CS725: Assignment 2

10 Marks, Due on September 9^{th} in class

1. Solve the optimization problem (11) on page 47 of `http://www.cse.iitb.ac.in/~cs725/notes/classNotes/lecturenotes_cs725_aut11.pdf`. Every step of the derivation and the justification of the solution (the corresponding theorem from `http://www.cse.iitb.ac.in/~cs725/notes/classNotes/BasicsOfConvexOptimization.pdf`) should be clearly specified.

   Now systematically solve problem (12). Does it have a global optimal? Can you derive the global optimal in closed form? Please provide precise explanations. How does this solution (if one exists) compare with the MAP or Bayes estimates for the multinomial Naive Bayes classifier? These estimates were specified in the class. The general definition of MAP and Bayes estimates are also given in the table on page 53 of `http://www.cse.iitb.ac.in/~cs725/notes/classNotes/lecturenote_2010.pdf`.

   (5 Marks)

2. Suppose we believe our one-dimensional samples come from a uniform distribution

$$p(x|\mu) = U(0,\mu) = \frac{1}{\mu}\ 0 \leq x \leq \mu \ \text{ and } \ 0 \ \text{ otherwise}$$

   $U(0,\mu)$ is the uniform distribution.

   In a Bayesian Inferencing setting, let $p(\mu) = U(0, 200)$.

   Now consider the first data point (random sample) $x_1 = 50$. What is the posterior distribution $p(\mu\,|x_1)$? Consider now a second data point $x_2 = 75$. What is the new posterior distribution $p(\mu\,|x_1, x_2)$? In general, what is the distribution $p(\mu\,|x_1, x_2, \ldots, x_m)$ after having observed m data points $\{x_1, x_2, \ldots, x_m\}$?

   What is the maximum likelihood estimate of μ for the distribution $p(x|\mu)$ specified above, given data $\{x_1, x_2, \ldots, x_m\}$?

   (**5 Marks**)