CS761 Spring 2017 Homework 3

Assigned Apr. 6, due Apr. 20

Instructions:

- Homeworks are to be done individually.
- Typeset your homework in latex using this file as template (e.g. use pdflatex). Show your derivations.
- Hand in the compiled pdf (not the latex file) online. Instructions will be provided. We do not accept hand-written homeworks.
- Homework will no longer be accepted once the lecture starts.
- Fill in your name and email below.

Name:
Email:

(4 questions, 25 points each)

1. The Wisconsin State Climatology Office keeps a record on the number of days Lake Mendota was covered by ice at `http://www.aos.wisc.edu/~sco/lakes/Mendota-ice.html`. The article DETERMINING THE ICE COVER ON MADISON LAKES at `http://www.aos.wisc.edu/~sco/lakes/msn-lakes_instruc.html` serves as a fine example of the Wisconsin tradition to integrate science with beer.

   (a) As with any real problems, the data is not as clean nor as organized as one would like for machine learning. Produce a clean data set starting from 1855-56 and ending in 2016-17 for the output variable DAYS. You do not need to attach your data set, but please produce a scatter plot of year vs. DAYs. Show us the sample mean and sample variance (round to 5 digits after decimal point).

   (b) Perform ordinary least squares to estimate a linear model

   $$y = \alpha + \beta x$$

   where y is DAYS and x is the year. For example, for 1855-56 the year is 1855. Show us $\hat{\alpha}, \hat{\beta}$, and an estimate of the standard error on β: $\widehat{s.e}(\hat{\beta})$.

   (c) Perform nonparametric kernel regression using the Nadaraya-Watson estimator on this data set (input: year, output: days). Use the Gaussian kernel. Write your own code for the Nadaraya-Watson estimator. Show us the leave-one-out score (Equation 23 in lecture notes `http://pages.cs.wisc.edu/~jerryzhu/cs761/kde.pdf`) for bandwidth $h = 10^{-1}, 10^{-0.9}, 10^{-0.8}, \ldots, 10^{2}$, respectively.

   (d) For $h = 10^{-1}, 10^{2}$ and the optimal h you found, respectively, plot the function estimated by Nadaraya-Watson.

2. Consider a Gaussian Process $f \sim GP(m, k)$ over $\mathbb{R}$ with mean function

   $$m(x) = \sin(\frac{\pi x}{100}) + \frac{x}{100}$$

   and kernel function

   $$k(x, x') = \frac{1}{16} \exp\left(-\frac{(x - x')^2}{2\sigma^2}\right).$$

   (a) Let $\sigma = 40$ (note: this is the standard deviation, not variance). Approximate the random function f by drawing $f(1), f(2), \ldots, f(100)$ from the appropriate marginal distribution. Plot the curve by connecting the dots. Show six such random functions on the same plot, together with the mean function m.

   (b) Do the same with $\sigma = 10$.

(c) Do the same with $\sigma = 1$.

(d) Let $\sigma = 40$. Now let us observe $f(40) = 0$ and $f(120) = 1$. Now draw f from the posterior Gaussian Process conditioned on these two observations. Again, show six such f from the posterior on the same plot.

(e) Do the same with $\sigma = 10$.

(f) Do the same with $\sigma = 1$.

3. Imagine a stick of length a. On the ground, draw parallel lines a apart. Randomly throw the stick to the ground. Each time, the stick may or may not intersect with a line.

   (a) What is the probability that the stick intersects with a line? Show your work.

   (b) Propose a Monte Carlo method for estimating π based on this.

   (c) Actually perform the experiment. Tell us about it.

4. Consider an undirected graphical model on a binary tree with 15 nodes. Each node takes value in $\{-1, 1\}$. All edges share the same potential function $\psi(u, v) = \exp(\alpha uv)$, where u, v are a pair of parent-child nodes.

   (a) Write down the joint probability distribution defined by this graphical model.

   (b) Let $\alpha = 1$. Let r be the root node and s be the left-most leaf node. Use brute force (enumerating all trees) to compute $p(r \mid s = 1)$.

   (c) Implement Gibbs sampling to estimate $p(r \mid s = 1)$. Start with the all-minus-1 tree except for $s = 1$. Go over levels in top-down order, left-to-right within each level. Discard a burn-in of 10^4 samples. Use the next 10^5 samples for estimation. Do not perform thinning.

   (d) Implement Metropolis-Hastings sampling to estimate $p(r \mid s = 1)$. Clearly define and discuss your proposal distribution (which has to be different than Gibbs). Use the same burn-in and number of samples as above.