Assignment10 – Learning

Given: July 10 Due: July 15

Problem 10.1 (Passive Reinforcement Learning)

Consider the example on Passive Learning in 4×3 world from the slides.

1. Give the transition model to the extent that it can be learned from these trials.
2. How could we learn the entire model?
3. How would we proceed to learn the utilities of the states?

Problem 10.2 (Active Reinforcement Learning)

Consider reinforcement learning in an unknown non-deterministic environment.

1. Explain the difference between a passive and an active agent.
2. What is the critical trade-off in designing an actively learning agent?

Problem 10.3 (Logical Formulation of Learning)

Some people with different attributes go sunbathing, the result class is whether they get sunburned:

#	Hair	Height	Weight	Lotion	Sunburned
1	Blonde	Short	Light	No	Yes
2	Blonde	Short	Average	Yes	No
3	Brown	Short	Average	Yes	No
4	Blonde	Short	Average	No	Yes
5	Blonde	Tall	Heavy	No	Yes
6	Brown	Tall	Heavy	No	No

We want to model these examples in first-order logic.

1. Explain which predicates are needed. For each predicate give the arity and the domain of each argument. Give the formal representation of the description and the classification of Example 1.
2. For attributes A, B, we write $A_1, \ldots, A_n \succ B$ if any two examples that agree on all of the attributes A_i also agree on the attribute B.
   Explain whether Height$\succ$ Weight and Weight $\succ$ Height hold or do not hold.
3. Give a minimal subset $\mathcal{A} \subseteq \{Hair, Height, Weight, Lotion\}$ such that $\mathcal{A} \succ Sunburned$ holds.
   Using the predicates introduced above, give a logical formula that captures the rule $\mathcal{A} \succ Sunburned$.

Problem 10.4 (Inductive Learning)

Consider the family tree given by the following relations:

couple	children
A, B	E, F
C, D	G, H
F, G	I, J

Assume we already know the predicate $\texttt{par}(x, y)$ for x being a parent of y.

Our goal is to learn the predicate $\texttt{gp}(x, y)$ for x being a grandparent of y. That means to find a formula D such that $\forall x, y.\texttt{gp}(x, y) \Leftrightarrow D(x, y)$.

We do not know D, but we have the following examples for gp:

person-pair	grandparent
A, I	yes
B, I	yes
A, J	yes
A, E	no
A, F	no
F, A	no
A, A	no
C, J	yes
D, H	no
I, A	no

1. Give the intended formula D_1, i.e., the correct definition of grandparent.
2. Give a formula D_2 that covers exactly the positive examples.
3. Explain the pros and cons of learning the formula D as D_1 vs. D_2.
4. We want to learn algorithmically the formula $D(x, y) = \exists u_1, \ldots, u_l.L_1 \wedge \ldots \wedge L_k$ where each L_i is a literal of the form $P(x, y, u_1, \ldots, u_l)$ or $\neg P(x, y, u_1, \ldots, u_l)$ for some predicate symbol P including the equality predicate, i.e., $P \in \{\texttt{par}, \texttt{gp}, =\}$. We do so by building the set $\{L_1, \ldots, L_k\}$ of literals gradually (with the understanding that each L_i can have any free variables in addition to x and y, which we will collect at the end as the u_i).
   (a) If we start with the empty set of literals, give all useful choices that we can make for the first literal.
   (b) For each choice L that is non-recursive (i.e., P is not the target predicate gp), positive (i.e., the literal is not negated), not an equality (i.e., P is not the equality predicate $=$), and does not introduce a new variable (i.e., only uses x and y), which examples are falsely classified?