Discussion 07

Probability, Monte Carlo & Global Illumination

Computer Graphics and Imaging
UC Berkeley CS 184/284A

Discussion 7 Announcements

- Homework 2 was just due
- Homework 3 checkpoint due 7/22, full assignment due 7/30
    - We will not accept extension requests after 7/30 at 11:59PM
- Final project groups and proposals due by 7/27

Probability Review

Discrete Random Variable

- Random variable, *X*.

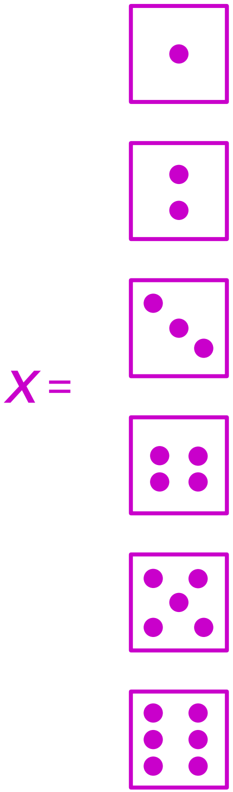

Discrete Random Variable

- Random variable, *X*.
- Probability Mass Function, $P[X = x]$.

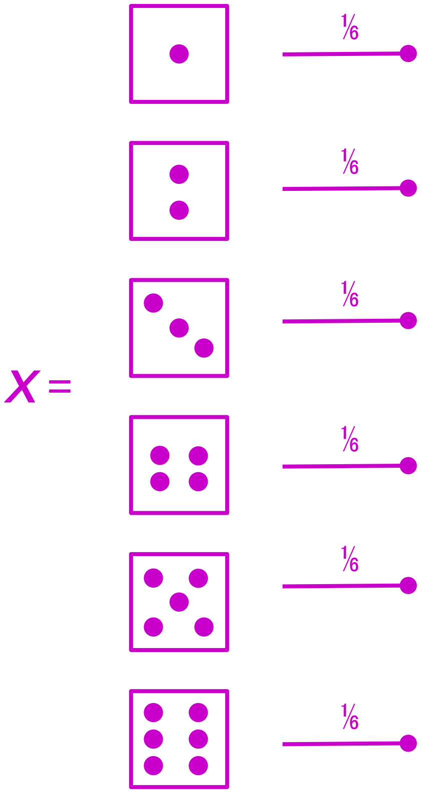

Discrete Random Variable

- Random variable, *X*.
- Probability Mass Function, *P[X = x]*.
- Expectation, *E[X]*.

$$E[X] = \sum_i x_i p_i$$

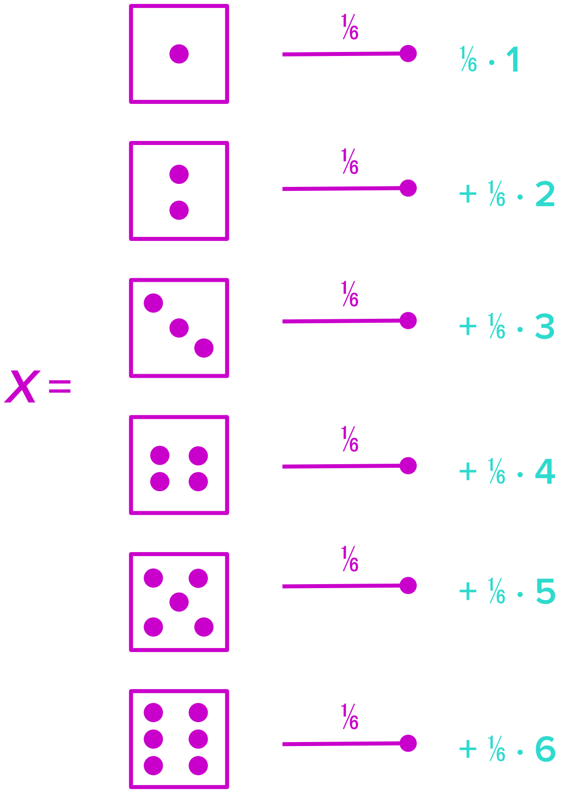

Discrete Random Variable

- Random variable, *X*.
- Probability Mass Function, *P[X = x]*.
- Expectation, *E[X]*.

$$E[X] = \sum_i x_i p_i$$

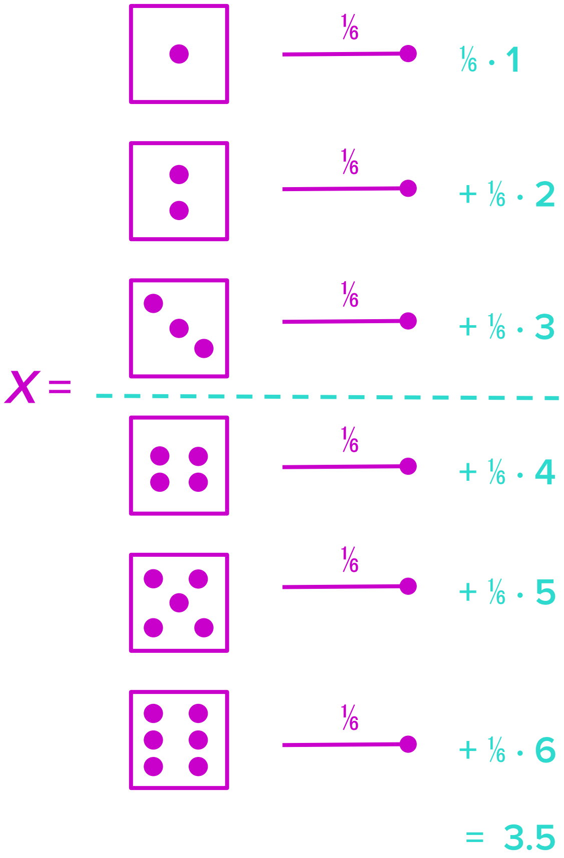

Continuous Random Variable

- Random variable, *X*.

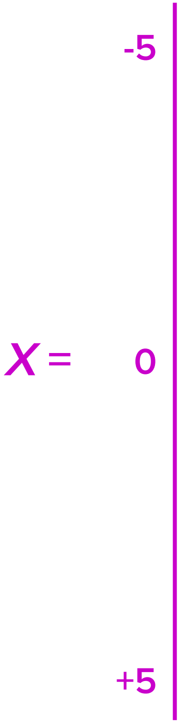

Continuous Random Variable

- Random variable, *X*.
- Probability Distribution Function, *p(x)*.

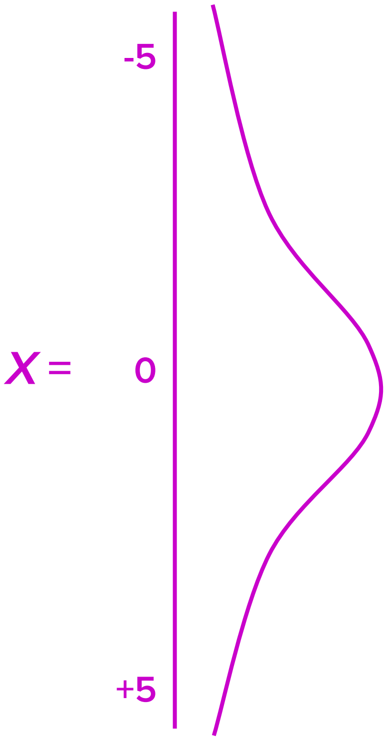

Continuous Random Variable

- Random variable, *X*.
- Probability Distribution Function, *p(x)*.
- Cumulative Distribution Function, ***F(x)***.

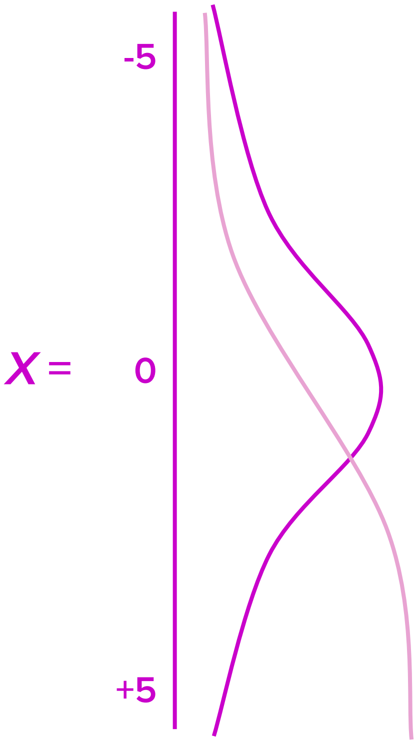

Continuous Random Variable

- Random variable, *X*.
- Probability Distribution Function, *p(x)*.
- Cumulative Distribution Function, ***F(x)***.

$$F(x) = \int_{-\infty}^{x} p(t)\mathrm{d}t$$

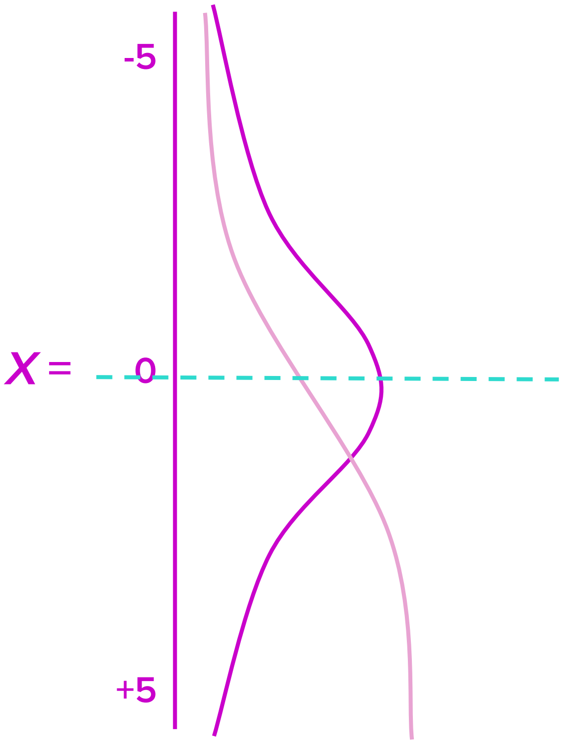

Continuous Random Variable

- Random variable, *X*.
- Probability Distribution Function, *p(x)*.
- Cumulative Distribution Function, ***F(x)***.

$$F(x) = \int_{-\infty}^{x} p(t)\mathrm{d}t$$

- Expectation, *E[X]*.

$$E[X] = \int xp(x)dx$$

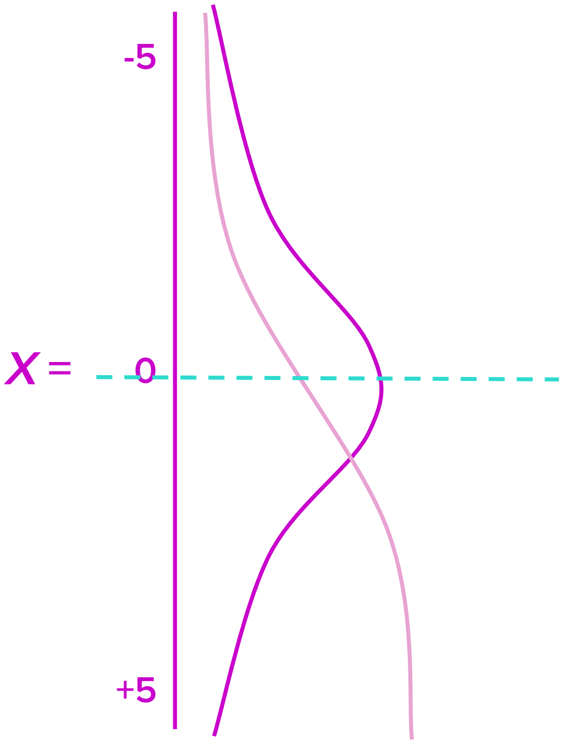

Inversion Method

Inversion Method

- Say our computer can sample values for *U*, which is *Uniform[0, 1]*.
- How to sample values from any PDF, *p(x)*?

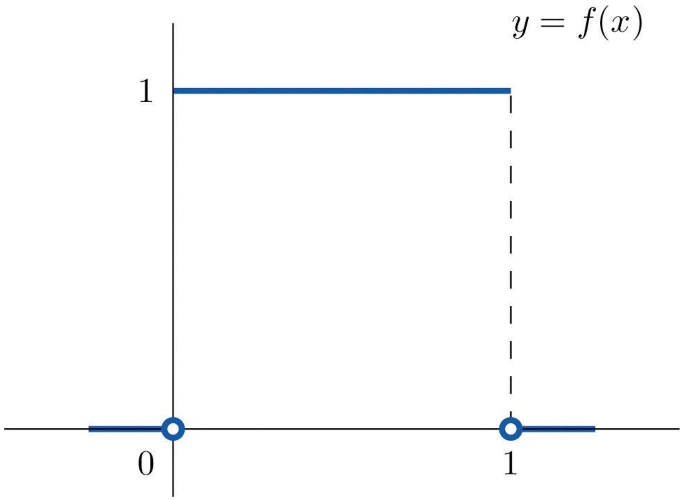

Inversion Method

- Say our computer can sample values for *U*, which is *Uniform[0, 1]*.
- How to sample values from any PDF, *p(x)*?

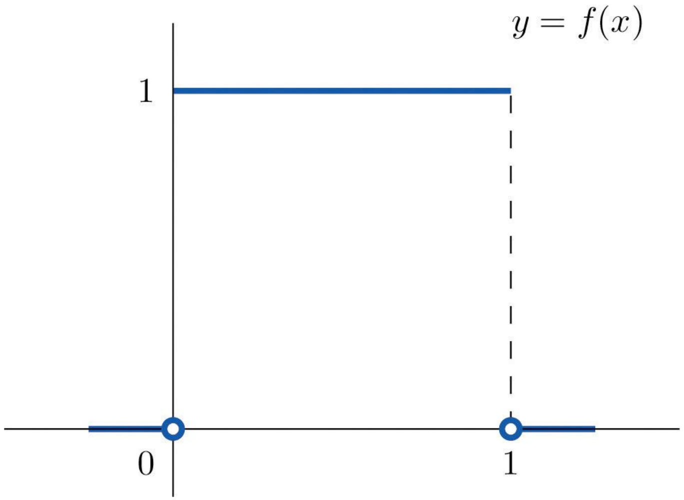

1. Calculate the CDF, *F(x)*, by integrating *p(x)*.
2. Invert the CDF → $F^{-1}(x)$.
3. Sampling *X* according to *p(x)* is achieved by sampling *U* → $X = F^{-1}(U)$.

Inversion Method

We want to create a cumulative distribution function (CDF) to sample from! The benefit of the inversion method (especially for Monte Carlo 👀 – wink wonk!!) is that we can generate a probability distribution using uniform random numbers.

Steps

1. Calculate the CDF, $F(x)$, by integrating $p(x)$.
2. Invert the CDF ➔ $F^{-1}(x)$.
3. Sampling X according to $p(x)$ is achieved by sampling U ➔ $X = F^{-1}(U)$.

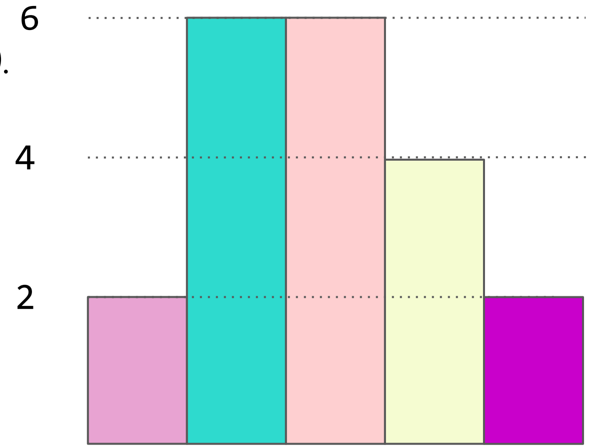

Inversion Method

Compute PDF (probability density function) by normalizing

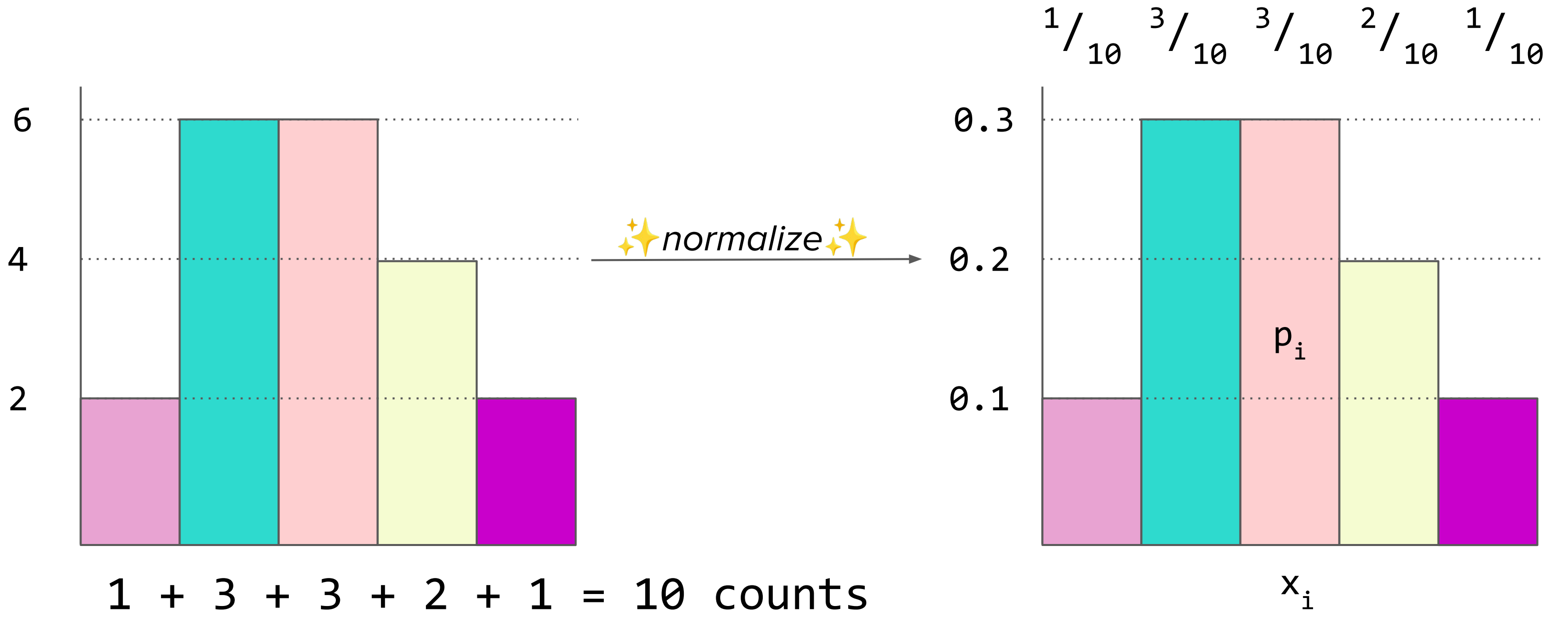

Inversion Method

Compute CDF (cumulative density function)

`P(x) = Pr(X < x)`

As a hint, you can think of the CDF as the "area under the curve" of a PDF! This means to get a CDF, you can commonly integrate the PDF.

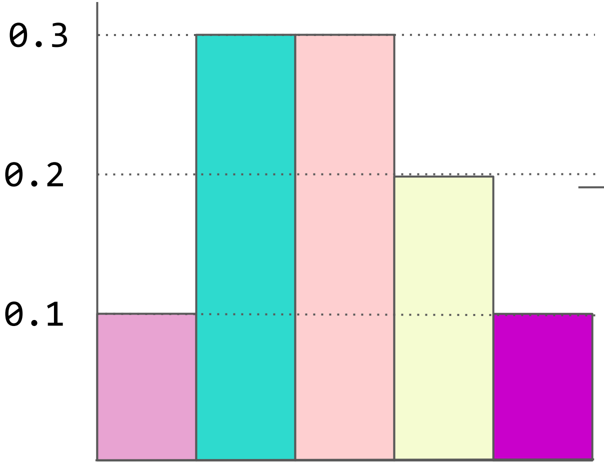

compute CDF from PDF

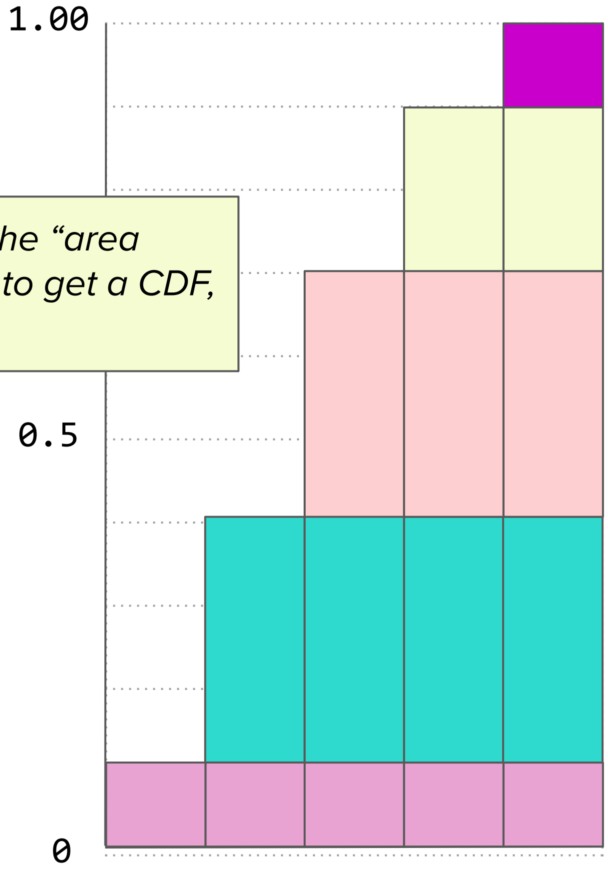

Inversion Method

Step 3 - Invert the CDF, and solve for $F^{-1}(\xi)$

$F^{-1}(\xi_1) =$

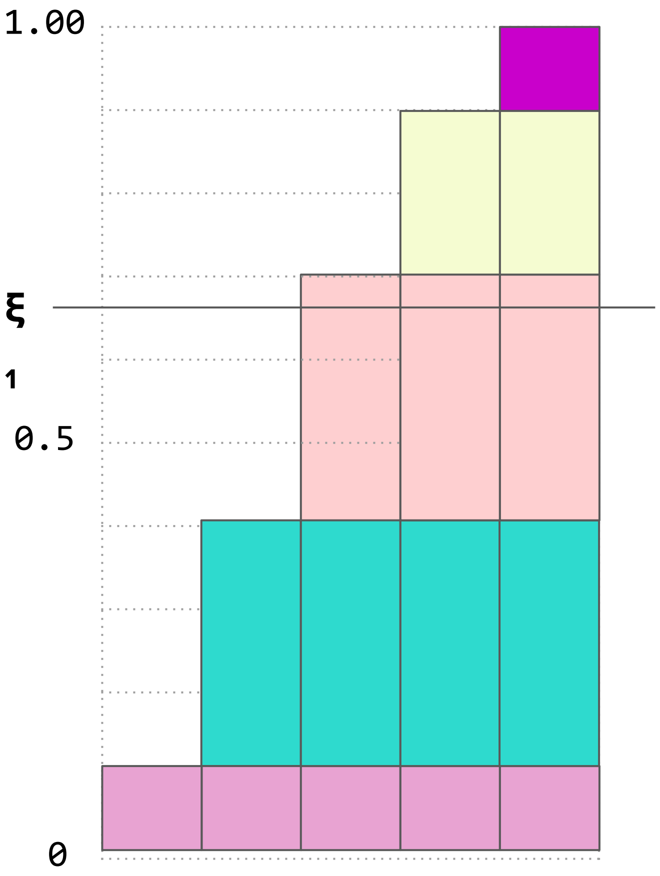

Inversion Method

Step 3 - Invert the CDF, and solve for $F^{-1}(\xi)$

$F^{-1}(\xi_1)$ =

$F^{-1}(\xi_2)$ =

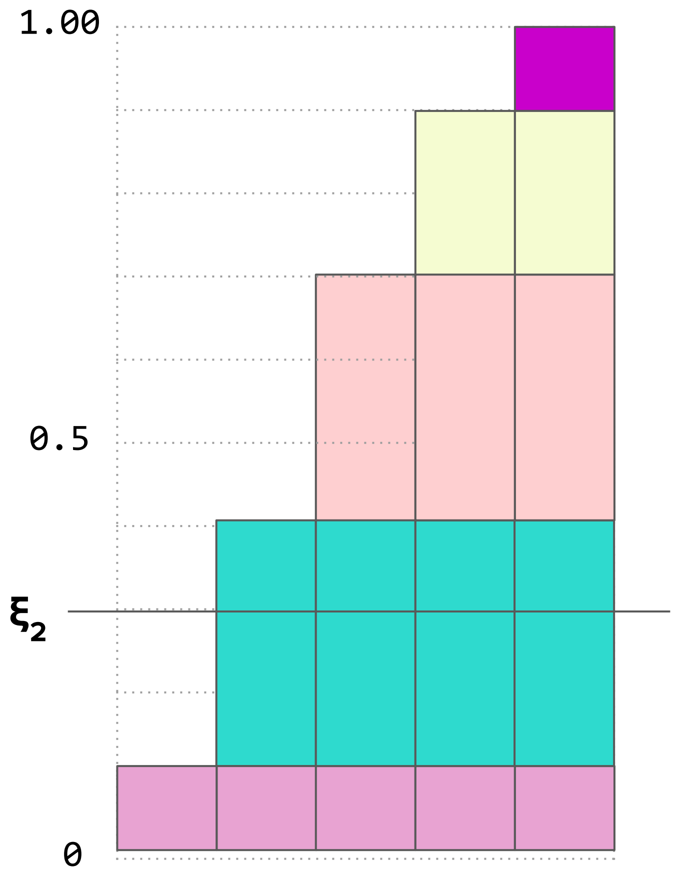

Inversion Method - Review

1. Compute PDF by normalizing
2. Compute CDF
3. Sample from p(x)

Called the "inversion method"

Cumulative probability distribution function

$$P(x) = \Pr(X < x)$$

Construction of samples:

Solve for $x = P^{-1}(\xi)$

Must know the formula for:

1. The integral of $p(x)$
2. The inverse function $P^{-1}(x)$

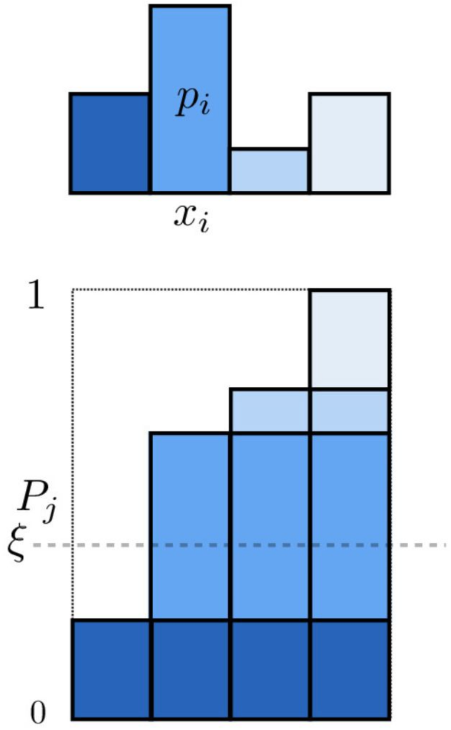

Worksheet Question 1

1. Inversion Method

Given a uniform random variable U in the interval $[0, 1]$, we can generate a random variable from any other one dimensional distribution using its cumulative distribution function: $X = F^{-1}(U)$. This is how we choose sample points when running a ray tracing algorithm.

1. What function of U will return a sample from the exponential distribution (with parameter λ)? This distribution has density $p_\lambda(x) = \lambda e^{-\lambda x}$, and is defined for $x \geq 0$.

As a reminder, the steps of the inversion method are

1. Compute CDF (cumulative density function), by integrating the PDF
2. Invert the CDF → $F^{-1}(x)$.
3. Sampling X according to $p(x)$ is achieved by sampling U → $X = F^{-1}(U)$.

1. Inversion Method

Given a uniform random variable U in the interval $[0, 1]$, we can generate a random variable from any other one dimensional distribution using its cumulative distribution function: $X = F^{-1}(U)$. This is how we choose sample points when running a ray tracing algorithm.

1. What function of U will return a sample from the exponential distribution (with parameter λ)? This distribution has density $p_\lambda(x) = \lambda e^{-\lambda x}$, and is defined for $x \geq 0$.

Solution: First, we need to calculate the CDF:

$$F_\lambda(x) = \int_0^x p_\lambda(t)dt = \int_0^x \lambda e^{-\lambda t}dt = -e^{-\lambda t}\Big|_0^x = 1 - e^{-\lambda x}$$

Set the CDF equal to U:

$$U = 1 - e^{-\lambda x}$$

Solving for x:

$$e^{-\lambda x} = 1 - U$$

$$-\lambda x = \ln(1 - U)$$

$$x = -\frac{\ln(1 - U)}{\lambda}$$

Thus, the inverse function is given by:

$$F_\lambda^{-1}(x) = -\frac{\ln(1 - x)}{\lambda}$$

and we can return:

$$-\frac{\ln(1 - U)}{\lambda}$$

to sample from the exponential distribution.

1. Inversion Method

2. For $\lambda = 1$, the plot below shows the PDF and CDF. If we sample $U = 0.75$, determine the corresponding sampled value using the inverse method. Mark the corresponding point on the CDF curve in the graph.

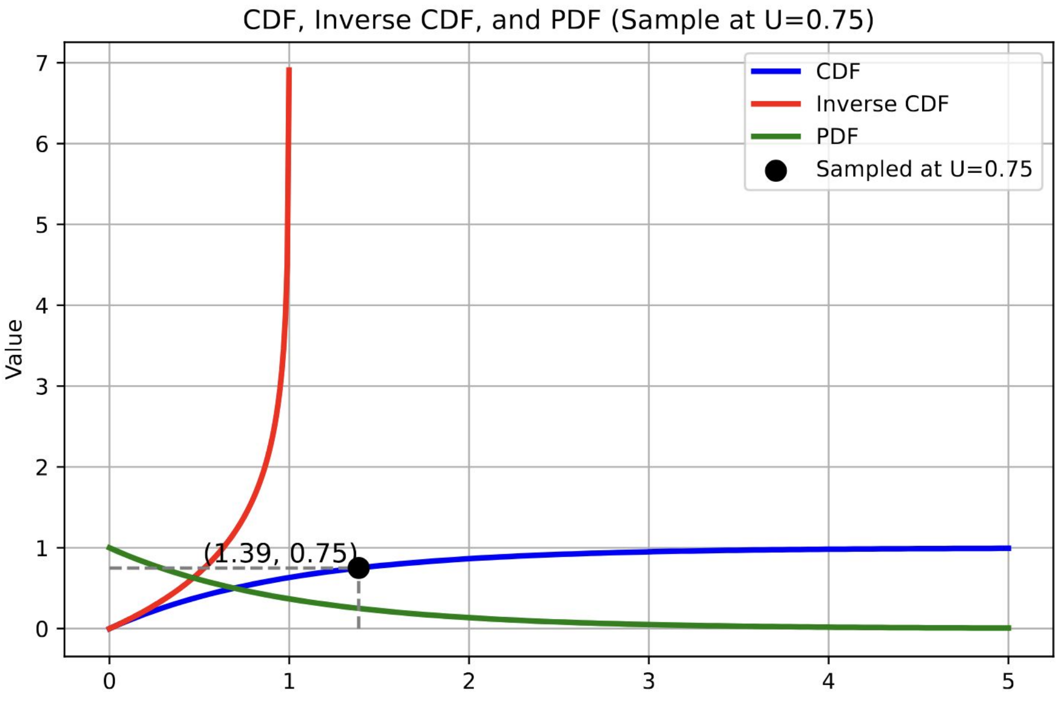

1. Inversion Method

3. What does the x-axis represent for the blue CDF curve?

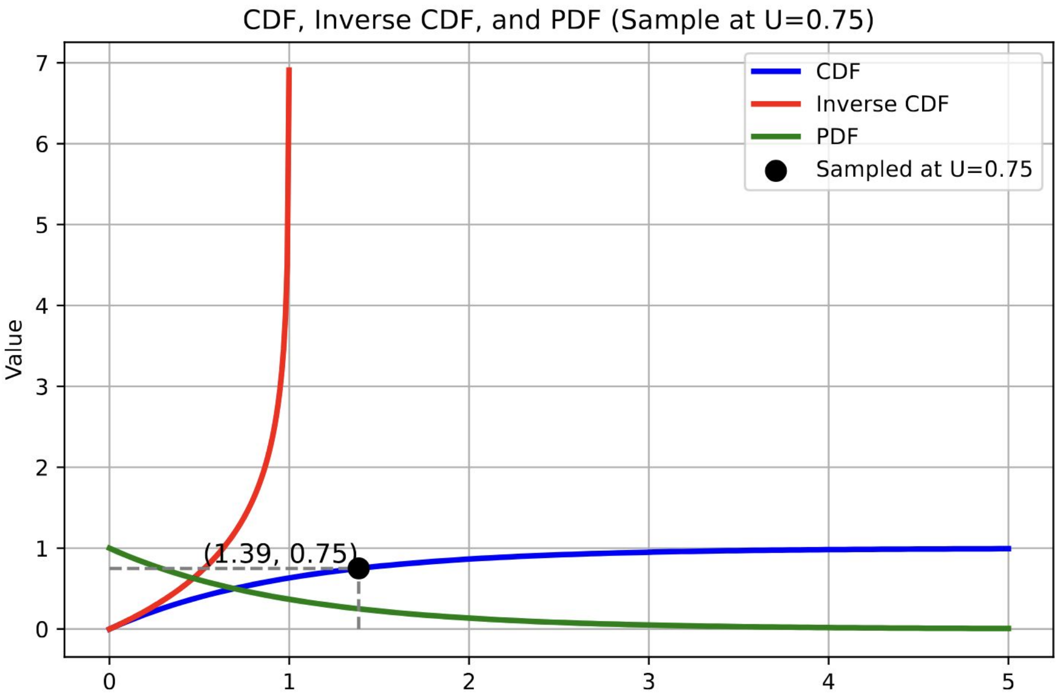

1. Inversion Method

3. What does the x-axis represent for the blue CDF curve?

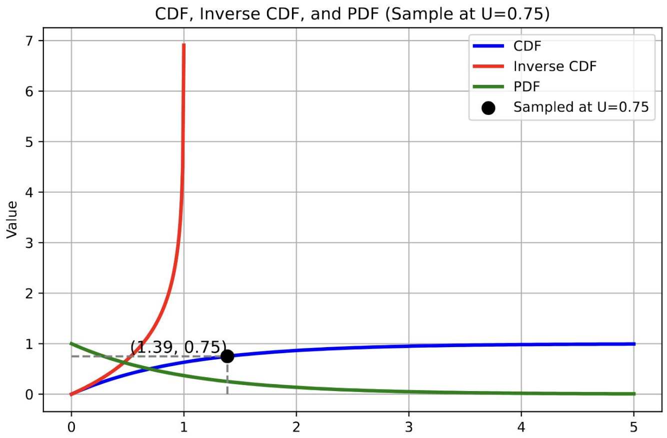

Solution: The x-axis represents the possible values that can be sampled from the exponential distribution. Given a uniform random variable U, the CDF maps it to a corresponding sampled value x from the distribution. The x-axis, therefore, represents the range of values we could obtain when transforming a uniform sample using the inverse CDF.

Monte Carlo Integration

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

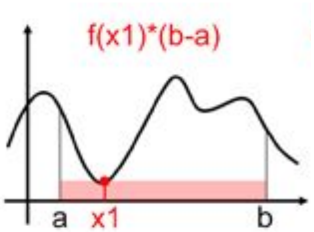

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

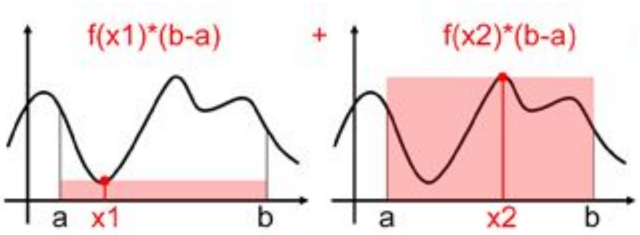

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

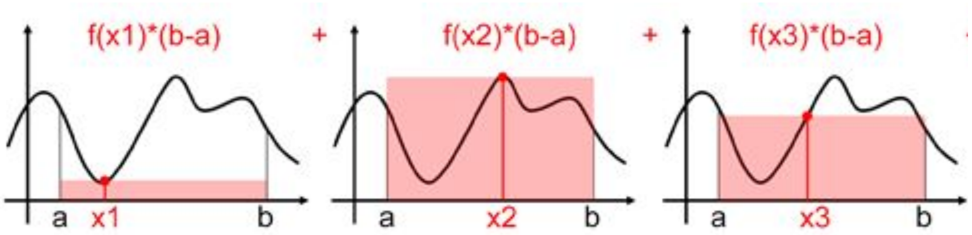

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

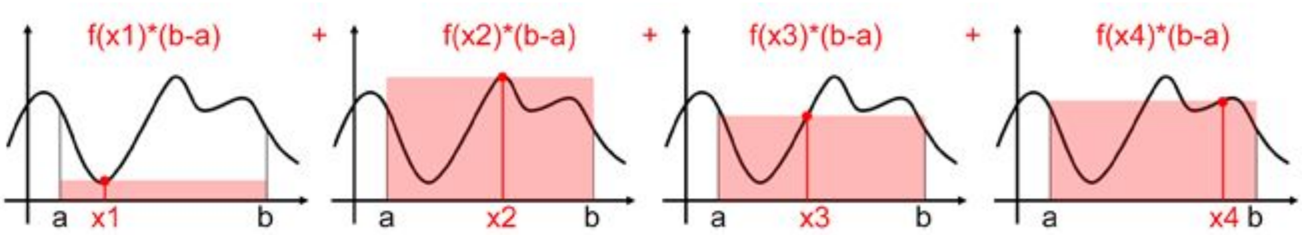

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

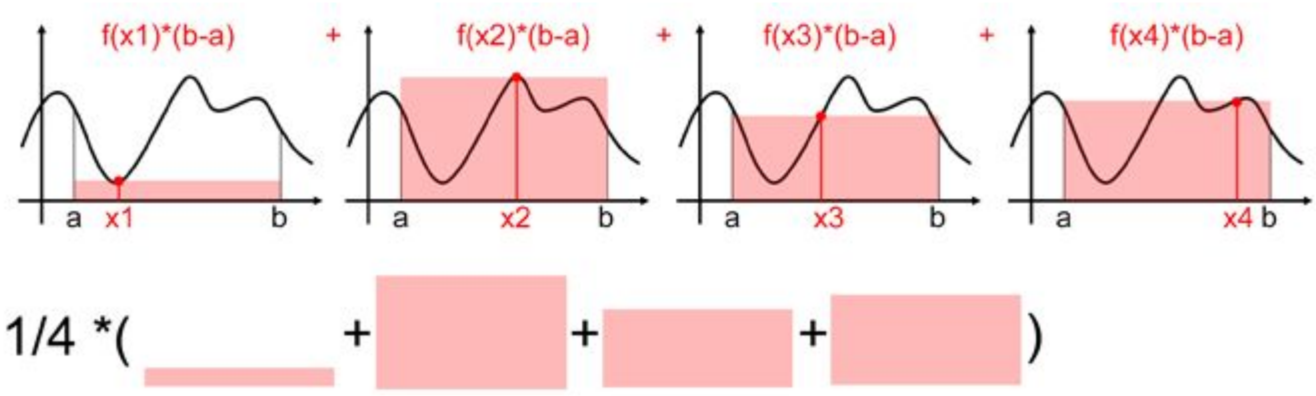

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.

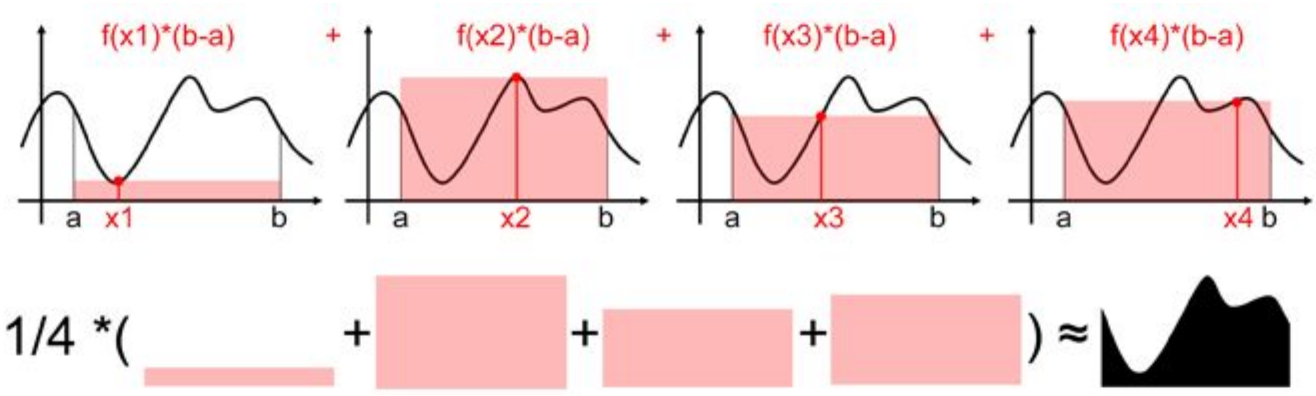

Monte Carlo Integration

- Estimate the value of an integral: $\int_a^b f(x)dx$
- Take *n* samples of *f(x)*, where *x* is randomly chosen from between *a* and *b*.
- This estimate is unbiased ➔ $E[F_n] = \int_a^b f(x)dx$

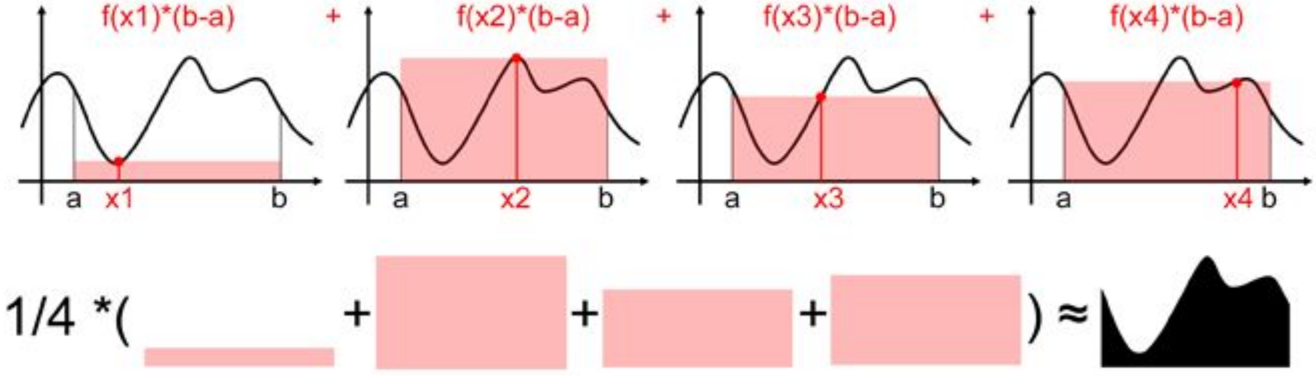

Why not Riemann sums?

- Hard to approximate complex shapes using rectangles.

Why not Riemann sums?

- Hard to approximate complex shapes using rectangles.
- Need lots of samples to get a good estimate!
  (Curse of dimensionality.)

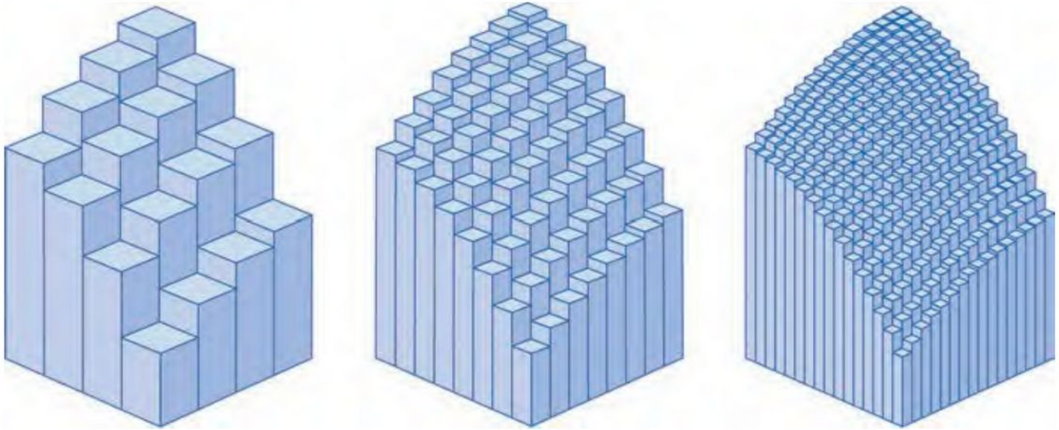

Why not Riemann sums?

- Hard to approximate complex shapes using rectangles.
- Need lots of samples to get a good estimate! (Curse of dimensionality.)
- Monte Carlo sampling is easy to implement.

```
random_uniform()
```

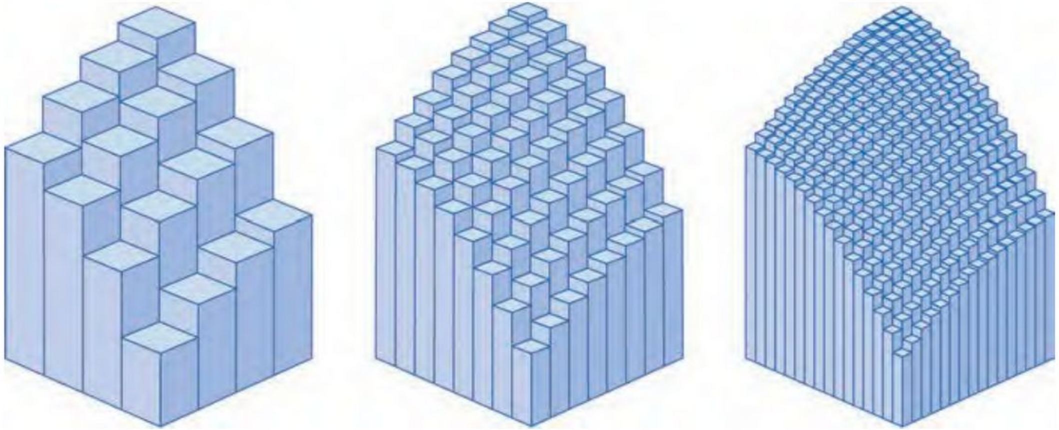

Monte Carlo Integration

Let us define the Monte Carlo estimator for the definite integral of given function $f(x)$

Definite integral $$\int_a^b f(x)dx$$

Random variable $$X_i \sim p(x)$$

Monte Carlo estimator $$F_N = \frac{1}{N}\sum_{i=1}^{N}\frac{f(X_i)}{p(X_i)}$$

p(x) can be any distribution! (Not just uniform.)

Question 2

2. Unbiased Estimators

1. Let $f : [-2, 2] \times [-2, 2] \to \mathbb{R}$ be a function. You have a machine that allows you to sample $2n$ values independently and uniformly from the interval $[-2, 2]$. Construct an unbiased Monte Carlo estimator for

$$F = \int_{-2}^{2} \int_{-2}^{2} f(x, y) dx dy.$$

2. Unbiased Estimators

1. Let $f : [-2,2] \times [-2,2] \to \mathbb{R}$ be a function. You have a machine that allows you to sample $2n$ values independently and uniformly from the interval $[-2,2]$. Construct an unbiased Monte Carlo estimator for

$$F = \int_{-2}^{2} \int_{-2}^{2} f(x,y)dxdy.$$

REMINDER: **Monte Carlo estimator** $$F_N = \frac{1}{N} \sum_{i=1}^{N} \frac{f(X_i)}{p(X_i)}$$

Where N = number of samples, X_i: ith random variable sampled from distribution, p(x) is the PDF of distribution for sampling values for the domain, and f(x) is our function (our range) we would like to estimate

2. Unbiased Estimators

1. Let $f : [-2,2] \times [-2,2] \to \mathbb{R}$ be a function. You have a machine that allows you to sample $2n$ values independently and uniformly from the interval $[-2,2]$. Construct an unbiased Monte Carlo estimator for

$$F = \int_{-2}^{2} \int_{-2}^{2} f(x,y)\,dx\,dy.$$

Solution: Draw n random samples $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$. Then we take

$$\langle F_n \rangle := \frac{1}{n} \sum_{i=1}^{n} \frac{f(X_i, Y_i)}{p(X_i, Y_i)}, \text{ where } p(X_i, Y_i) = p(X_i)p(Y_i) = \left(\frac{1}{4}\right)\left(\frac{1}{4}\right) = \frac{1}{16}$$

p(x):

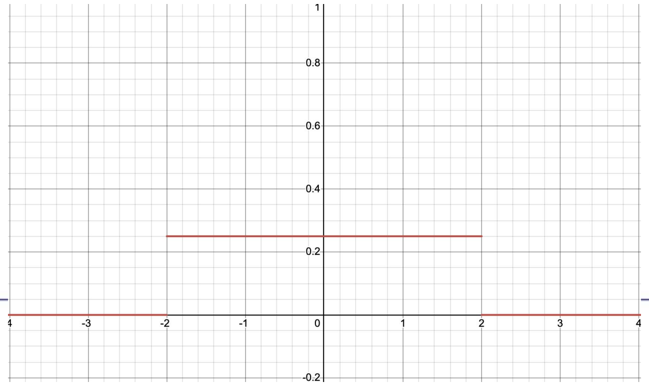

p(x, y):

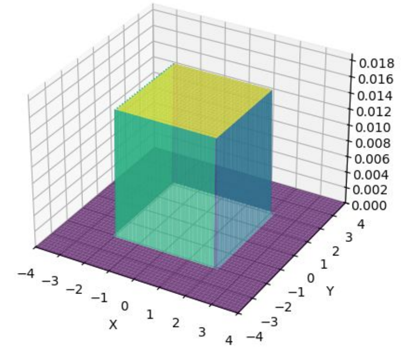

2. Unbiased Estimators

1. Let $f : [-2, 2] \times [-2, 2] \to \mathbb{R}$ be a function. You have a machine that allows you to sample $2n$ values independently and uniformly from the interval $[-2, 2]$. Construct an unbiased Monte Carlo estimator for

$$F = \int_{-2}^{2} \int_{-2}^{2} f(x, y) dx dy.$$

Solution: Draw n random samples $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$. Then we take

$$\begin{aligned} \langle F_n \rangle &:= \frac{1}{n} \sum_{i=1}^{n} \frac{f(X_i, Y_i)}{p(X_i, Y_i)}, \text{ where } p(X_i, Y_i) = p(X_i)p(Y_i) = \left(\frac{1}{4}\right)\left(\frac{1}{4}\right) = \frac{1}{16} \\ &= \frac{1}{n} \sum_{i=1}^{n} \frac{f(X_i, Y_i)}{\frac{1}{16}} \\ &= \frac{16}{n} \sum_{i=1}^{n} f(X_i, Y_i). \end{aligned}$$

Light Transport

Light Transport

- ω_i is the "direction" of <u>incoming</u> light.
- ω_o is the "direction" of <u>outgoing</u> light.

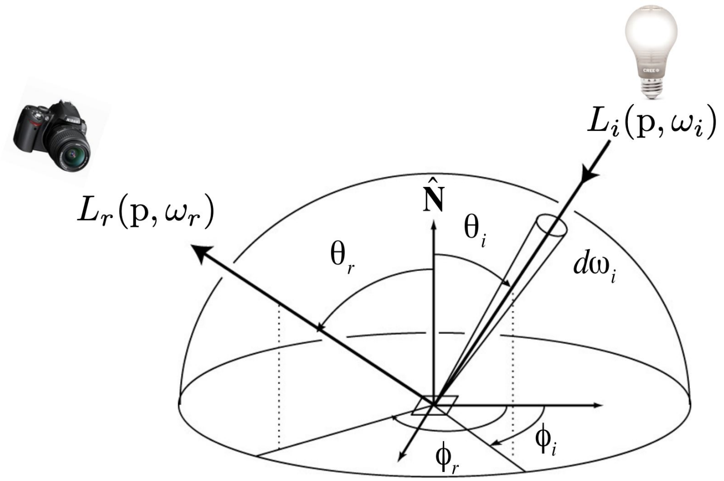

- The BRDF, $f_r(p, \omega_i, \omega_o)$, gives the amount of light from ω_i that is reflected in the direction ω_o.

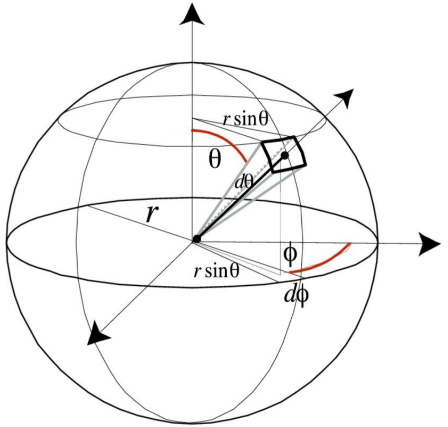

$$\mathrm{d}A = (r\,\mathrm{d}\theta)(r\sin\theta\,\mathrm{d}\phi) = r^2\sin\theta\,\mathrm{d}\theta\,\mathrm{d}\phi$$

$$\mathrm{d}\omega = \frac{\mathrm{d}A}{r^2} = \sin\theta\,\mathrm{d}\theta\,\mathrm{d}\phi$$

Solid Angle Review

Solid Angle

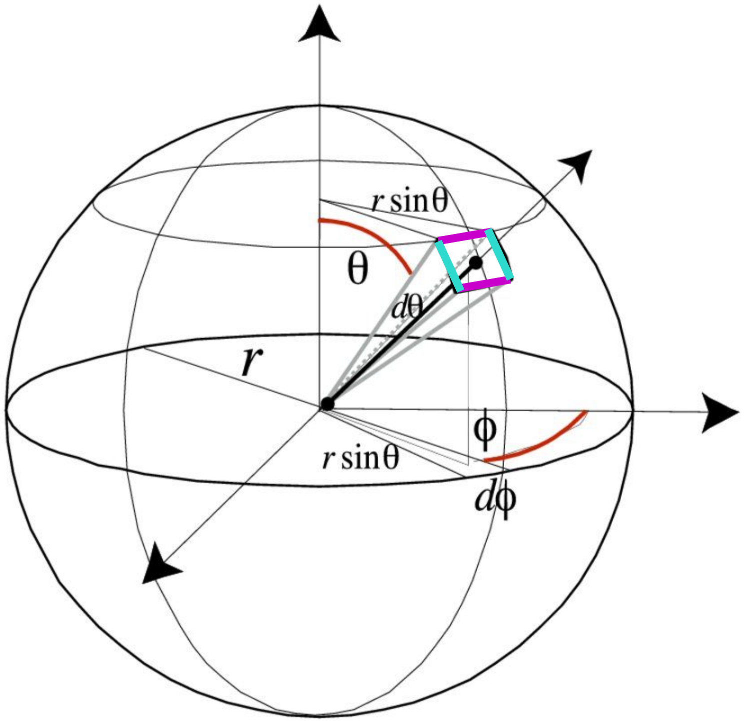

$$dA = (r\, d\theta)(r \sin\theta\, d\phi)$$
$$= r^2 \sin\theta\, d\theta\, d\phi$$

$$d\omega = \frac{dA}{r^2} = \sin\theta\, d\theta\, d\phi$$

Question 3

6. Conceptually, why does cosine-weighted hemisphere sampling outperform uniform sampling over a hemisphere?

6. Conceptually, why does cosine-weighted hemisphere sampling outperform uniform sampling over a hemisphere?

- **Importance sampling → want PDF of *X*'s to resemble integrand.**
- **Why does this work?**

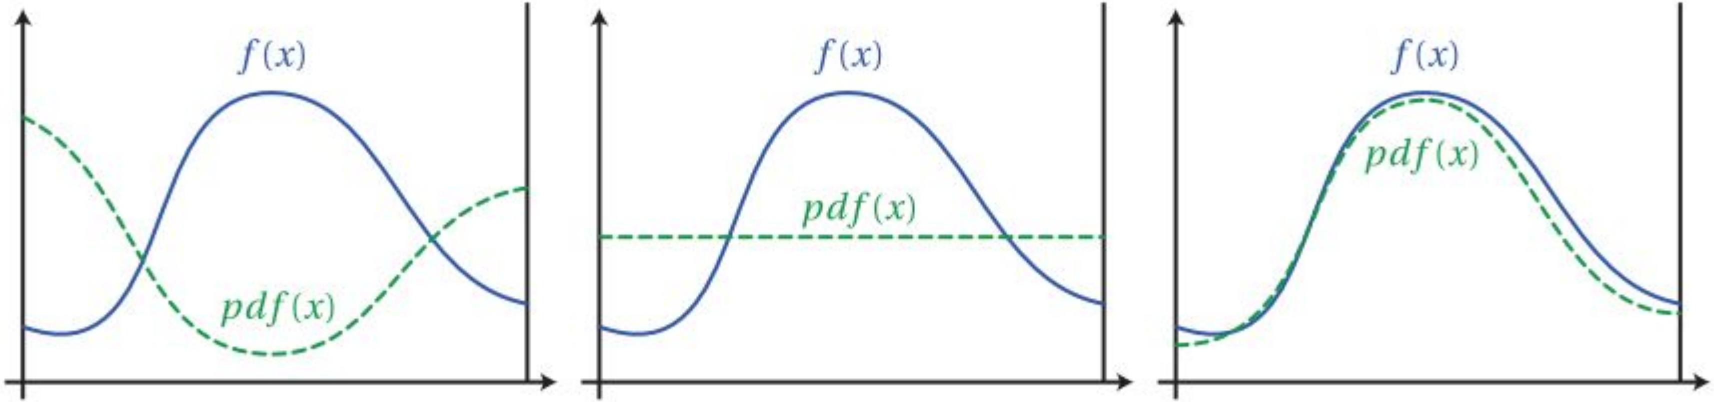

Global Illumination

Global Illumination

- Recursively calculate how much light falls onto point *p*.

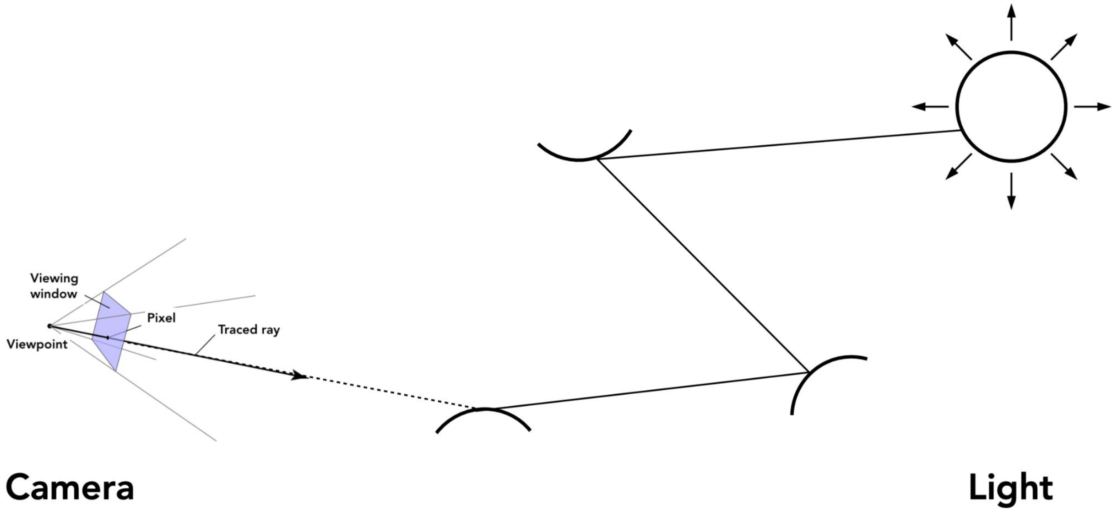

Global Illumination

- Recursively calculate how much light falls onto point *p*.
- Radiance from light source → *p* → camera (1-bounce).

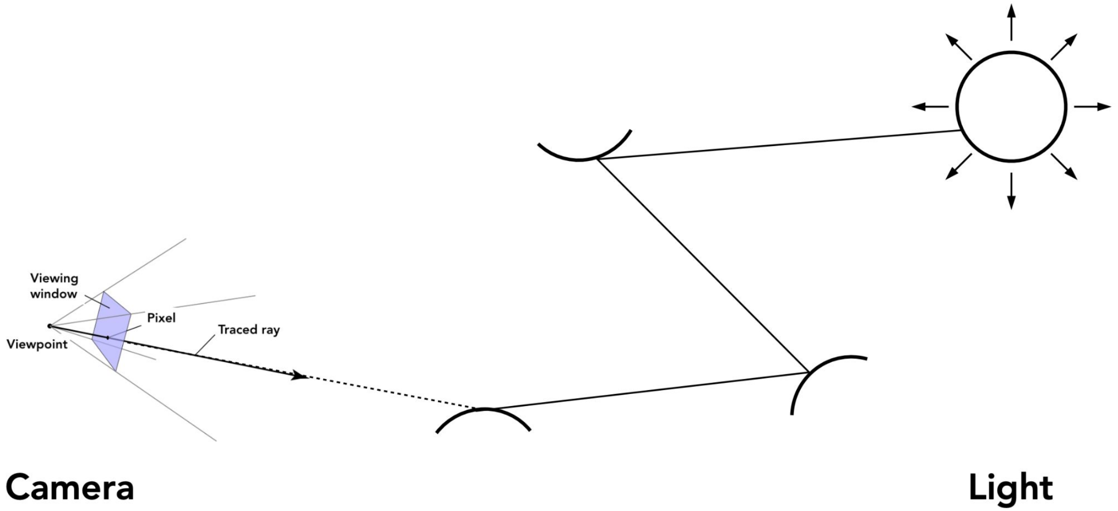

Global Illumination

- Recursively calculate how much light falls onto point *p*.
- Radiance from light source ➔ *p* ➔ camera (1-bounce).
- Radiance from light source ➔ *p*’ ➔ *p* ➔ camera (2-bounce).

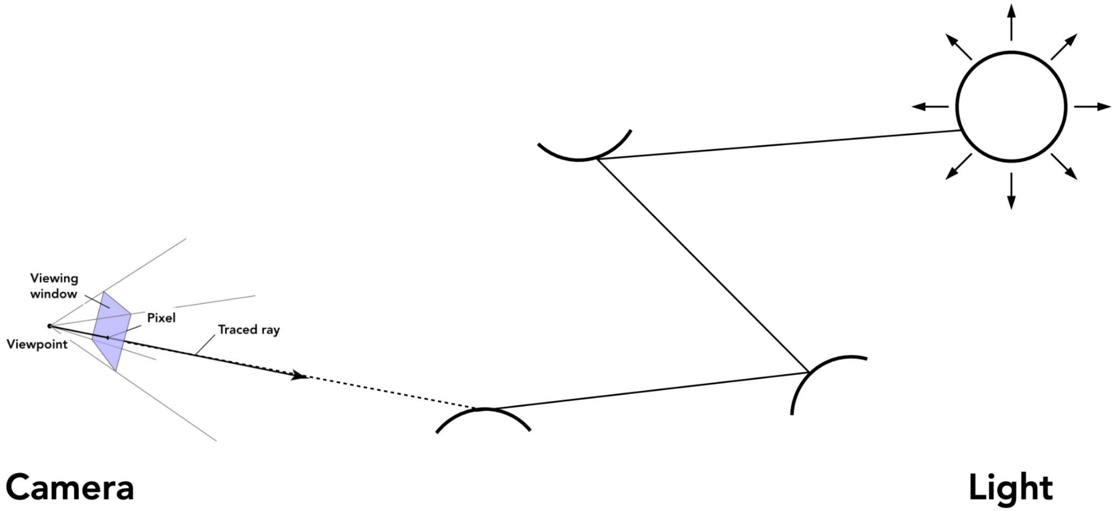

Global Illumination

- Recursively calculate how much light falls onto point *p*.
- Radiance from light source ➔ *p* ➔ camera (1-bounce).
- Radiance from light source ➔ *p*’ ➔ *p* ➔ camera (2-bounce).
- And so forth...

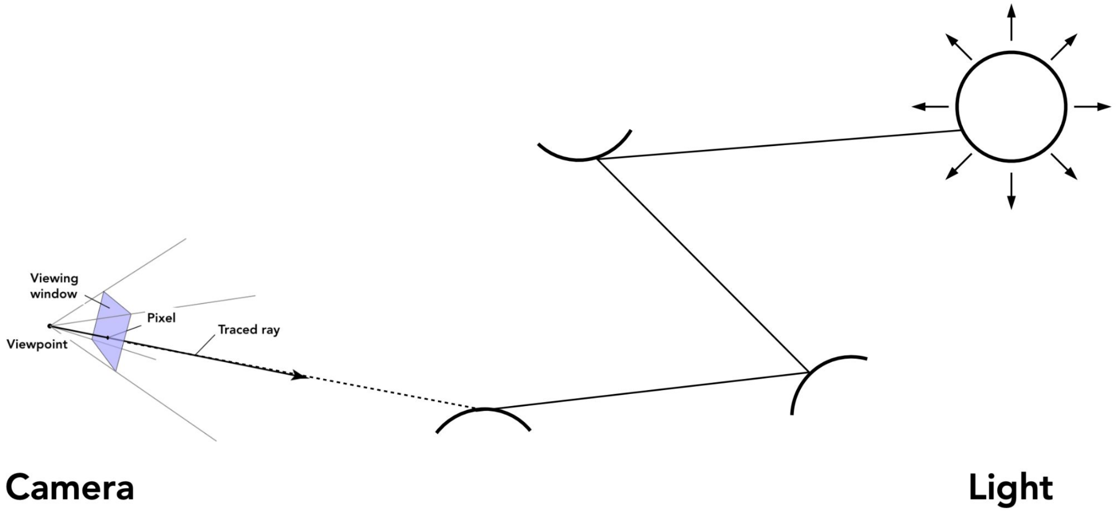

Global Illumination

- Recursively calculate how much light falls onto point *p*.
- Radiance from light source ➔ *p* ➔ camera (1-bounce).
- Radiance from light source ➔ *p'* ➔ *p* ➔ camera (2-bounce).
- And so forth...

Question 4 correction:

$$f_r(\mathrm{p}, \omega_i, \omega_o) = \frac{\rho}{\pi}$$

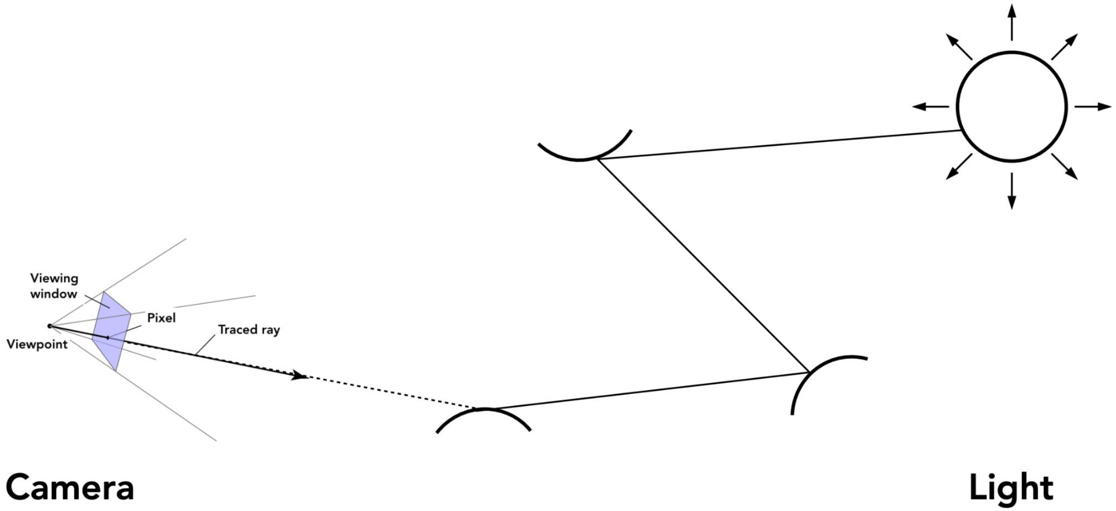

Question 4

Importance Sampling

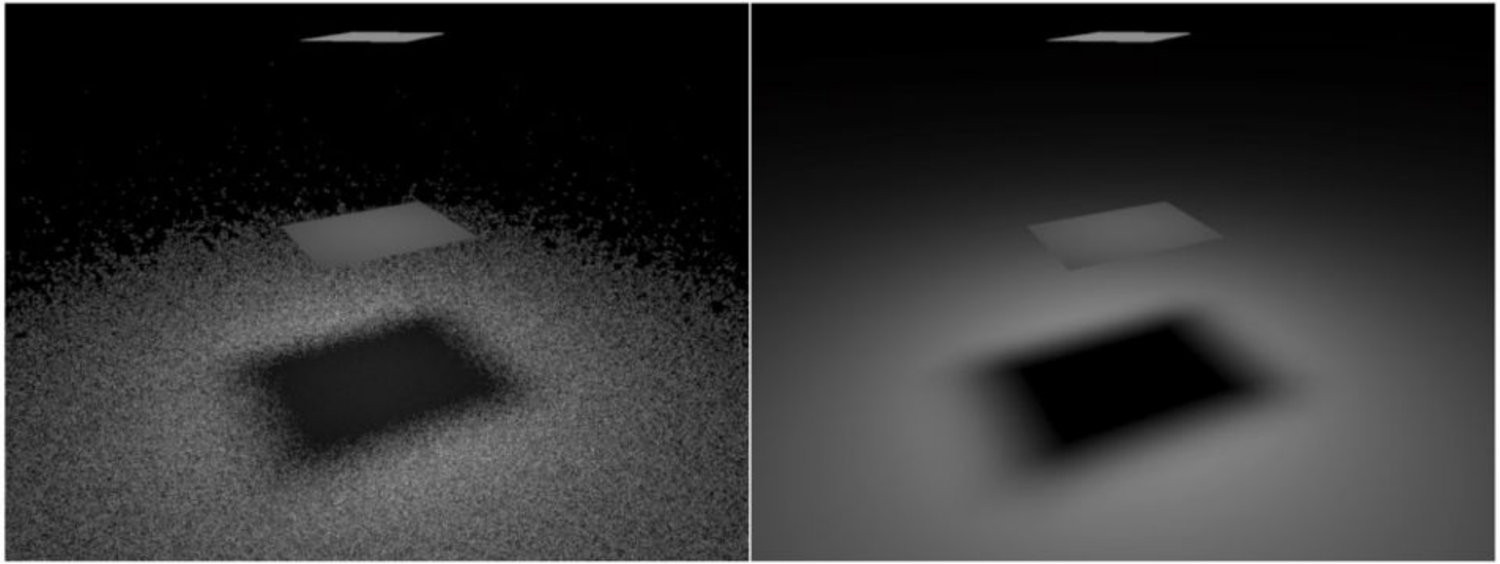

Sampling solid angle

100 random directions on hemisphere

Sampling light source area

100 random points on area of light source

2-Bounce Path-Tracing

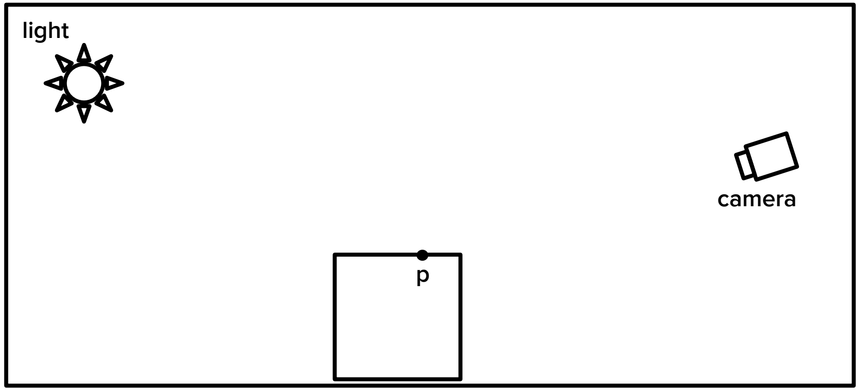

2-Bounce Path-Tracing

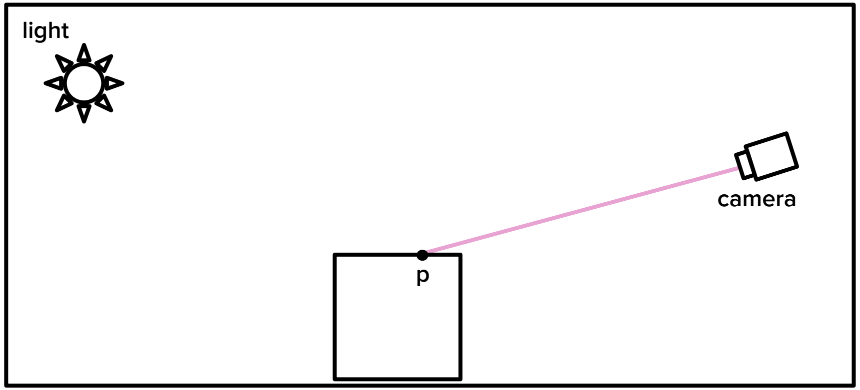

2-Bounce Path-Tracing

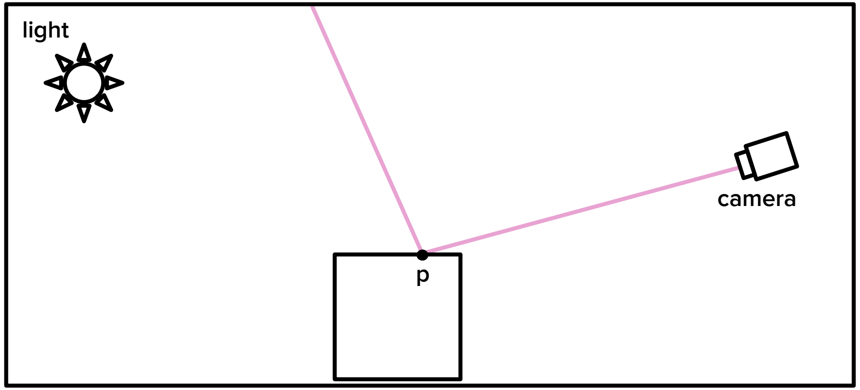

2-Bounce Path-Tracing

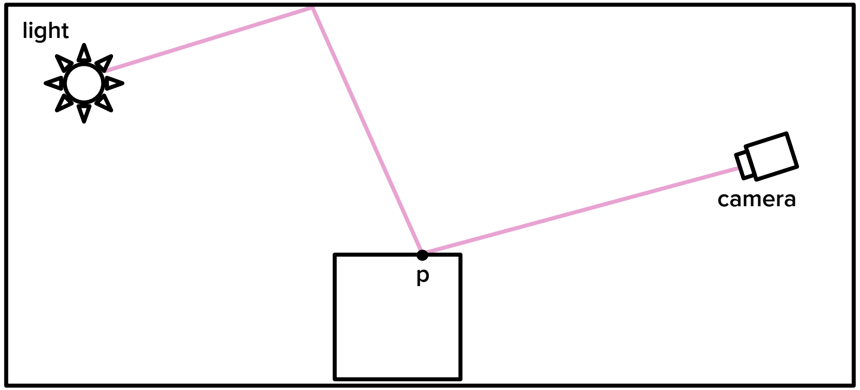

2-Bounce Path-Tracing

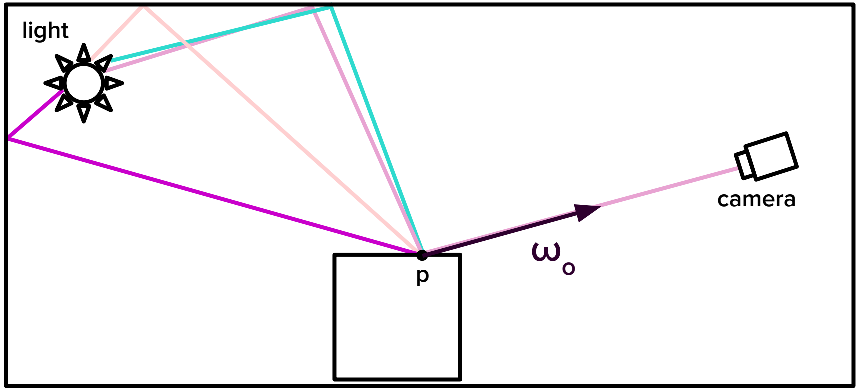

3-Bounce Path-Tracing

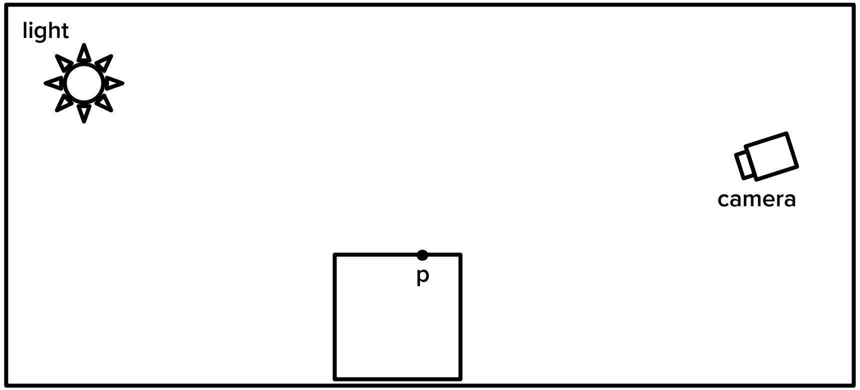

3-Bounce Path-Tracing

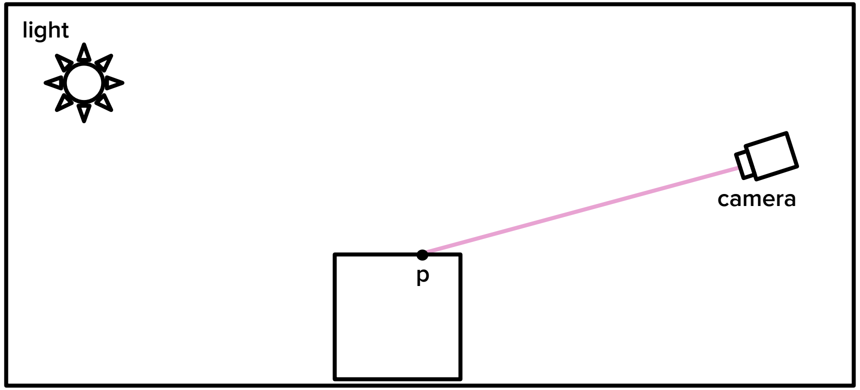

3-Bounce Path-Tracing

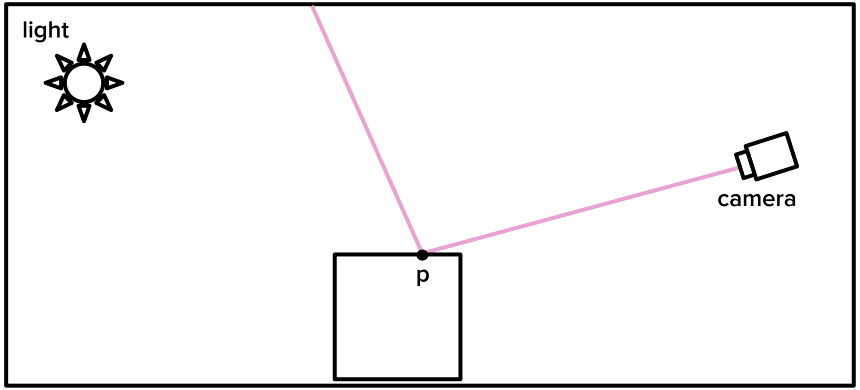

3-Bounce Path-Tracing

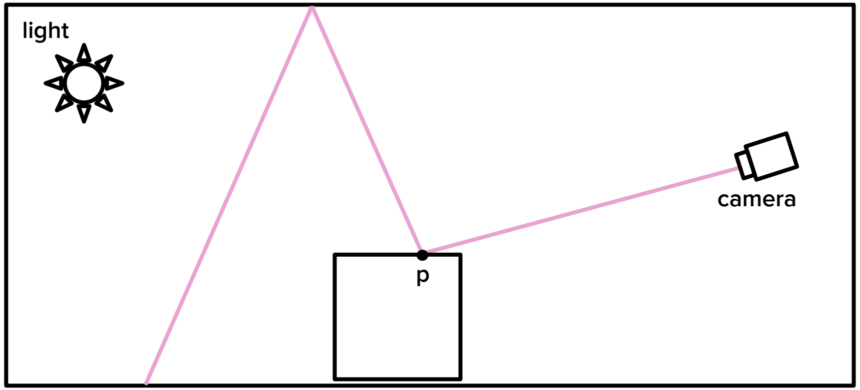

3-Bounce Path-Tracing

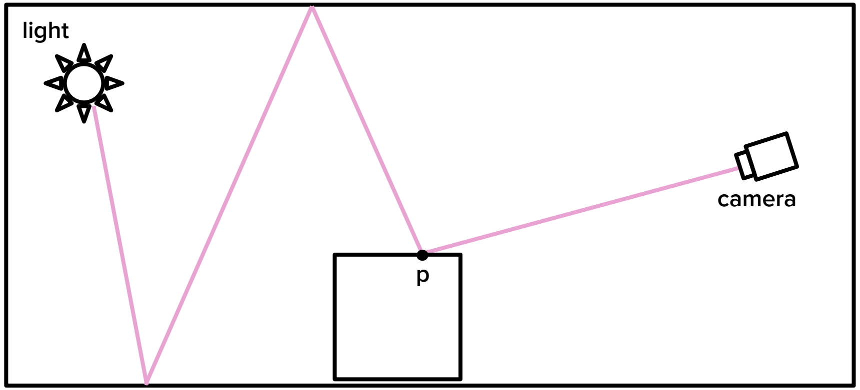

3-Bounce Path-Tracing

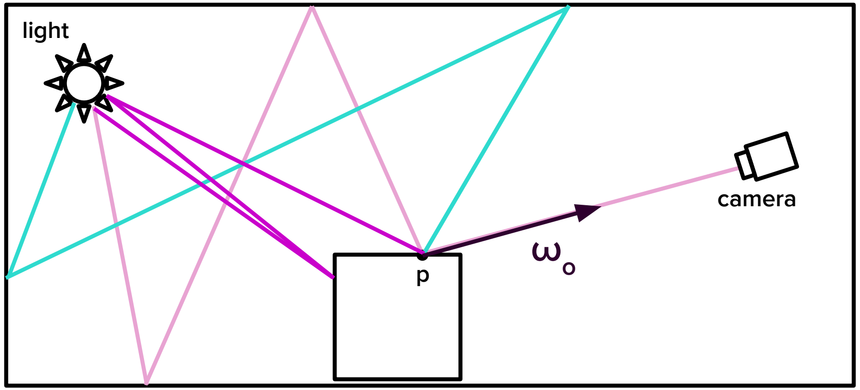

Termination Conditions

Termination Conditions

- 0-bounce, 1-bounce, 2-bounce… ➔ how to prevent ∞-bounce?

Termination Conditions

- 0-bounce, 1-bounce, 2-bounce... → how to prevent ∞-bounce?
- **Russian roulette:** at each bounce, randomly terminate current ray with probability $1 - p_{rr}$.

Termination Conditions

- 0-bounce, 1-bounce, 2-bounce... ➔ how to prevent ∞-bounce?
- **Russian roulette:** at each bounce, randomly terminate current ray with probability 1 - p_{rr}.
- Unbias the estimator (proof in lecture):

$$X_{\mathrm{rr}} = \begin{cases} \dfrac{X}{p_{\mathrm{rr}}}, & \text{with probability } p_{\mathrm{rr}} \\ 0, & \text{otherwise} \end{cases}$$

Question 5

Let's Take Attendance.

- Be sure to select Discussion 7 and input your TA's secret word 😉
- Any feedback? Let us know!