

Planning with a Known Model

(INF8250AE: Introduction to Reinforcement Learning)

Amir-massoud Farahmand

Polytechnique Montréal & Mila

POLYTECHNIQUE
MONTRÉAL

UNIVERSITÉ
D'INGÉNIERIE



Adaptive Agents Lab
(Adage)



Table of Contents

- 1 Some Initial Attempts
- 2 Value Iteration
- 3 Policy Iteration
 - Convergence of Policy Iteration
- 4 Linear Programming

Goal

We study some algorithms to solve the Planning problem: computing the (optimal) value function for a given MDP.

- Know the fundamental planning algorithms:
 - ▶ Value Iteration (VI): $V_{k+1} \leftarrow TV_k$.
 - ▶ Policy Iteration (PI): (1) Compute V^{π_k} ; (2) Compute $\pi_{k+1} \leftarrow \pi_g(V^{\pi_k})$
 - ▶ Linear Programming
- Learn the intuition behind them
- Learn the mathematics justifying them

We refer to these frequently in studying and analyzing RL/Planning algorithms.

Learning Objectives


You need to

- Remember: Value Iteration, Policy Iteration, and LP-based algorithms
- Understand: Why VI, PI, LP work
- Apply: VI and PI to solve an MDP

How to Compute the Optimal Policy π^* ?

- **Question:** How can we find the optimal policy?
- **Assumption:** MDP is known, that is, we know $\mathcal{R}$ and $\mathcal{P}$.
- The assumption of knowing the MDP does not hold in the RL setting.
- But designing methods for finding the optimal policy with known model provides the foundation for developing methods for the RL setting.
- We use concepts and properties such as
 - ▶ Value function for a policy π (V^π) and optimal value function (V^*)
 - ▶ Relation between V^* (or Q^*) and π^* through the greedy policy
 - ▶ Contraction and Monotonicity

Different Approaches to Find π^*



- **Value-based:** Compute Q^* (or V^*) and then $\pi^* \leftarrow \pi_g(Q^*)$.
- **Direct policy search:** Search in the space of policies without explicitly constructing the optimal value function.
- **Hybrid:** Explicitly construct value function to guide the search in the policy space.

Policy Evaluation vs. Control Problems

- **Policy Evaluation (PE)**: Problem of computing the value function of a given policy π , i.e., V^π or Q^π .
 - ▶ Not the ultimate goal of an RL agent (finding the optimal policy is), but is often needed as an intermediate step in finding the optimal policy.
- **Control**: Problem of finding the optimal value function V^* or Q^* or optimal policy π^* .


Dynamic Programming (DP): Methods that benefit from the structure of the MDP, such as the recursive structure encoded in the Bellman equation, in order to compute the value function.

Policy Evaluation

Problem Statement: Given an MDP $(\mathcal{X}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ and a policy π , we would like to compute V^π or Q^π .

$$V^\pi(x) = \mathbb{E} \left[\sum_{t=1}^{\infty} \gamma^{t-1} R_t \mid X_1 = x \right].$$

Policy Evaluation: A Naive Approach



Idea: Expand the tree of all possible futures!

Example: the expected reward at time $t = 2$ is

$$\sum_a \pi(a|x) r(x,a) + \gamma \sum_{a, x', a'} \pi(a|x) \mathcal{P}(x'|x, a) \pi(a'|x') r(x', a') + \gamma^2 \sum \dots$$

Remark

- This is inefficient. The size of the tree grows very fast.
- The sample-based approximate version of this idea works.

Policy Evaluation: Linear System of Equations

Q: Can we improve the efficiency?

Key Idea: Benefit from the recursive structure of the value function

$$Q^\pi(x, a) = r(x, a) + \dots \quad V^\pi = T^\pi V^\pi. \quad \forall (x, a) \in \mathcal{X} \times \mathcal{A}$$

$$V(x) = r^\pi(x) + \gamma \sum_{x' \in \mathcal{X}} \mathcal{P}^\pi(x'|x) V(x'), \quad \forall x \in \mathcal{X}$$

In the discrete state-action case:

- $n = |\mathcal{X}|$ equations
- $|\mathcal{X}|$ unknowns $(V(x_1), \dots, V(x_n))$

$$V = T^\pi V$$

$$V \neq V^\pi \quad ? \quad \text{X}$$

Policy Evaluation: Linear System of Equations

We have n equations in the form of:

$$V(x) - \gamma \sum_{x' \in \mathcal{X}} \mathcal{P}^\pi(x'|x) \underline{V(x')} = \underline{r^\pi(x)},$$

More compactly in the matrix form:

$$\rightarrow \underline{(\mathbf{I} - \gamma \mathcal{P}^\pi)} \underline{V} = \underline{r^\pi},$$

which is the same form of a generic linear system of equations:

$$\underline{A_{n \times n}} \underline{x_{n \times 1}} = \underline{b_{n \times 1}}.$$

Policy Evaluation: Linear System of Equations

Solving $A_{n \times n} x_{n \times 1} = b_{n \times 1}$:

- Compute A^{-1} and then calculate $A^{-1}b$.
- Better: Use various linear solvers.

Remark

To solve the control problem of finding V^* , we need to solve $V = T^*V$,
i.e.,

$$V(x) = \max_{a \in A} \left\{ r(x, a) + \gamma \sum_{x' \in \mathcal{X}} \mathcal{P}(x'|x, a) V(x') \right\}.$$

Handwritten notes: A blue arrow points from the V in the equation to the V^* in the text below. Another blue arrow points from the V in the equation to the $V = V^*$ below. A blue arrow points from the V in the equation to the V^* in the text below.

This is *not* a linear system of equations anymore!

Value Iteration

Value Iteration (PE)

Starting from $\underline{V_0} \in \mathcal{B}(\mathcal{X})$, we compute a sequence of $(\underline{V_k})_{k \geq 0}$ by

$$\underline{V_{k+1}} \leftarrow T^\pi \underline{V_k} [= \underline{r^\pi + \gamma \mathcal{P}^\pi V_k}].$$

By the contraction property of the Bellman operator:

$$\lim_{k \rightarrow \infty} \|\underline{V_k} - \underline{V^\pi}\|_\infty = \underline{0}.$$

Remark


Similar procedure to compute Q^π , i.e., $\underline{Q_{k+1}} \leftarrow T^\pi \underline{Q_k}$.

Value Iteration (Control)


$$\begin{aligned}V_{k+1} &\leftarrow T^*V_k, \\Q_{k+1} &\leftarrow T^*Q_k.\end{aligned}$$

By the contraction property of the Bellman optimality operator, it is guaranteed that $V_k \rightarrow V^*$ (or $Q_k \rightarrow Q^*$).


Value Iteration



Value Iteration



Value Iteration



Convergence of Value Iteration

RL Acceleration

We can quantify the converge rate of the VI algorithm

Theorem (Convergence of Value Iteration)

The VI algorithm generates a sequence of (V_k) that converges to $\frac{V^\pi}{G^\pi}$ (for PE) or V^* (for Control) with the rate of $O(\gamma^k)$.

G^*

$$\|V_k - V^*\| \leq O(\gamma^k)$$

$$\begin{aligned} \gamma = 0.9 &\rightarrow 0.9^{10} = 0.34 \\ \gamma = 0.99 &\rightarrow 0.99^{10} = 0.90 \end{aligned}$$

$$\begin{aligned} 0.9^{100} &= 0.000026 \\ 0.99^{100} &= 0.36 \\ 0.99^{1000} &= 0.000043 \end{aligned}$$

Convergence of Value Iteration (Proof)

We prove it for the Control problem. The proof for the PE problem is very similar.

Since $V^* = T^*V^*$ and $V_k = T^*V_{k-1}$, we have

$$V_k - V^* = \underline{V_k} - \underline{T^*V^*} = \underline{T^*V_{k-1}} - T^*V^*.$$

We take the supremum norm from both sides, and use the contraction property of the Bellman operator:

$$\|V_k - V^*\|_{\underline{\infty}} = \|T^*V_{k-1} - T^*V^*\|_{\underline{\infty}} \leq \underline{\gamma} \|V_{k-1} - V^*\|_{\underline{\infty}}.$$

Therefore, we get that

$$\begin{aligned} \|V_k - V^*\|_{\infty} &\leq \underline{\gamma} \|V_{k-1} - V^*\|_{\infty} \leq \underline{\gamma}^2 \|V_{k-2} - V^*\|_{\infty} \leq \dots \\ &\leq \underline{\gamma}^k \|V_0 - V^*\|_{\infty}. \end{aligned}$$

$O(\gamma^k)$

Value Iteration


- VI uses an existing approximation V_k of V^* (or V^π) to get a better approximation of V^* (or V^π). This idea is called bootstrapping in the RL literature.
 - ▶ Note: This is different from the bootstrapping method in statistics.
- VI is one of the fundamental algorithms for Planning.
- The current formulation of VI is a special case of an operator/matrix splitting-based formulation.
- We also have approximate value iteration algorithms too, where the new value function $V_{k+1} \approx TV_k$.
- Many RL algorithms are essentially the sample-based variants of (approximate) VI too.

Policy Iteration

Policy Iteration

A different approach is based on the iterative application of the following two steps:

- (Policy Evaluation) Given a policy π_k , compute V^{π_k} (or Q^{π_k}).
- (Policy Improvement) Find a new policy π_{k+1} that is better than π_k , i.e., $V^{\pi_{k+1}} \geq V^{\pi_k}$ (with a strict inequality in at least one state, unless at convergence).



Policy Iteration



Q: How to perform Policy Evaluation and Policy Improvement?

- Policy Evaluation: This is clear. We can either solve a linear system of equations or perform VI (PE) to compute the value of a policy π_k .
- Policy Improvement: Choose the greedy policy, i.e.,

$$\underline{\pi_{k+1}}(x) \leftarrow \underline{\pi_g}(x; \underline{Q^{\pi_k}}) = \underset{a \in \mathcal{A}}{\operatorname{argmax}} Q^{\pi_k}(x, a), \quad \forall x \in \mathcal{X}.$$

Remark

The Policy Iteration (PI) algorithm refers to the specific case that we pick the new policy π_{k+1} as $\pi_g(Q^{\pi_k})$.

Why Greedy Policy for Policy Improvement? (Intuition)

Assume that at state x , we act according to $\pi_g(x; Q^{\pi_k})$, and afterwards, we follow π_k .

The value of this new policy is

$$Q^{\pi_k}(x, \pi_g(x; Q^{\pi_k})) = Q^{\pi_k}(x, \operatorname{argmax}_{a \in \mathcal{A}} Q^{\pi_k}(x, a)) = \max_{a \in \mathcal{A}} Q^{\pi_k}(x, a).$$

Comparing $\max_{a \in \mathcal{A}} Q^{\pi_k}(x, a)$ with $V^{\pi_k}(x) = Q^{\pi_k}(x, \pi_k(x))$, we see

$$Q^{\pi_k}(x, \pi_g(x; Q^{\pi_k})) = \max_{a \in \mathcal{A}} Q^{\pi_k}(x, a) \geq V^{\pi_k}(x).$$

So this new policy is equal to or better than π_k at state x .

Policy Iteration

Recall that:

- V^{π_k} is the unique fixed point of T^{π_k} .
- The greedy policy $\pi_{k+1} = \pi_g(Q^{\pi_k})$ satisfies $T^{\pi_{k+1}}Q^{\pi_k} = T^*Q^{\pi_k}$.

We can summarize each iteration of the Policy Iteration algorithm:

- (Policy Evaluation) Given π_k , compute Q^{π_k} , i.e., find a Q that satisfies $Q = T^{\pi_k}Q$.
- (Policy Improvement) Obtain π_{k+1} as a policy that satisfies $T^{\pi_{k+1}}Q^{\pi_k} = T^*Q^{\pi_k}$.

Policy Iteration

- 1: Initialize π_0 arbitrarily
- 2: $k \leftarrow 0$
- 3: **repeat**
- 4: $\underline{Q^{\pi_k}} \leftarrow \text{solution of } Q = T^{\pi_k} Q$ $\triangleright$ Policy Evaluation: compute Q^{π_k}
- 5: $\underline{\pi_{k+1}} \leftarrow \underline{\pi_g(Q^{\pi_k})}$ $\triangleright$ Policy Improvement
- 6: $k \leftarrow k + 1$
- 7: **until** $\underline{\pi_k} = \underline{\pi_{k-1}}$

Approximate Policy Iteration

We also have **approximate** policy iteration algorithms too, where policy evaluation or improvement steps are performed approximately:

- $Q \approx T^{\pi_k} Q$
- $T^{\pi_{k+1}} Q^{\pi_k} \approx T^* Q^{\pi_k}$

We discuss this later when we get to function approximation.

$$Q^*(x, \cdot) = [1, \underbrace{2}_{a_2}, -1, \underbrace{2}_{a_4}, 0]$$

$$\pi_g(x; G^*) = \underset{a}{\operatorname{argmax}} \underbrace{Q^*(x, a)}_{a_2}$$

Convergence of Policy Iteration

- The Policy Iteration algorithm converges to the optimal policy.
- For finite MDPs, the convergence happens in a finite number of iterations.

Policy Improvement Theorem

Theorem (Policy Improvement)

If for policies π and π' , it holds that $T^{\pi'} Q^\pi = T^* Q^\pi$, we have that $\underline{Q^{\pi'} \geq Q^\pi}$.

In other words, the greedy policy is a proper policy improvement step.

Policy Improvement Theorem (Proof)

We first show that $T^{\pi'} Q^{\pi} \geq Q^{\pi}$. Notice that $T^{\pi'} Q^{\pi} = T^* Q^{\pi}$, by the assumption.

We have $T^* Q^{\pi} \geq T^{\pi} Q^{\pi} = Q^{\pi}$ because for any $(x, a) \in \mathcal{X} \times \mathcal{A}$, it holds that

$$\begin{aligned} r(x, a) + \gamma \int \mathcal{P}(\mathrm{d}x' | x, a) \max_{a' \in \mathcal{A}} Q^{\pi}(x', a') &\geq \\ r(x, a) + \gamma \int \mathcal{P}(\mathrm{d}x' | x, a) Q^{\pi}(x', \pi(x')) &= T^{\pi} Q^{\pi} \end{aligned}$$

Therefore, $T^{\pi'} Q^{\pi} = T^* Q^{\pi} \geq T^{\pi} Q^{\pi} = Q^{\pi}$.

Policy Improvement Theorem (Proof)

The second step is to use $T^{\pi'} Q^{\pi} \geq Q^{\pi}$ to show that $Q^{\pi'} \geq Q^{\pi}$.
 Apply $T^{\pi'}$ to both sides of $T^{\pi'} Q^{\pi} \geq Q^{\pi}$, and use the monotonicity property of the Bellman operator to conclude

$$T^{\pi'}(T^{\pi'} Q^{\pi}) \geq T^{\pi'} Q^{\pi} = T^* Q^{\pi} \geq Q^{\pi}$$

So we also have $(T^{\pi'})^2 Q^{\pi} \geq Q^{\pi}$.

By repeating this argument, we get that for any $m \geq 1$,

$$(T^{\pi'})^m Q^{\pi} \geq T^* Q^{\pi} \geq Q^{\pi} \tag{1}$$

Policy Improvement Theorem (Proof)

$$Q_\pi \rightarrow Q_{k+1} = T^\pi Q_k \Rightarrow Q_k \rightarrow Q^{\pi'}$$

$$(T^{\pi'})^m Q^\pi \geq T^* Q^\pi \geq \underline{Q^\pi}.$$

Take the limit of $m \rightarrow \infty$.

Because of the contraction property of the Bellman operator $T^{\pi'}$:

$$\lim_{m \rightarrow \infty} (T^{\pi'})^m \underline{Q^\pi} = \underline{Q^{\pi'}}. \quad (2)$$

By combining (1) and (2), we get that

$$\underline{Q^{\pi'}} = \lim_{m \rightarrow \infty} (T^{\pi'})^m \underline{Q^\pi} \geq \underline{T^* Q^\pi} \geq \underline{Q^\pi}, \quad (3)$$

$$Q^{\pi'} \geq Q^\pi$$

Convergence of Policy Iteration

- The Policy Improvement theorem shows that if we are given π_k , the new policy π_{k+1} is at least as good as the previous one.
- We can show that the PI algorithm converges to an optimal policy. We shall prove this.
- If $|\mathcal{X} \times \mathcal{A}| < \infty$, this happens in a finite number of iterations.

Convergence of Policy Iteration

Theorem (Convergence of the Policy Iteration Algorithm)

Let $(\pi_k)_{k \geq 0}$ be the sequence generated by the PI algorithm.

• For all k , we have that $V^{\pi_{k+1}} \geq V^{\pi_k}$, with equality if and only if $V^{\pi_k} = V^*$.

$$V^{\pi_{k+1}} = V^{\pi_k} \Rightarrow V^{\pi_k} = V^*$$

• $\lim_{k \rightarrow \infty} \|V^{\pi_k} - V^*\|_{\infty} = 0$.

• If the set of policies is finite, the PI algorithm converges in a finite number of iterations.

Remark

We follow the line of proof of Proposition 2.4.1 of Bertsekas 2018.

Convergence of Policy Iteration (Proof)

The basic idea behind the proof is that either we can strictly improve the policy, or if we cannot, we are already at the optimal policy.

Convergence of Policy Iteration (Proof)

Proof of $V^{\pi_{k+1}} \geq V^{\pi_k}$:

By the Policy Improvement Theorem (Theorem 2), we have that $V^{\pi_{k+1}} \geq V^{\pi_k}$.

Proof of $V^{\pi_{k+1}} = V^{\pi_k} \Rightarrow V^{\pi_k} = V^*$:

Suppose that instead of a strict inequality, we have an equality of

$$V^{\pi_{k+1}} = V^{\pi_k}.$$

Apply $T^{\pi_{k+1}}$ to both side to get

$$T^{\pi_{k+1}} V^{\pi_{k+1}} = T^{\pi_{k+1}} V^{\pi_k}.$$

As $T^{\pi_{k+1}} V^{\pi_k} = T^* V^{\pi_k}$ by the definition of the PI algorithm, we get that

$$T^{\pi_{k+1}} V^{\pi_{k+1}} = T^{\pi_{k+1}} V^{\pi_k} = T^* V^{\pi_k} = T^* V^{\pi_{k+1}},$$

where in the last step we used $V^{\pi_{k+1}} = V^{\pi_k}$ again.

Convergence of Policy Iteration (Proof)

By these equalities, we have

$$\Rightarrow T^{\pi_{k+1}} V^{\pi_{k+1}} = T^* V^{\pi_{k+1}}.$$

As $V^{\pi_{k+1}}$ is the value function of π_{k+1} , it satisfies the Bellman equation $T^{\pi_{k+1}} V^{\pi_{k+1}} = V^{\pi_{k+1}}$. Therefore, we also have

$$\underline{V^{\pi_{k+1}}} = \underline{T^*} V^{\pi_{k+1}}.$$

This means that $V^{\pi_{k+1}}$ is a fixed point of T^* .

But the fixed point of T^* is unique and is equal to V^* .

So we must have that

$$\underline{V^{\pi_{k+1}}} = \underline{V^*}.$$

Convergence of Policy Iteration (Proof)

Proof of $V^{\pi_k} = V^* \Rightarrow V^{\pi_{k+1}} = V^{\pi_k}$:

If $V^{\pi_k} = V^*$, then π_k is an optimal policy. The greedy policy of $V^{\pi_k} = V^*$ is still an optimal policy, hence $V^{\pi_{k+1}} = V^* = V^{\pi_k}$.



Convergence of Policy Iteration (Proof)

Proof of $\lim_{k \rightarrow \infty} \|V^{\pi_k} - V^*\|_{\infty} = 0$.

To prove the convergence, recall from (3) that

$$\underline{Q^{\pi_{k+1}}} \geq T^* \underline{Q^{\pi_k}} \geq Q^{\pi_k}.$$

By induction,

$$Q^{\pi_{k+1}} \geq T^* Q^{\pi_k} \geq T^* (T^* Q^{\pi_{k-1}}) \geq \dots \geq (T^*)^k Q^{\pi_0}.$$

By the definition of the optimal policy, we have $\underline{Q^{\pi}} \leq \underline{Q^*}$ for any π , including all π_k generated during the iterations of the PI algorithm.

So $\underline{Q^{\pi_{k+1}}}$ is sandwiched between Q^* and $(T^*)^k Q^{\pi_0}$, i.e.,

$$\underline{Q^*} \geq \underline{Q^{\pi_{k+1}}} \geq \underline{(T^*)^k Q^{\pi_0}}.$$

This entails that $\underline{Q^{\pi_{k+1}}} - Q^* \leq \underline{(T^*)^k Q^{\pi_0} - Q^*}$.


$$G^{\pi_{k+1}} > G^{\pi_k}$$

$$\Pi = \{\pi^{(1)}, \pi^{(2)}, \pi^{(3)}, \pi^{(4)}\}$$

$\uparrow \quad \uparrow \quad \uparrow \quad \uparrow$
 $\pi_3 \quad \pi_1 \quad \pi_2 \quad \pi_4$

(4)

$G^{\pi_2} > G^{\pi_1}$



Convergence of Policy Iteration (Proof)

We show that the RHS of $\|Q^{\pi_{k+1}} - Q^*\|_\infty \leq \|(T^*)^k Q^{\pi_0} - Q^*\|_\infty$ converges to zero. 

By the contraction property of the Bellman optimality operator, we have that

$$\lim_{k \rightarrow \infty} \|(T^*)^k Q^{\pi_0} - Q^*\|_\infty = 0.$$

Therefore,

$$\lim_{k \rightarrow \infty} \|Q^{\pi_k} - Q^*\|_\infty = 0.$$

This implies the convergence of V^{π_k} too.

Convergence of Policy Iteration (Proof)

Proof of finite convergence:

If the number of policies is finite, the number of times (4) can be a strict inequality is going to be finite too.

Convergence of Policy Iteration: Some Remarks

- The PI algorithm converges to the optimal policy in a finite number of iterations whenever the number of policies is finite.
- If the state space $\mathcal{X}$ and the action space $\mathcal{A}$ are finite, the number of deterministic policies are finite and is $|\mathcal{A}|^{|\mathcal{X}|}$.
- Even though finite, this can be very large.
 - ▶ Example: A 10×10 grid world problem with 4 actions at each state has $4^{100} \approx 1.6 \times 10^{60}$ possible policies.
- In practice, PI converges much faster.
- This suggest that the previous analysis might be crude.

Fast Convergence of Policy Iteration

It can be shown that the PI algorithm converges in

$$O\left(\frac{|\mathcal{X}||\mathcal{A}|}{1-\gamma} \log\left(\frac{1}{1-\gamma}\right)\right)$$

$$\frac{100 \times 4}{1-0.9} = 4000$$

iterations.

The proof is in the Foundations of Reinforcement Learning [Farahmand, 2025].

This is a significant quantitative improvement over the previous result.

Remark

This is a relatively recent result, which in various forms have been proven by Ye [2011]; Hansen et al. [2013]; Scherrer [2016].

$$\min_x f(x) \rightarrow \min_x f(x) \rightarrow \left| \begin{array}{l} \min c^T x \\ \bar{a}_i^T x + b_i \leq 0 \quad i=1, \dots \end{array} \right.$$

$g_i(x) \leq 0$

$\uparrow$

Linear Programming

Linear Programming for Finding V^*

We can find V^* by solving a Linear Program (LP) too.
Consider the set of all V that satisfy $V \geq T^*V$, i.e.,

$$\rightarrow C = \{ V : V \geq T^*V \}.$$

Interesting property: For any $V \in C$, we have

$$V \geq T^*V \Rightarrow T^*V \geq T^*(T^*V) = (T^*)^2V.$$

Repeating this argument, we get that for any $m \geq 1$,

$$V \geq (T^*)^m V.$$

Linear Programming for Finding V^*

From

$$\rightarrow V \geq (T^*)^m V.$$

we get that

$$\underline{V} \geq \lim_{m \rightarrow \infty} (T^*)^m \underline{V} = \underline{V^*}.$$

Also note that $\underline{V^*} \in C$ because $\underline{V^*} \geq T^* \underline{V^*}$ (and in fact, $\underline{V^*} = T^* \underline{V^*}$).

Interpretation:

- • Any $\underline{V} \in C$ is lower bounded by $\underline{V^*}$.
- $\underline{V^*}$ is the function in C that is smaller or equal to any other function in C (pointwise sense).

Linear Programming for Finding V^*

Choose a strictly positive vector $\mu > 0$ with the dimension of $\mathcal{X}$.
Solve

$$\min_{V \in C} \mu^\top V,$$

Can be written as

$$\begin{aligned} \min_V \quad & \underline{\mu^\top V}, \\ \text{s.t.} \quad & V(x) \geq (T^*V)(\underline{x}), \quad \forall x \in \underline{\mathcal{X}}. \end{aligned}$$

Linear objective; nonlinear constraints.

Linear Programming for Finding V^*

$$V(x) \geq \max \left\{ r(x, a_1) + \dots, r(x, a_2) + \dots \right\}$$

Each nonlinear constraint:

$$V(x) \geq \max_{a \in \mathcal{A}} \left\{ r(x, a) + \gamma \sum_y \mathcal{P}(y|x, a) V(y) \right\}$$

is equivalent to

$$V(x) \geq r(x, a) + \gamma \sum_y \mathcal{P}(y|x, a) V(y), \quad \forall a \in \mathcal{A}.$$

$$5 \geq \max \{1, 4, 3\} \Rightarrow \begin{array}{l} 5 \geq 1 \\ 5 \geq 4 \\ 5 \geq 3 \end{array}$$

Linear Programming for Finding V^*

$$\begin{array}{ll} \min_V & \mu^\top V, \\ \text{s.t.} & V(x) \geq r(x, a) + \gamma \sum_y \mathcal{P}(y|x, a) \underline{V(y)}, \quad \forall \underline{(x, a)} \in \underline{\mathcal{X} \times \mathcal{A}}. \end{array}$$

This is a linear program with $|\mathcal{X} \times \mathcal{A}|$ constraints.

Summary

Feedback:

- Add concrete (numerical) examples.
- In the Lab, try to show how the theory is practically used.

- Three methods for computing the optimal value function
 - ▶ Value Iteration
 - ▶ Policy Iteration
 - ▶ Linear Programming
- Established convergence of VI and PI
- These methods have variants for the RL setting.

References

- Dimitri P. Bertsekas. *Abstract dynamic programming*. Athena Scientific Belmont, 2nd edition, 2018.
- Amir-massoud Farahmand. *Foundations of Reinforcement Learning (draft)*. 2025.
- Thomas Dueholm Hansen, Peter Bro Miltersen, and Uri Zwick. Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor. *Journal of the ACM (JACM)*, 60(1): 1–16, 2013.
- Bruno Scherrer. Improved and generalized upper bounds on the complexity of policy iteration. *Mathematics of Operations Research*, 41(3): 758–774, 2016.
- Yinyu Ye. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Mathematics of Operations Research*, 36(4):593–603, 2011.