CS 530 — Principles of AI
Assignment One
Out: September 9, 2003
Due: September 30, 2003

About This Assignment

The problem in this assignment is to write an AI argument using a computer program that you implement. The majority of the text of this assignment describes the program that you should write and the context in which you should try to understand it. You will have to hand in the program you write electronically, and some of the credit for the assignment will depend on this program. However, the reason that you implement the program is so that you can characterize when and how it works. The most important part of this assignment is your English description of your experience and conclusions. The requirements for your writeup are at the end of the assignment.

Programming may be completed in the programming language of your choice. You will have to run your program on specific sample data. The data is available on the web in formats that can be read in easily into a Lisp or Prolog interpreter. Of course, it is also available in simple text files that can be conveniently read into a C or Java program – or preprocessed by a Perl script into whatever format you find convenient.

Program Description

Overall, the program constitutes a "development environment" for AI programming using the decision models described in "Agents in the Real World". You should be able to:

- Solve decision models: compute the optimal policy for a model.
- Train decision models: fit parameters in schematic decision models from sample interactions with the environment.
- Estimate expected performance: characterize the uncertainty of model results and simulate sample runs of a model.

A specific outline for this program follows.

Part 1. Describe concrete data structures for models and policies as described in "Agents in the Real World", and implement the function OPTIMUM to compute the best policy in a model.

Describe is a cover term I use for the different ways you create data structures in different languages. In C or Java, you define the data structures (e.g., in .h or class files); in languages like Prolog where you can introduce any mnemonic structured terms to hold data, you should add a comment saying what terms you're assuming; in a language like Lisp where all complex structures are stored the same way, you describe data structures by defining functions that build and access particular instances of those structures to represent your data. Describing the data structures is not asking for something special, it's just what you would do anyway. And of course your data structures can work any way you see fit to organize your program.

Your implementation of OPTIMUM should work in the general case. However, the dog model from "Agents in the Real World" would provide a good test case for your implementation.

Part 2. Now we introduce two new representations.

- *Model schemas* are representations that give the structure of a model but not the numerical parameters for probabilities and utilities.
- *Histories* are representations that describe a sample run of an agent, including a sequence of actions and observations and the utility outcome that the agent achieves as a result.

Describe concrete data structures for model schemas and a collection of histories, and implement a function TRAIN that estimates a model from a model schema and a collection of histories. In this case the histories form *training data* for your program. One simple way to write TRAIN is to traverse the model schema recursively, and keep track of the corresponding histories. At each stage, you can estimate probabilities for a node the model by the relative frequency with observed values actually occur in the training data. You can estimate utilities for the model by the average utility that you obtain in the training data. For possible outcomes that never occur in the training data—and you must handle this possibility—you can use a default utility of 0.5.

One way to test this function as you develop it is to use the data for Problem 3. There's also a sample data set of 100 trials generated according to the dog model from "Agents in the Real World", together with the actual parameter sets that you get from this data set.

Part 3. Enrich your program so that it can represent *performance models*. The performance model should include the standard deviation of observed utilities as well as the mean of observed utilities. (You can use a single set of data structures in your program as the models for Parts 1 and 2 and the performance model for Part 3.) To complete this problem you need to implement two algorithms that work on performance models in a general way.

First, implement a function that will SIMULATE a performance model by sampling. Choose actions at random; choose observations with the probability predicted by the model. Generate outcomes with a normal distribution with the specified mean and standard deviations (truncating negative values to 0 and large values to 1). The following frankly mysterious algorithm generates such normally distributed random variables, if the function $r()$ generates a random variable uniformly distributed between 0 and 1, μ is your mean and σ is your standard deviation:

repeat
  $x \leftarrow 2 * r() - 1$
  $y \leftarrow 2 * r() - 1$
  $r \leftarrow x^2 + y^2$
until $0 < r < 1$
return $\mu + \sigma * x * \sqrt{-2\ln(r)/r}$

Second, extend your training function so that it learns a performance model from a model schema and history data. In this case, compared to problem 2, you need to also estimate the standard deviation of performance. If u_i is the utility on the ith trial, there are n trials, and u is your estimate of the average utility, use the estimate S defined here:

$$S = \sqrt{\frac{(\sum_i u_i^2) - nu^2}{(n-1)}}$$

Again, do not worry if the formula is mysterious, either. Note that it can be computed iteratively.

Written work.

Problem Statement. With this assignment, the class web site contains three sets of test data files. Each data set contains five example files. The first four files in each set involve a small amount of training data (50–100 instances), while the last file involves a larger amount of training data (500-1000 instances). Your task is to explain the variability in performance of your program on this data, using additional results with your program to create the most convincing argument you can.

Each data file comes in three formats. In the vanilla format, each line of the file describes a history. The history is reported as a sequence of observations or actions (character strings separated by spaces) followed by the utility obtained on the history (a real number, as a string of digits). In the lisp format, the file itself defines a variable that holds the training data as a list of lists. In other words, each history is a list containing symbols for the observations and actions encountered and ending with the utility obtained, and the histories appear consecutively in one big list. Finally, in the prolog format, the file defines a unary predicate history; history(L) is true if L is a list describing an element of the training set; the Prolog list takes the same form as the lisp list (modulo Prolog syntax).

All of the histories are instances of this model schema:

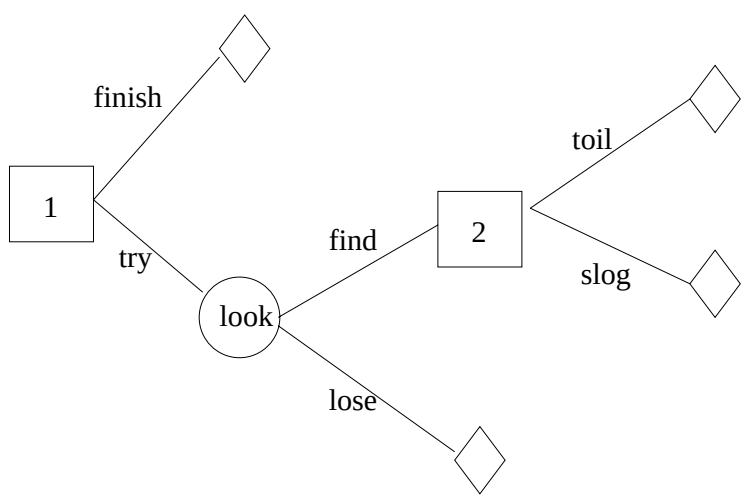

So sample elements might look like this

version	statement
vanilla	`try find slog 0.4`
lisp	`(setq t1 '((try find slog 0.4) ...))`
prolog	`history([try, find, slog, 0.4]).`

All of the files in each set come from the same model with the same parameters. However, different model parameters were used across the sets.

Use your answers to Problems 1 and 2 to automatically train a model from each file of data and compute the optimal policy for each trained model.

Note the variability in the answers and think about why you are finding this variability. Try to develop English arguments, including "back-of-the-envelope" calculations and tabulations of results obtained using your code, to give an analysis of the probability of determining the optimal

policy correctly for each of the three models of Problem 3 as the amount of training data gradually increases. Describe the features of the model that determine how variable the results are, and develop evidence that these features matter.

Hand In On Paper a five-page double-spaced writeup, documenting your experience and conclusions. Use this assignment to practice research writing. Set the context for your investigation by summarizing what your program does. Describe what you found in the sample data (qualitatively and numerically, perhaps using a chart or graph). Explain why you got these results, using the arguments that you have developed. Speculate about the general implications for cobining modeling, learning and analysis to design robust agents. Use your own words throughout and take the opportunity to frame the ideas behind your assignment and your results precisely for yourself.