# CPSC 322, Practice Exercise
# Solutions to Single-Stage Decision Networks

## 1 Directed Questions

- What is meant by a *one-off* decision? How can this be applied in the delivery robot example? **Answer:** The agent knows which actions are available, has preferences expressed by utilities of outcomes, and makes all the decisions before any action is carried out. In the delivery robot example, the decisions on wearing pads and taking the long or short route are made *before* the robot goes anywhere. Multiple decisions can be considered as a single macro decision.
- Define utility in a decision problem. **Answer:** The utility is a measure of desirability of possible worlds to an agent, i.e. indicates the agent's preferences. Let $U$ be a real-valued function such that $U(w)$ represents an agent's degree of preference for world $w$. The value of a utility is typically between 0 and 100.
- How do we calculate the *expected utility* of a decision? **Answer:** The expected utility is derived by summing over the possible worlds that select that decision, for each world $w$ multiplying $U(w)$ by $P(w)$.
- How do we compute an optimal one-off decision? **Answer:** If we calculate the expected utility for each decision as per the last question, we choose the decision that maximizes the expected utility.
- What are the three types of nodes in a single-stage decision network? **Answer:** Decision nodes, random variables (chance nodes), and utility nodes
- What is a policy for a single-stage decision network? What is an optimal policy? **Answer:** A policy for a single-stage decision network is an assignment of a value to each decision variable. The optimal policy is the policy whose expected utility is maximal.
- Describe the variable elimination steps for finding an optimal policy for a single-stage decision network. **Answer:** Prune all the nodes that are not ancestors of the utility node. Sum out all the chance nodes. There will be a single factor $F$ remaining that represents the expected utility for each combination of decision variables. If $v$ is the maximum value in $F$, return the assignment $d$ that gives that maximum value $v$.

## 2 A One-Off Decision

You are preparing to go for a bike ride and are trying to decide whether to use your thin road tires or your thicker, knobbier tires. You know from previous experience that your road tires are more likely to go flat during a ride. There's a 40% chance your road tires will go flat but only a 10% chance that the thicker tires will go flat.

Because of the risk of a flat, you also have to decide whether or not to bring your tools along on the ride (a pump, tire levers and a puncture kit). These tools will weigh you down.

The advantage of the thin road tires is that you can ride much faster. The table below gives the utilities for these variables:

| bringTools | flatTire | bringRoadTires | Satisfaction |
|---|---|---|---|
| T | T | T | 50.0 |
| T | T | F | 40.0 |
| T | F | T | 75.0 |
| T | F | F | 65.0 |
| F | T | T | 0.0 |
| F | T | F | 0.0 |
| F | F | T | 100.0 |
| F | F | F | 75.0 |

- Create the decision network representing this problem, using AISpace. **Answer:** An example is given in xml file `bikeride_tires_flat_tools.xml`.

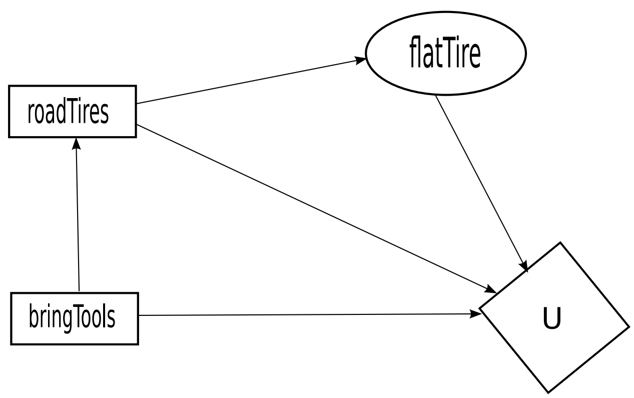


Figure 1: A decision problem.

- Use variable elimination to find the optimal policy.
  - What are the initial factors? **Answer:** There are two factors to begin with, one representing $p(flatTire|roadTires)$ and one representing the utilities.
  - Specify your elimination ordering and give each step of the VE algorithm. **Answer:** We sum out our chance node $flatTire$ first. This results in a new factor on the decisions. We eliminate $roadTires$ by maximizing that decision variable for each value of $bringTools$. This leaves one factor on $bringTools$. We maximize $bringTools$ in that final factor to get our answer.
- What is the optimal policy? What is the expected value of the optimal policy? **Answer:** The optimal policy is to take the thicker tires and leave the tools at home. The expected utility of this policy is 67.5
- Try changing the utilities and the probabilities in this problem, and identify which changes result in a different optimal policy. **Answer:** There are many possibilities here, e.g. changing the probability of a flat tire given the tire type, or decreasing the utilities for the two possible worlds FFT (currently 100) and FFF (currently 75).

# 3 Learning Goals

You can:

- Compare and contrast stochastic single-stage (one-off) decisions vs. multistage decisions.
- Define a utility function on possible worlds.
- Define and compute optimal one-off decision (max expected utility).
- Represent one-off decisions as a single-stage decision network and compute optimal decisions using variable elimination.