

RECALL: REINFORCE WITH BASELINE

Want: X, Y positively correlated;

FOR REINFORCE:

Base line

$$\theta \leftarrow \theta + \alpha \sum_{t=0}^{L-1} \delta^t \left(\overbrace{g(s_t, a_t)}^{\text{Base line}} - \underline{f(s_t)} \right) \nabla \log \pi_{\theta}(a_t | s_t)$$

Here: $f: S \rightarrow \mathbb{R}$ is a function indep of τ a traj

For any function $g(s)$,

$$\begin{aligned} & \mathbb{E} \left[\nabla \log \pi(a|s) g(s) \right] \\ &= \sum_a \pi(a|s) \nabla \log(\pi(a|s)) g(s) \\ &= \sum_a \frac{\pi(a|s)}{\pi(a|s)} \nabla \pi(a|s) g(s) \\ &= g(s) \sum_a \nabla \pi(a|s) = g(s) \nabla \sum_a \pi(a|s) \\ &= g(s) \nabla (1) = \underline{\underline{0}} \end{aligned}$$

Recall - Reinforce algorithm.

Algorithm: Stochastic Gradient Ascent on J .

REINFORCE -

- Initialize θ arbitrarily.
- for each episode k_0

Generate $S_0, A_0, R_0, S_1, A_1, \dots, S_{L-1}, A_{L-1}, R_{L-1}$

using θ .

For each t compute $G_t = \overbrace{g^{\pi_\theta}(s_t, a_t)}$

$$\nabla J(\theta) = \sum_{t=0}^{L-1} \delta^t \underbrace{G_t}_{\text{circled}} \frac{\partial \ln \pi(s_t | A_t; \theta)}{\partial \theta}$$

$$\theta = \theta + \alpha \nabla J(\theta)$$

end
→

The modification with baseline:

$$D_{t+1}^a = \theta_t + \alpha (r_t - b(s_t)) \frac{\nabla \pi(A_t | s_t; \theta)}{\pi(A_t | s_t; \theta)}$$

A natural choice of the baseline is an estimate of the state value $v(s_t, \omega)$, where $\omega \in \mathbb{R}^d$ is a wt vector learned by a method to be described!

Softmax policies: 1) $\pi_{\theta}(a|s) = \frac{\exp(\theta_{s,a})}{\sum_{a'} \exp(\theta_{s,a'})}$ | $\theta = \mathbb{R}^{|S||A|}$

2) Linear softmax policies:

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

For each s, a we have a feature map $\phi_{s,a} \in \mathbb{R}^d$; then $\theta \in \mathbb{R}^d$.

3) Neural softmax policies:

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

$f_{\theta}(s,a)$ parametrized by a neural network with $\theta \in \mathbb{R}^d$

Recall: in policy gradient, given a trajectory τ ,
we define $R(\tau) = (1-\gamma) \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$

↑ discounted total reward

and $V^{\pi_{\theta}}(\mu) = \mathbb{E}_{z \sim P_{z|\mu}^{\pi_{\theta}}} [R(z)]$

and we were optimizing $V^{\pi_{\theta}}(\mu)$ (i.e.) maximizing expected discounted total reward.

Using softmax policies,

$$\frac{\partial V^{\pi_{\theta}}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} A^{\pi_{\theta}}(s) \pi_{\theta}(a|s) A^{\pi_{\theta}}(s,a)$$

↑ check.
easy calculation

- However the gradients with respect to $\theta_{s,a}$ may be small even if $A(s,a)$ is large because $\pi_{\theta}(s|a)$ is small;
- To prevent probabilities from getting too small

add a regularizer and optimize

$$L_\lambda(\theta) := V^{\pi_\theta}(\mu) - \lambda \mathbb{E}_{s \sim \text{Units}} \left[\text{KL}(u_{\text{inf}, A}, \pi_\theta(\cdot|s)) \right]$$

$$\text{KL}(p, q) := \mathbb{E}_{z \sim p} -\log \left(\frac{q(z)}{p(z)} \right)$$

Policy gradient for $L_\lambda(\theta) =$

$$\underline{\theta^{t+1} \leftarrow \theta^t + \eta \nabla_{\theta} L_\lambda(\theta^t)}$$

Thm 1 Starting with any initial θ^0 and using

$$\lambda = \frac{\varepsilon(1-\gamma)}{2 \left\| \frac{d_{\pi^*}}{\mu} \right\|_{\infty}} \quad \text{and} \quad \eta = 1/\beta, \quad \beta := \frac{8\sigma}{(1-\gamma)^2} + \frac{2\lambda}{|S|}$$

we have, for all starting distributions ρ ,

$$\max_{t < T} \left\{ V^*(s) - V^t(s) \right\} \leq \varepsilon, \quad \text{for } T \geq O\left(\frac{|S|^2 |A|}{(1-\gamma)^2 \varepsilon^2} \left\| \frac{d_{\pi^*}}{\mu} \right\|_{\infty}^2 \right)$$

Natural gradients:

Idea: $f(x+y) = f(x) + \nabla f(x)^T y + \frac{1}{2} y^T H(f)_x y.$

A stationary pt in the nbd satisfies:

$$\nabla f(x) + H(f)_x y = 0.$$

$$y = -H(f)_x^{-1} \nabla f(x).$$

System matrix.

- The metric used to measure distances in the nbd is the one induced by the Hessian of f at x .
- Amari - argued that for policy gradient algorithms, we shouldn't use the standard Euclidean metric;

If the function being optimized is the loss function of a parametrized distribution $f(\pi_\theta)$, f : loss, π_θ the distribution, Amari

suggested using **FISCHER INFORMATION MATRIX**, as the choice for the Symm matrix giving us the optimal form;

• for each $s \in S$, $\pi_\theta(\cdot | s) \in$ a distribution.

$$F_\mu^\theta = \mathbb{E}_{s \sim d_\mu^{\pi_\theta}} \mathbb{E}_{a \sim \pi_\theta(\cdot | s)} \left[\nabla \log \pi_\theta(a | s) \nabla \log \pi_\theta(a | s)^\top \right]$$

• Natural policy gradient: NPG

$$\theta \leftarrow \theta + \eta (F_\mu^\theta)^\top \nabla V^{\pi_\theta}(\mu).$$

• Motivation for using that? \uparrow

For a distribution v over $S \times A$, and $\omega \in \mathbb{R}^d$ (the # parameters), define the compatible function approximator error L_v and minimal compatible function approximator error L_v^* as

$$L_v(\omega; \theta) = \mathbb{E}_{s, a \sim v} \left[\left(A^{\pi_\theta}(s, a) - \omega^T \nabla_\theta \log \pi_\theta(a|s) \right)^2 \right]$$

$$g_v: v(s, a) = d_{\mu}^{\pi_\theta} \pi_\theta(a|s)$$

$$\left(A^{\pi_\theta}(s, a) - \omega^T \nabla_\theta \log \pi_\theta(a|s) \right)^2$$

$$= \left(A^{\pi_\theta}(s, a) - \omega^T \nabla_\theta \log \pi_\theta(a|s) \right) \left(A^{\pi_\theta}(s, a) - \nabla_\theta \log \pi_\theta(a|s)^T \cdot \omega \right)$$

$$\nabla_\omega: -2 \nabla_\theta \log(\pi_\theta(a|s)) \cdot A^{\pi_\theta}(s, a)$$

$$+ 2 \nabla_\theta \log \pi_\theta(a|s) \cdot \left(\nabla_\theta \log \pi_\theta(a|s) \right)^T \omega$$

For a stationary point:

∇_{ω}

$$\mathbb{E}_{s \sim d^{\pi_{\theta}}} \mathbb{E}_{a \sim \pi_{\theta}(\cdot|s)} - A(s,a) \nabla_{\theta} \log \pi_{\theta}(a|s) + \nabla_{\theta} \log \pi_{\theta}(a|s) \nabla_{\theta} \log \pi_{\theta}(a|s)^T \omega$$

$= 0$ iff. ↖ $\nabla V^{\pi_{\theta}}(\mu)$

$$\mathbb{E}_{s \sim d^{\pi_{\theta}}} \mathbb{E}_{a \sim \pi_{\theta}(\cdot|s)} A^{\pi_{\theta}}(s,a) \nabla_{\theta} \log \pi_{\theta}(a|s) = \mathbb{E}_{s \sim d^{\pi_{\theta}}} \mathbb{E}_{a \sim \pi_{\theta}(\cdot|s)} \nabla_{\theta} \log \pi_{\theta}(a|s) \nabla_{\theta} \log \pi_{\theta}(a|s)^T \omega$$

$\therefore \omega = (\mathcal{F}_{\theta}^{\mu})^T \nabla V^{\pi_{\theta}}(\mu)$

Letting $\theta \in \mathbb{R}^{s \times A}$ and using softmax policy and NPG

$$(\mathcal{F}_{\theta}^{\mu})^T \nabla V^{\pi_{\theta}}(\mu) = \frac{1}{1-\gamma} A^{\pi_{\theta}}(s,a) - c(s)$$

↑
fn of state

A simple update

$$\pi^{(tr)}(a|s) = \frac{\pi^t(a|s) \exp(\eta A^t(s, a) / (1-\gamma))}{Z_\pi(s)}$$

$$= \sum_{a \in \mathcal{A}} \pi^t(a|s) \exp\left(\frac{\eta A^t(s, a)}{(1-\gamma)}\right)$$

Trick: For softmax policy class, perform the update $\theta \leftarrow \theta + \eta (F_\mu^\theta)^\top \nabla V^{\pi_\theta}(\mu)$

using $\mu \in \Delta(S)$

$$V^t(\mu) \approx V^*(\mu) - \frac{\log(A)}{\eta T} - \frac{1}{(1-\gamma)^2 T}$$

Set $\eta = (1-\gamma)^2 \log(A)$, then we get to within ϵ in $\frac{2}{(1-\gamma)^2 \epsilon}$ steps

- Recall Eq 1

$$\frac{1}{1-\gamma} \sum_s d^\pi(s) \sum_a \frac{\partial \pi_\theta(s,a)}{\partial \theta} \cdot Q^\pi(s,a)$$

$$\nabla V^{\pi_\theta}(\mu) = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^\pi} \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} \left[Q^\pi(s,a) \nabla \log \pi_\theta(a|s) \right]$$

or $\nabla \log \pi_\theta(s,a)$

Suppose we have $f_w: S \times A \rightarrow \mathbb{R}$ is an approximator to Q^π , with parameter w .

Natural: learn f_w , by following π & updating

$$w \text{ by } \Delta w_t \propto \frac{\partial}{\partial w} \left[\hat{Q}^\pi(s_t, a_t) - f_w(s_t, a_t) \right]^2$$

$$\propto \left[\hat{Q}^\pi(s_t, a_t) - f_w(s_t, a_t) \right] \frac{\partial f_w(s_t, a_t)}{\partial w}$$

↑
some unbiased estimator

So when we converge to a local minimum,

$$\sum_s d^\pi(s) \sum_a \pi(s,a) \left[Q^\pi(s,a) - f_w(s,a) \right] \frac{\partial f_w(s,a)}{\partial w} = 0$$

Suppose we find $f_w(s,a)$ s.t.:

$$\nabla V^{\pi_\theta}(\mu) = \frac{1}{1-\gamma} \sum_s d^\pi(s) \sum_a \frac{\partial \pi_\theta(s,a)}{\partial \theta} f_w(s,a) \quad **$$

It's clear that deriving f_w w.r.t

$$\frac{\partial f_w(s, a)}{\partial w} = \frac{1}{\pi(s, a)} \frac{\partial \pi(a, s)}{\partial \theta}, \quad \text{we can simplify}$$


* assuming **;

Proof: From ** and $\frac{\partial f_w(s, a)}{\partial w} = \frac{1}{\pi(s, a)} \frac{\partial \pi(a, s)}{\partial \theta}$

we get

$$\sum_s d^\pi(s) \sum_a \frac{\partial \pi(s, a)}{\partial \theta} [Q^\pi(s, a) - f_w(s, a)] = 0$$

Eq (2)



error is \perp gradient of policy;

$$\text{Eq (1)} - \text{Eq (2)}$$

$$= \sum_s d^\pi(s) \sum_a \frac{\partial \pi(s, a)}{\partial \theta} [Q^\pi(s, a) - Q^\pi(s, a) + f_w(s, a)]$$

$$= \sum_s d^\pi(s) \sum_a \frac{\partial \pi(s, a)}{\partial \theta} f_w(s, a)$$

The gradient has the form as in the red expression.

So a good approximator is

$$f_w(s, a) = \underbrace{\omega^T \nabla_{\theta} \ln \pi_{\theta}(s, a)}$$

$$\text{then } \frac{df_w(s, a)}{d\omega} = \frac{1}{\pi_{\theta}(s, a)} \nabla_{\theta} \pi_{\theta}(s, a)$$

• Is the motivation for compatible function approximator definition:

Then: If the Q value function approximator is compatible to the policy (i.e.) $\nabla_{\omega} f_w(s, a) = \nabla_{\theta} \log \pi_{\theta}(s, a)$

and the parameters ω minimize
$$\mathbb{E}_{\pi_{\theta}} \left[(Q^{\pi_{\theta}}(s, a) - f_w(s, a))^2 \right]$$
 then

policy gradient is exact
$$\nabla_{\theta} V^{\pi_{\theta}}(p) = \mathbb{E}_{\pi_{\theta}} \left[\nabla_{\theta} \log \pi_{\theta}(s, a) f_w(s, a) \right]$$

- Recall the general value function update

$$v(s_t) \leftarrow v(s_t) + \alpha [G_t - V(s_t)]$$

Comes from: Minimizing

$$\mathbb{E} \left[\frac{1}{2} (V^\pi(s) - v(s))^2 \right]$$

$$v \leftarrow v + \alpha \mathbb{E} \left[(V^\pi(s) - v(s)) \frac{\partial v(s)}{\partial v} \right]$$

- Do a stochastic gradient update at time t since we don't know $V^\pi(s)$;

Get an unbiased estimate - sample s & use G_t instead of $V^\pi(s)$

$$\therefore v \leftarrow v + \alpha (G_t - V(s_t)) \frac{\partial v(s_t)}{\partial v}$$

If w parameterizes values of states,

$$v_w: S \rightarrow \mathbb{R},$$

function approximator.

get update: $w \leftarrow w + \alpha (G_T - v_w(s_T)) \frac{\partial v_w(s_T)}{\partial w}$

TD Learning:

Bootstrapping method. bases its update on existing estimate.

An evaluation alg, estimates v^π

TD update: $v(s) \leftarrow v(s) + \alpha (r + \gamma v(s') - v(s))$

(ie) $v(s_t) \leftarrow v(s_t) + \alpha [R_t + \gamma v(s_{t+1}) - v(s_t)]$

• Instead of using G_T use $R_t + \gamma v(s_{t+1})$

$\delta_t = R_t + \gamma v(s_{t+1}) - v(s_t)$. TD error.

• Both TD & MC use estimates; $v(s_{t+1})$ in TD
 G_T in MC

• Note if we used v^π , $\mathbb{E}[\delta_t | S_t = s] = 0$.

TD update with function approximator:

$$w \leftarrow w + \alpha \left(R_{t+1} + \gamma v_w(S_{t+1}) - v_w(S_t) \right) \frac{\partial v_w(S_t)}{\partial w}$$

- TD is not really a gradient update algorithm
- Has good convergence properties; \rightarrow SHOW TO BE CONVERGENT TO v^π on an average.
- Converges when using linear function approximation; say $v_w(s) = w^T \underbrace{\phi(s)}_{\text{features extracted from } s}$;

- Comparison with Monte Carlo estimation of $R_s[s, a, s']$ & $R[s, a]$?

\updownarrow
TD converges to the same as v^π if every state is seen at least once;

- TD - only $|s|$ values;
- Dynamic / Tabular $|s|^2 |A|$ values;
- SARSA - TD for control:

$$q(s, a) \leftarrow q(s, a) + \alpha [r + \gamma q(s', a') - q(s, a)]$$

Stochastic update;

Using function approximator:

$$w \leftarrow w + \alpha (r + \gamma q_w(s', a') - q_w(s, a)) \frac{\partial q_w}{\partial w} \Big|_{s, a}$$

Greedy:

- Initialize $q(s, a)$;

for each episode do (ϵ greedy / softmax)

$\pi \leftarrow \text{Policy}(q)$

$s \sim \text{do}$

Choose a

for each time step until s is absorbing do

Take action a , get r & get s'

Choose a' from s' using π

$$q(s, a) \leftarrow q(s, a) + \alpha [r + q(s', a') - q(s, a)]$$

$s \leftarrow s'; a \leftarrow a'$

- Similarly SARSA with function approximator;
REINFORCE with baseline for estimating $\pi_{\theta} \approx \pi_{*}$.

Input: A differentiable policy parametrization $\pi(a|s, \theta)$ & state value parametrization $\hat{V}(s, \omega)$

Parameters: $\alpha^{\theta} > 0, \alpha^{\omega} > 0$

Initialize: θ & ω

Loop (for each episode):

$s_0, A_0, R_1, \dots, s_{T-1}, A_{T-1}, R_T$ following $\pi(\cdot|\cdot, \theta)$

for each step $t = 0, \dots, T-1$

$$G \leftarrow \sum_{k=t+1}^T \gamma^{k-t-1} R_k$$

$$\delta \leftarrow G - \hat{V}(s_t, \omega)$$

$$\omega \leftarrow \omega + \alpha^{\omega} \delta \nabla_{\omega} \hat{V}(s_t, \omega)$$

$$\theta \leftarrow \theta + \alpha^{\theta} \delta \nabla_{\theta} \ln \pi(A_t | s_t, \theta)$$

end

• Q learning : Off policy TD-control.

• gets an estimate of q^* ;

$$q(s,a) \leftarrow q(s,a) + \alpha \left[r + \delta \max_{a'} q(s',a') - q(s,a) \right]$$

Off policy - update done to get to q^*
and really independent of π used to generate
the data;

• Algorithm as before;

Next TD(x):

Combines the goodness of MC & TD(s)

updates after the
entire episode is over;

↑
immediate update

n -step bootstrapping.

- Has the benefits of MC & one-step; Actions cause updates instantly, in $TD(0)$, but variance is large because of this;
- MC- variance is less \therefore updates are done at the end.
- In MC updates are performed for each state based on the entire sequence of rewards from that state to the end of an episode.
- Natural idea: Perform an update based on an intermediate number of rewards.

Ex: A two step update, would look at first two rewards and estimated value of the state two steps later.

$$\bullet G_{t,t+2} := R_{t+1} + \gamma R_{t+2} + \gamma^2 \underbrace{V_{t+1}(S_{t+2})}$$

the estimate for S_{t+2} takes place of
 $\gamma^2 R_{t+3} + \gamma^3 R_{t+4} + \dots + \gamma^{T-t-1} R_T;$

• target for n -step update is the n -step return

$$G_{t:t+n} := R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots + \gamma^{n-1} R_{t+n-1} + \gamma^n V_{t+n-1}(S_{t+n}).$$

= (if $t+n \geq T$, missing terms are treated as zero)

• Involves future rewards; to use it we need to R_{t+n} and need to compute V_{t+n-1}

$$\bullet \boxed{V_{t+n}(S_t) \stackrel{a}{=} V_{t+n-1}(S_t) + \alpha [h_{t:t+n} - V_{t+n-1}(S_t)]} \\ 0 \leq t \leq T$$

$$V_{t+n}(s) = V_{t+n-1}(s) \quad \forall s \neq s_t;$$

• No changes during first $n-1$ steps of an episode

TD [n].

Input: π

Parameters: α, n

Initialize $V(s)$ arbitrarily $\forall s \in S$.

Loop for each episode:

Initialize & store $s_0 \neq$ terminal

$T \leftarrow \infty$

Loop for $t=0, 1, \dots$

If $t < T$

Take action as per $\pi(\cdot | s_t)$

observe & store reward as R_{t+1} , next state s_{t+1}

If s_{t+1} is terminal $T \leftarrow t+1$

$\tau \leftarrow t-n+1$ $\min(\tau+n, T)$ $i=\tau-1$

If $\tau \geq 0$, $G \leftarrow \sum_{i=\tau+1}^{\min(\tau+n, T)} \gamma^i R_i$

If $\tau+n < T$, $G \leftarrow G + \delta^n V(s_{t+n})$

$V(s_t) \leftarrow V(s_t) + \alpha [G - V(s_t)]$

until $\tau = \pi-1$

$$\max_s \left[\mathbb{E}_{\pi} \left[G_{t:t+n} \mid S_t = s \right] - V_{\pi}(s) \right] \leq$$

$$\gamma^n \max_s \left| V_{t+n-1}(s) - v_{\pi}(s) \right|$$