

A detailed line art illustration of a robot arm holding a pen. The robot has a helmet-like head with a circular sensor on the side. The arm is articulated with various joints and segments, ending in a hand that firmly grips a pen. The drawing is composed of clean, black outlines on a white background.

RL Part 3.1: Policy Search Methods using Policy Gradients

Jan Peters
Gerhard Neumann

Motivation

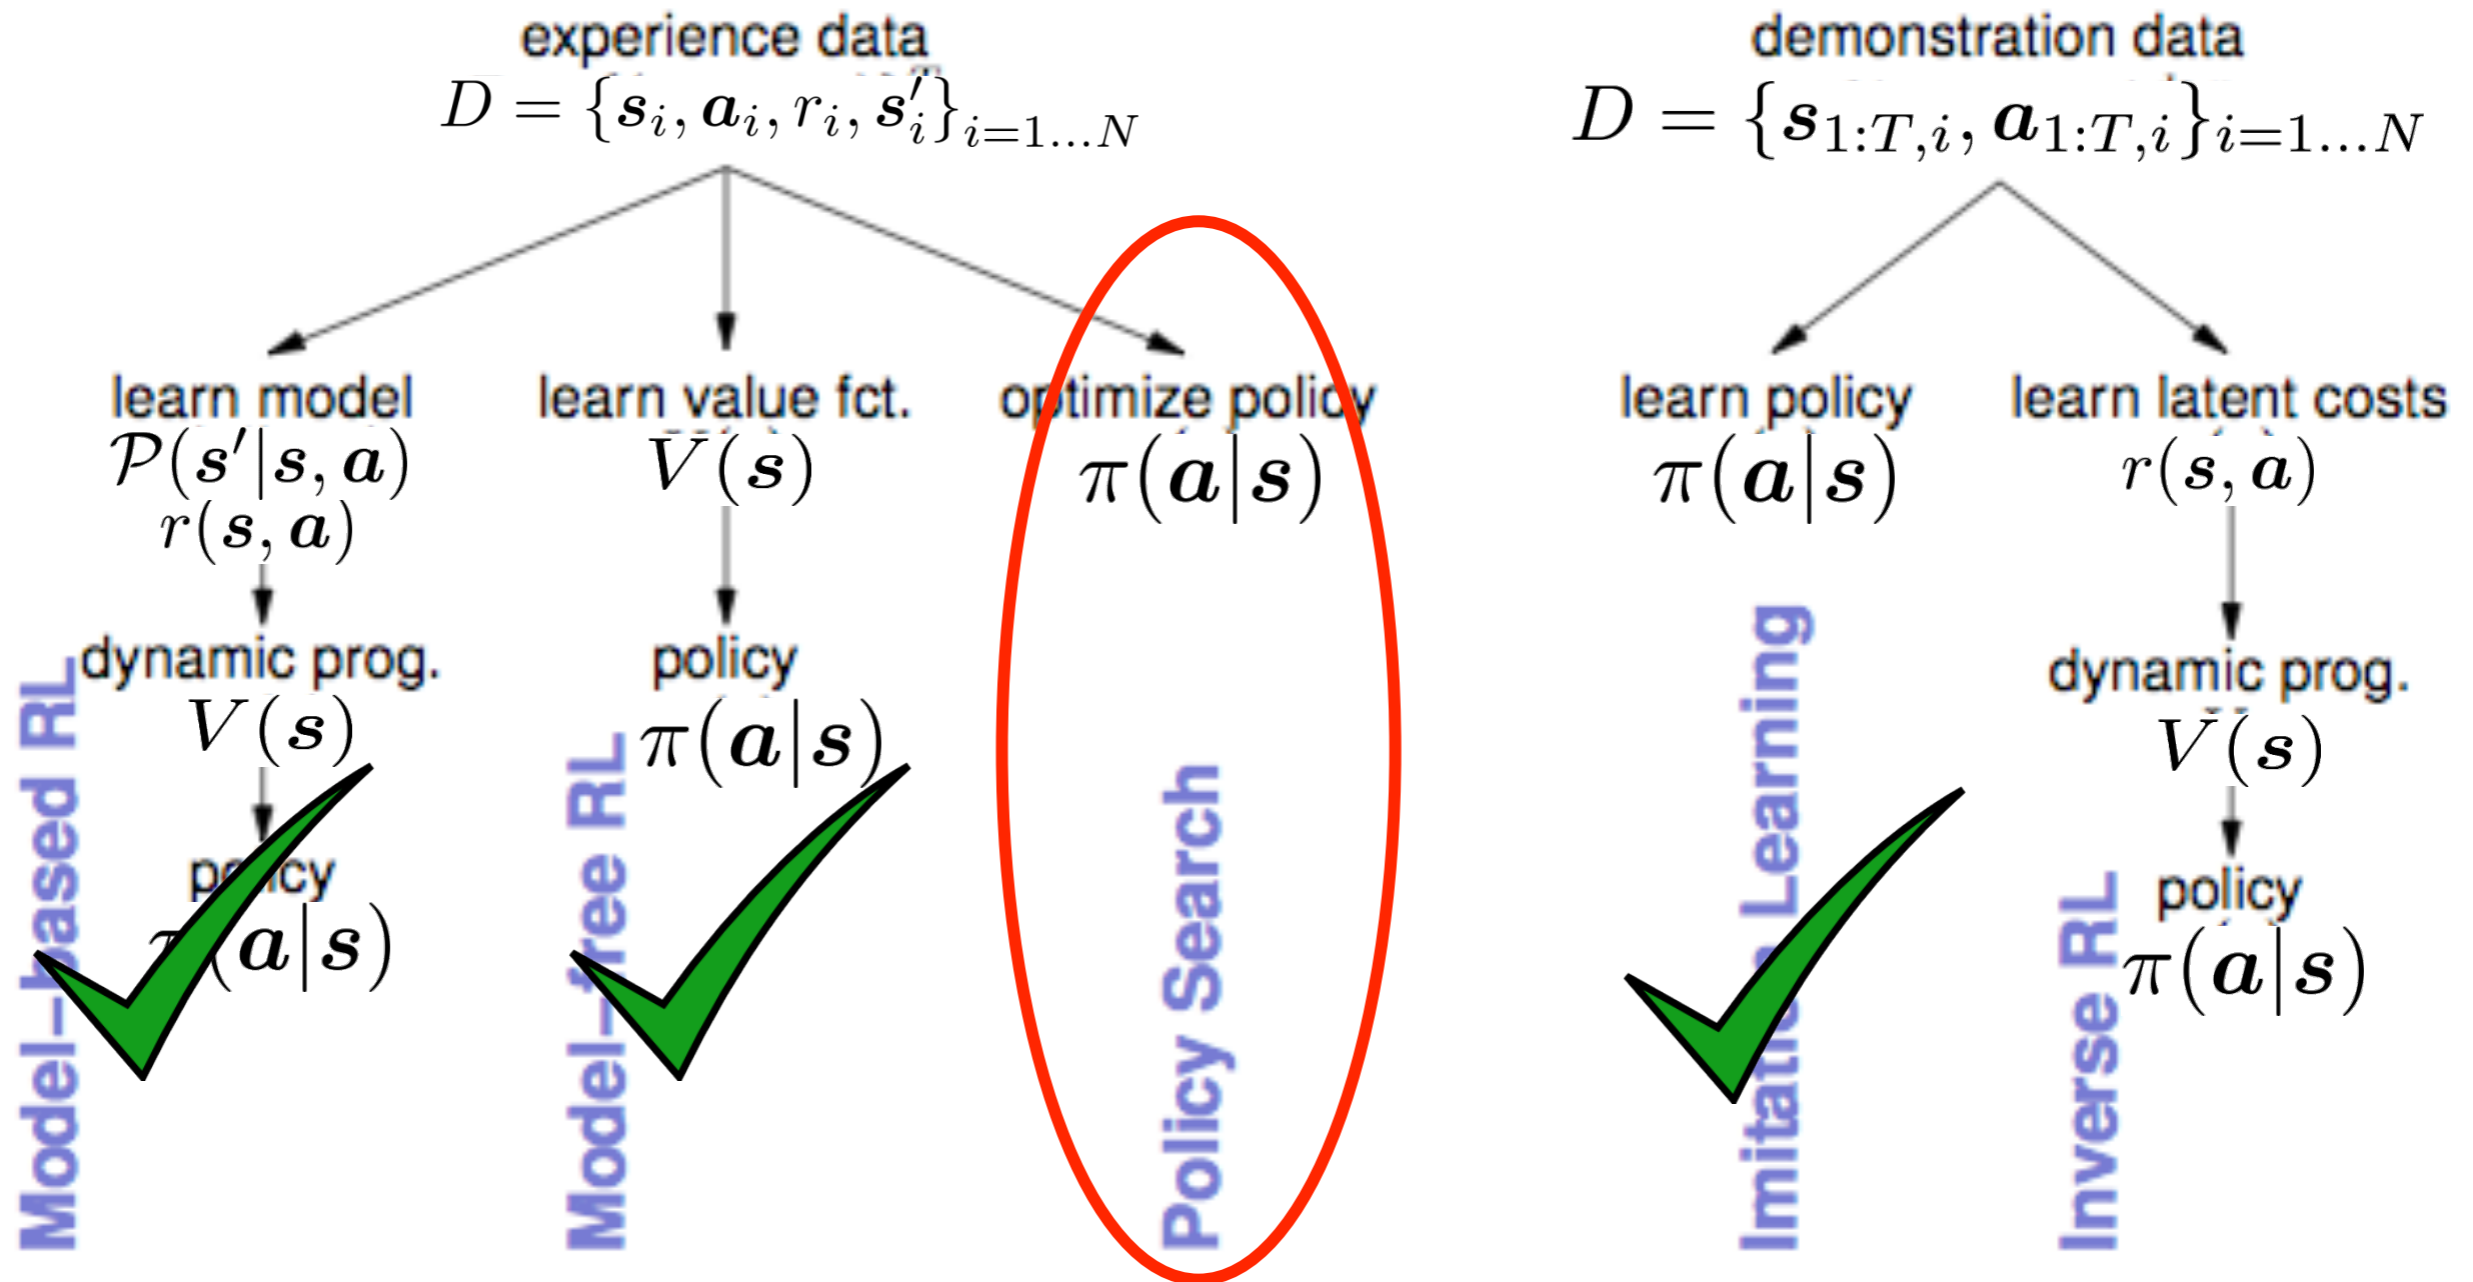
Limits of Value Functions:

- Fill-up state-space: Exponential explosion with the number of dimensions
- **Continuous actions?**
- **Value Function Approximation Error** might propagate and arbitrarily distort the policy update!
- **Exploration** on the real system?

Many of these problems can be fixed by **using parametric policies** and policy search

- Improving upon demonstrations
- Value function is not (always) needed
- Using task-appropriate policies is possible

Bigger Picture



Outline of the Lecture



1. Categorization of Policy Search

- I. Episode-Based versus Step-Based Policy Search

2. Policy Gradients

- I. Episode-Based Policy Gradients
- II. Step-Based Policy Gradients

3. Relative Entropy and Natural Gradients



Action Selection

... in value-based algorithms:

Greedy or soft-max policy: $\pi(\mathbf{a}|\mathbf{s}) = \frac{\exp(\beta Q(\mathbf{s}, \mathbf{a}))}{\sum_{\mathbf{a}'} \exp(\beta Q(\mathbf{s}, \mathbf{a}'))}$

Difficult in **continuous action spaces**

Alternatively, we can use parametrized policies for action selection

For example: Gaussian Policies

$$\pi(\mathbf{a}|\mathbf{s}; \boldsymbol{\theta}) = \mathcal{N}(\mathbf{a} | f_{\mathbf{w}}(\mathbf{s}), \sigma^2 \mathbf{I}), \quad \boldsymbol{\theta} = \{\mathbf{w}, \sigma^2\}$$

Continuous actions **can be easily incorporated**

Policy Search: How to find good parameters $\boldsymbol{\theta}$?

Model-free policy search



Pseudo-Algorithm:

Repeat

- 1. Explore:** Generate trajectories $\tau^{[i]}$ following the current policy π_k
- 2. Policy Evaluation:** Assess quality of trajectory or actions
 - ➔ **Episode-Based Policy Evaluation**
 - ➔ **Step-Based Policy Evaluation**
- 3. Policy Update:** Compute new policy π_{k+1} from trajectories and evaluations

Episode-based evaluation strategy



Evaluation Strategy:

- We directly assess the quality of a **parameter vector** $\theta^{[i]}$ **by the returns**

$$R_{[i]} = \sum_{t=1}^T r_t^{[i]}$$

- **High variance in returns** (sum of T random variables)

Data-set used for policy update

$$\mathcal{D}_{\text{episode}} = \{ \theta^{[i]}, R^{[i]} \}_{i=1 \dots N}$$

- One data-point per trajectory
- Works for a **moderate number of parameters**

Explore **in parameter space at each episode**

Step-based evaluation strategy



Evaluation Strategy:

We assess the quality of **single state-action pairs by using the reward to come**

$$Q_t^{[i]} = \sum_{h=t}^T r_h^{[i]}$$

Less variance in Q_t (sum of $T-t$ random variables)

Data-set used for policy update:

$$\mathcal{D}_{\text{step}} = \left\{ \mathbf{s}_t^{[i]}, \mathbf{a}_t^{[i]}, Q_t^{[i]} \right\}_{i=1 \dots N, t=1 \dots T}$$

One data-point per state-action pair

Explore **in action space at each time step** with **stochastic low-level policy**

$$\mathbf{a}_t \sim \pi(\cdot | \mathbf{s}; \boldsymbol{\theta})$$

Summary



Step-based:

Exploration in Action Space

Less variance in quality assessment.

More data-points to fit policy

Less likely to create unstable policies

Uses the structure of the RL problem

decomposition in single timesteps

Episode-based:

Exploration in Parameter Space

Allows for more sophisticated exploration strategies

Is often very efficient for a small amount of parameters

Generalization and multi-task learning

E.g. open loop policies such as DMPs

Structure-less optimization

„Black-Box Optimizer“

Episode-based evaluation strategy



Episode Based:

Algorithm 3 Episode-Based Policy Evaluation for Learning an Upper-Level Policy

repeat

Exploration:

Sample parameter vector $\theta^{[i]} \sim \pi_{\omega_k}(\theta)$, $i = 1 \dots N$

Sample trajectory $\tau^{[i]} \sim p_{\theta^{[i]}}(\tau)$

Policy Evaluation:

Evaluate policy parameters $R^{[i]} = \sum_{t=0}^T r_t^{[i]}$

Compose data set $\mathcal{D}_{\text{ep}} = \left\{ \theta^{[i]}, R^{[i]} \right\}_{i=1 \dots N}$

Policy Update:

Compute new policy parameters ω_{k+1} using \mathcal{D}_{ep}

Algorithms: Episode-based REPS, Episode-based PI²

PEPG, NES, CMA-ES, RWR

until Policy converges $\omega_{k+1} \approx \omega_k$

Step-based evaluation strategy



Step Based:

Algorithm 2 Policy Search with Step-Based Policy Evaluation

repeat

Exploration:

Create samples $\tau^{[i]} \sim \pi_{\theta_k}(\tau)$ following $\pi_{\theta_k}(\mathbf{u}|\mathbf{x})$, $i = 1 \dots N$

Policy Evaluation:

Evaluate actions: $Q_t^{[i]} = \sum_{h=t}^T r_h^{[i]}$ for all t and all i

Compose data set: $\mathcal{D}_{\text{step}} = \left\{ \mathbf{x}_t^{[i]}, \mathbf{u}_t^{[i]}, Q_t^{[i]} \right\}_{i=1 \dots N, t=1 \dots T-1}$

Policy Update:

Compute new policy parameters θ_{k+1} using \mathcal{D} .

Algorithms: REINFORCE, G(PO)MDP, NAC, eNAC,
PoWER, PI²

until Policy converges $\theta_{k+1} \approx \theta_k$

Episode-based Policy Search



We learn a search distribution $\pi(\theta; \omega)$ over the parameters of the low-level control policy $\pi(a|s; \theta)$

$\pi(\theta; \omega)$ is called **upper-level policy**

For example, $\pi(\theta; \omega) = \mathcal{N}(\mu, \Sigma)$

$\omega = \{\mu, \Sigma\}$... parameters of upper level policy

To reduce variance in the returns, $\pi(a|s; \theta)$ is often modelled as deterministic policy, i.e.,

$$\pi(a|s; \theta) \rightarrow a = \pi(s)$$

Episode-based Policy Search



Search for policy $\pi(\theta; \omega)$ that maximizes the expected return

$$J_{\omega} = \int \pi(\theta; \omega) R_{\theta} d\theta$$

Upper-Level Policy $\pi(\theta; \omega)$:

Stochastic, chooses parameters of low-level policy / movement primitive
Implements **exploration in parameter space for information gathering**

Return R_{θ} : Expected long-term reward for the trajectory τ
that corresponds to θ

$$R_{\theta} = \mathbb{E} \left[\sum_{t=1}^T r_t \mid \theta \right]$$

Episode-based Policy Search Algorithms



Policy Search Algorithms:

Given: initial upper level policy $\pi(\theta; \omega_0)$

Repeat until convergence

Exploration:

Sample from **stochastic policy**:

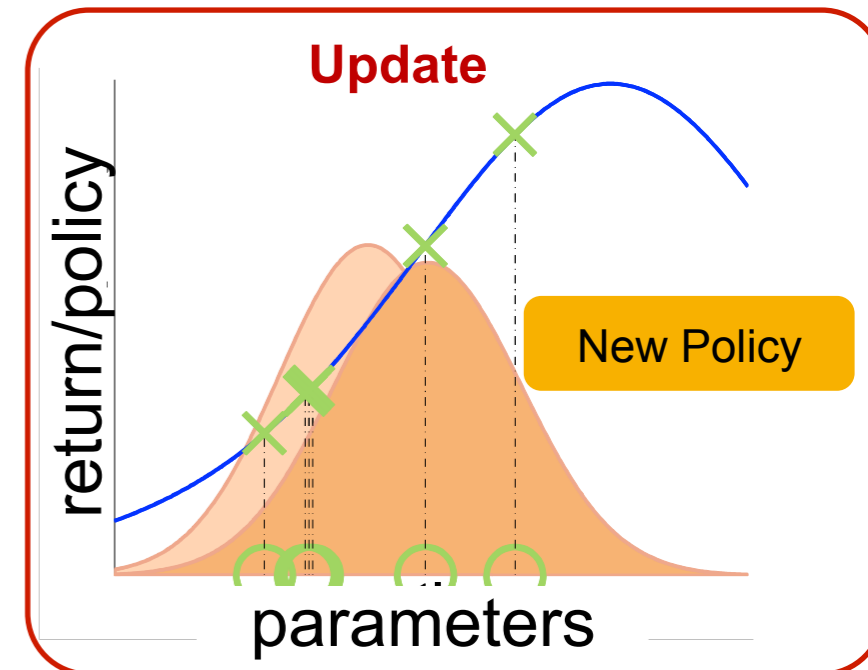
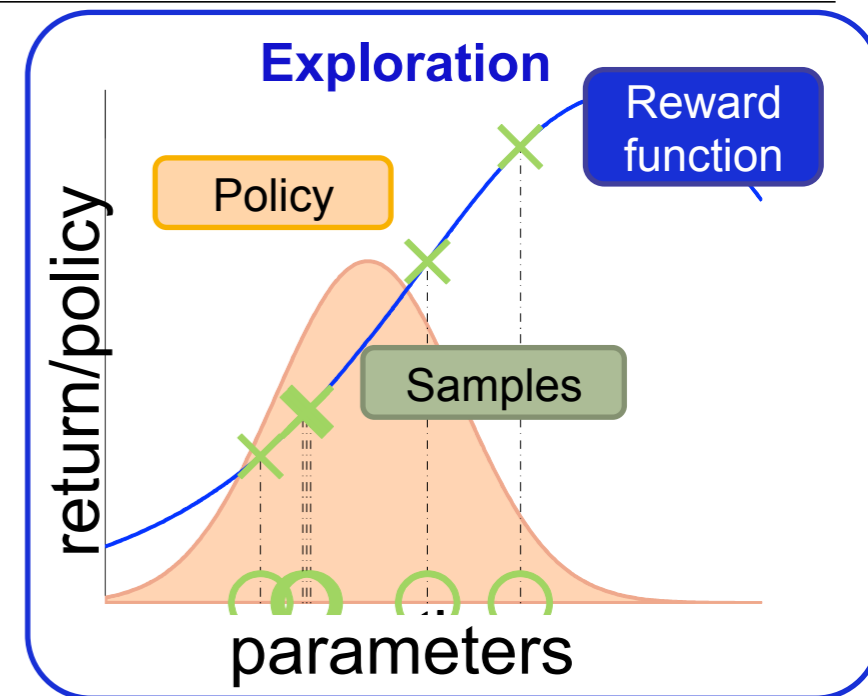
$$\theta_i \sim \pi(\theta; \omega_k), i = 1 \dots N$$

Collect returns by executing θ_i

$$R_i = R_{\theta_i}$$

Update:



Obtain new policy $\pi(\theta; \omega_{k+1})$ from samples



Exploration versus Exploitation



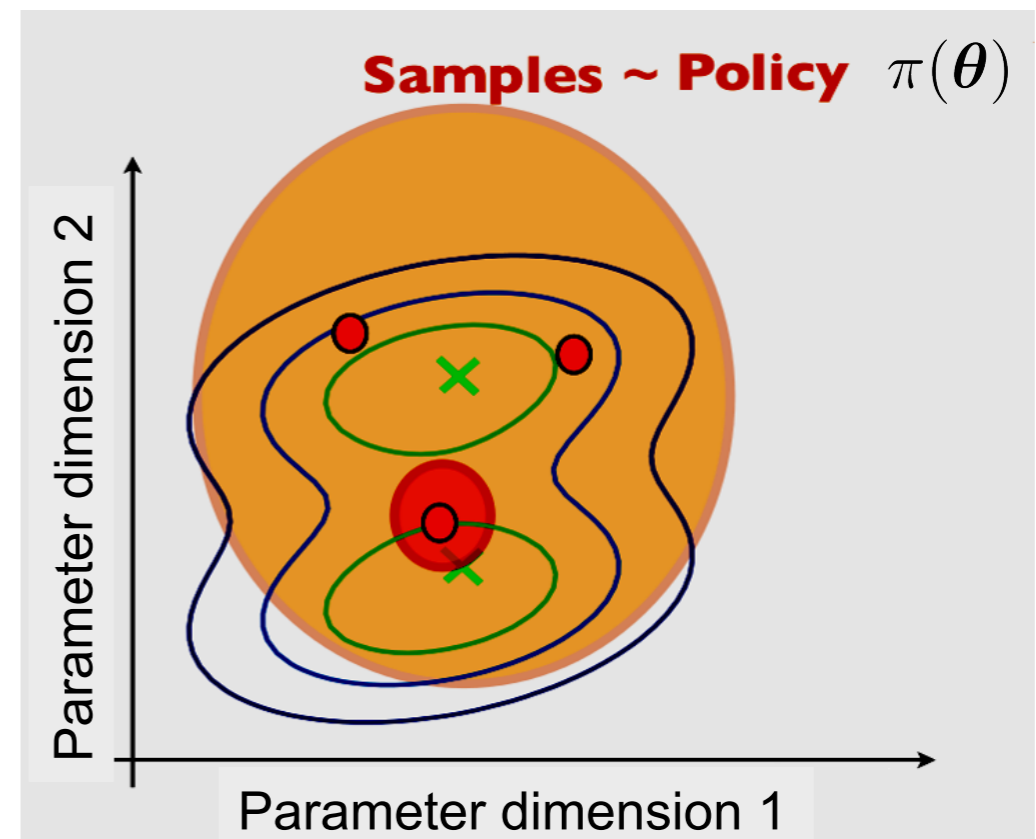

How should we update the policy?



Exploration versus Exploitation





How **greedily** should we update the policy?



Exploration versus Exploitation

How **greedily** should we update the policy?

→ How can we **control this update**?



→ We need to find a metric to measure the „**distance**“ between two policies



Exploration versus Exploitation


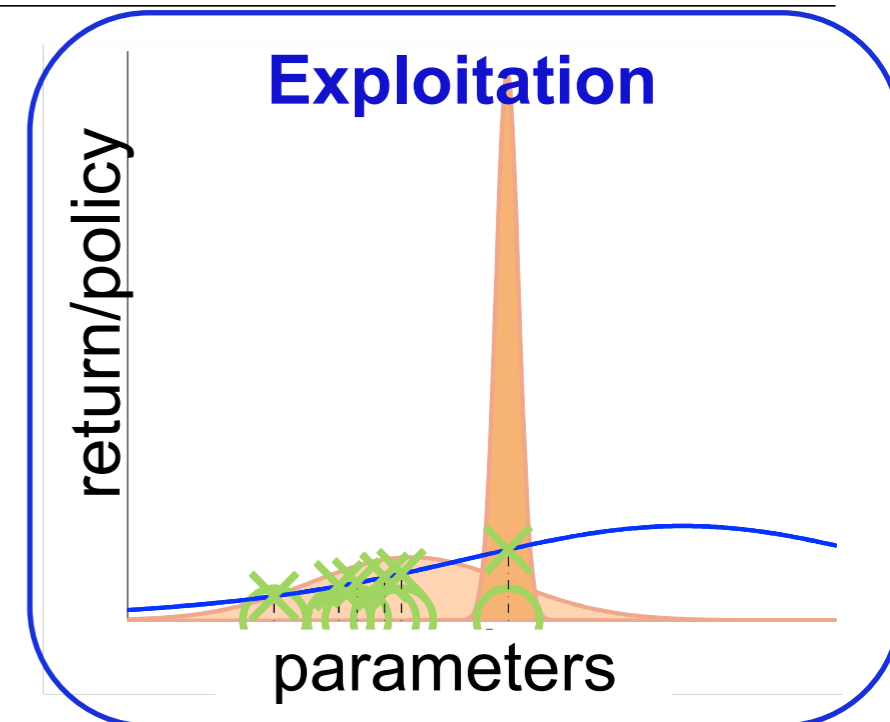
We have to choose a tradeoff between

- ➔ **Exploitation:** Maximizes reward on the samples
- ➔ **Exploration:** Continue to explore in the next iteration

Fundamental Question in Policy Search

- ➔ How can we control the trade-off between exploration and exploitation?
- ➔ We need to quantify the difference between two policies
- ➔ We will get to know **different metrics for policies**

Typically, we want to **limit the distance between two subsequent policies** for the update



Greedy vs Incremental



Why is it useful to control the step-width of the policy update?

Greedy Updates:

$$\theta_{\pi'} = \operatorname{argmax}_{\tilde{\theta}} E_{\pi_{\tilde{\theta}}} \{Q^{\pi}(x, u)\}$$



**potentially
unstable learning
process with large
policy jumps**

Policy Search Updates:



**stable learning
process with
smooth policy
improvement**

Outline of the Lecture



1. Categorization of Policy Search

- I. Episode-Based versus Step-Based Policy Search

2. Policy Gradients

- I. Episode-Based Policy Gradients
- II. Step-Based Policy Gradients


3. Relative Entropy and Natural Gradients

Gradient-based Policy Updates



Gradient computation


*Estimate
Gradient*
 $\nabla J(\boldsymbol{\theta})$



Policy Improvement

*Update
Parameters*


$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k + \alpha \nabla J(\boldsymbol{\theta}_k)$$



Finite Differences



1. Perturb the parameters of your policy:



2. Approximate J by first order Taylor approximation

$$J(\theta + \delta\theta) = J(\theta) + \frac{\partial J(\theta)}{\partial \theta} \delta\theta$$

3. Solve for $\frac{\partial J(\theta)}{\partial \theta}$ in a least squares sense (linear regression):

$$\nabla_{\theta}^{\text{FD}} J = \frac{\partial J(\theta)}{\partial \theta} = (\Delta \Theta^T \Delta \Theta)^{-1} \Delta \Theta^T \Delta J$$

Can be used to update a single parameter estimate (e.g. mean)

Likelihood Policy Gradients



How can we update a distribution $\pi(\theta; \omega)$ over the parameter vector (including variance)? **Log-ratio**

trick

$$\nabla \log f(x) = \frac{1}{f(x)} \nabla f(x) \quad \Rightarrow \quad \nabla f(x) = f(x) \nabla \log f(x)$$

Gradient of the expected return

$$\nabla_{\omega} J_{\omega} = \nabla_{\omega} \int \pi(\theta; \omega) R_{\theta} d\theta = \int \nabla_{\omega} \pi(\theta; \omega) R_{\theta} d\theta$$

$$= \int \pi(\theta; \omega) \nabla_{\omega} \log \pi(\theta; \omega) R_{\theta} d\theta$$

$$\approx \sum_{i=1}^N \nabla_{\omega} \log \pi(\theta_i; \omega) R_i$$

Only needs
samples!

23 **This gradient is called **Parameter Exploring Policy Gradient (PGPE)****

Baselines...



We can always **subtract a baseline** from the gradient...

$$\nabla_{\omega} J_{\omega} = \sum_{i=1}^N \nabla_{\omega} \log \pi(\boldsymbol{\theta}_i; \omega) (R_i - b)$$

Why?

The gradient estimate can have a high variance

Subtracting a baseline can reduce the variance

Its still unbiased...

$$\mathbb{E}_{p(\boldsymbol{x}; \omega)} [\nabla_{\omega} \log p(\boldsymbol{x}; \omega) b] = b \int \nabla_{\boldsymbol{x}} p(\boldsymbol{x}; \omega) = b \nabla_{\boldsymbol{x}} \int p(\boldsymbol{x}; \omega) = 0$$

Good baseline: Average reward

but there are optimal baselines for each alg. that minimize the variance



Step-based Policy Gradient Methods

The returns can still have **a lot of variance**

$$R_{\theta} = \mathbb{E} \left[\sum_{t=1}^T r_t | \theta \right]$$

➔ It is the sum over T random variables

There is less variance in the rewards to come: $Q_t^{[i]} = \sum_{h=t}^T r_h^{[i]}$

➔ Step-based algorithms can be more efficient when estimating the gradient

➔ For step-based algorithms, we have to compute the gradient $\nabla_{\theta} J$ for the low-level policy $\pi(a|s; \theta)$

Outline of the Lecture



1. Categorization of Policy Search

- I. Episode-Based versus Step-Based Policy Search

2. Policy Gradients

- I. Episode-Based Policy Gradients
- II. Step-Based Policy Gradients

3. Relative Entropy and Natural Gradients



Step-based Policy Gradient Methods

Some more basic notation

Trajectory distribution: $p(\boldsymbol{\tau}; \boldsymbol{\theta}) = p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi(\mathbf{a}_t | \mathbf{s}_t; \boldsymbol{\theta}) p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$

Return for a single trajectory: $R(\boldsymbol{\tau}) = \sum_{t=1}^{T-1} r_t + r_T$

Expected long term reward $J(\boldsymbol{\theta})$ can be written as **expectation over the trajectory distribution**

$$J(\boldsymbol{\theta}) = \mathbb{E}_{p(\boldsymbol{\tau}; \boldsymbol{\theta})} [R(\boldsymbol{\tau})] = \int p(\boldsymbol{\tau}; \boldsymbol{\theta}) R(\boldsymbol{\tau}) d\boldsymbol{\tau}$$



Step-Based Likelihood Ratio Gradient

Instead of computing the gradient of the upper-level policy, we compute the **gradient of the trajectory distribution**

$$\nabla_{\theta} J_{\theta} = \sum_{i=1}^N \nabla_{\theta} \log p(\tau^{[i]}; \theta) R(\tau^{[i]})$$

How do we compute $\nabla_{\theta} \log p(\tau^{[i]}; \theta)$?

$$p(\tau; \theta) = p(\mathbf{s}_1) \prod_{t=1}^{T-1} \pi(\mathbf{a}_t | \mathbf{s}_t; \theta) p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$$

$$\log p(\tau; \theta) = \sum_{t=1}^{T-1} \log \pi(\mathbf{a}_t | \mathbf{s}_t; \theta) + \text{const}$$

Model-dependent terms do not depend on parameters, derivative is now easy

$$\nabla_{\theta} \log p(\tau; \theta) = \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t | \mathbf{s}_t; \theta)$$

Lets plug it in...



Result:

$$\begin{aligned}\nabla_{\theta} J &= \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \theta) R(\tau^{[i]}) \\ &= \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \theta) \left(\sum_{h=1}^{T-1} r_h^{[i]} + r_T^{[i]} \right)\end{aligned}$$

This algorithm is called the **REINFORCE Policy Gradient**

- ➔ Wait... we still use the returns $R(\tau) = \sum_{t=1}^{T-1} r_t + r_T$ (high variance)
- ➔ What did we gain with our step-based version? Not too much yet...



Using the rewards to come...

Simple Observation:

Rewards in the past are not correlated with actions in the future

$$\mathbb{E}_{p(\tau)} [r_t \log \pi(\mathbf{a}_h | \mathbf{s}_h)] = 0, \forall t < h$$

This observation leads to the **Policy Gradient Theorem**

$$\begin{aligned} \nabla_{\theta}^{\text{PG}} J &= \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \theta) \left(\sum_{h=t}^{T-1} r_h^{[i]} + r_T^{[i]} \right) \\ &= \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \theta) Q_h^{[i]} \end{aligned}$$

- ➔ The rewards to come have less variance
- ➔ We can do it again with a baseline...



Metric in standard gradients

How can we choose the step-size to control our policy update?

Simple (naive) idea:

- ➔ Use **distance in parameter space as metric**
- ➔ Episode-based: $L_2(\pi_{k+1}, \pi_k) = \|\omega_{k+1} - \omega_k\|$
- ➔ Step-based: $L_2(\pi_{k+1}, \pi_k) = \|\theta_{k+1} - \theta_k\|$

Choose step size, such that $L_2(\pi_{k+1}, \pi_k) \leq \epsilon$

$$\alpha_k = \frac{1}{\|\nabla J\|} \epsilon$$

Metric in standard gradients



Is the distance in parameter space a good idea?

Consider the following policy:

$$\pi(a|s; \theta) = \mathcal{N}(a|\theta_1 s_1 + \theta_2 s_2, \sigma^2)$$

with $s_1 \in [0, 1]$ and $s_2 \in [0, 1000]$

Lets consider the distances of $\theta_1 = [1, 1]^T$, $\theta_2 = [1.1, 1]^T$, $\theta_3 = [1, 1.1]^T$

The distances $\|\theta_1 - \theta_2\|$ and $\|\theta_1 - \theta_3\|$ are the same

Policy $\pi(a|s, \theta_3)$ is much more different from $\pi(a|s, \theta_1)$ than $\pi(a|s, \theta_2)$

The euclidian metric is **not invariant to scaling of the variables!**

Metric in standard gradients



Can we define a metric that is invariant to transformation of the parameters?

Idea:

- ➔ Define a matrix M that captures the „influence“ of the parameters on the policy
- ➔ Use matrix M to define a new metric that incorporates this influence
- ➔ $L_M(\pi_{k+1}, \pi_k) = \|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k\|_M = (\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k)^T \mathbf{M} (\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k)$
- ➔ Large change in parameters are more expensive in directions with large influence



Metric in standard gradients

How to use such metric for gradient ascent?

Find an update direction $\Delta\theta$ that is **most similar to the standard gradient** $\nabla_{\theta} J = \left[\frac{dJ}{d\theta_1}, \dots, \frac{dJ}{d\theta_n} \right]$

$$\Delta\theta^* = \operatorname{argmax}_{\Delta\theta} \nabla_{\theta} J \Delta\theta$$

with limited distance, i.e.,

$$\text{s.t.: } L_M(\pi_{k+1}, \pi_k) = \Delta\theta^T M \Delta\theta \leq \epsilon$$

Solution to this constraint optimization problem (see lecture notes)

$$\Delta\theta^* = \lambda M^{-1} \nabla_{\theta} J \propto M^{-1} \nabla_{\theta} J$$

➡ Now we „only“ have to find a proper matrix M

Outline of the Lecture



1. Categorization of Policy Search

- I. Episode-Based versus Step-Based Policy Search

2. Policy Gradients

- I. Episode-Based Policy Gradients
- II. Step-Based Policy Gradients

3. Relative Entropy and Natural Gradients



We need to find a better metric...

What do we want?

1. Invariance to the representation of the policy (e.g. parameter transformations)
2. Invariance to transformations of the rewards

Alternative way to **measure the distance between two policies**

Policies are probability distributions

We can measure „distances“ of distributions

For example, **Relative Entropy** or **Kullback-Leibler divergence**

$$\text{KL}(p||q) = \sum_{\mathbf{x}} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})}$$

Information-theoretic „distance“ measure between distributions



Kullback-Leibler Divergence

Properties:
$$\text{KL}(p||q) = \sum_{\mathbf{x}} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})}$$

$$\text{KL}(q||p) \geq 0, \quad \text{KL}(q||p) = 0 \Rightarrow p = q$$

Not symmetric, **so not a real distance**

$$\text{KL}(q||p) \neq \text{KL}(p||q)$$

KL for Gaussians:

$$\text{KL}(p||q) = \log \frac{|B|}{|A|} + \text{tr}(\mathbf{B}^{-1}\mathbf{A}) + (\mathbf{b} - \mathbf{a})^T \mathbf{B}^{-1}(\mathbf{b} - \mathbf{a}) - n$$

with $p(\mathbf{x}) = \mathcal{N}(\mathbf{a}, \mathbf{A})$ and $q(\mathbf{x}) = \mathcal{N}(\mathbf{b}, \mathbf{B})$

... compare with euclidian metric:

Distance scales with inverse covariance of q

Kullback-Leibler Divergence




2 types of KL:

Moment projection: $\operatorname{argmin}_q \text{KL}(p||q) = \sum_{\mathbf{x}} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})}$

q is large wherever p is large

Same as **Maximum Likelihood** estimate (blackboard)!



Kullback-Leibler Divergence





2 types of KL:

Information projection: $\operatorname{argmin}_q \operatorname{KL}(q||p) = \sum_{\mathbf{x}} q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x})}$

q is zero wherever p is zero (zero forcing)

not unique for most distributions





KL divergences and the Fisher information matrix

The Kullback Leibler divergence can be **approximated by the Fisher information matrix (2nd order Taylor approximation)**

$$\text{KL}(p_{\theta+\Delta\theta}||p_{\theta}) \approx \Delta\theta^T \mathbf{G}(\theta)\Delta\theta$$

where $\mathbf{G}(\theta)$ is the **Fisher information matrix (FIM)**

$$\mathbf{G}(\theta) = \mathbb{E}_p[\nabla_{\theta} \log p_{\theta}(\mathbf{x}) \nabla_{\theta} \log p_{\theta}(\mathbf{x})^T]$$

- ➔ Captures information how the **single parameters influence the distribution**



Properties of the Fisher information matrix

$$\mathbf{G}(\boldsymbol{\theta}) = \mathbb{E}_p[\nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{x}) \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{x})^T]$$

If the distribution is Gaussian, i.e., $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_{\boldsymbol{\alpha}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}})$
with $\boldsymbol{\alpha} \in \mathbb{R}^n$ and $\boldsymbol{\beta} \in \mathbb{R}^m$

then the FIM is a $(n + m) \times (n + m)$ matrix and is given by

$$\mathbf{G}(\boldsymbol{\theta}) = \text{diag}(\mathbf{G}_1(\boldsymbol{\alpha}), \mathbf{G}_2(\boldsymbol{\beta})) \text{ with } \boldsymbol{\theta} = [\boldsymbol{\alpha}^T, \boldsymbol{\beta}^T]^T$$

$$\mathbf{G}_1(\boldsymbol{\alpha})_{i,j} = \frac{\partial \boldsymbol{\mu}_{\boldsymbol{\alpha}}}{\partial \alpha_i} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\mu}_{\boldsymbol{\alpha}}}{\partial \alpha_j}^T$$

$$\mathbf{G}_2(\boldsymbol{\beta}) = 0.5 \text{tr} \left(\boldsymbol{\Sigma}_{\boldsymbol{\beta}}^{-1} \frac{\partial \boldsymbol{\Sigma}_{\boldsymbol{\beta}}}{\partial \beta_i} \boldsymbol{\Sigma}_{\boldsymbol{\beta}}^{-1} \frac{\partial \boldsymbol{\Sigma}_{\boldsymbol{\beta}}}{\partial \beta_j} \right)$$

Homework: Check \mathbf{G}_1 for $\boldsymbol{\mu}_{\boldsymbol{\alpha}} = \boldsymbol{\alpha}$ and $\boldsymbol{\mu}_{\boldsymbol{\alpha}} = \boldsymbol{\phi}(s)^T \boldsymbol{\alpha}$

Kullback Leibler divergences



The Natural gradient uses the Fisher information matrix as metric

$$\nabla_{\theta}^{\text{NG}} J = \operatorname{argmax}_{\Delta\theta} \Delta\theta^T \nabla_{\theta} J$$

$$\text{s.t.: } \text{KL}(p_{\theta+\Delta\theta} || p_{\theta}) \approx \Delta\theta^T G(\theta) \Delta\theta \leq \epsilon$$

The solution to this optimization problem is given as:

$$\nabla_{\theta}^{\text{NG}} J \propto G(\theta)^{-1} \nabla_{\theta} J$$

As every parameter has the same influence under metric \mathbf{M} , the natural gradient is **invariant to linear transformations of the parameter space!**

Are they useful?




Linear Quadratic Regulation

$$x_{t+1} = Ax_t + Bu_t$$


$$u_t \sim \pi(u|x_t) = \mathcal{N}(u|kx_t, \sigma)$$

$$r_t = -x_t^T Q x_t - u_t^T R u_t$$

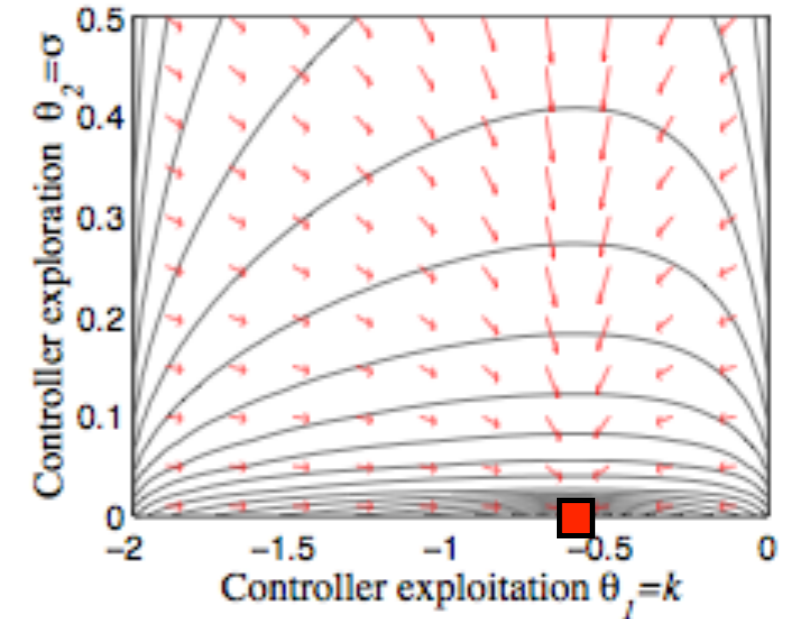
Two-State Problem




(a) LQR policy gradient




(b) LQR natural gradient



(c) Two state policy gradient



(d) Two state natural gradient




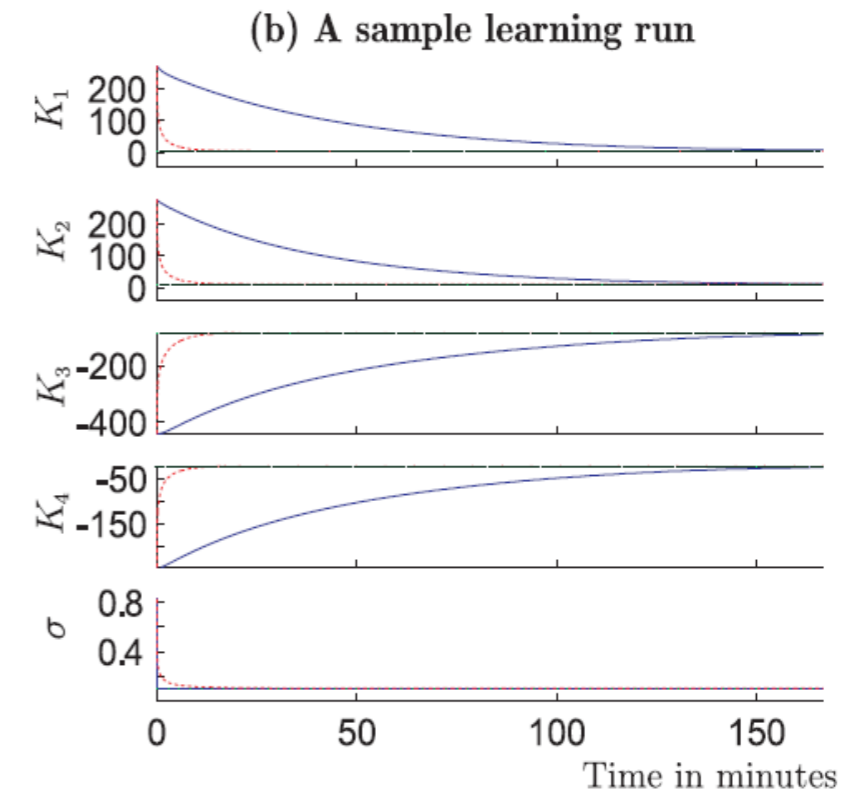

(Peters et al. 2003, 2005)

Comparison



Cart-Pole Balancing

- ➔ Learn 4 gains and 1 variance parameter
- ➔ Blue: standard gradient
- ➔ Red: natural gradient





Computing the Natural Gradient

2-Variants:

Episode-Based:

Closed Form for Gaussian distributions

Step-Based:

Can be avoided due to compatible value function approximation

Computing the NG (episode-based)



In the episode-based case, if the distribution is Gaussian, the Fisher Information matrix can be computed in closed form

Used by the **Natural Evolution Strategy (NES)**

Initialize: μ_0, Σ_0

For $k = 0$ **to** L

Create evaluate samples: $\theta_i \sim \pi(\theta | \mu_k, \Sigma_k)$ $R_i = \sum_{t=1}^T r_{i,t}$

Compute Gradient: $\nabla_{\omega} J = \sum_{i=1}^N \nabla_{\omega} \log \pi(\theta_i | \omega) R_i$

Compute FIM $G(\omega)$ **in closed form for Gaussians**

Compute Natural Gradient: $\nabla_{\omega}^{\text{NES}} J = G(\omega)^{-1} \nabla_{\omega} J$

Update Parameters: $\omega_{k+1} = \omega_k + \eta \nabla_{\omega}^{\text{NES}} J$

end

Computing the NG (step-based)



Policy Gradient Theorem with baseline

$$\nabla_{\theta}^{\text{PG}} J = \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \theta) (Q_h^{[i]} - b_h(\mathbf{s}))$$

To further improve the gradient estimate we can try to estimate **the reward to come**

$$f_w(\mathbf{s}, \mathbf{a}) = \psi(\mathbf{s}, \mathbf{a})^T \mathbf{w} \approx (Q_h^{[i]} - b_h(\mathbf{s}^{[i]}))$$

and use $\nabla_{\theta}^{\text{FA}} J = \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\theta} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \theta) f_w(\mathbf{s}^{[i]}, \mathbf{a}^{[i]})$ as gradient

It can be shown that this **gradient is still unbiased if:**

$$\psi(\mathbf{s}, \mathbf{a}) = \nabla_{\theta} \log \pi(\mathbf{a} | \mathbf{s})$$

Compatible Function Approximation



Compatible Function Approximation:

$$f_{\mathbf{w}}(\mathbf{s}, \mathbf{a}) = \psi(\mathbf{s}, \mathbf{a})^T \mathbf{w} \approx (Q_h^{[i]} - b_h(\mathbf{s}^{[i]})) \quad \psi(\mathbf{s}, \mathbf{a}) = \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}|\mathbf{s})$$

Basis functions of $Q(\mathbf{s}, \mathbf{a})$ are compatible to the policy

The compatible function approximation is mean-zero!

$$\mathbb{E}_{p(\boldsymbol{\tau})} [\nabla \log \pi(\mathbf{a}|\mathbf{s}; \boldsymbol{\theta})^T \mathbf{w}] = 0$$

Thus, it can only represent the Advantage Function:

Baseline

$$f_{\mathbf{w}}(\mathbf{s}, \mathbf{a}) = \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}|\mathbf{s}; \boldsymbol{\theta})^T \mathbf{w} = Q^{\pi}(\mathbf{s}, \mathbf{a}) - V^{\pi}(\mathbf{s})$$

The advantage function tells us, how much better an action is in comparison to the expected performance




Can the Compatible FA be learned?

The compatible function approximation represents an advantage function


$$f_w^\pi(\mathbf{x}, \mathbf{u}) = Q^\pi(\mathbf{x}, \mathbf{u}) - V^\pi(\mathbf{x}) = A^\pi(\mathbf{x}, \mathbf{u}).$$

The advantage function is very different from the value functions

Value Function $Q^\pi(\mathbf{x}, \mathbf{u})$



Advantage Function $A^\pi(\mathbf{x}, \mathbf{u})$



In order to learn $f_w(s, a)$ we need to learn $V^\pi(s)$

Combatible Function Approximation



Gradient with **Combatible Function Approximation**:

$$\nabla_{\boldsymbol{\theta}}^{\text{FA}} J = \sum_{i=1}^N \sum_{t=1}^{T-1} \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta})^T \mathbf{w}$$

$$\nabla_{\boldsymbol{\theta}}^{\text{FA}} J = \mathbb{E}_{p(\tau)} \left[\nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta})^T \right] \mathbf{w}$$

$$\nabla_{\boldsymbol{\theta}}^{\text{FA}} J = \mathbf{F}(\boldsymbol{\theta}) \mathbf{w}$$

We showed [Peters & Schaal, 2008]:

$$\begin{aligned} \mathbf{F}(\boldsymbol{\theta}) &= \mathbb{E}_{p(\tau)} \left[\nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta})^T \right] \\ &= \mathbb{E}_{p(\tau)} \left[\nabla_{\boldsymbol{\theta}} \log p(\tau; \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \log p(\tau; \boldsymbol{\theta})^T \right] = \mathbf{G}(\boldsymbol{\theta}) \end{aligned}$$



Connection to V-Function approximation

Compatible Function Approximation:

$$\nabla_{\theta}^{\text{FA}} J = \mathbf{F}(\theta) \mathbf{w}$$

We showed: \mathbf{F} is the Fisher information matrix!

$$\mathbf{F}(\theta) = \mathbf{G}(\theta)$$

That makes the natural gradient very simple !

$$\nabla_{\theta}^{\text{NG}} J = \mathbf{G}(\theta)^{-1} \nabla_{\theta}^{\text{FA}} J = \mathbf{G}(\theta)^{-1} \mathbf{F}(\theta) \mathbf{w} = \mathbf{w}$$

So we just have to learn \mathbf{w}



What about this additional FA?

In many cases, we don't have a good basis functions for $V^\pi(\mathbf{s})$

For one rollout i , if we sum up the Bellman Equations

$$Q_1^\pi(\mathbf{s}_1^{[i]}, \mathbf{a}_1^{[i]}) = r(\mathbf{s}_1^{[i]}, \mathbf{a}_1^{[i]}) + V_2^\pi(\mathbf{s}_2^{[i]})$$

$$V_1^\pi(\mathbf{s}_1^{[i]}) + f_w(\mathbf{s}_1^{[i]}, \mathbf{a}_1^{[i]}) = r(\mathbf{s}_1^{[i]}, \mathbf{a}_1^{[i]}) + V_2^\pi(\mathbf{s}_2^{[i]})$$

$$V_1^\pi(\mathbf{s}_1^{[i]}) + \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_1^{[i]} | \mathbf{s}_1^{[i]}; \boldsymbol{\theta}) \mathbf{w} = r(\mathbf{s}_1^{[i]}, \mathbf{a}_1^{[i]}) + V_2^\pi(\mathbf{s}_2^{[i]})$$

for each time step

$$V_1^\pi(\mathbf{s}_1^{[i]}) + \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_1^{[i]} | \mathbf{s}_1^{[i]}; \boldsymbol{\theta}) \mathbf{w} = r(\mathbf{s}_1^{[i]}, \mathbf{a}_1^{[i]}) + V_2^\pi(\mathbf{s}_2^{[i]}) \quad | + \text{both sides}$$

$$V_2^\pi(\mathbf{s}_2^{[i]}) + \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_2^{[i]} | \mathbf{s}_2^{[i]}; \boldsymbol{\theta}) \mathbf{w} = r(\mathbf{s}_2^{[i]}, \mathbf{a}_2^{[i]}) + V_3^\pi(\mathbf{s}_3^{[i]}) \quad | + \text{both sides}$$

\vdots

| + both sides

$$V_{T-1}^\pi(\mathbf{s}_{T-1}^{[i]}) + \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_{T-1}^{[i]} | \mathbf{s}_{T-1}^{[i]}; \boldsymbol{\theta}) \mathbf{w} = r(\mathbf{s}_{T-1}^{[i]}, \mathbf{a}_{T-1}^{[i]}) + V_T^\pi(\mathbf{s}_T^{[i]})$$



What about this additional FA?

We can now eliminate the values $V^\pi(\mathbf{s})$ of the intermediate states, we obtain

$$\underbrace{V^\pi(\mathbf{s}_1^{[i]})}_J + \underbrace{\left(\sum_{t=1}^{T-1} \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta}) \right)}_{\boldsymbol{\varphi}^T} \mathbf{w} = \sum_{t=1}^T r(\mathbf{s}_t^{[i]}, \mathbf{a}_t^{[i]})$$

ONE offset parameter J suffices as additional function approximation!

at least if we only have one initial state



Episodic Natural Actor-Critic

In order to get \mathbf{w} , we can use linear regression

$$\underbrace{V^\pi(\mathbf{s}_1^{[i]})}_J + \underbrace{\left(\sum_{t=1}^{T-1} \nabla_{\boldsymbol{\theta}} \log \pi(\mathbf{a}_t^{[i]} | \mathbf{s}_t^{[i]}; \boldsymbol{\theta}) \right)}_{\boldsymbol{\varphi}^T} \mathbf{w} = \sum_{t=1}^T r(\mathbf{s}_t^{[i]}, \mathbf{a}_t^{[i]})$$

Policy
Evaluation

$$\mathbf{\Phi} = \begin{bmatrix} \varphi_1 & \varphi_2 & \dots & \varphi_N \\ 1 & 1 & \dots & 1 \end{bmatrix}^T$$

$$\mathbf{R} = [R_1, R_2^T, \dots, R_N^T]^T$$

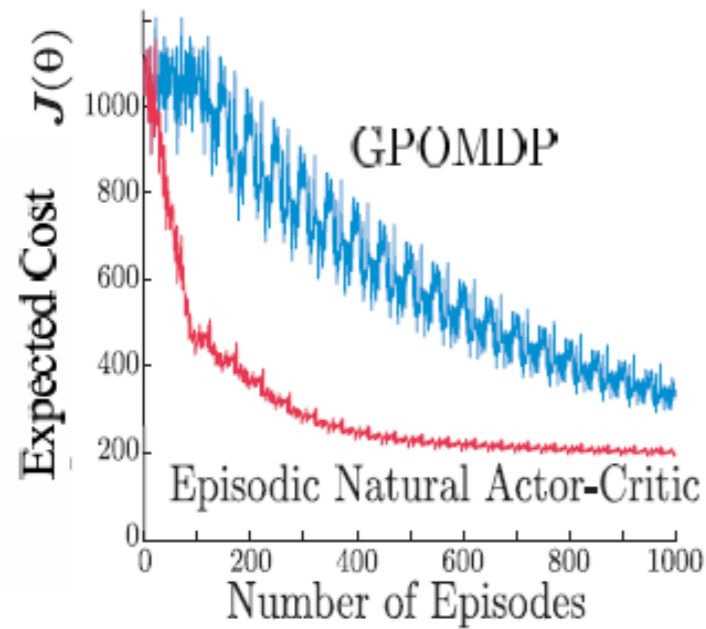
Linear
Regression

$$\begin{bmatrix} \mathbf{w} \\ J \end{bmatrix} = \left(\mathbf{\Phi}^T \mathbf{\Phi} \right)^{-1} \mathbf{\Phi}^T \mathbf{R}$$


**Actor: Natural
Policy Gradient
Improvement**

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t \mathbf{w}_t.$$


Results...



(a) Expected Cost



(b) Position of motor primitives



(c) Controls of motor primitives

Toy Task: Optimal point to point movements with DMPs

GPOMP: Standard Gradient (Equivalent to Policy Gradient Theorem)

Learning T-Ball



- 1) Teach motor primitives by imitation
- 2) Improve movement by Episodic Natural-Actor Critic

*Good
performance
often after
150-300 trials.*



Important Points



Points worth highlighting:

- ➔ The metric really matters in policy search
- ➔ Natural policy gradients are *independent* of the chosen policy parameterization!
- ➔ They correspond to *steepest descent in policy space* and not in the parameter space.
- ➔ *Convergence to a local minimum* is guaranteed!

Conclusion



- ➔ Policy Search is a powerful and practical alternative to value function and model-based methods.
- ➔ Policy gradients have dominated this area for a long time and solidly working methods exist.
- ➔ Say still need a lot of samples and we need to tune the learning rate
- ➔ Learning the exploration rate is still an open problem
- ➔ Newer methods focus on probabilistic policy search approaches.