

Policy Gradient ²

DOROZHKO Anton

Novosibirsk State University

June 6, 2020

²SpinningUP RL research : Part 3

Outline

- 1 Recap Q-learning
- 2 Policy Gradient

Q-learning

MDP $(S, A, \cancel{R}, \cancel{R}, \gamma)$

goal $\max_{\pi} \mathbb{E}[G_t]$

$$\hat{Q}(s, a) = \mathbb{E}_{\pi} [G_t | a, s]$$

$$\downarrow$$

$$\pi(a|s) = \underset{a}{\operatorname{argmax}} \hat{Q}(s, a) \quad | \quad \epsilon\text{-greedy}$$

Action-Value approximation

- Find w that minimizes loss between $Q^\pi(s, a)$ and $\hat{Q}(s, a; w)$

target *NN*
- MSE

$$J(w) = \mathbb{E}_\pi[(Q^\pi(s, a) - \hat{Q}(s, a; w))^2]$$

- SGD samples the gradient

$$-\frac{1}{2} \nabla_w J(w) = \mathbb{E}_\pi[(Q^\pi(s, a) - \hat{Q}(s, a; w)) \nabla_w \hat{Q}^\pi(s, a; w)]$$

$$\Delta w = -\frac{1}{2} \alpha \nabla_w J(w)$$

Incremental Model-Free Control

$$r_0 + \gamma r_1 + \dots + \gamma^i r_i$$

- MC
 - 200 apm
 - 5 minutes
 - 1000 steps
 - $\Delta w = \alpha (G_t - \hat{Q}^\pi(s_t, a_t; w)) \nabla_w \hat{Q}(s_t, a_t; w)$



- SARSA

$$\Delta w = \alpha (r + \gamma \hat{Q}(s', a'; w) - \hat{Q}^\pi(s_t, a_t; w)) \nabla_w \hat{Q}(s_t, a_t; w)$$

- Q-learning

$$G_t = r_t + \gamma R_{t+1} + \dots$$

$$\Delta w = \alpha (r + \gamma \max_a \hat{Q}(s', a'; w) - \hat{Q}^\pi(s_t, a_t; w)) \nabla_w \hat{Q}(s_t, a_t; w)$$

$$\gamma = 0$$

$$Q \sim r_t$$

$$\gamma \rightarrow 1$$

$$Q \approx r_t + \gamma r_{t+1} + \dots + \gamma^i r_i$$

What if ?

$$\hat{Q}(s, a; w) \rightarrow \pi$$

What if we can directly optimize our expected reward w.r.t. the parameters of our policy ?

$$\pi_{\theta}(a|s) \quad J_{\theta}(\pi) = \mathbb{E}[G_t]$$

$$\nabla_{\theta} J_{\theta}(\pi)$$

Policy Optimization

- Stochastic, parametrized policy $\pi_\theta \rightarrow f(\theta) : \mathcal{S} \rightarrow \Delta A$
- maximize expected return

- SGD with **policy gradient**

$$J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [G_t(\tau)]$$

$$\theta_{k+1} = \theta_k + \alpha \nabla_{\theta} J(\pi_\theta)|_{\theta_k}$$

NN : $x(s) \rightarrow$
logits
over
actions

τ trajectory
($s_0, a_0, s_1, a_1, \dots$)


r_0, r_1, \dots

How to compute policy gradient

- Derive the analytical gradient
- Compute sample estimate

1) Probability of a Trajectory

Trajectory



$$\tau = (s_0, a_0, \dots, s_{T+1})$$

Probability of a Trajectory

$$P(\tau|\theta) = \underbrace{\rho_0(s_0)}_{\text{Markov property}} \prod_{t=0}^T \underbrace{P(s_{t+1}|s_t, a_t)}_{\text{Markov property}} \underbrace{\pi_\theta(a_t, s_t)}_{\text{Markov property}}$$

Take log

$$\log P(\tau|\theta) = \log \rho_0(s_0) + \sum_t \log P(s_{t+1}|s_t, a_t) + \sum_t \log \pi_\theta(a_t, s_t)$$

2) The Log-Derivative Trick

$$(\log x)' = \frac{1}{x}$$

$$\nabla_x \log f(x) = \frac{1}{f(x)} \underbrace{\nabla_x f(x)}_{\text{? hard}}$$

easy →

$$\nabla_{\theta} \log P(\tau | \theta) = \frac{\nabla_{\theta} P(\tau | \theta)}{P(\tau | \theta)}$$

Rearrange

$$\nabla_{\theta} P(\tau | \theta) = \underbrace{P(\tau | \theta)}_{f(x)} \nabla_{\theta} \log f(x)$$

3) Gradients of Environment Functions

$$\begin{aligned} \underline{\nabla_{\theta} J(\pi_{\theta})} &= \nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] \\ &= \underline{\nabla_{\theta}} \int R(\tau) \cdot P(\tau | \theta) \end{aligned}$$

- $\rho_0(s_0)$, $P(s_{t+1}|s_t, a_t)$ and $G_t(\tau)$
- no dependence on θ
- Gradients w.r.t θ are zero

$$\mathbb{E}_{x \sim p} f(x) = \int f(x) \cdot p(x) dx$$

Grad-Log-Prob of Trajectory

$$\begin{aligned}
 \nabla_{\theta} \log P(\tau|\theta) &= \cancel{\nabla_{\theta} \log \rho_0(s_0)} + \sum_{t=0}^T \left(\cancel{\nabla_{\theta} \log P(s_{t+1}|s_t, a_t)} \right. \\
 &\quad \left. + \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) \right) \\
 &= \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t|s_t).
 \end{aligned}$$

$$\underline{\nabla_{\theta} J(\pi_{\theta})} = \dots$$

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] = \nabla_{\theta} \int R(\tau) \cdot P(\tau|\theta)$$

$$= \int R(\tau) \underline{\nabla_{\theta} P(\tau|\theta)}$$

$$= \int R(\tau) \underline{P(\tau|\theta)} \nabla_{\theta} \log P(\tau|\theta)$$

$$= \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau) \nabla_{\theta} \log P(\tau|\theta)]$$

$$= \mathbb{E}_{\tau \sim \pi} [R(\tau) \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t)]$$

Basic Policy Gradient

Derivation for Basic Policy Gradient


$$\begin{aligned}
 \nabla_{\theta} J(\pi_{\theta}) &= \nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] && \text{def} \\
 &= \nabla_{\theta} \int_{\tau} P(\tau|\theta) R(\tau) && \text{Expand expectation} \\
 &= \int_{\tau} \nabla_{\theta} P(\tau|\theta) R(\tau) && \text{Bring gradient under integral} \\
 &= \int_{\tau} P(\tau|\theta) \nabla_{\theta} \log P(\tau|\theta) R(\tau) && \text{Log-derivative trick} \\
 &= \mathbb{E}_{\tau \sim \pi_{\theta}} [\nabla_{\theta} \log P(\tau|\theta) R(\tau)] && \text{Return to expectation form} \\
 \therefore \nabla_{\theta} J(\pi_{\theta}) &= \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) R(\tau) \right] && \text{Expression for grad-log-prob}
 \end{aligned}$$

policy gradient

REINFORCE

How to compute expectation ?

$$R(\tau) = r_0 + \gamma r_1 + \dots$$



$$\mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) R(\tau) \right]$$

$$\hat{g} = \frac{1}{N} \sum_{i=1}^N \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) R(\tau)$$

$$\nabla_{\theta} J$$

N games with \mathbb{J}_{θ}

$$\theta_{k+1} = \theta_k - \alpha \nabla_{\theta} J$$

Basic Policy Gradient in words

$$Q_w(s, a) : S \times A \rightarrow \mathbb{R}$$

$$\pi_\theta(a|s) : S \rightarrow \Delta A$$

- Define policy as parametrized function
- Write the gradient of the loss function
- Play games collect trajectories (1 or more)
- Compute cumulative discounted rewards for each t
- Compute gradients $\nabla_\theta J(\theta)$
- Update parameters of your policy
- **Comment: throw out trajectories**

$$R(r)$$

after each step of SGD

$$\theta_{k+1} = \theta_k - \alpha \nabla_\theta J$$

Loss is not a normal loss

$$\text{MSE loss}$$

$$\frac{1}{N} \sum (f(x) - y)^2$$

- ① Data is dependent on the parameters π_θ
- ② Is not connected to performance $J(\pi_\theta)$ after one step
- ③ So, minimization of this "loss" itself for several steps on the same batch will not give better rewards.

Consider only future rewards



for the whole game

$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) R(\tau) \right]$$

Lemma

$$\mathbb{E}_{x \sim P_\theta} [\nabla_\theta \log P_\theta(x)] = 0$$

Proof: distributions are normalized \rightarrow take gradient \rightarrow log derivative trick

$$\int_x P_\theta(x) = 1$$

$$\nabla_\theta \int_x P_\theta(x) = \nabla_\theta 1 = 0$$

$$\nabla_\theta \int_x P_\theta(x) = \int_x P_\theta(x) \nabla_\theta \log P_\theta(x) = \underline{\mathbb{E}_{x \sim P_\theta} [\nabla_\theta \log P_\theta(x)]} \quad \square$$

log deriv
trick

Proof of go-to-reward usage


$$\mathbb{E}_{\tau} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta} \left[\sum_{t'=1}^T R \right] \right] = \mathbb{E}_{\tau} \left[\sum_t \nabla_{\theta} \log \pi_{\theta} \left[\sum_{t'=1}^{t-1} R + \sum_{t'=t}^T R \right] \right]$$

$$\nabla_{\theta} J_{\theta}(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \sum_{t'=t}^T R(s_{t'}, a_{t'}, s_{t'+1}) \right]$$

$$= \sum_t \mathbb{E}_{\tau} \nabla_{\theta} \log \pi_{\theta} \left[\sum_{t'=1}^{t-1} R \right] + \dots$$

doesn't depend on τ after t

Consider only future rewards



$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \sum_{t'=t}^T R(s_{t'}, a_{t'}, s_{t'+1}) \right]$$

Why it's better ?

Recall the sum of random variables with 0 mean and some variance.

$$\sum_{t=0}^T R \sim \text{i.i.d.} \quad \text{Var}(R)$$

Improvement: Baseline

Consequence of lemma

For $\forall b(s)$:

$$\mathbb{E}_{a_t \sim \pi_\theta} [\nabla_\theta \log \pi_\theta(a_t | s_t) b(s_t)] = 0$$

const w.r.t. a_t

This allows:

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \nabla_\theta \log \pi_\theta(a_t | s_t) \left(\underbrace{\sum_{t'=t}^T R(s_{t'}, a_{t'}, s_{t'+1})}_{\hat{Q}} - \underbrace{b(s_t)}_{\frac{\partial \text{Var}}{\partial b} = 0} \right) \right]$$

Var

We can choose b to reduce variance. In practice $b(s_t) = V^\pi(s_t)$ value function works well.

Basic Policy Gradient

Algorithm 1 Vanilla Policy Gradient Algorithm

- 1: Input: initial policy parameters θ_0 , initial value function parameters ϕ_0
- 2: **for** $k = 0, 1, 2, \dots$ **do**
- 3: Collect set of trajectories $\mathcal{D}_k = \{\tau_i\}$ by running policy $\pi_k = \pi(\theta_k)$ in the environment.
- 4: Compute rewards-to-go \hat{R}_t .
- 5: Compute advantage estimates, \hat{A}_t (using any method of advantage estimation) based on the current value function V_{ϕ_k} .
- 6: Estimate policy gradient as

$$\hat{g}_k = \frac{1}{|\mathcal{D}_k|} \sum_{\tau \in \mathcal{D}_k} \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) |_{\theta_k} \hat{A}_t$$

- 7: Compute policy update, either using standard gradient ascent,

$$\theta_{k+1} = \theta_k + \alpha_k \hat{g}_k,$$

or via another gradient ascent algorithm like Adam.

- 8: Fit value function by regression on mean-squared error:

$$\phi_{k+1} = \arg \min_{\phi} \frac{1}{|\mathcal{D}_k| T} \sum_{\tau \in \mathcal{D}_k} \sum_{t=0}^T (V_{\phi}(s_t) - \hat{R}_t)^2,$$

typically via some gradient descent algorithm.

- 9: **end for**

policy



$$Q(s, a) = V(s) + A(a)$$

const $\forall a$

101
102
99

+100

+1
+2
-1

Policy Gradient with automatic differentiation

$$\nabla_{\theta} J(\theta) \approx \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \nabla_{\theta} \log \pi_{\theta}(a_{i,t} | s_{i,t}) \hat{G}_{i,t}$$

How can we compute policy gradient with autodiff ?
 Create a function s.t. it's gradient is policy gradient.
 You will implement "pseudo-loss"

$$J(\theta) \approx \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \log \pi_{\theta}(a_{i,t} | s_{i,t}) \hat{G}_{i,t}$$

cross entropy (discrete) or squared error (gaussian)

It is good if you can ...

- 1 Define Markov Decision Process
- 2 Describe some task in terms of MDP
- 3 Understand value and action-value functions
- 4 Be able to apply Policy Iteration and Value Iteration
- 5 Understand model-based vs model-free
- 6 Be able to apply MC, TD(0) algorithms to Q-function
- 7 Be able to apply Q-learning, SARSA, Expected value SARSA
- 8 Describe the "Exploitation / Exploration" dilemma
- 9 Understand ϵ -Greedy and "Optimism in the face of uncertainty" (UCB) ideas
- 10 Understand the idea of Policy Optimization (Policy Gradient)