Lecture 9

Module Three: Probabilistic Reasoning

Topics: Probability, Bayes' theorem, Markov Processes, Bayesian Inference
Lab exercises: Localization and mapping

Probability Theory

Probability theory provides a framework for modeling and reasoning about uncertainty making precise statements about uncertain situations and drawing reliable inferences from unreliable observations
Probability theory provides a framework for designing systems that are robust to uncertainty

Probability Theory: Events

Probabilities are assigned to events, which are possible outcomes of an experiment
Example: flip three coins in succession
There are eight atomic events: HHH, HHT, HTH, HTT, THH, THT, TTH, TTT
Atomic events are mutually exclusive
Set of all atomic events is collectively exhaustive (cover all cases)
Set of all possible atomic events is called the sample space U

Probability Theory: Axioms of Probability

The probabilities that are assigned to events must obey three axioms:
Non-negativity: Pr[A] ≥ 0 for all events A
Scaling: Pr[U] = 1
Additivity: if A/B is empty, Pr[A∪B] = Pr[A] + Pr[B]

Conditional Probability

Computing probabilities sequentially using decision trees
Example: Assume the probability of rain on a given day is 0.1. However, if it rains today, the probability of rain tomorrow is 0.15. Similarly, if it does not rain today, the probability of rain tomorrow is 0.05

Check Yourself: Conditional Probability

Correct answer:
Oscar has lost his dog in either forest A with a prior probability 0.4 or in forest B with a prior probability 0.6
If the dog is in forest A and Oscar spends a day searching for it in forest A, the conditional probability that he will find the dog that day is 0.25. Similarly, if the dog is in forest B and Oscar spends the day looking for it there, he will find the dog that day with probability 0.15
In which forest should Oscar look to maximize the probability that he will find the dog on the first day of the search?
Oscar should look in forest A

Check Yourself: Conditional Probability

Correct answer:
In which forest should Oscar look to maximize the probability that he will find the god on the second day of the search?
Oscar should look in

Bayes' Rule

Decision trees are sequential, but set representation is symmetric
We can compute the probability of intersections two ways: $Pr(R_1 \cap R_2) = Pr(R_2|R_1)Pr(R_1) = Pr(R_1|R_2)Pr(R_2)$
Anti-sequential reasoning: what is the probability that it rained yesterday given that it rained today?
Bayes' Rule: $Pr(R_1|R_2) = \frac{Pr(R_2|R_1)Pr(R_1)}{Pr(R_2)}$

Oral AIDS Test: OraSure [OMT]

This test is performed by collecting a sample from your mouth and sending it to the lab for processing
The OraSure test has 98.6% sensitivity and 97.7% specificity
Sensitivity: Pr[posTest|AIDS] = 0.986
Specificity: Pr[negTest|noAIDS] = .977

Check Yourself: Bayes' Rule

Correct answer:
Given that the patient tests positive, what is the probability that the patient has AIDS?

Random Variables

A random variable is the probabilistic analog of a (deterministic) variable
While the value of a deterministic variable is a number, the value of a random variable is drawn from a distribution
Example: Let X represent the result of the toss of a die. Now we can write Pr[X=3] = 1/6

Joint Probability Distributions

Probability laws for multi-dimensional sample spaces are given by joint probability distributions
Let V represent the toss of the first die and W represent the toss of the second die
Pr[V, W] represents the joint probability distribution
Pr[v, w] represents the Pr[V = v and W = w]

Reducing Dimensionality

The dimensionality of a joint probability distribution can be reduced in two very different ways:
- Marginalizing refers to collapsing one or more dimensions by summing over all possible outcomes along those dimensions