortunately this often means claiming that some idea that people have worked on and are still working on will not form the basis for an intelligent system and so should not be included. We need people to make such judgments; however they could be wrong. It is much better to have a number of such resources, each of which gives a coherent view to the field. The diversity of what they choose to include and omit will make the discipline stronger.