T-61.3050 PROBLEMS 4/2007

In T1 on 5 October 2007 at 10 o'clock.

You should solve the problems before the problem session and give the solved problems to the assistant. Please write clearly and leave a wide (left or right) margin. The solutions should be stapled together **with a cover sheet** containing your name, student number and the numbers of problems you have solved.

For the problems where a "correct" solution exists (math and algorithm questions) the assistant will present one possible solution during the session. In some cases the questions do not have a single correct answer, but the idea is that you think about the problem and are prepared to discuss it with the assistant and other students during the session.

See `http://www.cis.hut.fi/Opinnot/T-61.3050/2007/problems` for up-to-date information of the problem session.

This problem sheet has two pages.

1. (Alpaydin (2004) Ch 4, Exercise 2) Write the log likelihood for a multinomial sample and show equation 4.6. (Equation 4.6 essentially reads $\hat{\theta}_{kML} = \frac{1}{N}\sum_t x_k^t$.)

2. (Modified from Alpaydin (2004) Ch 4, Exercise 3) Write the code that generates a normal sample with given μ and σ^2, and the code that calculates m, s^2 and the unbiased estimate $\hat{\sigma}^2$ from the sample. Do the same using the Bayes' estimator assuming a prior distribution for μ. Study the performance of the estimates as for different training set sizes.

3. A lighthouse is somewhere off a piece of straight coastline at a position y km along the shore and a distance of 1 km out at sea. The lighthouse has a rotating beacon that emits a series of short highly collimated flashes at random intervals and hence at random azimuths θ^t. The pulses are intercepted on the coast by photo-detectors that record only the fact that a flash aimed at certain position has occurred, but not the angle nor time at which it came. N flashes have so far been recorded at positions $\mathcal{X} = \{x^t\}_{t=1}^N$. See figure 1. (This setup is discussed in detail in "Data Analysis: a Bayesian Tutorial" by Sivia, section 2.4, pages 31–36.)

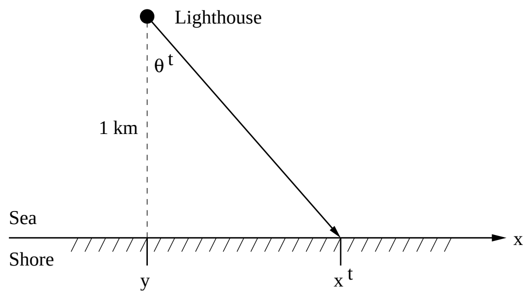

Figure 1: See problem 3 for explanation.

(a) The probability density (likelihood) $p(x \mid y)$ that the next flash is observed at position x, given that the lighthouse is located at y, obeys Cauchy distribution. Write down the Cauchy distribution. (Hint: `http://mathworld.wolfram.com/CauchyDistribution.html`)

(b) Compute analytically the expected position of a flash $E_{p(x|y)}[x]$ and the second moment $E_{p(x|y)}[x^2]$, given that the lighthouse is located at y. (Hint: you should notice something strange.)

(c) Choose a suitable prior probability density $p(y)$ for the position y of the lighthouse and write down a formula for the posterior probability density for the position of the lighthouse, given the prior and the fact that the flashes have been recorded at positions in $\mathcal{X}$.

(d) Download the locations of recorded flashes from the course web site at
`http://www.cis.hut.fi/Opinnot/T-61.3050/2007/problems#4`
What is the mean and median location of the flashes, and the maximum a posteriori (MAP) estimate of the location of the lighthouse y? How can you explain the differences between the three numbers (mean, median, MAP estimate)? (Hint: It may be easier to solve the MAP estimate numerically or plot the posterior distribution and inspect the plot.)

(e) What would you have obtained as an estimate of the location of the lighthouse if you had assumed that the positions of the flashes obey Gaussian distribution?

(f) Where is the lighthouse?