CS373: Glossary of Terms

Unit One

Normalized

A probability distribution is said to be normalized if the sum of the individual probabilities sums to one. Example: if all 5 cells have equal probability of 0.2, this is a normalized distribution since 0.2+0.2+0.2+0.2+0.2 = 1.

Convolution

Convolution is a mathematical operation that takes two functions and measures their overlap. To be more specific, it measures the amount of overlap as you slide one function over another. For example, if two functions have zero overlap, the value of their convolution will be equal to zero. If they overlap completely, their convolution will be equal to one. As we slide between these two extreme cases, the convolution will take on values between zero and one. The animations here and here should help clarify this concept.

Probability Distribution (PD)

A probability distribution is a function that assigns a probability to all possible outcomes in the sample space. For example, when we flip a fair coin, this function would assign the value of ½ to each of the possibilities in the sample space {heads,tails}.

Uniform PD

A probability distribution in which all possible outcomes are equally likely.

Prior PD
This is the name we give to the probability distribution we have before we make an observation. Compare with the posterior probability distribution.

Posterior PD
The name for a probability distribution after we have made an observation.

Sample Space

Sample Space refers to the set of all possible outcomes. The sample space for a distribution depends on what you are measuring. If you flip a coin the sample space is {heads, tails} because those are the two possible outcomes. If you are a robot car in one of 5 possible cells, the sample space is {cell 1, cell 2, cell 3, cell 4, cell 5}

Monte Carlo Localization

This is the technical name for the localization procedure used throughout Unit One. This procedure repeats the “sense” and “move” steps to determine a belief about the robot’s location. Sensing always results in an increase in the robot’s knowledge, while moving always decreases the robot’s knowledge about its position.

Stationary State

A stationary state, also known as a stationary probability distribution, is one in which the probability remains constant (stationary) over time. In Unit One, we found that repeated applications of the move function to a uniform probability distribution did not affect the distribution. This means we can call the uniform distribution a stationary state in this example.

Probability

Conditional Probability
Conditional probability is an updated belief based on new information. This new information usually comes from a measurement, which we label Z. We write conditional probabilities in the following form:

$$P(X|Z)$$

and this is read as “The probability of event X given Z.”

As an example: let’s say we have two coins. One is fair, and one has two heads. If we pick up the fair coin, the probability of flipping heads is one half. If we pick up the unfair coin, the probability of heads is one. Mathematically, we can write

$$P(Heads \mid Fair\ Coin) = 1/2$$
$$P(Heads \mid Unfair\ Coin) = 1$$

Bayes' Rule

Bayes' Rule gives us a way to use conditional probabilities to find their reverse forms. For example, in the previous definition we knew the probability of heads when we had a fair coin, but what if we flipped the question? What is the probability that the coin I just flipped is fair, given that it came up heads? For this we need Bayes' Rule, which can be written as follows:

$$P(X|Z) = \frac{P(Z|X)P(X)}{P(Z)}$$

Check out the Unit One course notes or the videos to see an example of Bayes' Rule worked out for you.

Total Probability Theorem

The Total Probability Theorem is a way of counting up all the possible ways for an event to happen. It says that the probability of an event occurring is equal to the sum of all the possible ways that event can occur. For example, let's say we want to know the probability of rolling two dice in such a way that the total is seven. We just have to count up all the ways this can happen and add together those probabilities. So, how can we get a 7? We could get a 1 and 6; a 2 and 5; a 3 and 4; a 4 and 3; a 5 and 2; or a 6 and 1. Each of these possibilities has probability $\frac{1}{36}$, and there are 6 of these possibilities thus making the **total probability** of rolling a 7 equal to $\frac{1}{36} * 6 = 1/6$.

The formal definition for the total probability theorem is:

$$P(X) = \sum_i P(X|Z_i)P(Z_i)$$

In our example before, X would represent rolling a seven, and each of the Z_i would be the different possible ways to roll a seven.