Project 2B Specification

Goals

The goal of this project is to develop—through experience—your ability to:

1. Make technical choices.
2. Bring multiple ideas together to solve a problem.
3. See how one problem can serve as a precedent for the solution of another.
4. Deal with problems that have less specification than that typical of a problem set.
5. Organize and perform work with a partner, if you choose to work with a partner.
6. Write about your work.

The Learning Problem

For Project 1's research option, you were to thing about translation from a change-oriented description to a trajectory-oriented description. For Project 2's research option, we want you to build a system that learns facts about the world based on pairs of descriptions.

Variation 1

Consider, for example, the following situation and description pair, expressed in trajectory-suggesting English:

1. Scene: Boy, rock, tree, and bird all at separate locations.
2. A boy walks to the rock and then the boy throws the rock at the bird. Example. Witnessed by learner.
3. A boy walks to the tree and then the boy throws the rock at the bird. Near miss. Oracle says cannot happen.

You would want your system to learn that you have to go to a throwable object before you can throw it.

You might also think about what sorts of supplementary representations, situations, examples, and near misses you would need to learn that birds can fly, but people cannot, or that a cats chase birds so that they can eat them.

The learning done should lay the foundation for a future system that critiques what is said to it. For example, once learning has occurred, the future system should say *What!* if told that *Patrick flew to the top of a tree*.

We believe the right approach is through a combination of Winston style near-miss learning (subject of a forthcoming lecture) and Jackendoff-style trajectory representation. You can learn more about near-miss learning from the following material:

1. Near-miss learning, Ch. 16 of *Artificial Intelligence, 3rd edition*, by Patrick Winston.
2. Explanation-based learning, Ch. 17 of *Artificial Intelligence, 3rd edition*, by Patrick Winston.

Variation 2

Same as variation 1, except that you decide to base your work on a transition-space representation of examples and near misses, rather than a trajectory-space representation.

Your job

1. Implement a W-learning mechanism for trajectory-space or transition-space learning.
2. Teach your system some set of interesting facts about the world using simple pairs of descriptions.

Check points

There is only one check point: you are to provide us with evidence of working code and your final report on the latest day MIT rules permit us to ask for it.

Report Length

The right length for a paper is always the shortest length that covers what you want to say clearly. As a rough guide, we do not want you to write more than **five pages**, exclusive of illustrations, code, run printout, and the like. We would be surprised if you can say what needs to be said in much less.