Assignment 7

1. Implement Hamiltonian Monte Carlo (HMC) method in which "momentum" variables are sampled and updated in the usual way (presented in lecture notes and Neal). Use your algorithm to sample from one-dimensional exponential distribution $X \sim Exp(\lambda)$, where $\lambda = 2$. Use $L = 25$ for the number of leapfrog steps. Try picking the step size of $\varepsilon = 0.1$ (see the note above). You must be sampling only for $x \geq 0$, because $Exp(\cdot)$ is only defined there. (Refer to Lecture 1, p.38, for the exponential distribution.) In order to sample only on the support of $Exp(\cdot)$ you must impose constraints: implement these in your HMC algorithm.

   (i) (w=2) First you need to determine $U(q)$. **Derive** the correct form $U(q) = \lambda q - \ln\,\lambda$.
   (ii) (w=1) What is the minimal form of $U(q)$ that does the job for HMC. Why?
   (iii) (w=1) Write down the constraints for this case. You can do this either by writing lines of pseudo code or by indicating clearly where the constraints are in your HMC algorithm.
   (iv) (w=3) Simulate $Exp(2)$ using your HMC. Use 10^6 samples. Plot the distribution as a histogram using 50 bins together with the pdf of $Exp(2)$ (linear scale/axes are fine).
   (v) (w=1) Report the value of ε you used (see below).

   If you succeed in your implementation, the histogram should follow the pdf quite well (of course the binning is rather sparse, so you'll see this in the graph) and that's it, you're done. **If you don't succeed in getting the pdf too well, you can try decreasing ε. Report the ε that will give a reasonable alignment.** (You'll get slightly less points for having to use $\varepsilon < 0.1$.)

   Hint. When deriving $U(q)$ you will find the identity
   $A\exp(-B) = \exp\left[\ln(A) - B\right]$ useful.

   Note. The usual warning of python functions: Check the parametrization of the library functions you decide to use.