

Homework 1: Chapters 1 - 5

The following exercises are due at the beginning of class on **Tuesday, October 1**. Some of these problems may take a while to solve, so I recommend that you work on this assignment over the course of multiple days.

1. **[10 pts.]** Prove the following:
 - a) For every purely reactive agent, there is a behaviorally equivalent standard agent.
 - b) There exist standard agents that have no behaviorally equivalent purely reactive agent.
2. **[10 pts.]** There are two ways of specifying tasks by utility functions: by associating utilities with either states ($u : E \rightarrow \mathcal{R}$) or with runs ($u : R \rightarrow \mathcal{R}$). The second type of utility function is strictly more expressive than the first. Give an example of a utility function over runs that cannot be defined simply by associating utilities with states.
3. **[45 pts. total]** Consider the vacuum-world example from Chapter 3 (pp. 52-54). Note: This is different from the environment we used for our programming assignment. In particular, there are no obstacles and the agent does not need to return home after it is done cleaning the room.
 - a) **[15 pts.]** Imagine that your agent must be able to perform the task in environments of arbitrary size. Assume that the agent's internal state includes a ground atom of the form `worldSize(width, height)` that is used to indicate the maximum x coordinate ($width - 1$) and maximum y coordinate ($height - 1$). Write a set of general rules that will allow the agent to function properly regardless of the world size. Try to keep this set of rules as small as possible. Hint: You may find it convenient to use predicates to represent some mathematical comparisons and to treat arithmetic expressions as terms (in fact, this could be considered shorthand for pre-defined functions that return the result of applying arithmetic operations to their terms). You may also assume the presence of other helpful mathematical predicates, as long as you are clear about how they should be defined.
 - b) **[15 pts.]** Now consider a solution to the problem that uses situation calculus.
 - i) Give precondition and effect axioms for the *forward* action. Note, you may need to modify some of the domain predicates so that they become fluents. As above, you may use arithmetic expressions as terms.
 - ii) Give a successor state axiom for the fluent `Facing(d,s)`, where d is the direction the agent is facing in situation s .
 - c) **[15 pts.]** Finally, use the book's version of the subsumption architecture approach to design a purely reactive agent for this environment. You can assume a "wall" percept which tells the agent that it is facing the boundary of the room. You do not need to implement this agent, only give a formal description of it. How does it compare with the logic-based example? Which do you think will perform better? Which is simpler?
4. **[10 pts.]** Consider the architecture presented by Bratman, Israel, and Pollack in "Plans and Resource-Bounded Practical Reasoning." Compare Fig. 1 of this paper to the pseudo-code for a practical reasoning agent presented in Fig. 4.3 of our textbook (p. 75). Among other things, discuss which modules correspond to which functions, how the inputs and outputs are similar or different, and how the architectures differ in capabilities.

5. **[15 pts.]** Consider the Jam system shown in Figure 4.6 (p. 82). Assume that PERFORM move \$OBJ1 \$OBJ2 has the following effects: add the fact that \$OBJ1 is on \$OBJ2, delete the fact that \$OBJ2 is CLEAR (unless it is the table), and delete the fact that \$OBJ1 is on whatever it used to be on. Trace the execution of this PRS agent until it has achieved its top-level goal “blocks_stacked”. Show beliefs, intention stack, and executed action (if any) at each step. Be sure that the intention stack includes information that allows the agent to resume executing a plan that was suspended in order to accomplish a subgoal. Explain any assumptions you need to make about the high-level algorithm described on pages 80-81 in order to ensure correct behavior. If you wish, you may use a more compact representation for the agent’s beliefs than that used by Jam (e.g., you might choose to use first-order logic).
6. **[10 pts.]** In “Unifying Control in a Layered Agent Architecture” by Fischer et. al (1994), the authors’ experiments find that behavior-based strategies tend to be more robust than plan-based, cooperative strategies. Based on your understanding of the various agent strategies tried, and your own analysis of the results, develop a hypothesis to explain this result.