

Ascent

Policy Gradient Descent

Lucas Janson and Sham Kakade

**CS/Stat 184: Introduction to Reinforcement Learning
Fall 2022**

Today

- Recap
- Today:
 - How do we learn/compute a good policy in an intractably large MDP?
 - Policy gradient descent is one of the most effective methods.

• QL was our first local search method.

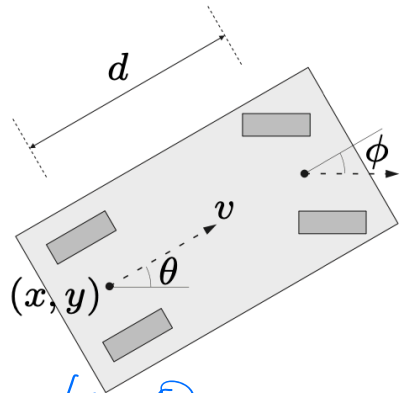
Recap

Example:

2-d car navigation

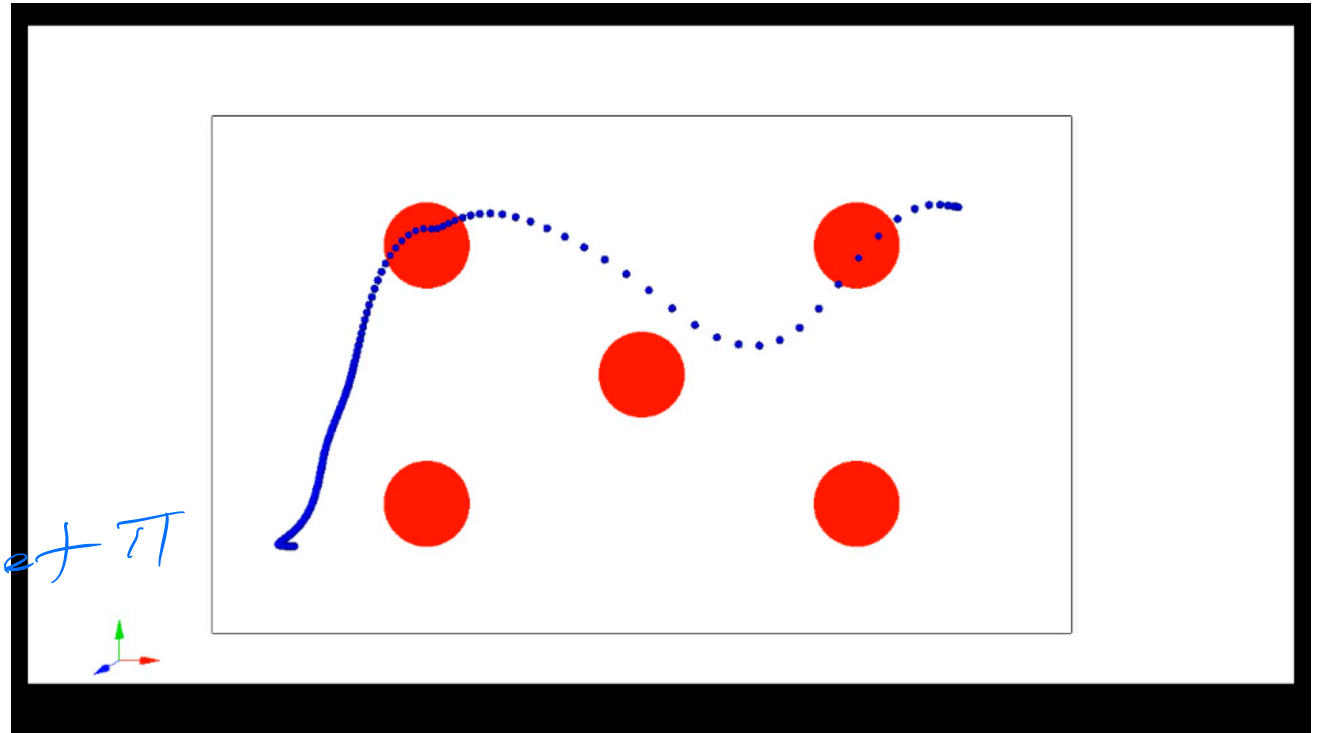
Cost function is designed such that it gets to the goal without colliding with obstacles (in red)

practice: Model Predictive Control (MPC)



for $t = 0, \dots$

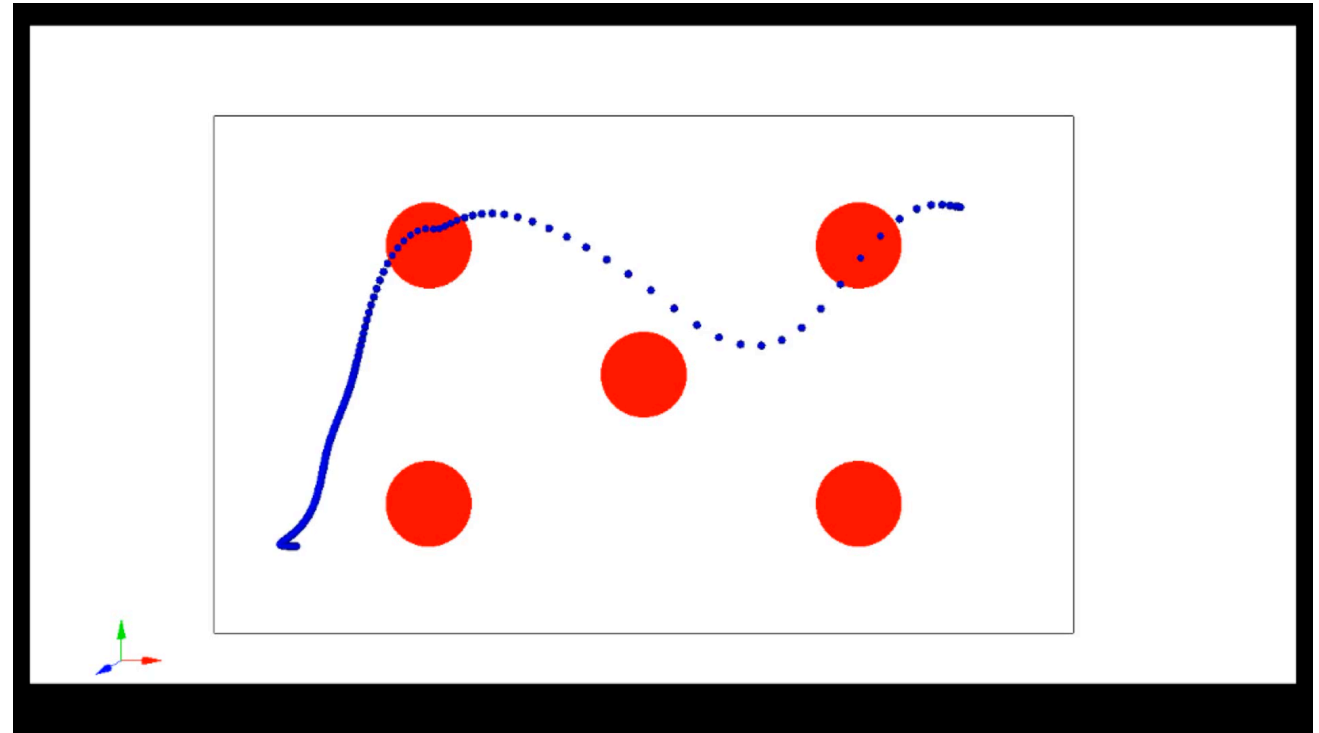
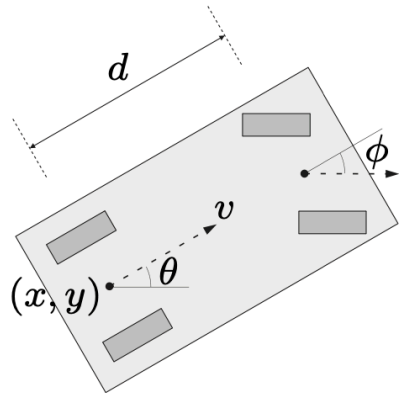
- run ILQR to get π
- $q_t = \pi(s_t)$
- repeat



Example:

2-d car navigation

Cost function is designed such that it gets to the goal without colliding with obstacles (in red)



Today:

Policy Gradient Descent

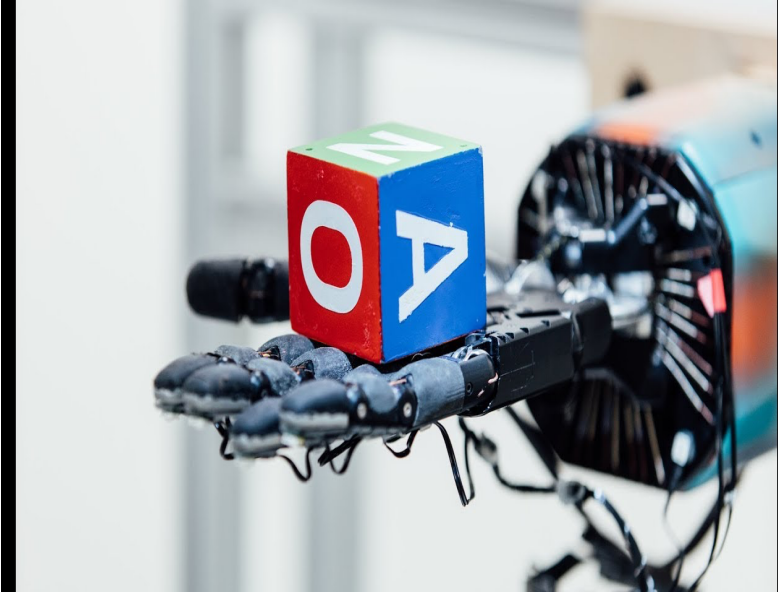
Policy Optimization



[AlphaZero, Silver et.al, 17]



[OpenAI Five, 18]



[OpenAI, 19]

Today: Policy Gradient Derivation

Consider parameterized policy:

$$\pi_{\theta}(a | s) = \pi(a | s; \theta)$$

Today: Policy Gradient Derivation

Consider parameterized policy:

$$\pi_{\theta}(a \mid s) = \pi(a \mid s; \theta)$$

$$\begin{aligned} J(\theta) &:= E_{s_0 \sim \mu_0} [V^{\pi_{\theta}}(s_0)] \\ &= E \left[\sum_{h=0}^{\infty} \gamma^h r_h \mid \mu_0, \pi_{\theta} \right] \end{aligned}$$

Today: Policy Gradient Derivation

Consider parameterized policy:

$$\pi_{\theta}(a | s) = \pi(a | s; \theta)$$

$$J(\theta) := E_{s_0 \sim \mu_0} [V^{\pi_{\theta}}(s_0)]$$

$$= E \left[\sum_{h=0}^{\infty} \gamma^h r_h \mid \mu_0, \pi_{\theta} \right]$$

$$\max_{\theta} J(\theta)$$

$$\theta_{t+1} = \theta_t + \eta \nabla_{\theta} J(\pi_{\theta})|_{\theta=\theta_t}$$



Today: Policy Gradient Derivation

Consider parameterized policy:

$$\pi_{\theta}(a | s) = \pi(a | s; \theta)$$

$$\begin{aligned} J(\theta) &:= E_{s_0 \sim \mu_0} [V^{\pi_{\theta}}(s_0)] \\ &= E \left[\sum_{h=0}^{\infty} \gamma^h r_h \mid \mu_0, \pi_{\theta} \right] \end{aligned}$$

$$\theta_{t+1} = \theta_t + \eta \nabla_{\theta} J(\pi_{\theta}) |_{\theta=\theta_t}$$

Main question for today's lecture:
how to compute the gradient?

Outline for today

1. Recap on Gradient Descent (GD) and Stochastic Gradient Descent (SGD)

2. Warm up: computing gradient using importance weighting

3. Policy Gradient formulations

Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

GD minimizes the above objective function as follows:

Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

GD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_{\theta}(x) - y)^2$)

GD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

$$\theta_{t+1} = \theta_t - \eta \nabla J(\theta_t)$$

Gradient Descent

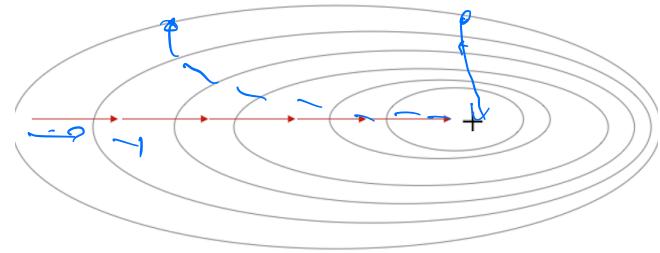
Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y} (f_{\theta}(x) - y)^2$)

GD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

$$\theta_{t+1} = \theta_t - \eta \nabla J(\theta_t)$$

Gradient Descent



Stochastic Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

SGD minimizes the above objective function as follows:

Stochastic Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

SGD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

Stochastic Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

SGD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

$$\theta_{t+1} = \theta_t - \eta_t g_t$$

Stochastic Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

SGD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

$$\theta_{t+1} = \theta_t - \eta_t g_t$$

where $\mathbb{E}[g_t] = \nabla_\theta J(\theta_t)$

Stochastic Gradient Descent

Given an objective function $J(\theta) : \mathbb{R}^d \mapsto \mathbb{R}$, (e.g., $J(\theta) = \mathbb{E}_{x,y}(f_\theta(x) - y)^2$)

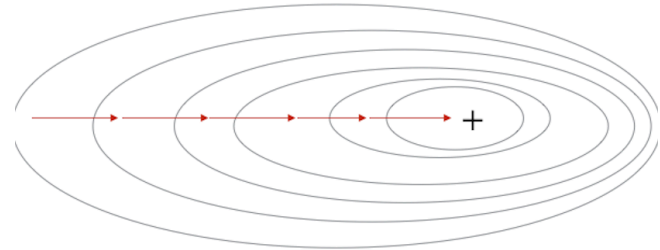
SGD minimizes the above objective function as follows:

Initialize θ_0 , for $t = 0, \dots$:

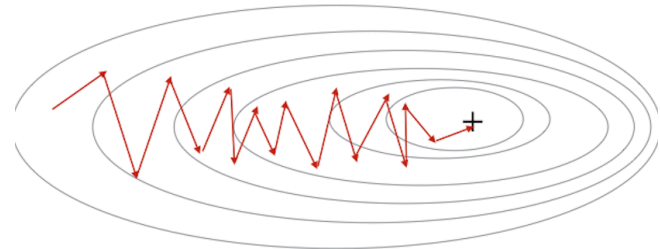
$$\theta_{t+1} = \theta_t - \eta_t g_t$$

where $\mathbb{E}[g_t] = \nabla_\theta J(\theta_t)$

Gradient Descent



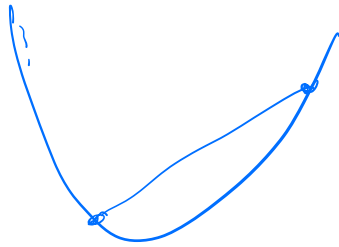
Stochastic Gradient Descent



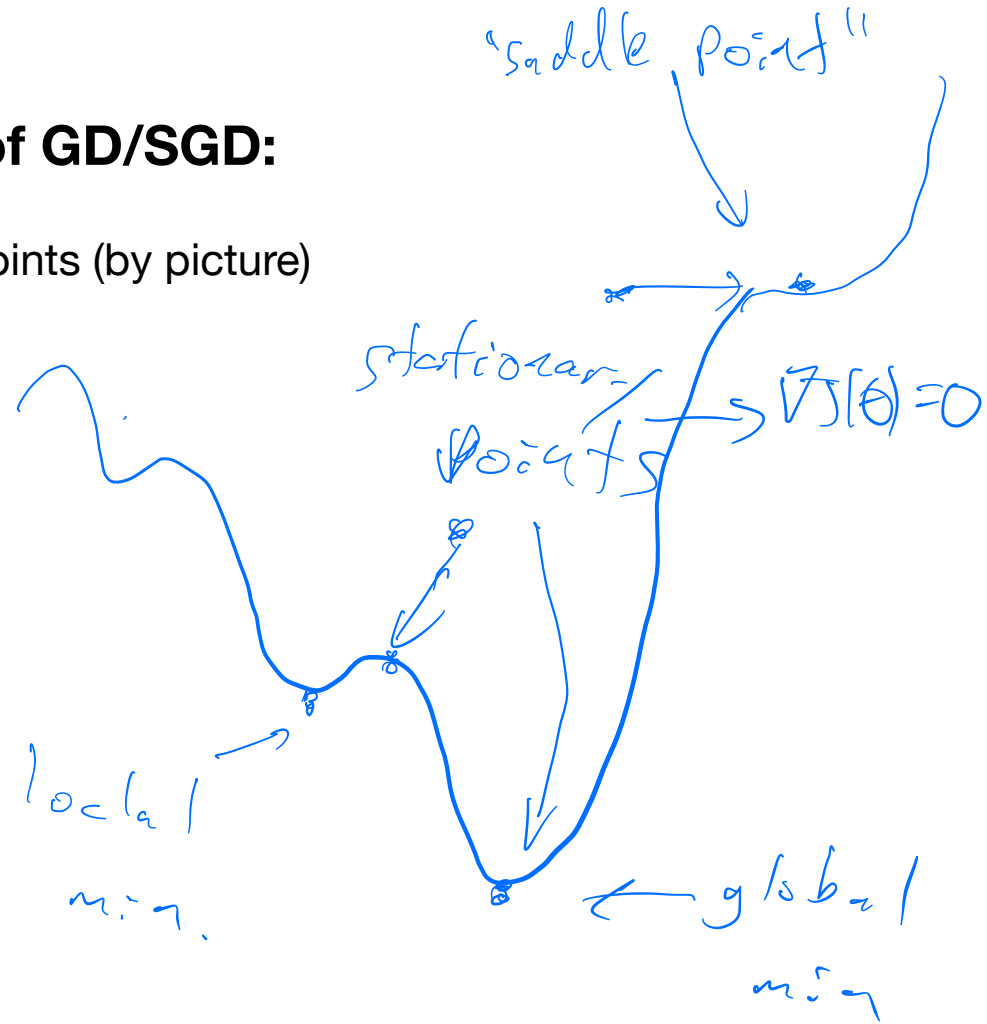
Brief overview of GD/SGD:

Brief overview of GD/SGD:

- Global optima, local optima, and saddle points (by picture)



← convex
function



"saddle point"

stationary
points

$$\nabla J(\theta) = 0$$

local
min.

← global
min

Brief overview of GD/SGD:

- Global optima, local optima, and saddle points (by picture)
- For convex functions (with certain regularity conditions, such as “smoothness”),
 - GD (with an appropriate constant learning rate) converges to the [global optima](#).
 - SGD (with an appropriately decaying learning rate) converges to the [global optima](#).

Brief overview of GD/SGD:

- Global optima, local optima, and saddle points (by picture)
- For convex functions (with certain regularity conditions, such as “smoothness”),
 - GD (with an appropriate constant learning rate) converges to the **global optima**.
 - SGD (with an appropriately decaying learning rate) converges to the **global optima**.
- For non-convex functions, we hope to find a **local minima**.

Brief overview of GD/SGD:

- Global optima, local optima, and saddle points (by picture)
- For convex functions (with certain regularity conditions, such as “smoothness”),
 - GD (with an appropriate constant learning rate) converges to the **global optima**.
 - SGD (with an appropriately decaying learning rate) converges to the **global optima**.
- For non-convex functions, we hope to find a **local minima**.
- What we can prove (under mild regularity conditions) is a little weaker:
 - GD (with an appropriate constant learning rate) converges to a **saddle point**.
 - SGD (with an appropriately decaying learning rate) converges to a **saddle point**.

$$\hat{J}(\theta) = \frac{1}{N} \sum_{i=1}^N (\hat{f}_{\theta}(x_i) - y_i)^2$$

some sort of “nice” stationary point.

SGD: Convergence to a Stationary Point for Nonconvex Functions

- Def of β -smooth: $\|\nabla_{\theta}J(\theta) - \nabla_{\theta}J(\theta_0)\|_2 \leq \beta\|\theta - \theta_0\|_2$

SGD: Convergence to a Stationary Point for Nonconvex Functions

- Def of β -smooth: $\|\nabla_{\theta}J(\theta) - \nabla_{\theta}J(\theta_0)\|_2 \leq \beta\|\theta - \theta_0\|_2$

- **[Theorem]** Suppose we run SGD: $\theta_{t+1} = \theta_t - \eta \widetilde{\nabla}_{\theta}J(\theta_t)$, for T steps, where $\mathbb{E} \left[\widetilde{\nabla}_{\theta}J(\theta_t) \right] = \nabla_{\theta}J(\theta_t)$ with $\eta = O(1/\sqrt{T})$. Assume:

SGD: Convergence to a Stationary Point for Nonconvex Functions

- Def of β -smooth: $\|\nabla_{\theta}J(\theta) - \nabla_{\theta}J(\theta_0)\|_2 \leq \beta\|\theta - \theta_0\|_2$

$$\eta_{\epsilon} = O\left(\frac{1}{\sqrt{\epsilon}}\right)$$

- [Theorem]** Suppose we run SGD: $\theta_{t+1} = \theta_t - \eta \widetilde{\nabla}_{\theta}J(\theta_t)$, for T steps, where $\mathbb{E} \left[\widetilde{\nabla}_{\theta}J(\theta_t) \right] = \nabla_{\theta}J(\theta_t)$ with $\eta = O(1/\sqrt{T})$. Assume:
 - $J(\theta)$ is β -smooth.

SGD: Convergence to a Stationary Point for Nonconvex Functions

- Def of β -smooth: $\|\nabla_{\theta}J(\theta) - \nabla_{\theta}J(\theta_0)\|_2 \leq \beta\|\theta - \theta_0\|_2$

- [Theorem]** Suppose we run SGD: $\theta_{t+1} = \theta_t - \eta \widetilde{\nabla}_{\theta}J(\theta_t)$, for T steps, where $\mathbb{E} \left[\widetilde{\nabla}_{\theta}J(\theta_t) \right] = \nabla_{\theta}J(\theta_t)$ with $\eta = O(1/\sqrt{T})$. Assume:
 - $J(\theta)$ is β -smooth.
 - $J(\theta)$ is bounded: $|J(\theta)| \leq M, \forall \theta$.

SGD: Convergence to a Stationary Point for Nonconvex Functions

- Def of β -smooth: $\|\nabla_{\theta}J(\theta) - \nabla_{\theta}J(\theta_0)\|_2 \leq \beta\|\theta - \theta_0\|_2$

- **[Theorem]** Suppose we run SGD: $\theta_{t+1} = \theta_t - \eta \widetilde{\nabla}_{\theta}J(\theta_t)$, for T steps, where $\mathbb{E} \left[\widetilde{\nabla}_{\theta}J(\theta_t) \right] = \nabla_{\theta}J(\theta_t)$ with $\eta = O(1/\sqrt{T})$. Assume:
 - $J(\theta)$ is β -smooth.
 - $J(\theta)$ is bounded: $|J(\theta)| \leq M, \forall \theta$.
 - $\widetilde{\nabla}_{\theta}J(\theta)$ has “bounded second moment”: $\mathbb{E} \left[\|\widetilde{\nabla}_{\theta}J(\theta_t)\|_2^2 \right] \leq \sigma^2$,

SGD: Convergence to a Stationary Point for Nonconvex Functions

- Def of β -smooth: $\|\nabla_{\theta}J(\theta) - \nabla_{\theta}J(\theta_0)\|_2 \leq \beta\|\theta - \theta_0\|_2$

• **[Theorem]** Suppose we run SGD: $\theta_{t+1} = \theta_t - \eta \widetilde{\nabla}_{\theta}J(\theta_t)$, for T steps, where $\mathbb{E} [\widetilde{\nabla}_{\theta}J(\theta_t)] = \nabla_{\theta}J(\theta_t)$ with $\eta = O(1/\sqrt{T})$. Assume:

- $J(\theta)$ is β -smooth.
- $J(\theta)$ is bounded: $|J(\theta)| \leq M, \forall \theta$.
- $\widetilde{\nabla}_{\theta}J(\theta)$ has “bounded second moment”: $\mathbb{E} [\|\widetilde{\nabla}_{\theta}J(\theta_t)\|_2^2] \leq \sigma^2$,

then, in T steps, SGD will find a θ such that: $\|\nabla_{\theta}J(\theta)\|^2 \leq O\left(\sqrt{M\beta\sigma^2/T}\right)$.

$$\|\nabla J_{\theta}\| \leq \left(\frac{1}{T}\right)^{1/4}$$

Outline for today



1. Recap on Gradient descent and stochastic gradient descent

2. Warm up: computing gradient using importance weighting

3. Policy Gradient formulations

Importance Sampling (and the Likelihood Ratio Method)

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:
 - $J(\theta)$ is “difficult” to compute.
 - P_θ is “easy” to compute.
 - We have a distribution ρ , that is easy to sample from and where $\max P_\theta(x)/\rho(x) < \infty$

$$\begin{aligned}\nabla J(\theta) &= \nabla \sum_x P_\theta(x) f(x) = \nabla \sum_x P(x) \frac{P_\theta(x)}{P(x)} f(x) \\ &= \sum_x P(x) \frac{\nabla P_\theta(x)}{P(x)} f(x) = \mathbb{E}_{x \sim P} \left[\frac{\nabla P_\theta(x)}{P(x)} f(x) \right]\end{aligned}$$

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:
 - $J(\theta)$ is “difficult” to compute.
 - P_θ is “easy” to compute.
 - We have a distribution ρ , that is easy to sample from and where $\max_x P_\theta(x)/\rho(x) < \infty$

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x) = \nabla_\theta \mathbb{E}_{x \sim \rho} \frac{P_\theta(x)}{\rho(x)} f(x)$$

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:
 - $J(\theta)$ is “difficult” to compute.
 - P_θ is “easy” to compute.
 - We have a distribution ρ , that is easy to sample from and where $\max_x P_\theta(x)/\rho(x) < \infty$

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x) = \nabla_\theta \mathbb{E}_{x \sim \rho} \frac{P_\theta(x)}{\rho(x)} f(x) = \mathbb{E}_{x \sim \rho} \frac{\nabla_\theta P_\theta(x)}{\rho(x)} f(x)$$

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:
 - $J(\theta)$ is “difficult” to compute.
 - P_θ is “easy” to compute.
 - We have a distribution ρ , that is easy to sample from and where $\max_x P_\theta(x)/\rho(x) < \infty$

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x) = \nabla_\theta \mathbb{E}_{x \sim \rho} \frac{P_\theta(x)}{\rho(x)} f(x) = \mathbb{E}_{x \sim \rho} \frac{\nabla_\theta P_\theta(x)}{\rho(x)} f(x) \approx \frac{1}{N} \sum_{i=1}^N \frac{\nabla_\theta P_\theta(x_i)}{\rho(x_i)} f(x_i)$$

$$x_i \sim \rho$$

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:
 - $J(\theta)$ is “difficult” to compute.
 - P_θ is “easy” to compute.
 - We have a distribution ρ , that is easy to sample from and where $\max_x P_\theta(x)/\rho(x) < \infty$

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x) = \nabla_\theta \mathbb{E}_{x \sim \rho} \frac{P_\theta(x)}{\rho(x)} f(x) = \mathbb{E}_{x \sim \rho} \frac{\nabla_\theta P_\theta(x)}{\rho(x)} f(x) \approx \frac{1}{N} \sum_{i=1}^N \frac{\nabla_\theta P_\theta(x_i)}{\rho(x_i)} f(x_i)$$

To compute gradient at θ_0 : $\nabla_\theta J(\theta_0)$ (in short of $\nabla_\theta J(\theta) |_{\theta=\theta_0}$)

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:

- $J(\theta)$ is “difficult” to compute.
- P_θ is “easy” to compute.
- We have a distribution ρ , that is easy to sample from and where $\max_x P_\theta(x)/\rho(x) < \infty$

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x) = \nabla_\theta \mathbb{E}_{x \sim \rho} \frac{P_\theta(x)}{\rho(x)} f(x) = \mathbb{E}_{x \sim \rho} \frac{\nabla_\theta P_\theta(x)}{\rho(x)} f(x) \approx \frac{1}{N} \sum_{i=1}^N \frac{\nabla_\theta P_\theta(x_i)}{\rho(x_i)} f(x_i)$$

To compute gradient at θ_0 : $\nabla_\theta J(\theta_0)$ (in short of $\nabla_\theta J(\theta) |_{\theta=\theta_0}$)

By setting the sampling distribution $\rho = P_{\theta_0}$

Importance Sampling (and the Likelihood Ratio Method)

For $J(\theta) = \mathbb{E}_{x \sim P_\theta} [f(x)]$, our goal is to accurately compute $\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x)$.

- Suppose:

- $J(\theta)$ is “difficult” to compute.
- P_θ is “easy” to compute.

- We have a distribution ρ , that is easy to sample from and where $\max_x P_\theta(x)/\rho(x) < \infty$

$$\frac{\nabla_\theta P_\theta(x)}{P_\theta(x)} = \nabla_\theta \log P_\theta(x)$$

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{x \sim P_\theta} f(x) = \nabla_\theta \mathbb{E}_{x \sim \rho} \frac{P_\theta(x)}{\rho(x)} f(x) = \mathbb{E}_{x \sim \rho} \frac{\nabla_\theta P_\theta(x)}{\rho(x)} f(x) \approx \frac{1}{N} \sum_{i=1}^N \frac{\nabla_\theta P_\theta(x_i)}{\rho(x_i)} f(x_i)$$

To compute gradient at θ_0 : $\nabla_\theta J(\theta_0)$ (in short of $\nabla_\theta J(\theta)|_{\theta=\theta_0}$)

By setting the sampling distribution $\rho = P_{\theta_0}$

$$\nabla_\theta J(\theta_0) = \mathbb{E}_{x \sim P_{\theta_0}} \left[\nabla_\theta \ln P_{\theta_0}(x) \cdot f(x) \right]$$

$$\approx \frac{1}{N} \sum_{i=1}^N \nabla_\theta \log P_{\theta_0}(x_i) f(x_i)$$

Importance Sampling (and the Likelihood Ratio Method)

Importance Sampling (and the Likelihood Ratio Method)

To compute gradient at θ_0 : $\nabla_{\theta} J(\theta_0)$ (in short of $\nabla_{\theta} J(\theta) |_{\theta=\theta_0}$)

Importance Sampling (and the Likelihood Ratio Method)

To compute gradient at θ_0 : $\nabla_{\theta} J(\theta_0)$ (in short of $\nabla_{\theta} J(\theta) |_{\theta=\theta_0}$)

By setting the sampling distribution $\rho = P_{\theta_0}$

Importance Sampling (and the Likelihood Ratio Method)

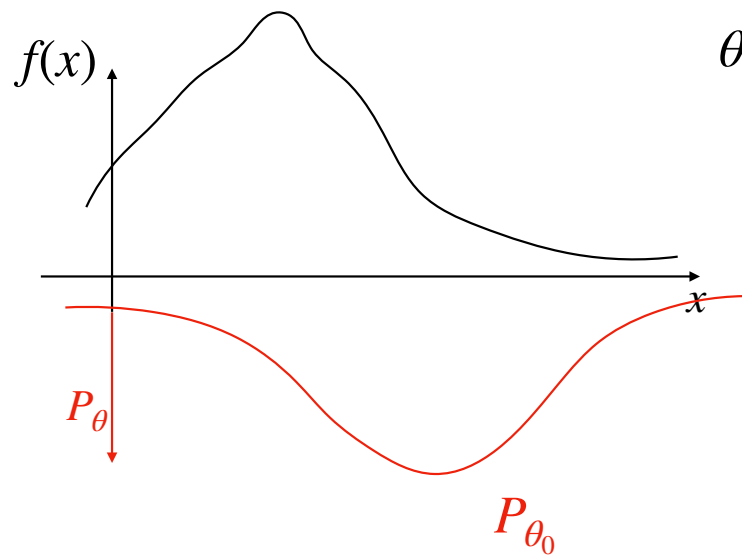
To compute gradient at θ_0 : $\nabla_{\theta} J(\theta_0)$ (in short of $\nabla_{\theta} J(\theta) |_{\theta=\theta_0}$)

By setting the sampling distribution $\rho = P_{\theta_0}$

$$\nabla_{\theta} J(\theta_0) = \mathbb{E}_{x \sim \rho} \frac{\nabla_{\theta} P_{\theta_0}(x)}{\rho(x)} f(x) = \mathbb{E}_{x \sim P_{\theta_0}} \left[\nabla_{\theta} \ln P_{\theta_0}(x) \cdot f(x) \right]$$

Example and Intuition

$$\nabla_{\theta} J(\theta) |_{\theta=\theta_0} = \mathbb{E}_{x \sim P_{\theta_0}} \nabla_{\theta} \ln P_{\theta_0}(x) \cdot f(x)$$

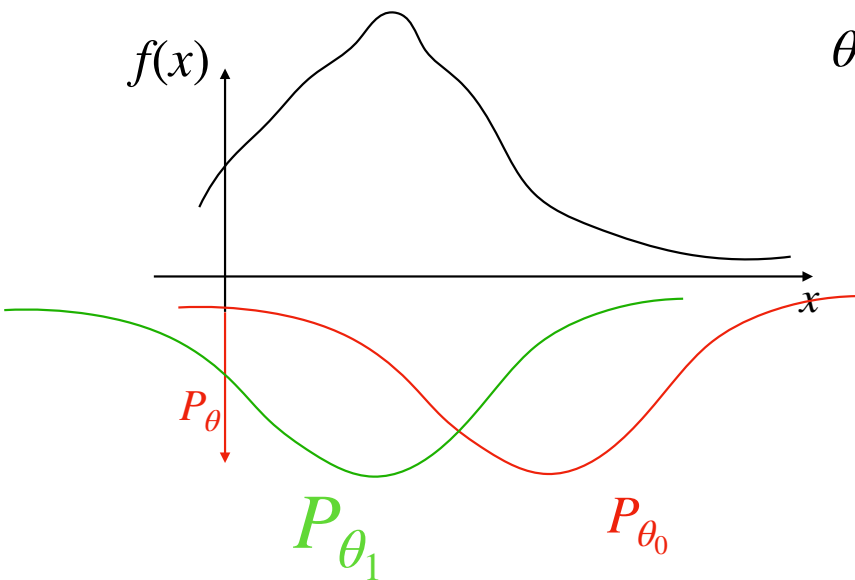


$$\theta_1 = \theta_0 + \eta \nabla_{\theta} J(\theta_0)$$

Example and Intuition

$$\nabla_{\theta} J(\theta) |_{\theta=\theta_0} = \mathbb{E}_{x \sim P_{\theta_0}} \nabla_{\theta} \ln P_{\theta_0}(x) \cdot f(x)$$

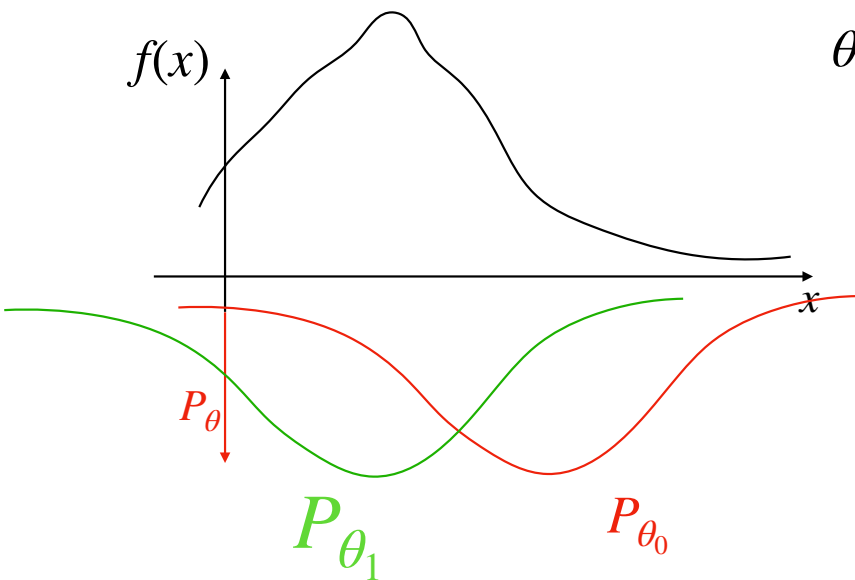
$$\theta_1 = \theta_0 + \eta \nabla_{\theta} J(\theta_0)$$



Example and Intuition

$$\nabla_{\theta} J(\theta) |_{\theta=\theta_0} = \mathbb{E}_{x \sim P_{\theta_0}} \nabla_{\theta} \ln P_{\theta_0}(x) \cdot f(x)$$

$$\theta_1 = \theta_0 + \eta \nabla_{\theta} J(\theta_0)$$

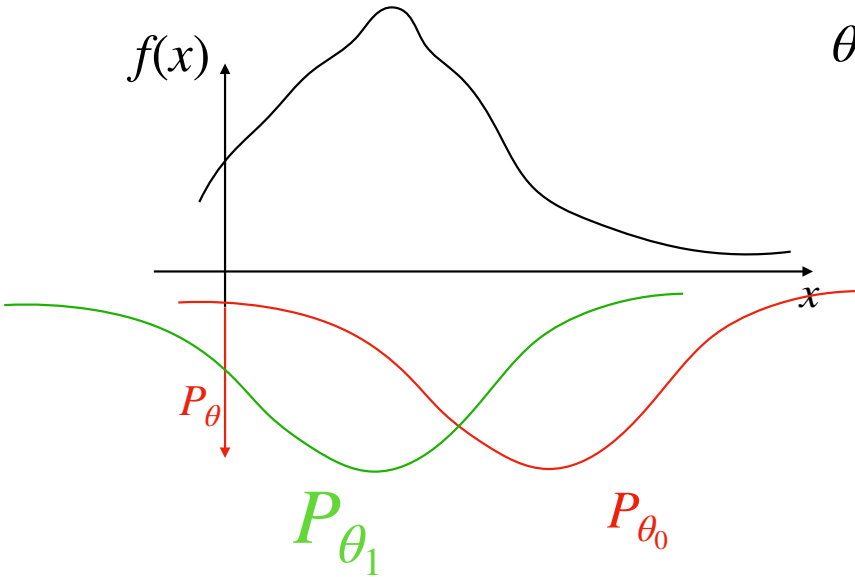


Update distribution (via updating θ) such that P_{θ} has high probability mass at regions where $f(x)$ is large

Example and Intuition

$$\nabla_{\theta} J(\theta) |_{\theta=\theta_0} = \mathbb{E}_{x \sim P_{\theta_0}} \nabla_{\theta} \ln P_{\theta_0}(x) \cdot f(x)$$

$$\theta_1 = \theta_0 + \eta \nabla_{\theta} J(\theta_0)$$



Update distribution (via updating θ) such that P_{θ} has high probability mass at regions where $f(x)$ is large

Using same idea, now let's move on to RL...

Outline for today

✓ 1. Recap on Gradient descent and stochastic gradient descent

✓ 2. Warm up: computing gradient using importance weighting

3. Policy Gradient formulations

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

2. Softmax linear Policy (We will try this in HW2)

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_\theta(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_\theta(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

2. Softmax linear Policy (We will try this in HW2)

Feature vector $\phi(s, a) \in \mathbb{R}^d$, and
parameter $\theta \in \mathbb{R}^d$

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

2. Softmax linear Policy (We will try this in HW2)

Feature vector $\phi(s, a) \in \mathbb{R}^d$, and
parameter $\theta \in \mathbb{R}^d$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta^{\top} \phi(s, a))}{\sum_{a'} \exp(\theta^{\top} \phi(s, a'))}$$

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

2. Softmax linear Policy (We will try this in HW2)

Feature vector $\phi(s, a) \in \mathbb{R}^d$, and
parameter $\theta \in \mathbb{R}^d$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta^{\top} \phi(s, a))}{\sum_{a'} \exp(\theta^{\top} \phi(s, a'))}$$

3. Neural Policy:

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

2. Softmax linear Policy (We will try this in HW2)

Feature vector $\phi(s, a) \in \mathbb{R}^d$, and
parameter $\theta \in \mathbb{R}^d$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta^{\top} \phi(s, a))}{\sum_{a'} \exp(\theta^{\top} \phi(s, a'))}$$

3. Neural Policy:

Neural network
 $f_{\theta} : S \times A \mapsto \mathbb{R}$

Policy Gradient: Examples of Policy Parameterization (discrete actions)

Recall that we consider parameterized policy $\pi_{\theta}(\cdot | s) \in \Delta(A), \forall s$

1. Softmax Policy for discrete MDPs:

$$\theta_{s,a} \in \mathbb{R}, \forall s, a \in S \times A$$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$$

2. Softmax linear Policy (We will try this in HW2)

Feature vector $\phi(s, a) \in \mathbb{R}^d$, and
parameter $\theta \in \mathbb{R}^d$

$$\pi_{\theta}(a | s) = \frac{\exp(\theta^{\top} \phi(s, a))}{\sum_{a'} \exp(\theta^{\top} \phi(s, a'))}$$

3. Neural Policy:

Neural network
 $f_{\theta} : S \times A \mapsto \mathbb{R}$

$$\pi_{\theta}(a | s) = \frac{\exp(f_{\theta}(s, a))}{\sum_{a'} \exp(f_{\theta}(s, a'))}$$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0|s_0)P(s_1|s_0, a_0)\pi_{\theta}(a_1|s_1)\dots$$

\swarrow $P_{\theta}(\tau) = \text{Prob of } \tau \text{ under } \pi_{\theta}$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0 | s_0)P(s_1 | s_0, a_0)\pi_{\theta}(a_1 | s_1)\dots$$

$$J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \underbrace{\left[\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h) \right]}_{R(\tau)}$$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0 | s_0)P(s_1 | s_0, a_0)\pi_{\theta}(a_1 | s_1)\dots$$

$$J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\underbrace{\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h)}_{R(\tau)} \right]$$

$$\nabla_{\theta} J(\pi_{\theta_0}) = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \ln \rho_{\theta_0}(\tau) R(\tau) \right]$$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0 | s_0)P(s_1 | s_0, a_0)\pi_{\theta}(a_1 | s_1)\dots$$

$$J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \underbrace{\left[\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h) \right]}_{R(\tau)}$$

$$\nabla_{\theta} J(\pi_{\theta_0}) = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \ln \rho_{\theta_0}(\tau) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \mu(s_0) + \ln \pi_{\theta_0}(a_0 | s_0) + \ln P(s_1 | s_0, a_0) + \dots \right) R(\tau) \right]$$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0 | s_0)P(s_1 | s_0, a_0)\pi_{\theta}(a_1 | s_1)\dots$$

$$J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \underbrace{\left[\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h) \right]}_{R(\tau)}$$

$$\nabla_{\theta} J(\pi_{\theta_0}) \Big|_{\theta = \theta_0} = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \ln \rho_{\theta_0}(\tau) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \rho(s_0) + \ln \pi_{\theta_0}(a_0 | s_0) + \ln P(s_1 | s_0, a_0) + \dots \right) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \pi_{\theta_0}(a_0 | s_0) + \ln \pi_{\theta_0}(a_1 | s_1) \dots \right) R(\tau) \right]$$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0 | s_0)P(s_1 | s_0, a_0)\pi_{\theta}(a_1 | s_1)\dots$$

$$J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \underbrace{\left[\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h) \right]}_{R(\tau)}$$

$$\nabla_{\theta} J(\pi_{\theta_0}) = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \ln \rho_{\theta_0}(\tau) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \rho(s_0) + \ln \pi_{\theta_0}(a_0 | s_0) + \ln P(s_1 | s_0, a_0) + \dots \right) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \pi_{\theta_0}(a_0 | s_0) + \ln \pi_{\theta_0}(a_1 | s_1) \dots \right) R(\tau) \right] = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta_0}(a_h | s_h) \right) R(\tau) \right]$$

Derivation of Policy Gradient: REINFORCE

$$\tau = \{s_0, a_0, s_1, a_1, \dots\}$$

$$\rho_{\theta}(\tau) = \mu(s_0)\pi_{\theta}(a_0 | s_0)P(s_1 | s_0, a_0)\pi_{\theta}(a_1 | s_1)\dots$$

$$J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \underbrace{\left[\sum_{h=0}^{\infty} \gamma^h r(s_h, a_h) \right]}_{R(\tau)}$$

Adjust policy such that
larger reward traj has
higher likelihood

$$\nabla_{\theta} J(\pi_{\theta_0}) = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \ln \rho_{\theta_0}(\tau) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \rho(s_0) + \ln \pi_{\theta_0}(a_0 | s_0) + \ln P(s_1 | s_0, a_0) + \dots \right) R(\tau) \right]$$

$$= \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\nabla_{\theta} \left(\ln \pi_{\theta_0}(a_0 | s_0) + \ln \pi_{\theta_0}(a_1 | s_1) \dots \right) R(\tau) \right] = \mathbb{E}_{\tau \sim \rho_{\theta_0}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta_0}(a_h | s_h) \right) R(\tau) \right]$$

Summary so far for Policy Gradients

We derived the most basic PG formulation:

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right]$$

Summary so far for Policy Gradients

We derived the most basic PG formulation:

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right]$$

Increase the likelihood of sampling an trajectory with high total reward

Obtaining a sample $\widetilde{\nabla}_{\theta} J(\theta)$ for REINFORCE (for this approach)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right], \text{ where } R(\tau) = \sum_{h=0}^{\infty} \gamma^h r(s_h, a_h)$$

• run π_{θ} on $s_0 \sim \mu$ to get τ

$$\bullet \quad \widetilde{\nabla}_{\theta} J(\theta) = \left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) \left(\sum_{h=0}^{\infty} \gamma^h r_h \right)$$

For finite horizon MDP (sometimes used with PG):

$$J(\theta) = \mathbb{E} \left[\sum_{h=0}^{H-1} r(s_h, a_h) \right]_{\rho_0, \pi_\theta}$$

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_\theta(\tau)} \left[\left(\sum_{h=0}^{H-1} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right]$$

$$\text{where } R(\tau) = \sum_{h=0}^{H-1} r(s_h, a_h)$$

For finite horizon MDP (sometimes used with PG):

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{H-1} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right]$$

$$\text{where } R(\tau) = \sum_{h=0}^{H-1} r(s_h, a_h)$$

Increase the likelihood of sampling an trajectory with high total reward

A improved PG formulation, for sampling (for the discounted setting)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right]$$

A improved PG formulation, for sampling (for the discounted setting)

$$\begin{aligned}\nabla J(\theta) &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \sum_{t=h}^{\infty} \gamma^t r_t \right) \right]\end{aligned}$$

A improved PG formulation, for sampling (for the discounted setting)

$$\begin{aligned}\nabla J(\theta) &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \sum_{t=h}^{\infty} \gamma^t r_t \right) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \gamma^h Q^{\pi_{\theta}}(s_h, a_h) \right) \right]\end{aligned}$$

A improved PG formulation, for sampling (for the discounted setting)

$$\begin{aligned}\nabla J(\theta) &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \sum_{t=h}^{\infty} \gamma^t r_t \right) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \gamma^h Q^{\pi_{\theta}}(s_h, a_h) \right) \right]\end{aligned}$$

Intuition: Change action distribution at h only affects rewards later on...)

A improved PG formulation, for sampling (for the discounted setting)

$$\begin{aligned}\nabla J(\theta) &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \sum_{t=h}^{\infty} \gamma^t r_t \right) \right] \\ &= \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \gamma^h Q^{\pi_{\theta}}(s_h, a_h) \right) \right]\end{aligned}$$

Intuition: Change action distribution at h only affects rewards later on...)

Exercise: Show this simplified version is equivalent to REINFORCE

A improved PG formulation, for sampling (for the discounted setting)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \sum_{t=h}^{\infty} \gamma^t r_t \right) \right]$$

Further simplification on PG (e.g., for finite horizon)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{H-1} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \cdot \sum_{\tau=h}^{H-1} r(s_{\tau}, a_{\tau}) \right) \right]$$

Further simplification on PG (e.g., for finite horizon)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{H-1} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \cdot \sum_{\tau=h}^{H-1} r(s_{\tau}, a_{\tau}) \right) \right]$$

(Change action distribution at h only affects rewards later on...)

Further simplification on PG (e.g., for finite horizon)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{H-1} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \cdot \sum_{\tau=h}^{H-1} r(s_{\tau}, a_{\tau}) \right) \right]$$

(Change action distribution at h only affects rewards later on...)

Exercise:

Show this simplified version is equivalent to REINFORCE

Summary for today

1. Importance Weighting (the likelihood ratio method)

2. The Policy Gradient:

REINFORCE (a direct application of the likelihood ratio method)

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim \rho_{\theta}(\tau)} \left[\left(\sum_{h=0}^{\infty} \nabla_{\theta} \ln \pi_{\theta}(a_h | s_h) \right) R(\tau) \right]$$

3. **SGAscent** With unbiased estimate of $\nabla_{\theta} J(\theta)$, SGA(hopefully) converges to a local optimal policy.

1-minute feedback form: <https://bit.ly/3RHtlxy>

