Monte Carlo Integration II & Sampling from PDFs

CS295, Spring 2017

Shuang Zhao

Computer Science Department

University of California, Irvine

Last Lecture

- Direct illumination
- Monte Carlo integration I

Today's Lecture

- Monte Carlo integration II
  - Convergence properties
  - Integrals over higher-dimensional domains

- Sampling from PDFs
  - Inversion method

Recap: Monte Carlo Integration

- Goal: to estimate $I = \int_a^b f(x)\,\mathrm{d}x$
  - Pick a probability density function $p(x)$
  - Generate *n* independent samples:

$$x_1, x_2, \ldots, x_n \sim p$$

  - Evaluate $\hat{I}_j := \dfrac{f(x_j)}{p(x_j)}$ for *j* = 1, 2, …, *n*
  - Return sample mean: $\bar{I} := \dfrac{1}{n}\sum_{j=1}^{n} \hat{I}_j$

How fast does this algorithm converge?

Central Limit Theorem

- Let $\{X_1, X_2, \ldots\}$ be a sequence of i.i.d. random variables with $\mathbb{E}[X_i] = \mu,\ \mathrm{Var}[X_i] = \sigma^2 < \infty$ for all *i*

- Let $S_n := \frac{1}{n}\sum_{i=1}^{n} X_i$, then $\frac{\sqrt{n}}{\sigma}(S_n - \mu) \xrightarrow[n\to\infty]{d} N(0,1)$

- It follows that for any *z* > 0:

$$\lim_{n\to\infty} \mathbb{P}\left[\frac{\sqrt{n}}{\sigma}(S_n - \mu) \le z\right] = \Phi(z)$$

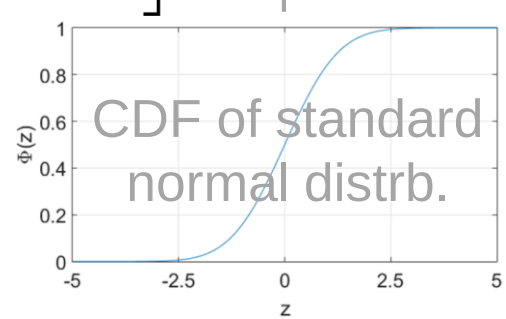

Central Limit Theorem

$$\lim_{n\to\infty} \mathbb{P}\left[\frac{\sqrt{n}}{\sigma}(S_n - \mu) \le z\right] = \Phi(z)$$

$$\Downarrow$$

$$\begin{aligned}
&\lim_{n\to\infty} \mathbb{P}\left[\left|\frac{\sqrt{n}}{\sigma}(S_n - \mu)\right| \le z\right] \\
= &\lim_{n\to\infty} \mathbb{P}\left[\mu \in \left(S_n - z\frac{\sigma}{\sqrt{n}},\ S_n + z\frac{\sigma}{\sqrt{n}}\right)\right] \\
= &\ 1 - 2\Phi(-z)
\end{aligned}$$

Confidence Interval

- $\left(S_n - z\frac{\sigma}{\sqrt{n}},\ S_n + z\frac{\sigma}{\sqrt{n}}\right)$ is called a $1 - 2\Phi(-z)$ confidence interval (CI) of μ
- **Interpretation:** there is a $1 - 2\Phi(-z)$ chance for μ to reside in this CI
- z is called the "critical value"
- Common critical values
  - z = 1.960: 95%-CIs
  - z = 2.576: 99%-CIs

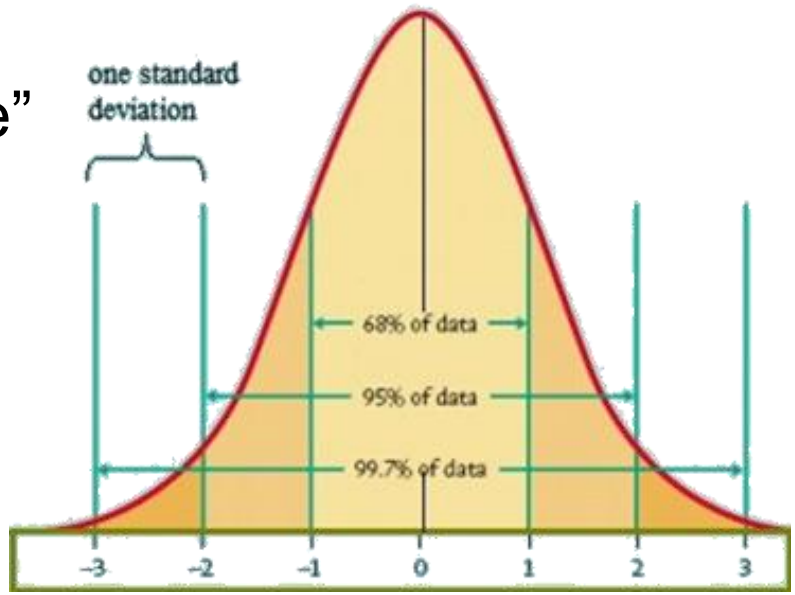

Confidence Interval

- $\left(S_n - z\frac{\sigma}{\sqrt{n}},\ S_n + z\frac{\sigma}{\sqrt{n}}\right)$ is called a $1 - 2\Phi(-z)$ confidence interval (CI) of *μ*
- **One more problem:** the standard deviation σ is generally *unknown*
- **Solution:** use (corrected) sample standard deviation

$$\sigma_n := \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(X_i - S_n)^2},$$

resulting in CIs of the form $\left(S_n - z\frac{\sigma_n}{\sqrt{n}},\ S_n + z\frac{\sigma_n}{\sqrt{n}}\right)$

Convergence of MC Methods

- The relative size of a CI, namely

$$\left(z \frac{\sigma_n}{\sqrt{n}} \right) / |S_n|$$

  is an important indicator for the reliability of estimated result S_n

- Monte Carlo methods generally have a convergence rate of $O(1/\sqrt{n})$

Today's Lecture

- Monte Carlo integration II
  - Convergence properties
  - Integrals over higher-dimensional domains

- Sampling from PDFs

Background:
Measures and Lebesgue Integration

$$I = \int_{\Gamma} f(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})$$

- Γ is a set specifying the domain of integration

- μ is a measure function capturing the "sizes" of Γ and its (measurable) subsets
  - Properties:

$$\mu(A) \geq 0 \quad \text{for all } A \subseteq \Gamma \text{ (that are measurable)}$$
$$\mu(A \cup B) = \mu(A) + \mu(B) \quad \text{if } A \cap B = \emptyset$$
$$\mu(\emptyset) = 0$$

Background:
Measures and Lebesgue Integration

$$I = \int_{\Gamma} f(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})$$

- Assuming Γ can be partitioned into n disjoint components $\Gamma_1, \Gamma_2, \ldots, \Gamma_n$ and $f(\boldsymbol{x}_i) \equiv a_i$ for all $\boldsymbol{x}_i$ in Γ_i
- Then, $I = \sum_{i=1}^{n} a_i\,\mu(\Gamma_i)$
- Special case: $\int_{\Gamma} \mathrm{d}\mu(\boldsymbol{x}) = \mu(\Gamma)$

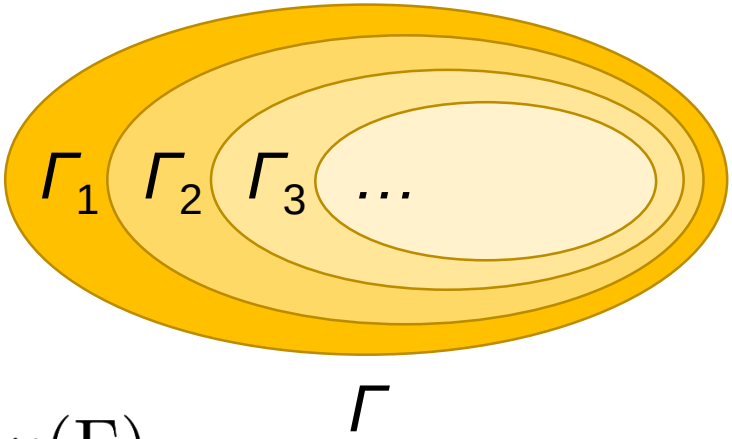

Background:
Measures and Lebesgue Integration

$$I = \int_{\Gamma} f(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})$$

- In general, Γ can be partitioned into many components Γ_1, Γ_2, … so that f becomes (approximately) constant in each component Γ_i

- In the limiting case,

$$I = \lim_{n\to\infty} \sum_{i=1}^{n} a_i\,\mu(\Gamma_i)$$

Background:
Borel Measure and 1D Integrals

$$\mu_{\mathrm{B}}((a,b]) = b - a$$

- The *Borel measure* captures the "length" of 1D intervals

$$\int_a^b f(x)\,\mathrm{d}x = \int_{(a,b]} f(x)\,\mathrm{d}\mu_{\mathrm{B}}(x)$$

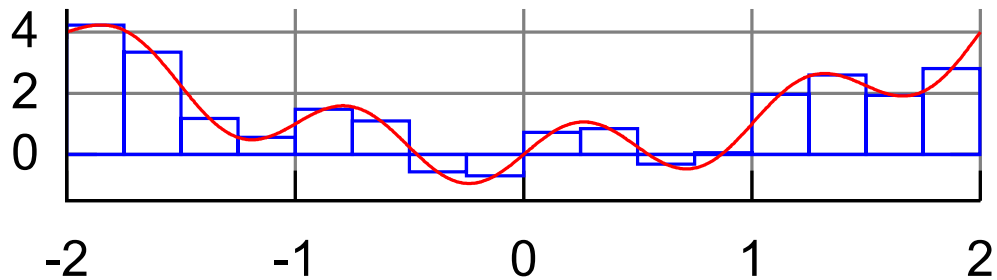

Background:
Lebesgue Measure

$$\lambda((a,b] \times (c,d]) = (b-a)(d-c)$$

- The *Lebesgue measure* captures the "volume" of hyper-cubes

$$\int_a^b \int_c^d f((x,y))\,\mathrm{d}y\,\mathrm{d}x = \int_{(a,b]\times(c,d]} f(\boldsymbol{r})\,\mathrm{d}\lambda(\boldsymbol{r})$$

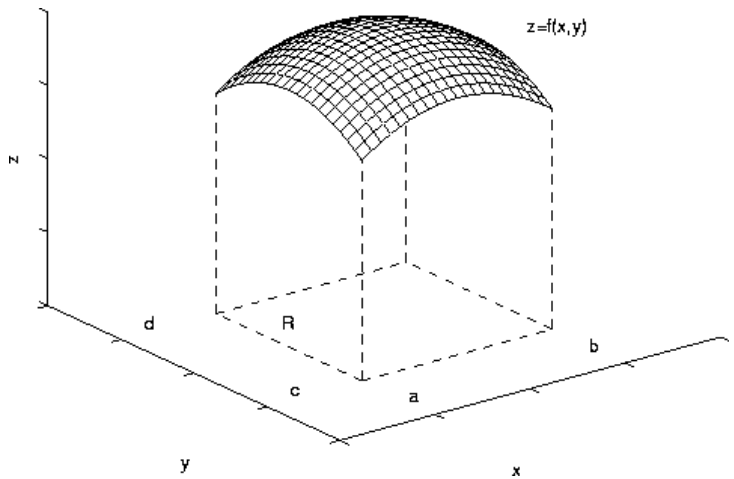

Background:
Solid Angle Measure

- For a region *A* on the unit sphere $\mathbb{S}^2$, the solid angle measure σ captures *A*'s surface area

Unit sphere

$$\sigma(\mathbb{S}^2) = 4\pi$$

$$\sigma(\Omega_+) = \sigma(\mathbb{H}^2) = 2\pi$$

Unit hemisphere

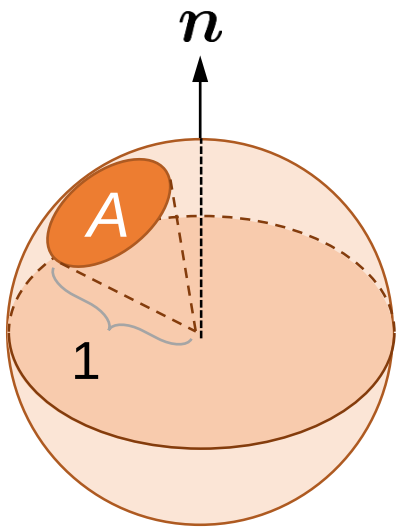

Background:
Projected Solid Angle Measure

- For a region *A* on the unit sphere $\mathbb{S}^2$, the projected solid angle measure $\sigma^\perp$ captures the surface area of $A^\perp$ (which is *A*'s projection on the plane perpendicular to the normal direction)

$$\sigma^\perp(\Omega_+) = \sigma^\perp(\mathbb{H}^2) = \pi$$

Unit hemisphere

$$\mathrm{d}\sigma^\perp(\boldsymbol{\omega}) = \cos\theta\,\mathrm{d}\sigma(\boldsymbol{\omega})$$

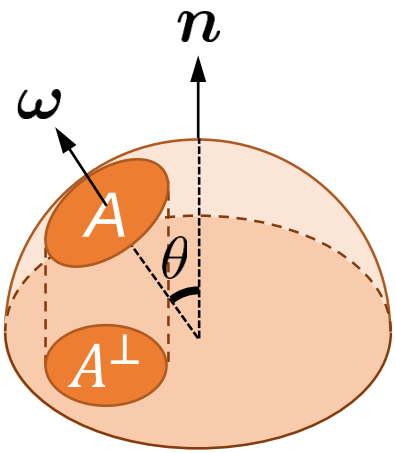

Background:
Probability Measures

- Given a sample space (i.e., all possible outcomes) Ω and $\mathcal{F} \subseteq \Omega$, $P(\mathcal{F})$ denotes the probability for an outcome in $\mathcal{F}$ to occur

- Example 1-1: fair coin
  - $\Omega = \{H, T\}$
  - $P(\{H\}) = P(\{T\}) = \frac{1}{2},\ P(\{H, T\}) = 1$

- Example 1-2: Poisson distribution (with parameter *λ*)
  - $\Omega = \{0, 1, 2, \ldots\} = \mathbb{N} \cup \{0\}$
  - $P(\{k\}) = \dfrac{\lambda^k \mathrm{e}^{-\lambda}}{k!} \quad \forall k \in \Omega$

Background:
Probability Measures

- Given a sample space (i.e., all possible outcomes) Ω and $\mathcal{F} \subseteq \Omega$, $P(\mathcal{F})$ denotes the probability for an outcome in $\mathcal{F}$ to occur

- Example 2-1: *U*(0, 1]
  - $\Omega = (0, 1]$
  - $P((a, b]) = b - a = \mu_B((a, b]) \quad \forall 0 \leq a < b \leq 1$

- Example 2-2: *N*(0, 1)
  - $\Omega = \mathbb{R}$
  - $P((a, b]) = \Phi(b) - \Phi(a) \quad \forall a < b$

Background:
Probability Measures

- Given a sample space (i.e., all possible outcomes) Ω and $\mathcal{F} \subseteq \Omega$, $P(\mathcal{F})$ denotes the probability for an outcome in $\mathcal{F}$ to occur

- Example 2-3: 1D continuous distrb. with CDF *F*
  - $\Omega = \mathbb{R}$
  - $P((a, b]) = F(b) - F(a) \quad \forall a < b$
  - $P((-\infty, a]) = F(a) \quad \forall a \in \mathbb{R}$
  - There is a one-to-one correspondence between *CDFs* and *probability measures*

Background:
High-Dimensional Distributions

- Consider a sample space $\Omega \subseteq \mathbb{R}^k$ associated with a measure μ
- E.g., $\Omega = \mathbb{R}^k, \mu = \lambda \text{ (Lebesgue)};$

  $\Omega = \mathbb{S}^2, \mu = \sigma \text{ (solid angle)}$
- Given a probability density function p satisfying

$$p(\boldsymbol{x}) \geq 0 \quad \text{and} \quad \int_\Omega p(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x}) = 1,$$

  the corresponding probability measure P is

$$P(A) = \int_A p(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})$$

  with $\mathrm{d}P(\boldsymbol{x}) = p(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})$

Background:
General Definition of Expected Values

$$\mathbb{E}[f(\boldsymbol{x})] := \int_{\Omega} f(\boldsymbol{x})\,\mathrm{d}P(\boldsymbol{x})$$

- Applies to both *discrete* and *continuous* cases

- Example 1: fair coin (discrete)
  - Consider *f*(H) = 2, *f*(T) = -1
    (i.e., you win 2 dollars for an H and lose 1 for a T)

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{x})] &= \textstyle\int_{\{H,T\}} f(\boldsymbol{x})\,\mathrm{d}P(\boldsymbol{x}) \\ &= f(H)P(\{H\}) + f(T)P(\{T\}) \\ &= 2\cdot\frac{1}{2} + (-1)\frac{1}{2} = \frac{1}{2}\end{aligned}$$

Background:
General Definition of Expected Values

$$\mathbb{E}[f(\boldsymbol{x})] := \int_{\Omega} f(\boldsymbol{x})\,\mathrm{d}P(\boldsymbol{x})$$

- Example 2: U(0, 1] (continuous)
  - Consider *f*(*x*) = *x* for all *x* in (0, 1]

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{x})] &= \int_{(0,1]} f(x)\,\mathrm{d}P(x) = \int_{(0,1]} x\,\mathrm{d}\mu_{\mathrm{B}}(x) \\ &= \int_0^1 x\,\mathrm{d}x = \frac{1}{2}\end{aligned}$$

Background:
General Definition of Expected Values

$$\mathbb{E}[f(\boldsymbol{x})] := \int_{\Omega} f(\boldsymbol{x})\,\mathrm{d}P(\boldsymbol{x})$$

- General continuous distribution with:
  - Sample space *Ω* and associated measure *μ*
  - Probability density *p*

"Classical" definition of expected value

$$\mathbb{E}[f(\boldsymbol{x})] = \int_{\Omega} f(\boldsymbol{x})\,\mathrm{d}P(\boldsymbol{x}) = \boxed{\int_{\Omega} f(\boldsymbol{x})\,p(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})}$$

Monte Carlo Integration over High-Dimensional Domains

$$I = \int_{\Gamma} f(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x})$$

- Let p be a pdf on Γ (with $f(\mathbf{x}) > 0$ implying $p(\mathbf{x}) > 0$)
- Then

$$\mathbb{E}\left[\frac{f(\boldsymbol{x})}{p(\boldsymbol{x})}\right] = \int_{\Gamma} \frac{f(\boldsymbol{x})}{p(\boldsymbol{x})}\,p(\boldsymbol{x})\,\mathrm{d}\mu(\boldsymbol{x}) = I$$

- Exactly the same as the 1D case!

Today's Lecture

- Monte Carlo integration II
  - Convergence properties
  - Integrals over higher-dimensional domains

- Sampling from PDFs

Sampling from PDFs

- Monte Carlo integration requires drawing independent samples from arbitrary probability density $p(\boldsymbol{x})$

- What we generally have in practice:
  - Generating random floating numbers from $U(0, 1]$

Inversion Method

- Given a 1D distribution with CDF *F*, then

$$X = F^{-1}(\xi) \quad \text{where } \xi \sim U(0,1]$$

follows this given distribution

- Proof: for all $a \in \mathbb{R}$,

$$\begin{aligned}\mathbb{P}[X < a] &= \mathbb{P}[F^{-1}(\xi) < a] \\ &= \mathbb{P}[\xi < F(a)] \\ &= F(a)\end{aligned}$$

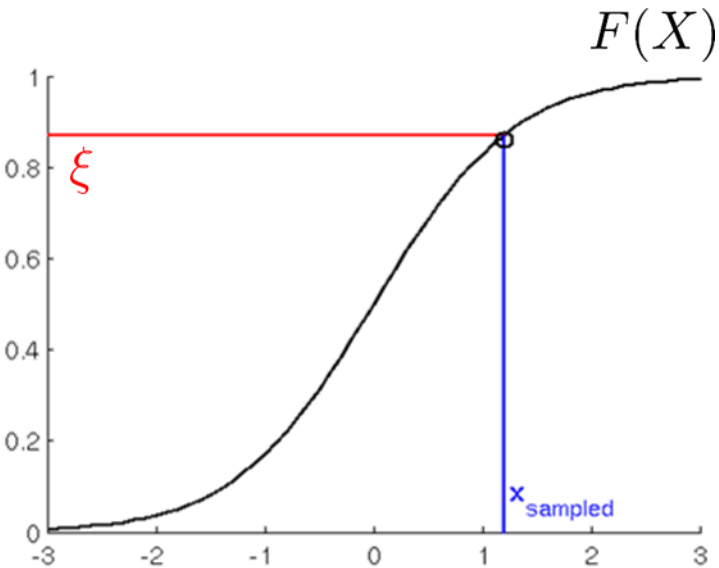

Inversion Method

- Also works for discrete distributions with finite outcomes: consider one with CMF *F*

- Sampling algorithm
  - Draw $\xi \sim U(0,1]$
  - Return the minimal integer *X* satisfying $F(X) \geq \xi$
  - Complexity: $O(\log N)$

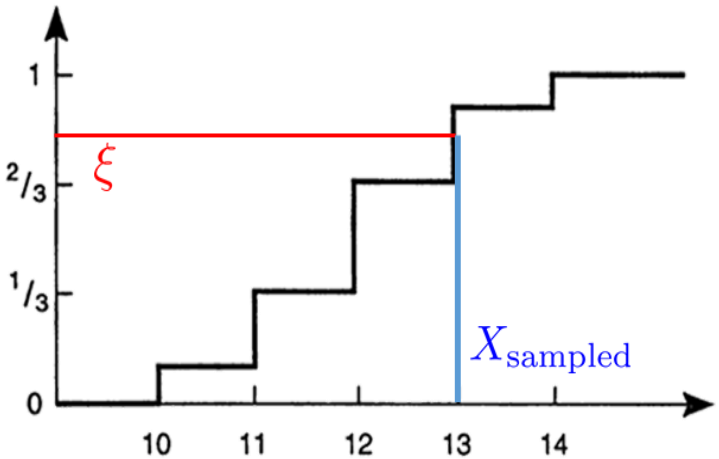

Next Lecture

- Sampling from PDFs II