CGR 2025 Lab 3: Numerical Integration

Duration: 1 hour
Tools: C++ compiler, any plotting tool

Objective

To estimate soft shadows in a given flatland scene via 1D numerical integration.

Task 1: Estimate irradiance

Given the following setup, write a function called Irradiance(x, O, l0, l1, ho, hl, N) with inputs

- O is a vector of pairs of floats (O_{i0}, O_{i1}) for i=1 to n obstacles
- Floats l0 and l1 as the starting and ending X-coordinates of the light
- Floats ho and hl which are the Y-coordinates of the occluders and the light
- Integer N is the number of Monte Carlo samples you will use

and which outputs a single float which is the irradiance at x. Estimate irradiance E(x) via naive Monte Carlo integration of the radiance L over directions (the integral is given below). That is, draw uniformly random samples of the angle and average the values of L along those directions at each x. O, l0, l1, h0 and hl will be provided by the tutor. They will be different for each group.

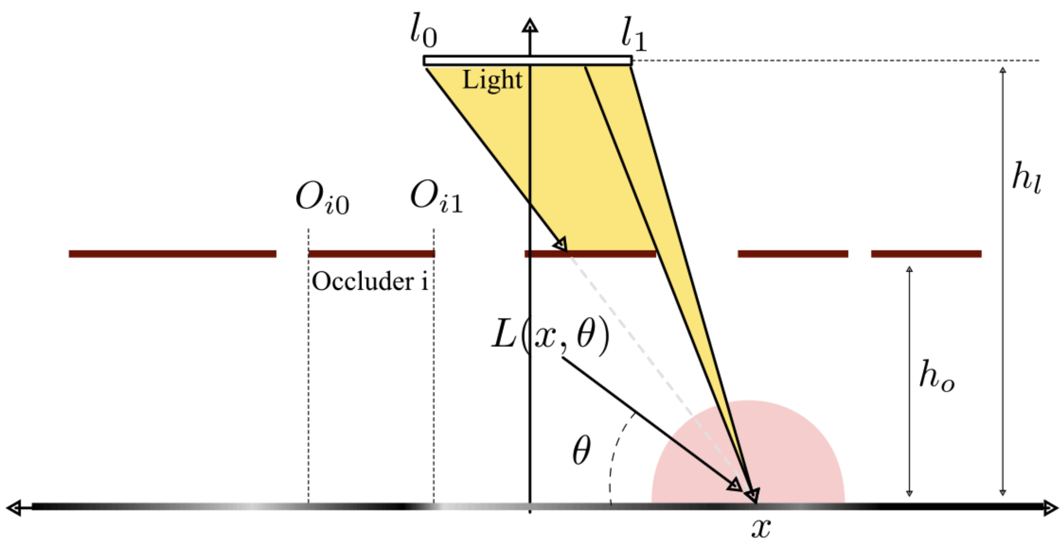

$$E(x) = \int_0^{\pi} L(x,\theta)\ \cos(\pi/2 - \theta)\ \mathrm{d}\theta$$

Once you have implemented the function, write the following to a file and plot using any tool:

a. plot E(x) vs x, for M (e.g. M=100) regularly sampled values of x spanning the minimum and maximum value in O. (use a fixed N, say 100)
b. plot V(x) vs x, where V(x) is the variance of estimates of E(x). Estimate the variance using K (e.g. K=50) samples of E(x) at each x.

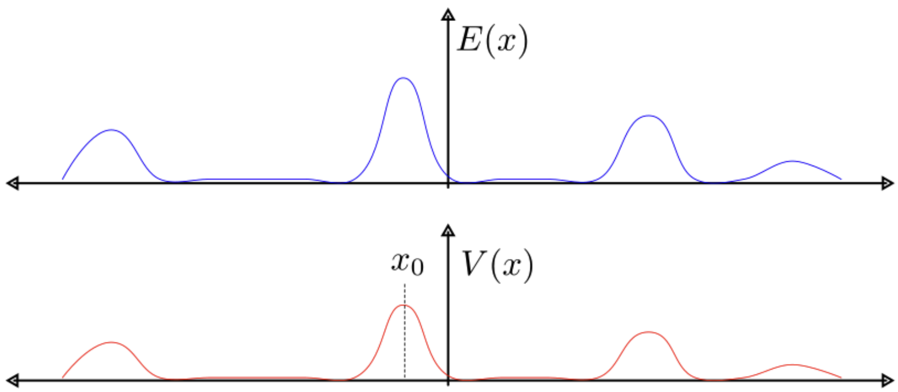

Task 2: Importance sampling

From the variance plot in 1.b, choose the corresponding x with the highest (or one of the highest) V(x) and design an importance sampling scheme just for that x location– let us call that x_0. That is, while the function Irradiance sampled uniformly over angles, write a new function ImportanceXIrradiance that will use a non-uniform pdf to use more samples with incoming light. You are free to manually craft this. For example, you can use a mixture of Gaussians where you place the centres to align with the gaps in the occluder as seen from x_0. ImportanceX0Irradiance will have the same inputs and outputs as Irradiance, including x (although you will call it with $x=x_0$). Your manually crafted pdf can be contained within the function.

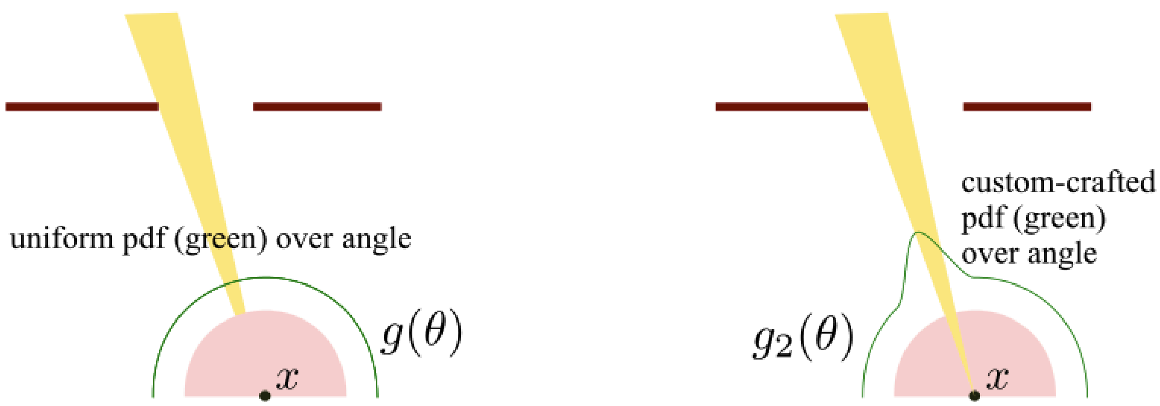

Repeat the plots in Task 1 and compare the plots produced using Irradiance vs ImportanceX0Irradiance.

For only one location x_0, plot the variance V(x0) vs N using a log-log scale. Compare the variance plot produced by the naive estimator against the importance sampled estimator. This is called a convergence plot, as described in class.

Marking

1. Correct plot of E(x) vs x using Irradiance – 1 mark
2. Correct plot of V(x) vs x using Irradiance – 1 mark
3. Correct plot of E(x) vs x using ImportanceX0Irradiance – 1 mark
4. Correct plot of V(x) vs x using ImportanceX0Irradiance – 1 mark
5. Explanation of the convergence plot – 1 mark