CS 520 — Artificial Intelligence
Final Report Assignment
Out: March 26, 2001
Due: April 30, 2001

This assigment asks you to prepare a written proposal describing a possible research prototype of some artificial intelligence application.

The requirements are straightforward. You should submit a paper of roughly 2000 words, written in your best academic style. The paper should

- motivate a constrained but significant problem that you will consider, making it clear that the problem involves some form of decision-making with real-world data from the environment, and thus potentially involves incomplete information, uncertainty, and noise;

- suggest a class of model that you can use to describe the environment (perhaps one of those studied in class), thereby articulating some assumptions about that environment which you conjecture to provide a good approximation that allows for tractable reasoning and learning;

- describe a potential source of ready data about the environment which will allow you to estimate the parameters of your model of the environment empirically, spelling out the real-world cost of obtaining the data (hopefully low) and the computational cost of fitting the model to the data (hopefully manageable);

- map out a feasible protocol of evaluation that you could run in a final system, if you constructed it as planned, in order to better understand the problem you have studied and the solutions that you have obtained for it;

- if possible, contrast the direction you propose with other strategies that you know of (from the research literature, for example) making it clear how your suggestion draws on but extends the insights of these other strategies.

Thus, the form of this writeup is relatively constrained—it's a skeleton for a description of an experimental study. The topic of the writeup, however, is completely unconstrained.

By the end of this course we will have seen a variety of abstractions for inference under uncertainty—categorization, tracking, parsing—and several ways to

connect inference to decision-making—explicit decision models, handcoded policies that depend on the agent's beliefs about the hidden state of the world, and full-scale planning. We will also have looked, in a general way, at a number of typical AI problems, in such domains as vision, language, robotics, information management, etc. Any of these abstractions, any of these decision-making strategies, any of these domains, provides a suitable topic for the paper.

You're welcome to explore these topics further, especially if you expect that you'd like to go on to research in a typical AI area. But I am just as curious to hear your suggestions about how the techniques that we've covered might apply to very different problems—problems that you're independently interested in and may even have discovered in work that you've done in a very different context from AI. As often as not, it's the friction of widely disparate ideas that provides the electricity of the intellectual world.

If you do succeed in connecting class material with your deeper interests, you need not think of this assignment as a dead-end to be handed in and forgotten. For those of you in the CS master's program, you will easily be able to revise and extend your paper into a Master's essay. For those of you in the CS PhD program, you might think seriously of actually carrying out the research you propose during a following independent study with an appropriate CS faculty member—perhaps even of writing up the results you obtain to fulfill the PhD program depth requirement. Others of you—less directly connected with computer science—might nevertheless think of your writeup as a springboard to infuse your research with new ideas and technology, and establish an interdiscipinary perspective on what you do already.