

16.410/413 Principles of Autonomy and Decision Making

Problem Set #4, Probability

Due by 5pm, Monday November 15, 2004

Turn in during class or give to Brian O'Conaill (33-336) before 5pm.

Objective

To refresh your probabilistic abilities, and to exercise your ability to perform some simple statistical inferences.

Problem 1: Mean and Variance

1. Let X be a random variable for which $E(X) = \mu$ and $\text{Var}(X) = \sigma^2$. Show that

$$E[(X - c)^2] = (\mu - c)^2 + \sigma^2.$$

2. Suppose that X and Y are independent random variables with finite variances such that $E(X) = E(Y)$. Show that

$$E[(X - Y)^2] = \text{Var}(X) + \text{Var}(Y).$$

3. Suppose that you have two path planners, A and B , and you want to generate the cheapest plan. Planner A generates a plan consisting of 100 moves. Each move is purely deterministic and has a cost of 1.

Planner B generates a plan consisting of 10 moves. Each of B 's moves has a cost of 1 with 85% probability, but a 15% probability of hitting an obstacle, which has a cost of 100. Which planner do you prefer and why?

Problem 2: Probabilistic Inference

A light spectrometer is capable of distinguishing terrestrial rocks from martian meteorites with 80% accuracy. Based on previous tests, NASA knows that 95% of all rocks in the Antarctic are terrestrial.


If the spectrometer takes a single measurement of a rock, and classifies it as a meteorite, what is the posterior probability that this rock is indeed a meteorite?

Let's assume that measurements are independent and identically distributed. How many times should the spectrometer measure the rock and classify it as a meteorite for us to have 95% confidence that this is a meteorite?

Problem 3: Hidden Markov Models

In this problem, you have to help Nick determine his location. Nick is travelling through MIT campus buildings, looking out the window at either Mass Ave. or Vassar St. as he goes. Cambridge streets tend to look very similar, so Nick has a small probability of looking at Mass Ave. and thinking it's Vassar St., and vice versa. Nick is making forward progress, but has some small probability of making random transitions due to absent-mindedness.


The figure below shows the navigation model Nick has of MIT. Each bold circle, labelled with a campus building, is a state. Each non-bold ellipse contains the probability of seeing each of the two streets.



Nick receives the following observations:

Mass Ave.
Vassar St.
Vassar St.

Use the Viterbi algorithm to fill in the lattice below with the α values:



What is the maximum likelihood sequence of states that Nick travelled through?

16.413 ONLY

Problem 4: Gaussian Models


Suppose that you have a robot that knows it is at pose $\mathbf{x} = (x_r, y_r, \theta_r)$, and it knows this quantity exactly. The robot has prior knowledge of the $\mathbf{b} = (x, y)$ position of a nearby beacon: this prior knowledge of the beacon is given by a Gaussian distribution of mean μ and covariance Σ :

$$p(\mathbf{b}) = \frac{1}{\sqrt{2\pi|\Sigma|}} e^{-\frac{1}{2}(\mathbf{b}-\mu)^T \Sigma^{-1}(\mathbf{b}-\mu)} \quad (1)$$

The robot has a range sensor that can measure the range r and bearing θ to a nearby beacon. Each measurement $\mathbf{m} = (r, \theta)$ can be expressed by some function of the robot's position and the true (but unknown) beacon position.

$$\mathbf{m} = h(\mathbf{x}, \mathbf{b})$$

This scenario is shown in the figure below.



Your problem is to determine the probability distribution of sensor measurements b given this model: that is, give an expression for $p(b)$.

- You can assume that $p(b)$ is Gaussian.
- You don't need to know anything about the sensor function h (although you can probably figure it out from the diagram).
- Consider linearizing h using a first-order Taylor series expansion. You can write your solution in terms of the derivatives of h .

Use the back page if necessary.

16.413 ONLY