CIS 730 Artificial Intelligence
CIS 530 Principles of Artificial Intelligence
Fall 2007

Homework 10 of 10: Machine Problem (MP10)
Term Project Experiments

Assigned: Sun 02 Dec 2007
Due: Fri 07 Dec 2007 (before midnight)

The purpose of this assignment is to help you with your final project implementation phase.

This homework assignment is worth a total of 20 points.
Upload a copy of your solution to your K-State drop box before the due date.

1. **(20 points for 530, 12 points for 730) Experiment Replication Script.** Write a script in any scripting language (e.g., Unix shell script, Tcl/Tk, Perl, Ruby, Python, Lua, or VB.net) to automate the collection of experimental results for your final project report. Alternatively, write a data generator or preprocessor using a scripting language or a high-level imperative language such as C++ or Java. Examples include:

   a) A Perl script to run *Angband* using different levels of training (the basic APWborg, trained with (for example) 1000, 2000, and 4000 turns of combat data.
   b) A Python driver for training a *TAC-Classic* agent and then bringing them into new games on a running server
   c) A Java program to generate test data for benchmarking the import module you write to map data from a particular protein ontology format (e.g., DIP) into the unified protein interaction format. Measure the ontology reasoner and I/O costs separately.

   Turn in a standalone table of results produced by this program along with your source code.

2. **(730 only, 8 points) Clustering.** Run the k-means clustering applet from Stanford University (http://nlp.stanford.edu/~danklein/demos/constrained-clustering-demo.shtml) and turn in a screenshot of the results for each of the three examples:

   a. first data set with 4 clusters, k-means. (Why does this sometimes produce imperfect cluster labels?)
   b. third data set with 2 clusters, spatial (Why don't any of the implemented algorithms work on this data set? What would be needed to separate the inner and outer clusters?)
   c. your own random data set with 3 clusters, 5 - 9 points each, instance-based

   Note: Include multiple screenshots if you get different results on different runs. Name them *mp10-2a-i*, where $i \geq 1$.

Class participation (required). Think about the following question and post your answer to CIS730-L by Fri 07 Dec 2007: What might be a potential pitfall of using a conversational Turing Test, *unrestricted by language*, to test the reasoning ability of a general-purpose AI system? Give an example of an evaluation difficulty: an ambiguity, subjective aspect, or problems that the *human* examinee might have with the problem (assuming he or she **does** speak all of the test languages).