


Reading Assignment: Read Sections 4.4–4.5 and 6.1–6.3.

Problems:

1. Suppose a 6-sided die is rolled. What is the expectation of the value showing? Suppose two 6-sided dice are rolled. What is the expectation of the sum of the values showing?
2. Rosen, Section 4.5, problem 32. Please compute the *net gain*, i.e. the expected amount of money you actually win. To compute this value, you must subtract the \$1 you pay for the ticket.
3. Rosen, Section 4.5, problem 44.
4. *Mobile robot localization*: Bayes' rule underlies most modern AI systems for probabilistic inference (see also Rosen, page 305, problems 40–42). One application of this rule is the update of a robot's position estimate based on new sensor information. To see, look at the example below.



The robot is placed in the hallway facing east and it does not know where it is (it only knows its orientation). The hallway is tessellated and the robot can be in any of the five squares (always facing east). Assume that the width of each square is 1 meter. Then the squares are numbered by their distance from the beginning of the hallway. Given the tessellation of the hallway, we can represent the position of the robot by a random variable L . This variable takes on values between 0 and 4, depending on the robot's current position.

At each point in time, the robot's position estimate is represented by a probability distribution over all possible locations. In the beginning, the robot does not know where it is. This can be represented by the following distribution:

$$P(L = 0) = P(L = 1) = P(L = 2) = P(L = 3) = P(L = 4) = 1/5 \quad (1)$$

The robot has a camera that can be pointed to the left or to the right. In order to determine where it is, the robot tries to detect doors by looking either to the left or to the right (note that the robot faces east). As you can see in the figure, there are four doors in the hallway.

Now we want to update the robot's position estimate based on an observation, i.e. we want to compute the probability distribution of L *given* an observation. In general, we denote observations by o (where o ranges over the possible observations). The probability distribution *after* the observation can be represented by $P(L = l|o)$ (l is a number between 0 and 4). However, it is not easy to compute this value directly. In mobile robot localization we use Bayes' rule to compute this value:

$$P(L = l|o) = \frac{P(o | L = l)P(L = l)}{\sum_{i=1}^n P(o | L = l_i)P(L = l_i)} \quad (2)$$

All we need to know are the values for $P(o | L = l)$ and $P(L = l)$. $P(L = l)$ is given by the distribution *before* the robot made the observation (given in Equation (1)).

The term $P(o | L = l)$ describes the probability of making an observation *given* the position of the robot. In our example, we have four possible observations: The robot can detect a door to its left, it can detect a door to its right, it can detect a wall to its left, or it can detect a wall to its right. We will denote these four observations by DL , DR , WL , and WR (i.e. these are the four possible values of o). To determine $P(o | L = l)$, the probability of making observation o at location l , we make the following assumption about the robot's ability to detect doors and walls. If the robot points the camera towards a door, then it successfully detects this door with probability 0.8. However, it sometimes confuses the door for a wall, i.e. with probability 0.2 it erroneously detects a wall, even though it looks at a door. Detecting walls is easier, therefore, the probability of detecting a wall if the robot is looking at a wall is 0.9. With probability 0.1 it thinks it detects a door even though it looks at a wall. There are four different types of positions in the hallway: Let l_{DL} denote any position in the hallway, at which there is a door to the left of the robot (these are positions 1 and 4 in our example). The other three types are given by l_{DR}, l_{WL} and l_{WR} . For example, l_{WR} are positions 0, 1 and 3. Now we can summarize the robot's ability to detect doors and walls as follows:

$$\begin{aligned}
 P(DL|L = l_{DL}) &= 0.8 \\
 P(DL|L = l_{WL}) &= 0.1 \\
 P(WL|L = l_{WL}) &= 0.9 \\
 P(WL|L = l_{DL}) &= 0.2 \\
 P(DR|L = l_{DR}) &= 0.8 \\
 P(DR|L = l_{WR}) &= 0.1 \\
 P(WR|L = l_{WR}) &= 0.9 \\
 P(WR|L = l_{DR}) &= 0.2
 \end{aligned}$$

For example, the probability that the robot detects a door to its left if it is at location 3, is determined by $P(DL|L = l_{WL})$, since there is actually a wall to the robot's left side at position 3. Now we can start to estimate the position of the robot.

- (a) Assume that the robot does not know where it is ($P(L)$ is given by equation (1)). The robot detects a wall to its right, *i.e.* it makes the observation WR . What is the distribution of the robot's position estimate after this observation? Use a calculator (or write a computer program) and Bayes' rule.
- (b) Assume the robot has already updated its position estimate based on the first observation WR . Now the robot turns the camera and looks to the left. What is the robot's position estimate, if it detects a door to the left? Use the distribution you computed in the previous step to represent $P(L = l_i)$. This probability distribution combines the fact that the robot has detected a wall to its right and a door to its left.
- (c) Compute the mean and the variance of the two previous distributions over possible robot positions.
- (d) After these two observations, the robot points the camera back to the right and, oops, it detects a door! Compute the new probability distribution and briefly discuss whether the new values make sense.
- (e) We might have assumed that the robot never fails when it detects a door or a wall. How could you represent such a perfect sensor model using the conditional probabilities? What would the distributions look like after the observations WR , DL and DR ?