

Tutorial on Reinforcement Learning

Marc Deisenroth




Department of Computing
Imperial College London

Department of Computer Science
TU Darmstadt

`m.deisenroth@imperial.ac.uk`




Machine Learning Summer School on Big Data
Hammamet, September 17, 2013

Machine Learning



- Unsupervised learning
- Supervised learning
- Reinforcement learning


Machine Learning



- Unsupervised learning
- Supervised learning
- Reinforcement learning


RL makes decisions!

RL Success Stories



- ▶ Games (e.g., G. Tesauro, D. Silver)
- ▶ Operations research and scheduling (e.g., W. Powell, P. Tadepalli)
- ▶ Recently: robotics (e.g., P. Abbeel, J. Peters, P. Stone, M. Riedmiller)

Motivation



- ▶ Learning system in an **unknown environment**
- ▶ Knowledge **only** through interacting with environment
- ▶ **Explores** the environment and receives **rewards**
- ▶ Find strategy/**policy**, which maximizes overall reward
 - ▶▶ Optimal behavior

Bayesian Decision Theory

- ▶ Make optimal decisions \mathbf{a}^* by maximizing an expected utility

$$\mathbf{a}^* \in \arg \max_{\mathbf{a}} \mathbb{E}[r(\mathbf{a})] = \arg \max_{\mathbf{a}} \sum_{j=1}^m r(\mathbf{s}_j, \mathbf{a}) p(\mathbf{s}_j)$$

\mathbf{a} : decision

\mathbf{s} : information about environment/[state](#)

Bayesian Decision Theory

- ▶ Make optimal decisions \mathbf{a}^* by maximizing an expected utility

$$\mathbf{a}^* \in \arg \max_{\mathbf{a}} \mathbb{E}[r(\mathbf{a})] = \arg \max_{\mathbf{a}} \sum_{j=1}^m r(\mathbf{s}_j, \mathbf{a}) p(\mathbf{s}_j)$$

\mathbf{a} : decision

\mathbf{s} : information about environment/[state](#)

- ▶ Bayesian sequential decision theory (statistics)
- ▶ Optimal control theory (engineering)
- ▶ Reinforcement learning (computer science, psychology)

Example: Winning the Lottery

Actions	Outcomes
a_1 : play	s_1 : Win the lottery
a_2 : don't play	s_2 : Don't win the lottery



Example: Winning the Lottery

Actions	Outcomes
a_1 : play	s_1 : Win the lottery
a_2 : don't play	s_2 : Don't win the lottery

Optimal action

$$a^* = \arg \max_{a_i} \sum_{j=1}^2 r_{ij} p(s_j | a_i)$$



Example: Winning the Lottery

Actions	Outcomes
a_1 : play	s_1 : Win the lottery
a_2 : don't play	s_2 : Don't win the lottery

Optimal action

$$a^* = \arg \max_{a_i} \sum_{j=1}^2 r_{ij} p(s_j | a_i)$$



$$\begin{array}{ll} p(s_1 | a_1) = 10^{-7} & r_{11} = 500,000 \text{ USD} \\ p(s_2 | a_1) = 1 - 10^{-7} & r_{12} = -1 \text{ USD} \\ p(s_1 | a_2) = 0 & r_{21} = 0 \text{ USD} \\ p(s_2 | a_2) = 1 & r_{22} = 0 \text{ USD} \end{array}$$

►► What is the optimal action for this decision problem?

From Bayesian Decision Theory to RL

- ▶ So far: single decisions. How do we make a **sequence of decisions** in order to achieve some **long-term rewards**?
- ▶ What about state-dependent actions $a_i|s_j$?


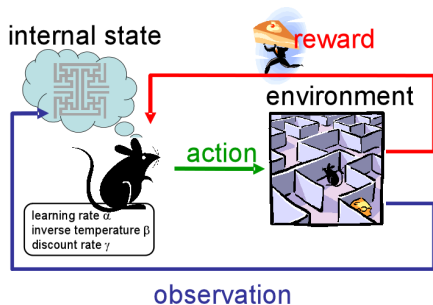
Motivation

Reinforcement Learning Set-Up

Value Function Methods


Policy Search

RL Set-up



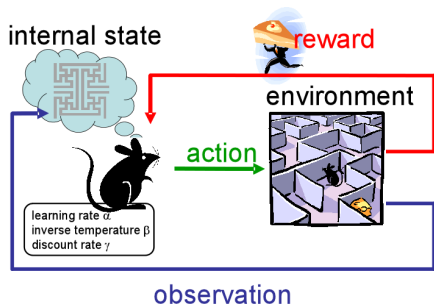
- ▶ **Agent** interacts with **environment** to gain knowledge
- ▶ **Explores** and receives **rewards**
- ▶ **Actions** change the **state** of the environment
- ▶ Choose actions to **maximize long-term reward**

RL Set-up



- ▶ **Agent** interacts with **environment** to gain knowledge
- ▶ **Explores** and receives **rewards**
- ▶ **Actions** change the **state** of the environment
- ▶ Choose actions to **maximize long-term reward**

RL Set-up




- ▶ **Agent** interacts with **environment** to gain knowledge
- ▶ **Explores** and receives **rewards**
- ▶ **Actions** change the **state** of the environment
- ▶ Choose actions to **maximize long-term reward**

▶▶ Markov Decision Process

Markov Decision Process: Definition

- \mathcal{S} : State space (finite)
- \mathcal{A} : Action space (finite)
- \mathcal{P} : Transition probability $p(s_{k+1}|s_k, a_k)$
- r : Reward function
- $\gamma \in [0, 1)$: Discount factor
- π : Policy
 - Deterministic: $a = \pi(s)$
 - Stochastic: $a \sim p_\pi(a|s)$


alternative notation: $p_\pi(a|s) = \pi(a|s)$



Markov Decision Process: Definition

- \mathcal{S} : State space (finite)
- \mathcal{A} : Action space (finite)
- \mathcal{P} : Transition probability $p(s_{k+1}|s_k, a_k)$
- r : Reward function
- $\gamma \in [0, 1)$: Discount factor
- π : Policy
 - Deterministic: $a = \pi(s)$
 - Stochastic: $a \sim p_\pi(a|s)$

alternative notation: $p_\pi(a|s) = \pi(a|s)$



Objective

Find a policy π^* that maximizes the expected long-term reward

$$V^\pi(s) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid s_0 = s, \pi\right], \quad r_{k+1} = r_{k+1}(s_k, a_k, s_{k+1})$$

Categorization of RL Algorithms

- Value function methods
 - ▶▶ Use structure of a value function to discover optimal policies
- Value function-free methods (e.g., policy search)
 - ▶▶ Search in policy space directly

Motivation

Reinforcement Learning Set-Up

Value Function Methods

Policy Search

Value Functions

- **State Value Function:** How good is it to be in a particular state s ?

Well, this depends on the current policy:

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[R | s_0 = s] = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, \pi\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(s_1) | s_0 = s, \pi] \quad \text{Self-consistency} \end{aligned}$$

Value Functions

- ▶ **State Value Function:** How good is it to be in a particular state s ?

Well, this depends on the current policy:

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[R | s_0 = s] = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, \pi\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(s_1) | s_0 = s, \pi] \quad \text{Self-consistency} \end{aligned}$$

- ▶ **State-Action Value Function:** How good is it to be in a particular state s and apply action a , and afterwards follow policy π ?

$$\begin{aligned} Q^\pi(s, a) &= \mathbb{E}[R | s_0 = s, a_0 = a, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, a_0 = a, \pi\right] \end{aligned}$$

Value Functions

- ▶ **State Value Function:** How good is it to be in a particular state s ?

Well, this depends on the current policy:

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[R | s_0 = s] = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, \pi\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(s_1) | s_0 = s, \pi] \quad \text{Self-consistency} \end{aligned}$$

- ▶ **State-Action Value Function:** How good is it to be in a particular state s and apply action a , and afterwards follow policy π ?

$$\begin{aligned} Q^\pi(s, a) &= \mathbb{E}[R | s_0 = s, a_0 = a, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, a_0 = a, \pi\right] \\ &= \mathbb{E}\left[r_1(s_0, a_0, s_1) + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | s_0 = s, a_0 = a, \pi\right] \\ &= \end{aligned}$$

Value Functions

- ▶ **State Value Function:** How good is it to be in a particular state s ?

Well, this depends on the current policy:

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[R|s_0 = s] = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, \pi\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(s_1) | s_0 = s, \pi] \quad \text{Self-consistency} \end{aligned}$$

- ▶ **State-Action Value Function:** How good is it to be in a particular state s and apply action a , and afterwards follow policy π ?

$$\begin{aligned} Q^\pi(s, a) &= \mathbb{E}[R | s_0 = s, a_0 = a, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, a_0 = a, \pi\right] \\ &= \mathbb{E}\left[r_1(s_0, a_0, s_1) + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | s_0 = s, a_0 = a, \pi\right] \\ &= \end{aligned}$$

Value Functions

- ▶ **State Value Function:** How good is it to be in a particular state s ?

Well, this depends on the current policy:

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[R | s_0 = s] = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, \pi\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(s_1) | s_0 = s, \pi] \quad \text{Self-consistency} \end{aligned}$$

- ▶ **State-Action Value Function:** How good is it to be in a particular state s and apply action a , and afterwards follow policy π ?

$$\begin{aligned} Q^\pi(s, a) &= \mathbb{E}[R | s_0 = s, a_0 = a, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | s_0 = s, a_0 = a, \pi\right] \\ &= \mathbb{E}\left[r_1(s_0, a_0, s_1) + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | s_0 = s, a_0 = a, \pi\right] \\ &= \mathbb{E}\left[r_1(s_0, a_0, s_1) + \gamma V^\pi(s_1) | s_0 = s, a_0 = a, \pi\right] \end{aligned}$$

Bellman Operator

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= T^\pi[V^\pi](\mathbf{s}) \\ T^\pi[V^\pi](\mathbf{s}) &:= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi] \end{aligned}$$

Bellman Operator

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= T^\pi[V^\pi](\mathbf{s}) \end{aligned}$$

$$T^\pi[V^\pi](\mathbf{s}) := \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi]$$

►► $V^\pi = T^\pi[V^\pi]$, T^π : **Bellman operator** (linear affine)

Bellman Operator

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= T^\pi[V^\pi](\mathbf{s}) \end{aligned}$$

$$T^\pi[V^\pi](\mathbf{s}) := \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi]$$

►► $V^\pi = T^\pi[V^\pi]$, T^π : **Bellman operator** (linear affine)

- **Fixed point equation** with a **unique** solution for V^π
(Banach's FP theorem)

Bellman Operator

$$\begin{aligned}V^\pi(\mathbf{s}) &= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi] \\&= T^\pi[V^\pi](\mathbf{s}) \\T^\pi[V^\pi](\mathbf{s}) &:= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi]\end{aligned}$$

►► $V^\pi = T^\pi[V^\pi]$, T^π : **Bellman operator** (linear affine)

- **Fixed point equation** with a **unique** solution for V^π
(Banach's FP theorem)

$$\begin{aligned}T^\pi[V^\pi] &= r^\pi + \gamma \mathcal{P}^\pi V^\pi \\ \implies V^\pi &= r^\pi + \gamma \mathcal{P}^\pi V^\pi\end{aligned}$$

for suitable representations $r^\pi, \mathcal{P}^\pi, V^\pi$

Optimal Policies and Value Functions

- Optimal policy π^* ensures that $V^{\pi^*}(s) \geq V^{\pi}(s) \quad \forall s \in \mathcal{S}, \pi$
- Existence of π^* ? Uniqueness?

Optimal Policies and Value Functions

- ▶ Optimal policy π^* ensures that $V^{\pi^*}(\mathbf{s}) \geq V^{\pi}(\mathbf{s}) \quad \forall \mathbf{s} \in \mathcal{S}, \pi$
- ▶ Existence of π^* ? Uniqueness?
- ▶ Optimal state value function:
$$\forall \mathbf{s} \in \mathcal{S} : V^*(\mathbf{s}) = V^{\pi^*}(\mathbf{s}) = \max_{\pi} V^{\pi}(\mathbf{s})$$

Optimal Policies and Value Functions

- ▶ **Optimal policy π^*** ensures that $V^{\pi^*}(\mathbf{s}) \geq V^{\pi}(\mathbf{s}) \quad \forall \mathbf{s} \in \mathcal{S}, \pi$
- ▶ Existence of π^* ? Uniqueness?
- ▶ **Optimal state value function:**
 $\forall \mathbf{s} \in \mathcal{S} : V^*(\mathbf{s}) = V^{\pi^*}(\mathbf{s}) = \max_{\pi} V^{\pi}(\mathbf{s})$
- ▶ **Optimal state-action value function:**
 $\forall \mathbf{s} \in \mathcal{S} : Q^*(\mathbf{s}, \mathbf{a}) = \max_{\pi} Q^{\pi}(\mathbf{s}, \mathbf{a})$
▶▶ Expected return of choosing action \mathbf{a} in state \mathbf{s} and afterwards following the optimal policy π^* . Note that

$$Q^*(\mathbf{s}, \mathbf{a}) = \mathbb{E}[r_{t+1} + \gamma V^*(\mathbf{s}_{t+1}) | \mathbf{s}_t = \mathbf{s}, \mathbf{a}_t = \mathbf{a}]$$

Bellman Optimality Equations

$$V^*(\mathbf{s}) = \max_a Q^*(\mathbf{s}, \mathbf{a})$$

Bellman Optimality Equations

$$\begin{aligned} V^*(\mathbf{s}) &= \max_{\mathbf{a}} Q^*(\mathbf{s}, \mathbf{a}) \\ &= \max_{\mathbf{a}} \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \end{aligned}$$

Bellman Optimality Equations

$$\begin{aligned} V^*(\mathbf{s}) &= \max_a Q^*(\mathbf{s}, \mathbf{a}) \\ &= \max_a \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \\ &= \max_a \mathbb{E} \left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \end{aligned}$$

Bellman Optimality Equations

$$\begin{aligned} V^*(\mathbf{s}) &= \max_{\mathbf{a}} Q^*(\mathbf{s}, \mathbf{a}) \\ &= \max_{\mathbf{a}} \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \\ &= \max_{\mathbf{a}} \mathbb{E} \left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \\ &= \max_{\mathbf{a}} \mathbb{E} \left[r_1 + \gamma V^*(\mathbf{s}_1) \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \end{aligned}$$

Bellman Optimality Equations

$$\begin{aligned} V^*(\mathbf{s}) &= \max_{\mathbf{a}} Q^*(\mathbf{s}, \mathbf{a}) \\ &= \max_{\mathbf{a}} \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \\ &= \max_{\mathbf{a}} \mathbb{E} \left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \\ &= \max_{\mathbf{a}} \mathbb{E} \left[r_1 + \gamma V^*(\mathbf{s}_1) \mid \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \mathbf{a}, \pi^* \right] \end{aligned}$$

►► $V^* = T^*[V^*]$ $T^* : \text{Bellman optimality operator (nonlinear)}$

Bellman Optimality Equations

$$\begin{aligned} V^*(s) &= \max_a Q^*(s, a) \\ &= \max_a \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid s_0 = s, a_0 = a, \pi^* \right] \\ &= \max_a \mathbb{E} \left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} \mid s_0 = s, a_0 = a, \pi^* \right] \\ &= \max_a \mathbb{E} \left[r_1 + \gamma V^*(s_1) \mid s_0 = s, a_0 = a, \pi^* \right] \end{aligned}$$

►► $V^* = T^*[V^*]$ $T^* : \text{Bellman optimality operator (nonlinear)}$

$$\begin{aligned} Q^*(s, a) &= \mathbb{E} [r_1 + \gamma \max_{a'} Q^*(s_1, a') \mid s_0 = s, a_0 = a] \\ &= \end{aligned}$$

Bellman Optimality Equations

$$\begin{aligned} V^*(s) &= \max_a Q^*(s, a) \\ &= \max_a \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} \mid s_0 = s, a_0 = a, \pi^* \right] \\ &= \max_a \mathbb{E} \left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} \mid s_0 = s, a_0 = a, \pi^* \right] \\ &= \max_a \mathbb{E} \left[r_1 + \gamma V^*(s_1) \mid s_0 = s, a_0 = a, \pi^* \right] \end{aligned}$$

►► $V^* = T^*[V^*]$ $T^* : \text{Bellman optimality operator (nonlinear)}$

$$\begin{aligned} Q^*(s, a) &= \mathbb{E} \left[r_1 + \gamma \max_{a'} Q^*(s_1, a') \mid s_0 = s, a_0 = a \right] \\ &= \mathbb{E} \left[r_1 + \gamma V^*(s_1) \mid s_0 = s, a_0 = a \right] \end{aligned}$$

Solving MDPs

- ▶ Assume you know V^*
- ▶ **One-step search:** be “greedy” with respect to V^*

$$\pi^*(s) = \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^*(s')]$$

Solving MDPs

- ▶ Assume you know V^*
- ▶ **One-step search:** be “greedy” with respect to V^*

$$\pi^*(s) = \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^*(s')]$$

- ▶ Assume you know Q^*
- ▶ **Zero-step search:**

$$\pi^*(s) = \max_a Q^*(s, a)$$

Solving MDPs

- ▶ Assume you know V^*
- ▶ **One-step search:** be “greedy” with respect to V^*

$$\pi^*(s) = \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^*(s')]$$

- ▶ Assume you know Q^*
- ▶ **Zero-step search:**

$$\pi^*(s) = \max_a Q^*(s, a)$$

Assumptions for solving the Bellman equations exactly:

- ▶ Know the transition probabilities $p(s'|s, a)$
- ▶ Sufficient computational resources available
- ▶ Markov property holds

Solving MDPs

- ▶ Assume you know V^*
- ▶ **One-step search:** be “greedy” with respect to V^*

$$\pi^*(s) = \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^*(s')]$$

- ▶ Assume you know Q^*
- ▶ **Zero-step search:**

$$\pi^*(s) = \max_a Q^*(s, a)$$

Assumptions for solving the Bellman equations exactly:

- ▶ Know the transition probabilities $p(s'|s, a)$
 - ▶ Sufficient computational resources available
 - ▶ Markov property holds
- ▶▶ Approximate solutions in practice

Solving MDPs (2)

- Exact: Dynamic programming
- Approximate: Monte Carlo, Temporal Difference Learning

Dynamic Programming

Assumptions:

- Perfect model $p(s'|s, a)$ is known
- Typically finite state spaces \mathcal{S} and action spaces \mathcal{A}
- Expected immediate rewards $\mathbb{E}[r(s, a, s')]$ are known
- ▶▶ Use value functions to structure the search for good policies

Policy Evaluation

Objective

For a given policy π , find the corresponding value function V^π

- ▶ Exploit the fixed-point property of the value function $V^\pi = T^\pi[V^\pi]$:
 - ▶ Initialize V_0^π arbitrarily
 - ▶ Find V^π as the limit of the sequence V_0^π, V_1^π, \dots

Policy Evaluation

Objective

For a given policy π , find the corresponding value function V^π

- ▶ Exploit the fixed-point property of the value function $V^\pi = T^\pi[V^\pi]$:
 - ▶ Initialize V_0^π arbitrarily
 - ▶ Find V^π as the limit of the sequence V_0^π, V_1^π, \dots

Update Rule

$$\forall s \in \mathcal{S} : V_{k+1}^\pi(s) \leftarrow \mathbb{E}[r(s, a, s') + \gamma V_k^\pi(s') | s, \pi]$$
$$a \sim p_\pi(a|s), \quad s' \sim p(s'|s, a)$$

► Bootstrapping

Policy Improvement

- So far: We know V^π , but we want V^*
- Find a better policy π'

Objective

Find a policy $\pi' \geq \pi$, i.e., $V^{\pi'} \geq V^\pi$

Policy Improvement Theorem

Policy Improvement Theorem

If π, π' are two (deterministic) policies with

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^{\pi}(\mathbf{s}, \pi'(\mathbf{s})) \geq Q^{\pi}(\mathbf{s}, \pi(\mathbf{s})) = V^{\pi}(\mathbf{s})$$

then $\pi' \geq \pi$ and $V^{\pi'} \geq V^{\pi}$, i.e., π' improves π .

Policy Improvement Theorem

Policy Improvement Theorem

If π, π' are two (deterministic) policies with

$$\forall s \in \mathcal{S} : \quad Q^\pi(s, \pi'(s)) \geq Q^\pi(s, \pi(s)) = V^\pi(s)$$

then $\pi' \geq \pi$ and $V^{\pi'} \geq V^\pi$, i.e., π' improves π .

- For stochastic policies:

$$\begin{aligned} Q^\pi(s, \pi'(s)) &= \mathbb{E}_a[Q^\pi(s, a)] & a &\sim p_{\pi'}(a|s) \\ &= \sum_a p_{\pi'}(a|s) Q^\pi(s, a) \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$V^\pi(\mathbf{s}) \stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s}))$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 Q^\pi(\mathbf{s}_2, \pi'(\mathbf{s}_2)) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 Q^\pi(\mathbf{s}_2, \pi'(\mathbf{s}_2)) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 Q^\pi(\mathbf{s}_2, \pi'(\mathbf{s}_2)) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 r_3 + \gamma^3 V^\pi(\mathbf{s}_3) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 Q^\pi(\mathbf{s}_2, \pi'(\mathbf{s}_2)) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 r_3 + \gamma^3 V^\pi(\mathbf{s}_3) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

Proof

$$\forall \mathbf{s} \in \mathcal{S} : \quad Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \geq V^\pi(\mathbf{s}) = Q^\pi(\mathbf{s}, \pi(\mathbf{s})) \quad (*)$$

$$\begin{aligned} V^\pi(\mathbf{s}) &\stackrel{(*)}{\leq} Q^\pi(\mathbf{s}, \pi'(\mathbf{s})) \\ &= \mathbb{E}[r_1(\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1) + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0)] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma Q^\pi(\mathbf{s}_1, \pi'(\mathbf{s}_1)) | \mathbf{s}_0 = \mathbf{s}, \mathbf{a}_0 = \pi'(\mathbf{s}_0), \mathbf{a}_1 = \pi'(\mathbf{s}_1)] \\ &= \mathbb{E}[r_1 + \gamma \mathbb{E}[r_2 + \gamma V^\pi(\mathbf{s}_2)] | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 V^\pi(\mathbf{s}_2) | \mathbf{s}_0 = \mathbf{s}] \\ &\stackrel{(*)}{\leq} \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 Q^\pi(\mathbf{s}_2, \pi'(\mathbf{s}_2)) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}[r_1 + \gamma r_2 + \gamma^2 r_3 + \gamma^3 V^\pi(\mathbf{s}_3) | \mathbf{s}_0 = \mathbf{s}] \\ &\dots \\ &\leq \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}\right] = V^{\pi'}(\mathbf{s}) \end{aligned}$$

Policy Improvement

- ▶ Easy to evaluate a change in the policy at a single state s

Policy Improvement

- ▶ Easy to evaluate a change in the policy at a single state s
- ▶ Extend this idea to **all** states:

$$\begin{aligned}\forall s \in \mathcal{S} : \quad \pi'(s) &= \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^\pi(s')] \\ &= \arg \max_a Q^\pi(s, a)\end{aligned}$$

Policy Improvement

- ▶ Easy to evaluate a change in the policy at a single state s
- ▶ Extend this idea to **all** states:

$$\begin{aligned}\forall s \in \mathcal{S} : \quad \pi'(s) &= \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^\pi(s')] \\ &= \arg \max_a Q^\pi(s, a)\end{aligned}$$

- ▶ When $\pi' = \pi$ then $V^{\pi'} = V^\pi$ ►► Convergence (Why?)

Policy Improvement

- ▶ Easy to evaluate a change in the policy at a single state s
- ▶ Extend this idea to **all** states:

$$\begin{aligned}\forall s \in \mathcal{S} : \quad \pi'(s) &= \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^\pi(s')] \\ &= \arg \max_a Q^\pi(s, a)\end{aligned}$$

- ▶ When $\pi' = \pi$ then $V^{\pi'} = V^\pi$ ►► Convergence (Why?)
- ▶ If V^π is known, we **need a model** $p(s'|s, a)$ for the **one-step look-ahead**

Policy Improvement

- ▶ Easy to evaluate a change in the policy at a single state s
- ▶ Extend this idea to **all** states:

$$\begin{aligned}\forall s \in \mathcal{S} : \quad \pi'(s) &= \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^\pi(s')] \\ &= \arg \max_a Q^\pi(s, a)\end{aligned}$$

- ▶ When $\pi' = \pi$ then $V^{\pi'} = V^\pi$ ►► Convergence (Why?)
- ▶ If V^π is known, we **need a model** $p(s'|s, a)$ for the **one-step look-ahead**
- ▶ If Q^π is known, we **don't need a model** (no prediction required)

Policy Improvement

- ▶ Easy to evaluate a change in the policy at a single state s
- ▶ Extend this idea to **all** states:

$$\begin{aligned}\forall s \in \mathcal{S} : \quad \pi'(s) &= \arg \max_a \mathbb{E}[r(s, a, s') + \gamma V^\pi(s')] \\ &= \arg \max_a Q^\pi(s, a)\end{aligned}$$

- ▶ When $\pi' = \pi$ then $V^{\pi'} = V^\pi$ ►► Convergence (Why?)
- ▶ If V^π is known, we **need a model** $p(s'|s, a)$ for the **one-step look-ahead**
- ▶ If Q^π is known, we **don't need a model** (no prediction required)
- ▶ Greedy policy update with respect to the value function (but look implicitly at **long-term** rewards)

Policy Iteration

$$\pi_0 \xrightarrow{E} V^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} V^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} V^* \xrightarrow{I} \pi^*$$

- E: policy evaluation
- I: policy improvement

Policy Iteration

$$\pi_0 \xrightarrow{E} V^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} V^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} V^* \xrightarrow{I} \pi^*$$

- E: policy evaluation
- I: policy improvement
- Strict policy improvement at each step (unless policy is already optimal)
- Converges often after a few iterations

Policy Iteration

$$\pi_0 \xrightarrow{E} V^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} V^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} V^* \xrightarrow{I} \pi^*$$

- ▶ E: policy evaluation
- ▶ I: policy improvement
- ▶ Strict policy improvement at each step
(unless policy is already optimal)
- ▶ Converges often after a few iterations
- ▶ Each policy evaluation is itself an iterative process
 - ▶▶ Can be really slow!

Value Iteration

$$\pi_0 \xrightarrow{E} V^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} V^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} V^* \xrightarrow{I} \pi^*$$

- Stop policy evaluation after a **single** update
 - ▶▶ No longer an iterative process

Value Iteration

$$\pi_0 \xrightarrow{E} V^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} V^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} V^* \xrightarrow{I} \pi^*$$

- ▶ Stop policy evaluation after a **single** update
 - ▶▶ No longer an iterative process

Update Rule

$$\begin{aligned} V_{k+1}(s) &= \max_a \mathbb{E}[r(s, a, s') + \gamma V_k(s') | s, a] \\ &= \max_a \sum_{s'} p(s' | s, a) (\mathbb{E}[r(s, a, s')] + \gamma V_k(s')) \end{aligned}$$

Value Iteration

$$\pi_0 \xrightarrow{E} V^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} V^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} V^* \xrightarrow{I} \pi^*$$


- ▶ Stop policy evaluation after a **single** update
- ▶▶ No longer an iterative process

Update Rule

$$\begin{aligned} V_{k+1}(s) &= \max_a \mathbb{E}[r(s, a, s') + \gamma V_k(s') | s, a] \\ &= \max_a \sum_{s'} p(s' | s, a) (\mathbb{E}[r(s, a, s')] + \gamma V_k(s')) \end{aligned}$$


- ▶ **Bootstrapping**
- ▶ **Bellman optimality equation as an update rule:** $V_{k+1} \leftarrow T^*[V_k]$
- ▶ Identical to policy evaluation backup if you add the max operator

Generalized Policy Iteration (GPI)




- Abstraction/generalization of policy iteration
- Two **interacting processes**: policy evaluation/improvement

Generalized Policy Iteration (GPI)



- Abstraction/generalization of policy iteration
- Two **interacting processes**: policy evaluation/improvement
- Update details abstracted away (policy needs to be greedy)
- Don't need to go full way to π or V^π , just "in the direction"

Generalized Policy Iteration (GPI)



- Abstraction/generalization of policy iteration
- Two **interacting processes**: policy evaluation/improvement
- Update details abstracted away (policy needs to be greedy)
- Don't need to go full way to π or V^π , just "in the direction"
- Both processes **converge to a single joint solution** (V^*, π^*)
- Value iteration (incomplete value function updates) is one example of GPI

Summary: Dynamic Programming

- ▶ Find optimal policies via **value functions** and **bootstrapping**
- ▶ **Exact** method (standard method in optimal control)

Summary: Dynamic Programming

- ▶ Find optimal policies via **value functions** and **bootstrapping**
- ▶ **Exact** method (standard method in optimal control)
- ▶ **Computationally expensive** (sweeps through state-action spaces)
 - ▶▶ Curse of dimensionality

Summary: Dynamic Programming

- ▶ Find optimal policies via **value functions** and **bootstrapping**
- ▶ **Exact** method (standard method in optimal control)
- ▶ **Computationally expensive** (sweeps through state-action spaces)
 - ▶▶ Curse of dimensionality
- ▶ Exponentially faster than any direct policy search method (if the policy space is not restricted)
- ▶ **Requires a model** $p(s'|s, a)$ and the knowledge of $\mathbb{E}[r(s, a, s')]$ **before** DP can be applied.
- ▶ 2 algorithms: **Policy iteration, value iteration**

Approximate Value Function Methods

Look trajectory samples


- Monte Carlo methods
- Temporal difference learning

Monte Carlo Methods

Key Idea

$$V^{\pi}(s) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \pi, s_0 = s\right]$$

Estimate value function by **averaging returns**
of **sampled trajectories**




Monte Carlo Methods

Key Idea

$$V^{\pi}(s) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \pi, s_0 = s\right]$$

Estimate value function by **averaging returns of sampled trajectories**




Properties:

- ▶ **Model free** (no knowledge of $p(s'|s, a)$ required) ► very general
- ▶ Learn from online experience (sampled trajectories of states, actions, rewards)
- ▶ Compute the same value function as DP (in the limit)

Episodic Set-up

- ▶ Consider a finite time horizon of length T
- ▶ Usually a fixed set of initial states s_0
- ▶ Observe rewards $r_1, r_2, \dots, r_T | s_0, \pi$
- ▶ Value estimates $V^\pi(s_0)$ updated at the end of an episode (not after each time step)



First-Visit Monte Carlo Policy Evaluation

- ▶ Generate trajectories with policy π
- ▶ Record reward after visiting state s , average at the end

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all states  $s \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $s$ 
5:      $R(s) \leftarrow [R(s), r]$  ▷ Append  $r$  to array
6:   end for
7:    $V^\pi(s) \approx \mathbb{E}[R(s)|\pi] \approx \frac{1}{|R(s)|} \sum_{j=1}^{|R(s)|} R(s)[j]$ 
8: end for
```

- ▶ Convergence to the correct value V^π in the limit

Properties

- ▶ Computational complexity **independent** of the size of the state space
- ▶ Very valuable if we are only interested in values starting from a small set of states s_0 ►► episodic set-up
- ▶ **No bootstrapping** (unlike DP)
- ▶ Learn from **actual** experience (DP only uses the model)

Properties

- ▶ Computational complexity **independent** of the size of the state space
- ▶ Very valuable if we are only interested in values starting from a small set of states s_0 ►► episodic set-up
- ▶ **No bootstrapping** (unlike DP)
- ▶ Learn from **actual** experience (DP only uses the model)

Next steps?

- ▶ Low-cost (computation) solution to computer V^π
- ▶ Is this useful for policy improvement? Why? (not?)

Properties

- ▶ Computational complexity **independent** of the size of the state space
- ▶ Very valuable if we are only interested in values starting from a small set of states s_0 ➡ episodic set-up
- ▶ **No bootstrapping** (unlike DP)
- ▶ Learn from **actual** experience (DP only uses the model)

Next steps?

- ▶ Low-cost (computation) solution to computer V^π
- ▶ Is this useful for policy improvement? Why? (not?)
- ▶ If we only know V^π , we need to perform a **one-step search** for policy improvement...
 - ➡ Need an **estimate of Q^π** if we don't have a model $p(s'|s, a)$

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all states  $s \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $s$ 
5:      $R(s) \leftarrow [R(s), r]$  ▷ Append  $r$  to array
6:   end for
7:    $V^\pi(s) \approx \mathbb{E}[R(s)|\pi] \approx \frac{1}{|R(s)|} \sum_{j=1}^{|R(s)|} R(s)[j]$ 
8: end for
```

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

- Any potential problems?

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

► Any potential problems?

► ∞ many trajectories (use GPI idea)

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

- ▶ Any potential problems?
- ▶ Deterministic policy?

▶ ∞ many trajectories (use GPI idea)

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

- ▶ Any potential problems?
- ▶ Deterministic policy? Not many state-action pairs are visited!

▶ ∞ many trajectories (use GPI idea)

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

- ▶ Any potential problems?
- ▶ Deterministic policy? Not many state-action pairs are visited!
- ▶ Ideally, estimate $Q^\pi(s, a)$ for **all** actions $a \in \mathcal{A}$

▶ ∞ many trajectories (use GPI idea)

Monte Carlo for Q Function

```
1: for  $i = 1$  to  $\infty$  do
2:   Generate trajectory  $\tau_i^\pi$  with policy  $\pi$ 
3:   for all state-action pairs  $(s, a) \in \tau_i$  do
4:      $r \leftarrow$  sum of rewards following the first occurrence of  $(s, a)$ 
5:      $R(s, a) \leftarrow [R(s, a), r]$  ▷ Append  $r$  to array
6:   end for
7:    $Q^\pi(s, a) \approx \mathbb{E}[R(s, a) | \pi] \approx \frac{1}{|R(s, a)|} \sum_{j=1}^{|R(s, a)|} R(s, a)[j]$ 
8: end for
```

- ▶ Any potential problems?
- ▶ Deterministic policy? Not many state-action pairs are visited!
- ▶ Ideally, estimate $Q^\pi(s, a)$ for **all** actions $a \in \mathcal{A}$

- ▶▶ ∞ many trajectories (use GPI idea)
- ▶▶ Maintain exploration

Exploration

- **Exploring starts:** non-zero probability that each state-action pair is chosen as the start (difficult in practice, not commonly used)

Exploration

- ▶ **Exploring starts:** non-zero probability that each state-action pair is chosen as the start (difficult in practice, not commonly used)
- ▶ **Stochastic policies** with non-zero probability on each action
 - ▶ $\forall s \in \mathcal{S}, a \in \mathcal{A} : p_{\pi}(a|s) > 0$
 - ▶ Example: ϵ -greedy policies, softmax policies
 - ▶ ϵ -greedy:

$$p_{\pi}(a|s) = \begin{cases} \pi(s) & \text{with probability } 1 - \epsilon \\ \mathcal{U}(\mathcal{A}) & \text{with probability } \epsilon \end{cases}$$

Exploration

- ▶ **Exploring starts**: non-zero probability that each state-action pair is chosen as the start (difficult in practice, not commonly used)
- ▶ **Stochastic policies** with non-zero probability on each action
 - ▶ $\forall s \in \mathcal{S}, a \in \mathcal{A} : p_{\pi}(a|s) > 0$
 - ▶ Example: ϵ -greedy policies, softmax policies
 - ▶ ϵ -greedy:

$$p_{\pi}(a|s) = \begin{cases} \pi(s) & \text{with probability } 1 - \epsilon \\ \mathcal{U}(\mathcal{A}) & \text{with probability } \epsilon \end{cases}$$

▶▶ **Trade off exploring** the world and **exploiting** current knowledge

On-Policy Monte Carlo Control


$$\pi_0 \xrightarrow{E} Q^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} Q^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} Q^* \xrightarrow{I} \pi^*$$

- ▶ Approximate optimal policies
- ▶ Follow GPI idea: ϵ -greedy policy

On-Policy Monte Carlo Control

$$\pi_0 \xrightarrow{E} Q^{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} Q^{\pi_1} \xrightarrow{I} \dots \xrightarrow{E} Q^* \xrightarrow{I} \pi^*$$

- ▶ Approximate optimal policies
- ▶ Follow GPI idea: ϵ -greedy policy
- ▶ MC version of policy iteration:
 - ▶ **Policy evaluation** via Monte Carlo estimates (stochastic policy)
▶▶ get Q^π
 - ▶ **Policy improvement**: select greedy policy with respect to Q^π



Off-Policy Monte Carlo Control

- So far **on-policy**: Evaluate Q^π and, subsequently, apply π when interacting with the environment

Off-Policy Monte Carlo Control

- ▶ So far **on-policy**: Evaluate Q^π and, subsequently, apply π when interacting with the environment
- ▶ **Off-policy**: Evaluate Q^π but when interacting with the environment, follow $\pi' \neq \pi$
 - ▶ Experience only generated from π' . Can we still estimate Q^π ?

Off-Policy Monte Carlo Control

- ▶ So far **on-policy**: Evaluate Q^π and, subsequently, apply π when interacting with the environment
- ▶ **Off-policy**: Evaluate Q^π but when interacting with the environment, follow $\pi' \neq \pi$
 - ▶ Experience only generated from π' . Can we still estimate Q^π ?
- ▶ Yes, if we ensure that every action under π is also (occasionally) taken under π' , i.e.,

$$\text{if } p_\pi(\mathbf{a}|\mathbf{s}) > 0 \text{ then } p_{\pi'}(\mathbf{a}|\mathbf{s}) > 0$$

Off-Policy Monte Carlo Control

- ▶ So far **on-policy**: Evaluate Q^π and, subsequently, apply π when interacting with the environment
- ▶ **Off-policy**: Evaluate Q^π but when interacting with the environment, follow $\pi' \neq \pi$
 - ▶▶ Experience only generated from π' . Can we still estimate Q^π ?
- ▶ Yes, if we ensure that every action under π is also (occasionally) taken under π' , i.e.,


$$\text{if } p_\pi(a|s) > 0 \text{ then } p_{\pi'}(a|s) > 0$$

- ▶ Learning can be slow if the policies are very explorative, i.e., $\epsilon \gg 0$

Summary: Monte Carlo Methods

- Learn optimal behavior from interaction
- Easy to focus them on a small set of start states
- Incremental implementation (updates after each episode) possible
- Exploration required!
- No bootstrapping (unlike DP)
- Model free (unlike DP)

Temporal-Difference Learning



- ▶ Between MC and DP
- ▶ Bootstrapping
- ▶ Model-free
- ▶ MC waits until the end of the episode to update V^π, Q^π
- ▶ TD only waits until the **next time step**. Update value functions based on observed reward and the current estimate of V^π, Q^π

TD Policy Evaluation

Generic Update Rule

$$V(\mathbf{s}) \leftarrow V(\mathbf{s}) + \alpha(\kappa - V(\mathbf{s})) = (1 - \alpha)V(\mathbf{s}) + \alpha\kappa$$

κ : “target”

TD Policy Evaluation

Generic Update Rule

$$V(\mathbf{s}) \leftarrow V(\mathbf{s}) + \alpha(\kappa - V(\mathbf{s})) = (1 - \alpha)V(\mathbf{s}) + \alpha\kappa$$

κ : “target”

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[\textcolor{red}{R}(\textcolor{red}{s}) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}\left[\sum_k \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}\right] \\ &= \mathbb{E}[\textcolor{blue}{r}_1 + \gamma V^\pi(\textcolor{blue}{s}_1) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

TD Policy Evaluation

Generic Update Rule

$$V(\mathbf{s}) \leftarrow V(\mathbf{s}) + \alpha(\kappa - V(\mathbf{s})) = (1 - \alpha)V(\mathbf{s}) + \alpha\kappa$$

κ : “target”

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[\textcolor{red}{R}(\mathbf{s}) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}\left[\sum_k \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}\right] \\ &= \mathbb{E}[\textcolor{blue}{r}_1 + \gamma V^\pi(\textcolor{blue}{s}_1) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

- ▶ **MC** uses an estimate of $\textcolor{red}{R}(\mathbf{s})$ as target κ (sample average)

TD Policy Evaluation

Generic Update Rule

$$V(\mathbf{s}) \leftarrow V(\mathbf{s}) + \alpha(\kappa - V(\mathbf{s})) = (1 - \alpha)V(\mathbf{s}) + \alpha\kappa$$

κ : “target”

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[\textcolor{red}{R}(\mathbf{s}) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}\left[\sum_k \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}\right] \\ &= \mathbb{E}[\textcolor{blue}{r}_1 + \gamma V^\pi(\textcolor{blue}{s}_1) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

- ▶ **MC** uses an estimate of $\textcolor{red}{R}(\mathbf{s})$ as target κ (sample average)
- ▶ **DP** uses an estimate of $\textcolor{blue}{r}_1 + \gamma V^\pi(\textcolor{blue}{s}_1)$ as target κ (V_k instead of V^π)

TD Policy Evaluation

Generic Update Rule

$$V(\mathbf{s}) \leftarrow V(\mathbf{s}) + \alpha(\kappa - V(\mathbf{s})) = (1 - \alpha)V(\mathbf{s}) + \alpha\kappa$$

κ : “target”

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[\mathbf{R}(\mathbf{s}) | \mathbf{s}_0 = \mathbf{s}] \\ &= \mathbb{E}\left[\sum_k \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}] \end{aligned}$$

- ▶ MC uses an estimate of $\mathbf{R}(\mathbf{s})$ as target κ (sample average)
- ▶ DP uses an estimate of $r_1 + \gamma V^\pi(\mathbf{s}_1)$ as target κ (V_k instead of V^π)
- ▶ TD target: $\kappa = r_1 + \gamma V^\pi(\mathbf{s}_1)$

TD Policy Evaluation

Generic Update Rule

$$V(s) \leftarrow V(s) + \alpha(\kappa - V(s)) = (1 - \alpha)V(s) + \alpha\kappa$$

κ : “target”

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[\textcolor{red}{R}(\textcolor{red}{s}) | s_0 = s] \\ &= \mathbb{E}\left[\sum_k \gamma^k r_{k+1} | s_0 = s\right] \\ &= \mathbb{E}[\textcolor{blue}{r}_1 + \gamma V^\pi(\textcolor{blue}{s}_1) | s_0 = s] \end{aligned}$$

- ▶ **MC** uses an estimate of $\textcolor{red}{R}(\textcolor{red}{s})$ as target κ (sample average)
- ▶ **DP** uses an estimate of $\textcolor{blue}{r}_1 + \gamma V^\pi(\textcolor{blue}{s}_1)$ as target κ (V_k instead of V^π)
- ▶ TD target: $\kappa = r_1 + \gamma V^\pi(s_1)$
 - Approximate V^π by V_k (same as DP)

TD Policy Evaluation

Generic Update Rule

$$V(s) \leftarrow V(s) + \alpha(\kappa - V(s)) = (1 - \alpha)V(s) + \alpha\kappa$$

κ : “target”

$$\begin{aligned} V^\pi(s) &= \mathbb{E}[\textcolor{red}{R}(\textcolor{red}{s}) | s_0 = s] \\ &= \mathbb{E}\left[\sum_k \gamma^k r_{k+1} | s_0 = s\right] \\ &= \mathbb{E}[r_1 + \gamma V^\pi(\textcolor{blue}{s}_1) | s_0 = s] \end{aligned}$$

- ▶ **MC** uses an estimate of $\textcolor{red}{R}(\textcolor{red}{s})$ as target κ (sample average)
- ▶ **DP** uses an estimate of $r_1 + \gamma V^\pi(\textcolor{blue}{s}_1)$ as target κ (V_k instead of V^π)
- ▶ TD target: $\kappa = r_1 + \gamma V^\pi(\textcolor{blue}{s}_1)$
 - ▶ Approximate V^π by V_k (same as DP)
 - ▶ Sample trajectories (same as MC)
 - ▶▶ Combine MC sampling with DP bootstrapping

TD(0)

- 1: **Init.:** Set $V(s)$ arbitrarily
- 2: **repeat** for each episode
- 3: $a \leftarrow p_{\pi}(a|s)$ ▷ Sample action in current state
- 4: Apply a , observe r, s' ▷ Transition to next state
- 5: $V(s) \leftarrow V(s) + \alpha(r + \gamma V(s') - V(s))$ ▷ Update V
- 6: $s \leftarrow s'$ ▷ Re-set current state
- 7: **until** s is terminal

TD(0)

- 1: **Init.:** Set $V(s)$ arbitrarily
- 2: **repeat** for each episode
- 3: $a \leftarrow p_{\pi}(a|s)$ ▷ Sample action in current state
- 4: Apply a , observe r, s' ▷ Transition to next state
- 5: $V(s) \leftarrow V(s) + \alpha(r + \gamma V(s') - V(s))$ ▷ Update V
- 6: $s \leftarrow s'$ ▷ Re-set current state
- 7: **until** s is terminal

- ▶ **Temporal difference error:** $r + \gamma V(s') - V(s)$ ($= V^{\text{new}} - V^{\text{old}}$)
- ▶ Difference to MC and DP:
 - ▶ MC waits until the end of the episode to update V
 - ▶ DP needs complete distribution $p(s'|s, a)$ of successor states to update V

TD(λ)

- TD(0) uses 1-step returns to update V^π
- MC uses full trajectories to update V^π
- TD(λ), $\lambda \in [0, 1]$, **blends** between them
 - $\lambda = 0$: TD(0), $\lambda = 1$: MC
 - TD(λ) update rule given as a **mixture of multi-step returns**
 - Mixing coefficients $(1 - \lambda)\lambda^k$, $k \geq 0$

SARSA: On-Policy TD Control

- Thus far, we only learned V^π using TD
- Not very useful without model when we want to do control
 - ▶▶ Learn Q^π

SARSA: On-Policy TD Control

- ▶ Thus far, we only learned V^π using TD
- ▶ Not very useful without model when we want to do control
 - ▶▶ Learn Q^π

TD Update for Q^π

$$Q(s, a) \leftarrow Q(s, a) + \underbrace{\alpha(r(s, a, s') + \gamma Q(s', a') - Q(s, a))}_{=(1-\alpha)Q^{\text{old}} + \alpha Q^{\text{new}}}$$

SARSA: On-Policy TD Control

- ▶ Thus far, we only learned V^π using TD
- ▶ Not very useful without model when we want to do control
 - ▶▶ Learn Q^π

TD Update for Q^π

$$Q(s, a) \leftarrow \underbrace{Q(s, a) + \alpha(r(s, a, s') + \gamma Q(s', a') - Q(s, a))}_{=(1-\alpha)Q^{\text{old}} + \alpha Q^{\text{new}}}$$

- ▶ Update rule needs (s, a, r, s', a') ▶▶ SARSA
- ▶ On-policy algorithm
- ▶ a' chosen from s' using policy derived from Q (e.g., ϵ -greedy)
- ▶ Convergence proofs for ϵ -greedy policies

SARSA: On-Policy TD Control

- ▶ Thus far, we only learned V^π using TD
- ▶ Not very useful without model when we want to do control
 - ▶▶ Learn Q^π

TD Update for Q^π

$$Q(s, a) \leftarrow \underbrace{Q(s, a) + \alpha(r(s, a, s') + \gamma Q(s', a') - Q(s, a))}_{=(1-\alpha)Q^{\text{old}} + \alpha Q^{\text{new}}}$$

- ▶ Update rule needs (s, a, r, s', a') ▶▶ SARSA
- ▶ On-policy algorithm
- ▶ a' chosen from s' using policy derived from Q (e.g., ϵ -greedy)
- ▶ Convergence proofs for ϵ -greedy policies

DEMO gridworld

Q-Learning: Off-Policy TD Control

TD Update for Q^*

$$Q(s, a) \leftarrow Q(s, a) + \alpha (r(s, a, s') + \gamma \max_{a'} Q(s', a') - Q(s, a))$$

Q-Learning: Off-Policy TD Control

TD Update for Q^*

$$Q(s, a) \leftarrow Q(s, a) + \alpha (r(s, a, s') + \gamma \max_{a'} Q(s', a') - Q(s, a))$$

- ▶ Off-policy TD control
- ▶ SARSA: learn Q^π , Q-learning: learn Q^*
- ▶ Update the value function Q **independent** of the policy the agent actually follows to generate the samples (max ...)
- ▶ For convergence: keep updating all state-action pairs

Q-Learning: Off-Policy TD Control

TD Update for Q^*



$$Q(s, a) \leftarrow Q(s, a) + \alpha (r(s, a, s') + \gamma \max_{a'} Q(s', a') - Q(s, a))$$

- Off-policy TD control
- SARSA: learn Q^π , Q-learning: learn Q^*
- Update the value function Q **independent** of the policy the agent actually follows to generate the samples (max ...)
- For convergence: keep updating all state-action pairs

DEMO gridworld 2

RL in Continuous Spaces: Function Approximation

- ▶ So far: discrete states and actions
 - ▶▶ Table representation sufficient
- ▶ In continuous spaces: **Function approximation** for better generalization:



Function Approximation

- Typically: (linear) basis function representation

$$V(\mathbf{s}) = \sum_i \theta_i \phi_i(\mathbf{s})$$

- Basis functions ϕ_i are fixed, only parameters θ_i need to be learned

Function Approximation

- Typically: (linear) basis function representation

$$V(s) = \sum_i \theta_i \phi_i(s)$$



- Basis functions ϕ_i are fixed, only parameters θ_i need to be learned
- Gradient descent to update the parameters. Example TD(0):

$$\boldsymbol{\theta}_{k+1} \leftarrow \boldsymbol{\theta}_k + \alpha (v_k - V_k(\mathbf{s}_k)) \frac{\partial V_k(\mathbf{s}_k)}{\partial \boldsymbol{\theta}_k}$$

v_k : approximation/estimate of $V^\pi(\mathbf{s}_k)$, e.g., MC estimate

α : learning rate


- Convergence if v_k is **unbiased**, i.e., $\mathbb{E}[v_k] = V^\pi(\mathbf{s}_k)$



Summary: RL with Value Functions

- Learn policies exploiting properties of the value functions V, Q
- Bellman equations/optimal principle
- Policy evaluation/improvement
- Exact solution: Dynamic programming
- Approximation solution: MC, TD
- Exploration
- Function approximation

Applications



- ▶ Board games (e.g., Tesauro, Silver, Riedmiller)
- ▶ Power systems (e.g., D. Ernst)
- ▶ Robocup (e.g., Stone, Riedmiller)
- ▶ Scheduling tasks (e.g., W. Powell)

Motivation

Reinforcement Learning Set-Up

Value Function Methods

Policy Search

RL without Value Functions: Policy Search

- ▶ Value function approximation is hard, especially in continuous domains
- ▶ Search directly in policy space?

RL without Value Functions: Policy Search

- ▶ Value function approximation is hard, especially in continuous domains
- ▶ Search directly in policy space?
- ▶ Only feasible in a **restricted class Π of policies**

$$\pi = \pi(\boldsymbol{\theta})$$

$$V^\pi(\boldsymbol{\theta}) = \mathbb{E}\left[\sum_{k=0}^T r_{k+1} | \boldsymbol{\theta}\right]$$

- ▶ Consider an episodic set-up, i.e., $p(s_0)$ is given

RL without Value Functions: Policy Search

- ▶ Value function approximation is hard, especially in continuous domains
- ▶ Search directly in policy space?
- ▶ Only feasible in a **restricted class Π of policies**

$$\pi = \pi(\boldsymbol{\theta})$$

$$V^\pi(\boldsymbol{\theta}) = \mathbb{E}\left[\sum_{k=0}^T r_{k+1} | \boldsymbol{\theta}\right]$$

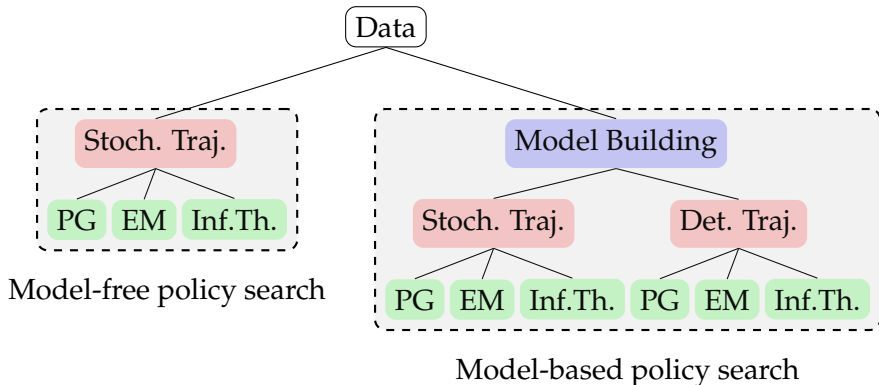
- ▶ Consider an episodic set-up, i.e., $p(s_0)$ is given

Objective (Policy Search)

For a given class $\Pi(\boldsymbol{\theta})$ of policies, find an optimal policy


$$\pi^* \in \arg \max_{\pi \in \Pi(\boldsymbol{\theta})} V^\pi(\boldsymbol{\theta})$$

Categorization of Policy Search Methods




- Data: $(s_i, a_i, r_i), i = 1, \dots, n$
- Model-based vs. model-free policy search
- Policy evaluation (red), policy improvement (green)

Model-based Set-up: Interaction and Simulation



interaction




simulation


Two alternating phases:

- **Interaction:** internal model is refined using experience from interacting with the real system

Model-based Set-up: Interaction and Simulation



interaction

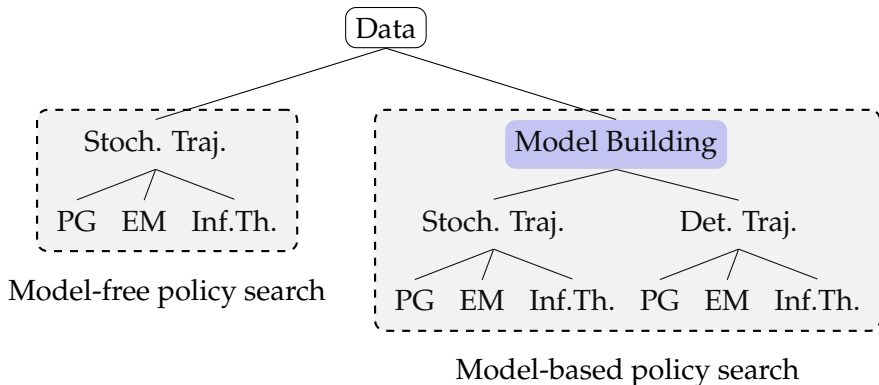


simulation



Two alternating phases:

- ▶ **Interaction:** internal **model is refined** using experience from interacting with the real system
- ▶ **Simulation:** internal model simulates consequences of actions in the real system, **policy is refined**

Model Building





Model Learning





► Pro:

Model Learning





- Pro: No “real” experiments with robot for policy evaluation and improvement (just simulate!) ►► Protect hardware
- Con:

Model Learning



- ▶ Pro: No “real” experiments with robot for policy evaluation and improvement (just simulate!) ►► Protect hardware
- ▶ Con: Model errors ►► Effects of “wrong” models?


Model Learning



- ▶ Pro: No “real” experiments with robot for policy evaluation and improvement (just simulate!) ►► Protect hardware
- ▶ Con: Model errors ►► Effects of “wrong” models?
- ▶ What are good models?

Model Learning


Model learning problem: Find a function $f : x \mapsto f(x) = y$



Observed function values

Model Learning


Model learning problem: Find a function $f : x \mapsto f(x) = y$



Plausible function approximators

Model Learning

Model learning problem: Find a function $f : x \mapsto f(x) = y$




Plausible function approximators

Predictions? Decision Making?

Model Learning

Model learning problem: Find a function $f : x \mapsto f(x) = y$




Plausible function approximators

Predictions? Decision Making? Model Errors!

Model Learning


Model learning problem: Find a function $f : x \mapsto f(x) = y$



Distribution over plausible functions

Model Learning

Model learning problem: Find a function $f : x \mapsto f(x) = y$



Distribution over plausible functions

- ▶ Express **uncertainty** about the underlying function
- ▶ **Bayesian models** (Bayesian linear regression, Gaussian process, ...)

Useful Models

- Probabilistic models!
- Reduce model errors and simulation/optimization bias
- Examples of probabilistic models
 - Bayesian linear regression
 - Gaussian process

Bayesian Linear Regression: Model

- ▶ Model: $y = \theta^\top \phi(x) = \sum_i \theta_i \phi_i(x)$
- ▶ $\phi(x)$: “Features”, e.g., $\phi(x) = [x, x^2]^\top$

Bayesian Linear Regression: Model

- ▶ Model: $y = \theta^\top \phi(x) = \sum_i \theta_i \phi_i(x)$
- ▶ $\phi(x)$: “Features”, e.g., $\phi(x) = [x, x^2]^\top$
- ▶ Distribution over model parameters θ :

$$p(\theta) = \mathcal{N}(\theta \mid m, S)$$

Bayesian Linear Regression: Model

- ▶ Model: $y = \theta^\top \phi(x) = \sum_i \theta_i \phi_i(x)$
- ▶ $\phi(x)$: “Features”, e.g., $\phi(x) = [x, x^2]^\top$
- ▶ Distribution over model parameters θ :

$$p(\theta) = \mathcal{N}(\theta | m, S)$$

- ▶▶ For any θ , one particular function is defined
- ▶▶ Distribution over θ induces **distribution over functions**

Bayesian Linear Regression: Model

- ▶ Model: $y = \theta^\top \phi(x) = \sum_i \theta_i \phi_i(x)$
- ▶ $\phi(x)$: “Features”, e.g., $\phi(x) = [x, x^2]^\top$
- ▶ Distribution over model parameters θ :

$$p(\theta) = \mathcal{N}(\theta \mid m, S)$$

- ▶▶ For any θ , one particular function is defined
- ▶▶ Distribution over θ induces **distribution over functions**

DEMO Bayesian regression

Bayesian Linear Regression: Predictions

- Model: $y = \phi^\top(x)\theta$, $\theta \sim \mathcal{N}(m, S)$

Bayesian Linear Regression: Predictions

- ▶ Model: $y = \phi^\top(x)\theta$, $\theta \sim \mathcal{N}(\mathbf{m}, \mathbf{S})$
- ▶ Predict function values $\mathbf{y} = [y_1, \dots, y_n]^\top$ at inputs $\mathbf{X} = [x_1, \dots, x_n]$

Define $\Phi = \phi(\mathbf{X}) \Rightarrow \mathbf{y} = \Phi^\top \theta$

$$\mathbb{E}[\mathbf{y}] =$$

$$\mathbb{V}[\mathbf{y}] =$$

Bayesian Linear Regression: Predictions

- ▶ Model: $y = \phi^\top(x)\theta$, $\theta \sim \mathcal{N}(m, S)$
- ▶ Predict function values $\mathbf{y} = [y_1, \dots, y_n]^\top$ at inputs $\mathbf{X} = [x_1, \dots, x_n]$

Define $\Phi = \phi(\mathbf{X}) \Rightarrow \mathbf{y} = \Phi^\top \theta$

$$\mathbb{E}[\mathbf{y}] = \mathbb{E}[\Phi^\top \theta] = \Phi^\top \mathbb{E}[\theta] = \Phi^\top m$$

$$\mathbb{V}[\mathbf{y}] =$$

Bayesian Linear Regression: Predictions

- ▶ Model: $y = \phi^\top(x)\theta$, $\theta \sim \mathcal{N}(m, S)$
- ▶ Predict function values $\mathbf{y} = [y_1, \dots, y_n]^\top$ at inputs $\mathbf{X} = [x_1, \dots, x_n]$

Define $\Phi = \phi(\mathbf{X}) \Rightarrow \mathbf{y} = \Phi^\top \theta$

$$\mathbb{E}[\mathbf{y}] = \mathbb{E}[\Phi^\top \theta] = \Phi^\top \mathbb{E}[\theta] = \Phi^\top m$$

$$\mathbb{V}[\mathbf{y}] = \mathbb{V}[\Phi^\top \theta] = \Phi^\top \mathbb{V}[\theta] \Phi = \Phi^\top S \Phi$$

Introduction to Gaussian Processes

- Generalization of Bayesian linear regression
- Nonparametric Bayesian regression method
- Probability distribution over functions
- Fully specified by
 - **Mean function** m (average function)
 - **Covariance function/kernel** k (assumptions on structure)

$$\text{Cov}[f(\mathbf{x}_p), f(\mathbf{x}_q)] = k(\mathbf{x}_p, \mathbf{x}_q)$$


Introduction to Gaussian Processes

- Generalization of Bayesian linear regression
- Nonparametric Bayesian regression method
- Probability distribution over functions
- Fully specified by
 - **Mean function** m (average function)
 - **Covariance function/kernel** k (assumptions on structure)

$$\text{Cov}[f(\mathbf{x}_p), f(\mathbf{x}_q)] = k(\mathbf{x}_p, \mathbf{x}_q)$$

- **Posterior predictive distribution** at \mathbf{x}_* is Gaussian:

$$p(f(\mathbf{x}_*) | \mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(f(\mathbf{x}_*) | m(\mathbf{x}_*), \sigma^2(\mathbf{x}_*))$$



Gaussian Process: Definition

Definition

A Gaussian process is a collection of random variables, any finite number of which has a joint Gaussian distribution.

Gaussian Process: Definition

Definition

A Gaussian process is a collection of random variables, any finite number of which has a joint Gaussian distribution.

- ▶ Look at Gaussian distributions of **function values** f_1, f_2, \dots
- ▶ All of them are jointly Gaussian distributed

$$\mathbb{E}[f(\mathbf{x})] = m(\mathbf{x})$$

$$\text{Cov}[f(\mathbf{x}_i), f(\mathbf{x}_j)] = k(\mathbf{x}_i, \mathbf{x}_j)$$

Gaussian Process: Predictions

- ▶ Given a training set $(\mathbf{x}_i, f(\mathbf{x}_i))_{i=1}^n$, we can predict function values f_{*j} at test inputs \mathbf{x}_{*j}
- ▶ First, compute the joint distribution:

$$p(f, f_* | \mathbf{X}, \mathbf{X}_*) = \mathcal{N} \left(\begin{bmatrix} m(\mathbf{X}) \\ m(\mathbf{X}_*) \end{bmatrix}, \begin{bmatrix} \mathbf{K} & k(\mathbf{X}, \mathbf{X}_*) \\ k(\mathbf{X}_*, \mathbf{X}) & \mathbf{K}_* \end{bmatrix} \right)$$
$$K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j) = \text{Cov}[f(\mathbf{x}_i), f(\mathbf{x}_j)]$$

Gaussian Process: Predictions

- ▶ Given a training set $(\mathbf{x}_i, f(\mathbf{x}_i))_{i=1}^n$, we can predict function values f_{*j} at test inputs \mathbf{x}_{*j}
- ▶ First, compute the joint distribution:

$$p(f, f_* | \mathbf{X}, \mathbf{X}_*) = \mathcal{N} \left(\begin{bmatrix} m(\mathbf{X}) \\ m(\mathbf{X}_*) \end{bmatrix}, \begin{bmatrix} \mathbf{K} & k(\mathbf{X}, \mathbf{X}_*) \\ k(\mathbf{X}_*, \mathbf{X}) & \mathbf{K}_* \end{bmatrix} \right)$$
$$K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j) = \text{Cov}[f(\mathbf{x}_i), f(\mathbf{x}_j)]$$

- ▶ Second, compute the conditional $p(f_* | \mathbf{X}, \mathbf{X}_*, f)$ by **plain Gaussian conditioning**:

Gaussian Process: Predictions


- ▶ Given a training set $(\mathbf{x}_i, f(\mathbf{x}_i))_{i=1}^n$, we can predict function values f_{*j} at test inputs \mathbf{x}_{*j}
- ▶ First, compute the joint distribution:

$$p(f, f_* | \mathbf{X}, \mathbf{X}_*) = \mathcal{N} \left(\begin{bmatrix} m(\mathbf{X}) \\ m(\mathbf{X}_*) \end{bmatrix}, \begin{bmatrix} \mathbf{K} & k(\mathbf{X}, \mathbf{X}_*) \\ k(\mathbf{X}_*, \mathbf{X}) & \mathbf{K}_* \end{bmatrix} \right)$$
$$K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j) = \text{Cov}[f(\mathbf{x}_i), f(\mathbf{x}_j)]$$

- ▶ Second, compute the conditional $p(f_* | \mathbf{X}, \mathbf{X}_*, f)$ by **plain Gaussian conditioning**:

$$p(f_* | \mathbf{X}, \mathbf{X}_*, f) = \mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*)$$
$$\boldsymbol{\mu}_* = m(\mathbf{X}_*) + k(\mathbf{X}_*, \mathbf{X})\mathbf{K}^{-1}(f - m(\mathbf{X}))$$
$$\boldsymbol{\Sigma}_* = \mathbf{K}_* - k(\mathbf{X}_*, \mathbf{X})\mathbf{K}^{-1}k(\mathbf{X}, \mathbf{X}_*)$$

Intuitive Introduction to Gaussian Processes




Prior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\emptyset] = m(\mathbf{x}_*) = 0$$

$$\mathbb{V}[f(\mathbf{x}_*)|\emptyset] = \sigma^2(\mathbf{x}_*) = \text{Cov}[f(\mathbf{x}_*), f(\mathbf{x}_*)] = k(\mathbf{x}_*, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Prior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\emptyset] = m(\mathbf{x}_*) = 0$$

$$\mathbb{V}[f(\mathbf{x}_*)|\emptyset] = \sigma^2(\mathbf{x}_*) = \text{Cov}[f(\mathbf{x}_*), f(\mathbf{x}_*)] = k(\mathbf{x}_*, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes




Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Intuitive Introduction to Gaussian Processes



Posterior belief about the function

Predictive (marginal) mean and variance:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = m(\mathbf{x}_*) = \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{y}$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}] = \sigma^2(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{X}, \mathbf{x}_*)^\top \mathbf{k}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}_*)$$

Properties

- ▶ Universal function approximator ►► extremely expressive
- ▶ Model gives “free” variance estimates


Properties

- ▶ Universal function approximator ►► extremely expressive
- ▶ Model gives “free” variance estimates
- ▶ Computationally involved:
 - ▶ Training: $\mathcal{O}(N^3)$ ►► Repeated inversion of $N \times N$ matrix
 - ▶ Mean prediction: $\mathcal{O}(N)$ ►► Scalar product
 - ▶ Variance prediction: $\mathcal{O}(N^2)$ ►► Matrix-vector multiplication

Properties

- ▶ **Universal function approximator** ▶ extremely expressive
- ▶ Model gives “free” variance estimates
- ▶ Computationally involved:
 - ▶ Training: $\mathcal{O}(N^3)$ ▶ Repeated inversion of $N \times N$ matrix
 - ▶ Mean prediction: $\mathcal{O}(N)$ ▶ Scalar product
 - ▶ Variance prediction: $\mathcal{O}(N^2)$ ▶ Matrix-vector multiplication
- ▶ Sparse approximations exist
- ▶ Code/book online: <http://www.gaussianprocess.org>

Policy Evaluation



Policy evaluation: Compute expected long-term reward

- Stochastic trajectory evaluation
- Deterministic trajectory evaluation


Policy Evaluation

- ▶ Stochastic inference (sampling) using either the learned model (simulator) or the real system
- ▶ Deterministic inference—only with a learned model

Stochastic Inference

- ▶ Sample trajectories (s_i, a_i, r_i) ►► Monte Carlo
- ▶ Conceptually very simple
- ▶ Requires a lot of “interactions” (if you don’t have a model or a good simulator)
 - Potentially impractical (e.g., in robotics)

Deterministic Inference



- ▶ Analytically propagate uncertainty through the model
- ▶ Computationally/mathematically more involved
- ▶ Can't do this for arbitrary systems, but for some.

Deterministic Inference: Example

Linear system

$$p(\mathbf{s}_t) = \mathcal{N}(\mathbf{s}_t \mid \mathbf{m}_t, \mathbf{S}_t)$$

$$\mathbf{s}_{t+1} = \mathbf{A}\mathbf{s}_t$$

Successor state distribution $p(\mathbf{s}_{t+1})$?

Deterministic Inference: Example

Linear system

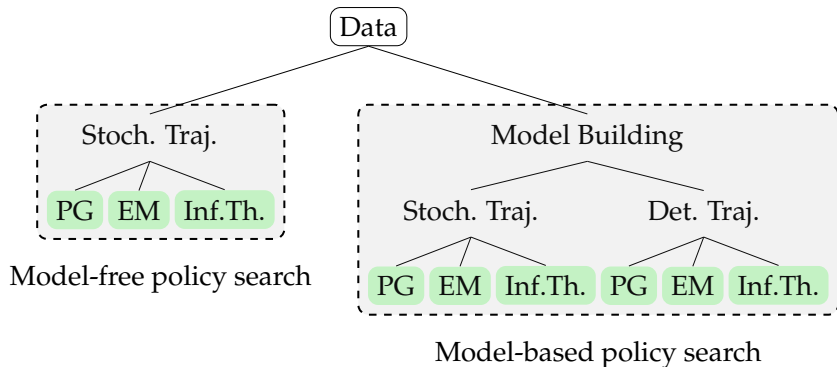
$$p(\mathbf{s}_t) = \mathcal{N}(\mathbf{s}_t \mid \mathbf{m}_t, \mathbf{S}_t)$$
$$\mathbf{s}_{t+1} = \mathbf{A}\mathbf{s}_t$$

Successor state distribution $p(\mathbf{s}_{t+1})$?

$$p(\mathbf{s}_{t+1}) = \mathcal{N}(\mathbf{s}_{t+1} \mid \mathbf{m}_{t+1}, \mathbf{S}_{t+1})$$
$$\mathbf{m}_{t+1} = \mathbf{A}\mathbf{m}_t, \quad \mathbf{S}_{t+1} = \mathbf{A}\mathbf{S}_t\mathbf{A}^\top$$

- In nonlinear/non-Gaussian systems, we need approximations (e.g., linearization, moment matching)

Policy Improvement



Policy improvement (green)

- [Policy gradients](#)
- Expectation Maximization
- Information theory

Policy Search: Policy Improvement

Objective

Find policy parameters θ^* , which maximize the expected long-term reward

$$V^\pi(\theta) = \mathbb{E}\left[\sum_{k=0}^T \gamma^k r_{k+1} | \theta\right], \quad s_0 \sim p(s_0)$$

- ▶ No global value function model V^π or Q^π
- ▶ Search directly in (policy) parameter space

Policy Search: Policy Improvement

Objective

Find policy parameters θ^* , which maximize the expected long-term reward

$$V^\pi(\theta) = \mathbb{E}\left[\sum_{k=0}^T \gamma^k r_{k+1} | \theta\right], \quad s_0 \sim p(s_0)$$

- No global value function model V^π or Q^π
- Search directly in (policy) parameter space
- ▶ One way: gradient-based optimization
 - Compute V^π with corresponding gradients $dV^\pi/d\theta$
 - Gradient-based optimizer for maximization (e.g., CG, BFGS)

Gradient Estimation for Stochastic Inference

Gradient Estimation for Stochastic Inference

- Finite (central) differences

$$\frac{dV^\pi(\boldsymbol{\theta})}{d\boldsymbol{\theta}} \approx \frac{V^\pi(\boldsymbol{\theta} + \epsilon) - V^\pi(\boldsymbol{\theta} - \epsilon)}{2\epsilon}$$

- Model $p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k)$ not required but useful
- **Large variance** of estimator ► many samples needed

Gradient Estimation for Stochastic Inference

- Finite (central) differences

$$\frac{dV^\pi(\boldsymbol{\theta})}{d\boldsymbol{\theta}} \approx \frac{V^\pi(\boldsymbol{\theta} + \epsilon) - V^\pi(\boldsymbol{\theta} - \epsilon)}{2\epsilon}$$

- Model $p(s_{k+1}|s_k, a_k)$ not required but useful
- **Large variance** of estimator ► many samples needed
- PEGASUS trick:
 - Fix the random seed and re-set
 - Smaller variance of the estimate of V^π and its gradient
 - No model $p(s_{k+1}|s_k, a_k)$ required
 - Only with simulator (where we can run exactly the same experiment)

Gradient Estimation for Deterministic Inference

- ▶ Finite differences and PEGASUS still work
- ▶ Analytic (=exact) gradients. Example (assume $r = r(\mathbf{s})$):

$$\begin{aligned}\frac{dV(\boldsymbol{\theta})}{d\boldsymbol{\theta}} &= \sum_t \gamma^t \frac{dr(\mathbf{s}_t)}{d\boldsymbol{\theta}} = \sum_t \gamma^t \frac{\partial r(\mathbf{s}_t)}{\partial \mathbf{s}_t} \frac{d\mathbf{s}_t}{d\boldsymbol{\theta}} \\ &= \sum_t \gamma^t \frac{\partial r(\mathbf{s}_t)}{\partial \mathbf{s}_t} \left(\frac{\partial \mathbf{s}_t}{\partial \mathbf{s}_{t-1}} \frac{d\mathbf{s}_{t-1}}{d\boldsymbol{\theta}} + \frac{\partial \mathbf{s}_t}{\partial \mathbf{a}_{t-1}} \frac{\partial \mathbf{a}_{t-1}}{\partial \boldsymbol{\theta}} \right)\end{aligned}$$

Gradient Estimation for Deterministic Inference

- ▶ Finite differences and PEGASUS still work
- ▶ Analytic (=exact) gradients. Example (assume $r = r(s)$):

$$\begin{aligned}\frac{dV(\boldsymbol{\theta})}{d\boldsymbol{\theta}} &= \sum_t \gamma^t \frac{dr(s_t)}{d\boldsymbol{\theta}} = \sum_t \gamma^t \frac{\partial r(s_t)}{\partial s_t} \frac{ds_t}{d\boldsymbol{\theta}} \\ &= \sum_t \gamma^t \frac{\partial r(s_t)}{\partial s_t} \left(\frac{\partial s_t}{\partial s_{t-1}} \frac{ds_{t-1}}{d\boldsymbol{\theta}} + \frac{\partial s_t}{\partial a_{t-1}} \frac{\partial a_{t-1}}{\partial \boldsymbol{\theta}} \right)\end{aligned}$$

- ▶ Requires
 - ▶ Forward model $s_t = f(s_{t-1}, a_{t-1})$
 - ▶ Differentiable policy $a = \pi(s, \boldsymbol{\theta})$

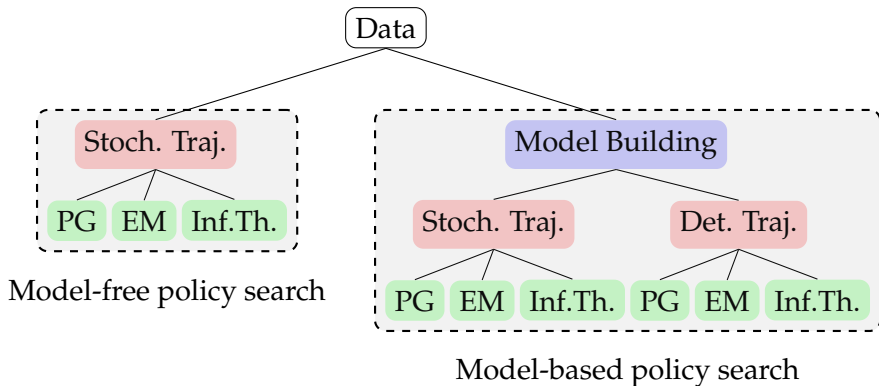
Gradient Estimation for Deterministic Inference

- ▶ Finite differences and PEGASUS still work
- ▶ Analytic (=exact) gradients. Example (assume $r = r(s)$):


$$\begin{aligned}\frac{dV(\boldsymbol{\theta})}{d\boldsymbol{\theta}} &= \sum_t \gamma^t \frac{dr(s_t)}{d\boldsymbol{\theta}} = \sum_t \gamma^t \frac{\partial r(s_t)}{\partial s_t} \frac{ds_t}{d\boldsymbol{\theta}} \\ &= \sum_t \gamma^t \frac{\partial r(s_t)}{\partial s_t} \left(\frac{\partial s_t}{\partial a_{t-1}} \frac{da_{t-1}}{d\boldsymbol{\theta}} + \frac{\partial s_t}{\partial s_{t-1}} \frac{ds_{t-1}}{d\boldsymbol{\theta}} \right)\end{aligned}$$

- ▶ Requires
 - ▶ Forward model $s_t = f(s_{t-1}, a_{t-1})$
 - ▶ Differentiable policy $a = \pi(s, \boldsymbol{\theta})$
- ▶ Mathematically more involved
- ▶ Gradients are exact (no variance): single trajectory evaluation

Policy Search



Applications in Robotics and Control



- ▶ Cart-pole: e.g., Riedmiller (2005), Deisenroth & Rasmussen (2011)
- ▶ Throttle valve control: Bischoff et al. (2013)
- ▶ Autonomous helicopter: e.g., Abbeel, Ng et al. (2003–2010), Bagnell & Schneider (2001)
- ▶ Pancake flipping (Kormushev et al., 2010)
- ▶ Throwing and catching balls (Kober et al., 2012)

Summary







- ▶ RL is a principled framework for sequential decision making under uncertainty
- ▶ Value functions V, Q
- ▶ Exact RL: Dynamic programming
- ▶ Approximate RL: Monte Carlo, TD
- ▶ Policy Search with applications in robotics

`m.deisenroth@imperial.ac.uk`

Thank you for your attention

Key References



- ▶ Sutton, Barto: *Reinforcement Learning: An Introduction* (online)
- ▶ Bertsekas: *Dynamic Programming and Optimal Control*, Vol. 1–2
- ▶ Szepesvári: *Algorithms for Reinforcement Learning* (online)
- ▶ Deisenroth et al.: *A Survey on Policy Search for Robotics* (online)

RL Software Packages

- ▶ RLGlue: <http://glue.rl-community.org/>
- ▶ RLPy: <http://acl.mit.edu/RLPy/>
- ▶ CLSquare: <http://www.ni.uos.de/index.php?id=70>
- ▶ PIQLE: <http://piqle.sourceforge.net/>
- ▶ RL Toolbox: <http://www.igi.tugraz.at/ril-toolbox/>
- ▶ LibPG: <http://code.google.com/p/libpgrl/>
- ▶ PILCO (policy search): <http://mlg.eng.cam.ac.uk/pilco>

References

- [1] P. Abbeel, A. Coates, M. Quigley, and A. Y. Ng. An Application of Reinforcement Learning to Aerobatic Helicopter Flight. In B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, volume 19, page 2007. The MIT Press, Cambridge, MA, USA, 2007.
- [2] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1 of *Optimization and Computation Series*. Athena Scientific, Belmont, MA, USA, 3rd edition, 2005.
- [3] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2 of *Optimization and Computation Series*. Athena Scientific, Belmont, MA, USA, 3rd edition, 2007.
- [4] B. Bischoff, D. Nguyen-Tuong, T. Koller, H. Markert, and A. Knoll. Learning Throttle Valve Control Using Policy Search. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases*, 2013.
- [5] M. P. Deisenroth, G. Neumann, and J. Peters. *A Survey on Policy Search for Robotics*, volume 2 of *Foundations and Trends in Robotics*. NOW Publishers, 2013.
- [6] M. P. Deisenroth and C. E. Rasmussen. PILCO: A Model-Based and Data-Efficient Approach to Policy Search. In *Proceedings of the International Conference on Machine Learning*, pages 465–472, New York, NY, USA, June 2011. ACM.
- [7] S. Gelly, M. Schoenauer, M. Sebag, O. Teytaud, L. Kocsis, D. Silver, and C. Szepesvári. The Grand Challenge of Computer Go: Monte Carlo Tree Search and Extensions. *Communications of the ACM*, 55(3), 2012.
- [8] Q. Gemine, E. Karangelos, D. Ernst, and B. Cornélusse. Active Network Management: Planning under Uncertainty for Exploiting Load Modulation. In *Proceedings of the 2013 IREP Symposium-Bulk Power System Dynamics and Control*, 2013.
- [9] J. Kober, K. Mülling, and J. Peters. Learning Throwing and Catching Skills. In *Proceedings of the IEEE/RSJ International Conference on Robot Systems*, 2012.
- [10] P. Kormushev, S. Calinon, and D. G. Caldwell. Robot Motor Skill Coordination with EM-based Reinforcement Learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010.
- [11] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, volume 703. John Wiley & Sons, 2007.
- [12] M. Riedmiller. Neural Fitted Q Iteration—First Experiences with a Data Efficient Neural Reinforcement Learning Method. In *Proceedings of the 16th European Conference on Machine Learning*, Porto, Portugal, 2005.
- [13] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning. The MIT Press, Cambridge, MA, USA, 1998.
- [14] C. Szepesvári. *Algorithms for Reinforcement Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2010.
- [15] G. J. Tesauro. Td-gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2):215–219, 1994.

Appendix

Self-Consistency of Value Functions

$$V^{\pi}(\mathbf{s}) = \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi]$$

Self-Consistency of Value Functions

$$V^{\pi}(\mathbf{s}) = \mathbb{E}[\textcolor{red}{R} | \mathbf{s}_0 = \mathbf{s}, \pi]$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[\textcolor{red}{R} | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \textcolor{red}{\gamma}^k \textcolor{red}{r}_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) \left(\mathbb{E}[r_1] + \gamma \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_1 = \mathbf{s}', \pi\right] \right) \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) \left(\mathbb{E}[r_1] + \gamma \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_1 = \mathbf{s}', \pi\right] \right) \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned}V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\&= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\&= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\&= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) \left(\mathbb{E}[r_1] + \gamma \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_1 = \mathbf{s}', \pi\right] \right) \\&= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) (\mathbb{E}[r_1(\mathbf{s}, a, \mathbf{s}')] + \gamma V^\pi(\mathbf{s}'))\end{aligned}$$

Self-Consistency of Value Functions


$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) \left(\mathbb{E}[r_1] + \gamma \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_1 = \mathbf{s}', \pi\right] \right) \\ &= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) (\mathbb{E}[r_1(\mathbf{s}, a, \mathbf{s}')] + \gamma V^\pi(\mathbf{s}')) \end{aligned}$$

Self-Consistency of Value Functions

$$\begin{aligned} V^\pi(\mathbf{s}) &= \mathbb{E}[R | \mathbf{s}_0 = \mathbf{s}, \pi] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+1} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \mathbb{E}\left[r_1 + \gamma \sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_0 = \mathbf{s}, \pi\right] \\ &= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) \left(\mathbb{E}[r_1] + \gamma \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{k+2} | \mathbf{s}_1 = \mathbf{s}', \pi\right] \right) \\ &= \sum_a p_\pi(a | \mathbf{s}) \sum_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{s}, a) (\mathbb{E}[r_1(\mathbf{s}, a, \mathbf{s}')] + \gamma V^\pi(\mathbf{s}')) \\ &= \mathbb{E}[r_1 + \gamma V^\pi(\mathbf{s}_1) | \mathbf{s}_0 = \mathbf{s}, \pi] \end{aligned}$$


Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Policy Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem




Demo: Value Iteration

Example: Shortest-path problem



Demo: Value Iteration

Example: Shortest-path problem



Reinforcement Learning

Marc Deisenroth

