

Data Science and Statistical Computing

Tutorial 6, Week 21

Q1

Let $f(x)$ be the Uniform(0, 1) distribution, $\tilde{f}(x)$ the Uniform(0, $\frac{1}{2}$) distribution, and consider the functional $g(x) = x^2$. Show that,

$$\mathbb{E}_{\tilde{f}} \left[\frac{g(X)f(X)}{\tilde{f}(X)} \right] \neq \mathbb{E}_f [g(X)]$$

Note: This is to illustrate the importance of the requirement that $\tilde{f}(\cdot)$ be a pdf such that $\tilde{f}(x) > 0$ whenever $g(x)f(x) \neq 0$ in order for importance sampling to be valid! This condition is clearly violated for the choices of $f(\cdot)$, $\tilde{f}(\cdot)$ and $g(\cdot)$ above.

Q2

We consider a very simple problem to keep the algebra easy(ish), but this is still a lengthy question. Please remember, you would not really need importance sampling in such a simple problem, we do it here to highlight some interesting points!

Let,

$$f(x) = \begin{cases} 2x & \text{for } x \in [0, 1] \\ 0 & \text{otherwise} \end{cases}$$

- Compute $\mu = \mathbb{E}_f[X]$ exactly by solving the relevant integral.
- Assume you have n Monte Carlo simulations $\{x_1, \dots, x_n\}$ from $f(\cdot)$. Write down the equation for the Monte Carlo estimator of μ , $\hat{\mu}_n$, as well as the variance of $\hat{\mu}_n$ as a function of n .

You decide to use a proposal distribution from the family of distributions having probability density function of the form:

$$\tilde{f}(x) = \begin{cases} \alpha x^{\alpha-1} & \text{for } x \in [0, 1] \\ 0 & \text{otherwise} \end{cases}$$

where $\alpha > 0$ is a parameter (we know this distribution is easy to inverse sample from tutorial 4, Q3).

- Assume you have n simulations $\{x_1, \dots, x_n\}$ from $\tilde{f}(\cdot)$. Write down the equation for the Importance Sampling estimator of μ , $\hat{\mu}_n$, and derive the variance of $\hat{\mu}_n$ as a function of α and n (Hint: Theorem 5.4).
- For what range of choices of $\alpha > 0$ does the Importance Sampling estimator have lower variance than the standard Monte Carlo estimator?
- By using the formula for $\tilde{f}_{\text{opt}}(x)$ in lectures, write down the *optimal* proposal distribution for this problem.

Does it belong to the family $\tilde{f}(\cdot | \alpha)$ above? If so, determine the value for α and use part (c) to write down the estimator variance. Does this variance make sense?!

Hint: take any single random draw $x_1 \in [0, 1]$ and compute $\hat{\mu}_1$ by importance sampling with the optimal proposal.

Q3

Let $f(x) = \lambda e^{-\lambda x}$, $x \geq 0$ for some fixed parameter value $\lambda > 0$. If you use proposal density $\tilde{f}(x) = \eta e^{-\eta x}$, $x \geq 0$, what are the weights for an importance sample, x_i , drawn from $\tilde{f}$? For what values of η are you guaranteed the weights are bounded?

Note: Think about what can happen if the weights are not bounded. It is often the case that if the weights are unbounded then the variance of the estimator is infinite, but requiring the weights to be bounded restricts us to using importance sampling in the same situations as rejection sampling (there are approaches to manage this, but it is beyond the scope of this second year course).