Homework #2: Monte Carlo Integration
Due Date: Thursday, November 13

Problem 1. Variance

Prove the identity

$$V[F] = E_X V_Y[F] + V_X E_Y[F] ,$$

where $F = f(X,Y)$ is a function of two random variables X and Y. Here E_X and E_Y denote the expected value with respect to X or Y only (that is, the other variable is treated as a constant), so that $E[F] = E_X E_Y[F]$. Recall that we used this identity in class to analyze the variance reduction due to stratified sampling and the method of expected values. It is very useful other Monte Carlo contexts as well.

Problem 2. Importance Sampling

In this problem we consider how to generate samples according to various distributions over the hemisphere. This is useful for evaluating light reflection from surfaces (part (e)).

2 (a). Using the inversion method, show how to sample from the density function

$$p(\omega) = C \cos^k(\theta) , \tag{1}$$

where ω is a point on the hemisphere with polar coordinates (θ, ϕ), and $0 \le \theta \le \pi/2$. Part of your job is to determine the normalization constant C. [Useful facts: $d\omega = \sin(\theta)\, d\theta\, d\phi$ and $p(\omega)\, d\omega = p(\theta, \phi)\, d\theta\, d\phi$.]

[Remark: this result can be used to sample from the *reciprocal Phong BRDF*,

$$f_{\mathrm{r}}(\omega_{\mathrm{i}} \to \omega_{\mathrm{o}}) = K_{\mathrm{s}}\, (\omega_{\mathrm{i}} \cdot \omega_{\mathrm{m}})^k .$$

Here K_{s} is a constant in the range $[0, 1]$, and ω_{m} is the *mirror direction* ω_{m} obtained by reflecting ω_{o} around the surface normal $\mathbf{N}$. Given the direction ω_{o} (toward the camera), the objective is to sample ω_{i} with probability proportional to $f_{\mathrm{r}}(\omega_{\mathrm{i}} \to \omega_{\mathrm{o}})$. This is equivalent to sampling from the density above, where θ is the angle between ω_{i} and ω_{m}.]

2 (b). The technique derived in part (a) does not work well for stratified sampling, since it leads to strata with elongated shapes. This problem investigates a better transformation for sampling the density (1), that yields strata with more uniform shapes.

Suppose that we start by using Shirley's technique to map a uniformly distributed point (U_1, U_2) in the unit square onto a uniformly distributed point (R, ϕ) in the unit disc. Given this, your task is to derive a transformation of the form $\theta = f(R)$ from the unit disc to the hemisphere, such that the transformed samples $\omega \equiv (\theta, \phi)$ have the density given by equation (1).

[Hint: if a random variable X has the cumulative distribution function $P(x)$, then $U = P(X)$ is a uniform random variable on $[0, 1]$. Thus you can solve for θ by equating the cumulative distribution functions of θ and R.]

2 (c). The rest of this problem deals with BRDF's that are based on *microfacet theory.* These BRDF's have the form

$$f_{\mathrm{r}}(\omega_{\mathrm{i}} \to \omega_{\mathrm{o}}) = K_{\mathrm{s}} \, \frac{D(\omega_{\mathrm{h}})}{4\cos(\theta_{\mathrm{i}})\cos(\theta_{\mathrm{o}})} \,, \tag{2}$$

where $\theta_{\mathrm{i}}, \theta_{\mathrm{o}}$ are the angles between $\omega_{\mathrm{i}}, \omega_{\mathrm{o}}$ and the normal $\mathbf{N}$, while ω_{h} is the *halfway vector*

$$\omega_{\mathrm{h}} = \frac{\omega_{\mathrm{i}} + \omega_{\mathrm{o}}}{\|\omega_{\mathrm{i}} + \omega_{\mathrm{o}}\|} \,.$$

The function $D(\omega_{\mathrm{h}})$ is the *microfacet distribution function*, which describes the microscale geometry of the surface. This function is normalized so that

$$\int_{\Omega} D(\omega_{\mathrm{h}}) \cos(\alpha) \, d\omega_{\mathrm{h}} = 1 \,,$$

where α is the angle between ω_{h} and $\mathbf{N}$, and Ω is the hemisphere of directions ω such that $\omega \cdot \mathbf{N} \geq 0$.

One common microfacet distribution is the *Beckmann distribution*,

$$D(\omega_{\mathrm{h}}) = C \frac{1}{m^2 \cos^4(\alpha)} e^{-[\tan(\alpha)/m]^2} \,.$$

This distribution can be integrated analytically, which allows the inversion method to be used for sampling. Your task for this part of the problem is to determine the normalization constant C, and show how to sample from the density

$$p_{\mathrm{h}}(\omega_{\mathrm{h}}) = D(\omega_{\mathrm{h}}) \cos(\alpha) \tag{3}$$

where $0 \leq \alpha \leq \pi/2$. [Useful fact: $d\tan(\alpha)/d\alpha = 1/\cos^2(\alpha)$.]

2 (d). We now consider how to sample the microfacet BRDF (2). Given ω_{o}, the goal is to randomly generate a direction ω_{i} according to a density function that is approximately proportional to $f_{\mathrm{r}}(\omega_{\mathrm{i}} \to \omega_{\mathrm{o}})$.

For a microfacet BRDF, this can be done by first sampling the halfway vector ω_{h}, according to some density $p_{\mathrm{h}}(\omega_{\mathrm{h}})$ that approximates $D(\omega_{\mathrm{h}})$. For example, the density (3) is a convenient approximation of this type (the extra factor of $\cos(\alpha)$ allows sampling using the inversion method). Next, the vector ω_{i} is obtained by reflecting ω_{o} around the vector ω_{h} (so that by construction, ω_{h} is the halfway vector for ω_{i} and ω_{o}).

Show that when ω_{i} is generated in this way, its density function is

$$p_{\mathrm{i}}(\omega_{\mathrm{i}}) = p_{\mathrm{h}}(\omega_{\mathrm{h}}) / (4\cos(\theta')) \,,$$

where $\theta' \in [0, \pi/2]$ is the angle between ω_{o} and ω_{h}.

[Hint: consider a differential solid angle $d\omega_{\mathrm{h}} = \sin(\theta') \, d\theta' \, d\phi'$ in the spherical coordinate system centered around ω_{o}. Work out the corresponding solid angle $d\omega_{\mathrm{i}}$ after the transformation above, and use the fact that $p_{\mathrm{i}}(\omega_{\mathrm{i}}) \, d\omega_{\mathrm{i}} = p_{\mathrm{h}}(\omega_{\mathrm{h}}) \, d\omega_{\mathrm{h}}$.]

2 (e). Finally, we put together the previous results to show how they can be used in a rendering system. The radiance $L_{\mathrm{o}}(\omega_{\mathrm{o}})$ reflected from a surface is given by the integral

$$L_{\mathrm{o}}(\omega_{\mathrm{o}}) = \int_{\Omega} L_{\mathrm{i}}(\omega_{\mathrm{i}}) f_{\mathrm{r}}(\omega_{\mathrm{i}} \to \omega_{\mathrm{o}}) \cos(\theta_{\mathrm{i}})\, d\omega_{\mathrm{i}} \;, \tag{4}$$

where $L_{\mathrm{i}}(\omega_{\mathrm{i}})$ is the incident radiance from direction ω_{i}. Let f_{r} be the Beckmann BRDF (2) from part (c), and suppose that a single sample ω_{i} is taken according to the method in part (d).

Your task is to write an expression for the sample value $Y = f(\omega_{\mathrm{i}})/p(\omega_{\mathrm{i}})$, where f is the integrand in equation (4), and p is the density function for ω_{i}. This is the quantity that would actually be computed and returned in a Monte Carlo calculation (recalling that Y is an unbiased estimator of the integral (4)).

Problem 3. Stratified Sampling

The goal of this problem is to determine the convergence rate of stratified sampling for discontinuous functions. Such functions are common in graphics, mainly due to visibility: for example, images are discontinuous in general at the boundaries between objects.

Consider a function $f(\mathbf{x})$ defined on $[0,1] \times [0,1]$, where $\mathbf{x} = (x, y)$. Suppose that we stratify this domain into an $n \times n$ grid of cells

$$C_{i,j} = [\frac{i-1}{n}, \frac{i}{n}] \times [\frac{j-1}{n}, \frac{j}{n}] \qquad \text{where } 1 \le i, j \le n \;,$$

and then we sample one point $\mathbf{X}_{i,j}$ uniformly within each of the $N = n^2$ cells. Finally, we compute the usual estimate

$$F_N = (1/N) \sum_{i=1}^{n} \sum_{j=1}^{n} f(\mathbf{X}_{i,j}) \;.$$

Answer all parts of this problem using the "Big-Oh" notation, e.g. $O(1)$ or $O(N^{-1})$. Rigorous justifications are not required.

3 (a). Let f be a piecewise constant function such that $f(x, y) = 1$ inside a finite number of polygons, and $f(x, y) = 0$ otherwise. Consider a cell $C_{i,j}$ that contains a discontinuity (i.e. one that is crossed by an edge). What is the variance of $Z_{i,j} = f(\mathbf{X}_{i,j})$?

3 (b). How many cells contain a discontinuity, as a function of N? [Hint: the total length of the discontinuity edges is a constant. What is the average length of a line segment crossing a cell?]

3 (c). What is the variance of F_N? What is the corresponding standard deviation? How does this compare to unstratified sampling?

3 (d). What are the corresponding results in s dimensions? [Start by thinking about $s = 3$.] How much does stratified sampling help as $s \to \infty$?

Problem 4. Latin Hypercube Sampling

The goal of this problem is to determine how well Latin Hypercube sampling works for pixels containing edges. Let the pixel region be $[0,1] \times [0,1]$, and suppose that it is crossed by a single edge passing through the pixel center. Specifically, consider the image function

$$f(x,y) = \begin{cases} 1 & \text{if } y \leq (x-h)/(1-2h), \\ 0 & \text{otherwise,} \end{cases}$$

where $0 \leq h \leq 1/2$. This corresponds to an edge passing through the points $(h,0)$ and $(1/2,1/2)$.

4 (a). Find the *anova decomposition* of $f(x,y)$, as described in class:

$$\begin{aligned} f(x,y) &= I + f_1(x) + f_2(y) + f_{12}(x,y) \ , && (5) \\ \text{where} \qquad I &= \iint f(x,y)\,dx\,dy \ , \\ f_1(x) &= \int [f(x,y) - I]\,dy \ , \\ f_2(y) &= \int [f(x,y) - I]\,dx \ , && (6) \end{aligned}$$

and $f_{12}(x,y)$ is determined by (5). Recall that these functions are orthogonal, and satisfy

$$\iint f^2\,dx\,dy = \iint (I^2 + f_1^2 + f_2^2 + f_{12}^2)\,dx\,dy \ .$$

4 (b). Let F_N be the estimator which uses N independent uniform samples, and let F'_N be the estimator which uses N Latin hypercube samples. Show that

$$\lim_{N\to\infty} \frac{V[F'_N]}{V[F_N]} = (1-4h^2)/3 \ .$$

[Hint: recall from class that the variance of Latin hypercube sampling in two dimensions is

$$V[F'_N] = \frac{1}{N} \iint f_{12}^2(x,y)\,dx\,dy + o(1/N) \ .$$

Here $f_{12}(x,y)$ is the component of $f(x,y)$ that is orthogonal to all *additive functions* (i.e. functions of the form $g_1(x) + g_2(y)$), and $o(1/N)$ denotes a function that decreases more rapidly than $1/N$.]

4 (c). Briefly interpret the result of part (b). How does Latin hypercube sampling compare to (i) uniform independent samples, and (ii) stratified sampling?