Assignment 1.

Given October 15, due October 29.

Objective: To explore simple sampling and Monte Carlo estimation.

(1) Write a program that uses the Box Muller algorithm to generate independent standard normal random variables. Test whether the program produces the density by generating lots of gaussians and putting them into bins. The bin size is Δx. The k^{th} bin is $B_k = (x_k - \Delta x/2, x_k + \Delta x/2)$. As always, $x_k = k\Delta x$. If we make N samples, the k^{th} "bin count", N_k, is the number of X_j that fall in B_k for $j = 1, \ldots, N$. The expected bin counts are

$$\mathbf{E}\left[N_k\right] = N \int_{B_k} f(x)dx \quad , \tag{1}$$

where $f(x)$ is the probability density function for a standard normal. If the bins are small enough, the integrals in (1) can be estimated using taylor series expansions of f about x_k. Probably, one or two terms suffice. Make a plot showing actual bin counts and expected bin counts. If Δx is fixed and N is made large enough, these plots should agree to within small statistical errors.

To test whether the Box Muller algorithm produces independent gaussians, make pairs (X_j, Y_j), make bins in 2D, and compare empirical bin counts to expected bin counts. Use the joint density function for independent standard normals. Bin B_{kl} will be a square centered at (x_k, y_l) with side Δx.

(2) Make a standard Monte Carlo estimate of $\mathbf{E}[X^{2m}]$ for $m = 1, 2, 3, 4$, where X is a standard normal. The exact answers are 1, 3, 15, and $105 = 7 \times 5 \times 3 \times 1$. Use standard normals made by Box-Muller. Make error bars. How large a sample size is needed so that the one standard deviation error bar is 1% of the exact answer? Why is it harder when m is larger? Hint: plot the integrand.

(3) Try computing $\mathbf{E}[X^8]$ by using importance sampling and gaussians with variance $\sigma > 1$. What value of σ works best?