

EECS 391: Introduction to Artificial Intelligence (Spring 2012) Written Homework 4 (Max points: 100)

Assigned Tuesday February 28, due 5pm Tuesday March 6. Write your answers neatly and remember to show all relevant work. Before turning in your work, staple your answer sheets together and write your name (or names if working as a pair) on each page.

1. From $Friends(John, Mary)$ we should be able to infer $\exists x Friends(John, x)$. Write an inference rule similar to resolution that allows this inference. Write the conditions that should be satisfied for the variables and terms for this rule to be used correctly. (10 points)
2. Derive the most general unifier if it exists: (i) $Q(y, G(A, B))$, $Q(G(x, x), y)$, (ii) $Knows(Father(y), y)$, $Knows(x, x)$. Here x and y are variables and A, B are constants. (10 points)
3. Suppose we omit the occurs check from the unification algorithm, so that it allows a literal like $P(x, f(x))$ to be unified with $P(r, r)$. Show that this allows the conclusion $\exists r P(r, r)$ to be inferred from $\forall x \exists y P(x, y)$. Give an intuitive example of a predicate P where this is not sound, i.e. $\forall x \exists y P(x, y)$ is true but $\exists r P(r, r)$ is false. (10 points)
4. Solve the riddle “Brothers and sisters have I none, but that man’s father is my father’s son” to find the identity of “that man.” Use the rules of the kinship domain in Section 8.3.2 of the book. (10 points)
5. Suppose a KB has just the following formulae:
 $Ancestor(Mother(x), x)$
 $Ancestor(x, y) \wedge Ancestor(y, z) \Rightarrow Ancestor(x, z)$
Can resolution prove $\neg Ancestor(John, John)$ from this KB? Explain how or why not. (10 points)
6. The monkey-and-bananas problem is the following: An AI monkey is in a room with a box and some bananas hanging from the ceiling. The box and the monkey have height *Short*, but if the monkey climbs on the box it will have height *Tall*. The bananas have the height *Tall*. The monkey can *Go* from one place to another, *Push* an object, *ClimbUp* or *ClimbDown* from an object and *Grab* or *Drop* an object. The result of a *Grab* is the monkey holds the object if they are at the same place at the same height. The monkey wants to get the bananas. Write this as a STRIPS planning problem. (10 points)
7. Explain why dropping negative effects from action schemas results in a relaxed problem. (10 points)
8. Figure 10.4 in the book shows a famous example of the “Sussman anomaly,” a very early AI planning problem that was difficult for many early automated planning systems because they were “non-interleaved.” Given a conjunctive goal $G_1 \wedge G_2$, a non-interleaved planner first

solves G_1 and then G_2 , or vice versa. (The resulting plan is the plan for G_1 concatenated with the plan for G_2 , or vice versa.) (i) Represent 10.4 as a planning problem. (ii) Solve it by hand. (iii) Explain why this could be difficult for noninterleaved planners. (10 points)

9. When we covered search, we talked about bidirectional search in class. Could bidirectional search be used in (i) state space planning and (ii) plan space planning? Describe how if so. (10 points)

10. Explain how forward state-space searching algorithms could also be considered plan-space algorithms, and describe the corresponding search space. (10 points)