ECE / CS 6501: Geometry of Data **Due 12:00 PM (noon), Sat 12/14**

Homework 5: Sampling Methods

Instructions: Submit a single Jupyter notebook (.ipynb) of your work to Canvas by 11:59pm on the due date. All code should be written in Python. **Be sure to show all the work involved in deriving your answers! If you just give a final answer without explanation, you may not receive credit for that question.**

You may discuss the concepts with your classmates, but write up the answers entirely on your own. Do not look at another student's answers, do not copy answers from the internet or other sources, and do not show your answers to anyone. **Remember to cite any sources you use, including any prompts given to AI chatbots.**

1. Download the Jupyter notebook `HW5-Starter.ipynb`. This defines an unormalized probability density function in $\mathbb{R}^2$ from which you will sample. First, implement an independence sampler with normal proposal density $g(y) = \mathcal{N}(y; 0, \sigma^2)$.

   (a) Run your sampler for 10,000 proposal samples. Keep track of the acceptance rate (proportion of proposed samples that are accepted), and report this.

   (b) Plot accepted samples over top of the density function (you can use the contour plot of the density provided in the starter code).

   (c) Repeat the experiment 3 times with different σ^2 values. What value do you find works best? Discuss the effect of σ^2 on the sampling process, that is, what happens if σ^2 is "too small" or "too large"?

2. Implement a random walk Metropolis sampler, with normal proposal density $g(y \mid x_{k-1}) = \mathcal{N}(y; x_{k-1}, \sigma^2)$. Repeat the same steps (a) - (c) from problem #1.

3. Implement a Langevin Monte Carlo sampler **or** a Hamiltonian Monte Carlo sampler and repeat steps (a) and (b) from problem #1. (You don't have to do step (c) - you can just tune the parameters for a single run of 10,000.)