

RL Theory

[Planning in MDPs](#) / 8. Approximate Policy Iteration

8. Approximate Policy Iteration

[PDF Version](#)

Note: On March 13, 2021, these notes were updated as follows:

- 1 Tighter bounds are derived; the old analysis was based on bounding $\|T^* - T\|$; the new analysis directly bounds $\|T^* - T\|$, which leads to a better dependence on the approximation error;
 - 2 Unbiased return estimates are introduced that use rollouts of random length.
-

One simple idea to use function approximation in MDP planning is to take a planning method that uses internal value functions and add a constraint that restrict the value functions to have a compressed representation.

As usual, two questions arise:

- Does this lead to an **efficient** planner? That is, can the computation be carried out in time polynomial in the relevant quantities, but not the size of the state space? In the case of linear functions the question is whether we can calculate the coefficients efficiently.
- Does this lead to an **effective** planner? In particular, how good a policy can we arrive at with a limited compute effort?

In this lecture, as a start into exploring the use of value function approximation in planning, we look at modifying policy iteration in the above described way. The resulting algorithm belongs to the family of **approximate policy iteration** algorithms, which consists of all algorithms derived from policy iteration by adding approximation to it.

We will work with linear function approximation. In particular, we will assume that the planner is given as a hint a feature-map ϕ . In this setting, since policy iteration hinges upon evaluating the policies obtained, the hint given to the planner is

considered to be “good” if the (action-)value functions of **all** policies are well-represented with the features.

This means, that we will work under assumption B2 from the previous lecture, which we copy here for convenience. In what follows we fix .

Assumption B2 (approximate universal value function realizability) The MDP and the featuremap are such that for any memoryless policy of the MDP,

Recall that here the notation means that can be approximated up to a uniform error of using linear combinations of the basis functions underlying the feature-map :

For any policy ,

One may question whether it is reasonable to expect that the value functions of all policies can be compressed. We will come back to this question later.


Approximate Policy Evaluation: Done Well

Recall that in phase of policy iteration, given a policy , the next policy is obtained as the policy that is greedy with respect to . If we found some coefficients such that

then when it comes to “using” policy , we could just use when an action is needed at state . Note that this action can be obtained at the cost of elementary operations, a small overhead compared to a table lookup (with idealized access times).

Hence, the main question is how to obtain this parameter in an efficient manner. To be more precise, here we want to control the uniform error committed in approximating .

To simplify the notation, let . A simple idea is **rolling out** with the policy from a fixed set to “approximately” measure the value of at the pairs in . For concreteness, let . Rolling out with policy this pair means using the simulator to simulate what would happen if we used policy for a number of consecutive time steps when the initial state is , the first



action a , but for subsequent time steps the actions are chosen using policy π for whatever states are encountered. If the simulation goes on for T steps, this way we get T trajectories starting in s_0 . For $t = 0, \dots, T-1$ let the trajectory obtained be τ_t . Thus,

,

where $\tau_t = (s_0, a_0, r_0, s_1, a_1, r_1, \dots)$, and for $t = 0, \dots, T-1$, τ_t is a trajectory, and τ_t is a trajectory.

. The figure on the right illustrates these trajectories.

Given these trajectories, the empirical mean of the discounted sum of rewards along these trajectories is used for approximating J_π :

—

Under the usual condition that the rewards are in the $[-\bar{r}, \bar{r}]$ interval, the expected value of J_π is in the $[-\bar{r}, \bar{r}]$ vicinity of the J_π and by averaging a large number of independent trajectories, we also achieve that the empirical means are tightly concentrated around their mean.

Using a randomization device, it is possible to remove the error (“bias”) introduced by truncating the trajectories at a fixed time. For this, just let τ_t be independent **geometrically distributed** random variables with parameter γ , which are also independently chosen from the trajectories. By definition τ_t is the number of $(1-\gamma)$ -parameter Bernoulli trials needed to get one success. With the help of these variables, define now τ_t by

—

Note that in the expression of J_π the discount factor is eliminated. To calculate J_π one can just perform a rollout with policy π as before, just in each time step t , after obtaining s_t , draw a Bernoulli variable with parameter γ to decide whether the rollout should continue.

To see why the above definition works, fix t and note that by definition, for $\tau_t = k$, and thus $\tau_t = k$. Therefore,

All in all, this means, that we expect that if we solve for the **least-squares problem**

we expect \hat{w} to be a good approximation to w^* . Or at least, we can expect this hold at the points of \mathcal{X} , where we are taking our measurements. The question is what happens **outside of \mathcal{X}** : That is, what guarantees can we get for **extrapolating** to points of \mathcal{X}' . The first thing to observe that unless we are choosing \mathcal{X} carefully, there is no guarantee about the extrapolation error will be kept under control. In fact, if the choice of \mathcal{X} is so unfortunate that all the feature vectors for points in \mathcal{X} are **identical**, the least-squares problem will have many solutions.

Our next lemma gives an explicit error bound on the extrapolation error. For the coming results we slightly generalize least-squares by introducing a weighting of the various errors in \mathcal{X} . For this, let $w(x)$ be a weighting function assigning a positive weight to the various error terms and let

\hat{w} be the minimizer of the resulting weighted squared-loss. A simple calculation gives that provided the (weighted) **moment matrix**

is nonsingular, the solution to the above weighted least-squares problem is unique and is equal to

From this expression we see that there is no loss of generality in assuming that the weights in the weighting function sum to one: $\sum_{i=1}^n w_i = 1$. We will denote this by writing $\mathbf{w} \in \Delta_n$ (here, Δ_n refers to the fact that we can see \mathbf{w} as an element of a n -simplex). To state the lemma recall the notation that for a positive definite, $n \times n$ matrix \mathbf{A} and vector \mathbf{b} ,

Lemma (extrapolation error control in least-squares): Fix any $\mathbf{w} \in \Delta_n$, $\mathbf{A} \in \mathbb{R}^{n \times n}$, and $\mathbf{b} \in \mathbb{R}^n$ such that the moment matrix $\mathbf{A} \mathbf{w} \mathbf{A}^T$ is nonsingular. Define

Then, for any $\mathbf{w} \in \Delta_n$ we have

Before the proof note that what this lemma tells us is that as long as we guarantee that the moment matrix is full rank, the extrapolation errors relative to predicting with some \mathbf{w} can be controlled by controlling

- 1 the value of $\|\mathbf{A} \mathbf{w} \mathbf{A}^T\|^{-1}$; and
- 2 the maximum deviation of the targets used in the weighted least-squares problem and the predictions with \mathbf{w} .

Proof: First, we relate \mathbf{w} to \mathbf{w} :

Then for a fixed ϵ ,

To get a sense of how to control the sum notice that if ϵ in the last sum was somehow replaced by ϵ^2 , using the definition of ϵ could greatly simplify the last expression. To get here, one may further notice that having the term in absolute value squared would help. Now, to get the squares, recall [Jensen's inequality](#), which states that for any convex function ϕ and probability distribution μ , $\phi(\int x d\mu) \leq \int \phi(x) d\mu$. Of course, this also works when μ is a finitely supported, which is the case here. Thus, applying Jensen's inequality with $\phi(x) = x^2$, we thus get

Plugging this back into the previous inequality gives the desired result.

It remains to be seen of whether ϵ can be kept under control. This is the subject of a classic result of Kiefer and Wolfowitz:

Theorem (Kiefer-Wolfowitz): Let \mathcal{X} be finite. Let μ be such that the underlying feature matrix Φ is rank d . There exists a set \mathcal{X}_ϵ and a distribution μ_ϵ over this set, i.e. $\mu_\epsilon(\mathcal{X}_\epsilon) = 1$, such that

$$\begin{aligned} 1 & \quad \|\mu - \mu_\epsilon\| \leq \epsilon; \\ 2 & \quad \|\Phi - \Phi_\epsilon\| \leq \epsilon; \end{aligned}$$

- 3 In the previous line, the inequality is achieved with equality and the value of $\bar{\mu}$ is best possible under all possible choices of μ and σ .

We will not give a proof of the theorem, but we give references at the end where the reader can look up the proof. When Σ is not full rank (i.e., Σ is not rank n), one may reduce the dimensionality (and the cardinality of \mathcal{X} reduces accordingly). The problem of choosing μ and σ such that $\bar{\mu}$ is minimized is called the $\bar{\mu}$ -optimal design problem in statistics. This is a specific instance of optimal experimental design.

Combining the Kiefer-Wolfowitz theorem with the previous lemma shows that least-squares amplifies the “measurement errors” by at most a factor of $\bar{\mu}$:

Corollary (extrapolation error control in least-squares via optimal design): Fix any Σ full rank. Then, there exists a set \mathcal{X} with at most n elements and a weighting function w such that for any μ and any σ ,

where $\bar{\mu}$ is given by

Importantly, note that μ and σ are chosen independently of Σ , that is, they are independent of the target. This suggests that in approximate policy evaluation, one should choose μ as in the Kiefer-Wolfowitz theorem and use the w weighted moment matrix. This leads to

where $\bar{\mu}$ is defined by Eq. 10.1 and $\bar{\sigma}$ is defined by Eq. 10.2. We call this procedure **least-square policy evaluation based on rollouts from $\bar{\mu}$ -optimal design points**, or LSPE- $\bar{\mu}$, for short. Note that we stick to the truncated rollouts, because this allows a simpler probabilistic analysis. That this properly controls the extrapolation error is as attested by the next result:

Lemma (LSPE- extrapolation error control): Fix any full-rank feature-map ϕ and take the set \mathcal{S} and the weighting function w as in the Kiefer-Wolfowitz theorem. Fix an arbitrary policy π and let ϵ and δ such that $\epsilon > \delta$ and assume that immediate rewards belong to the interval $[-R, R]$. Let β be as in Eq. (10.10). Then, for any n , with probability $1 - \delta$,

$$\| \hat{Q}_n^\pi - Q^\pi \|_\infty \leq \frac{2R}{\epsilon} \sqrt{\frac{\log(1/\delta)}{n}}$$

Notice that that from the Kiefer-Wolfowitz theorem, ϵ and therefore nothing in the above expression depends on the size of the state space. Now, say we want to make the above error bound at most ϵ with some value of n . From the above we see that it suffices to choose n and δ so that

$$\frac{2R}{\epsilon} \sqrt{\frac{\log(1/\delta)}{n}} \leq \epsilon$$

This, together with (10.10) gives

$$\| \hat{Q}_n^\pi - Q^\pi \|_\infty \leq \epsilon$$

Proof: In a nutshell, we use the previous corollary, together with [Hoeffding's inequality](#) and using that R , which follows since the rewards are bounded in $[-R, R]$.

► [Click here for the full proof.](#)

In summary, what we have shown so far is that if the features can approximate well the action-value function of a policy, then there is a simple procedure (Monte-Carlo rollouts and least-squares estimation based on an optimal experimental design) to produce an reliable estimate of the action-value function of the policy. The question remains whether if we use these estimates in policy iteration, the whole procedure will still give good policies after a sufficiently large number of iterations.

Progress Lemma with Approximation Errors

Here we give a refinement of the [geometric progress lemma](#) of policy iteration that allows for “approximate” policy improvement steps. This previous lemma stated that the value function of the improved policy is at least as large as the Bellman operator applied to the value function of the policy to be improved. Our new lemma is as follows:

Lemma (Geometric progress lemma with approximate policy improvement): Consider a memoryless policy π and its corresponding value function V^π . Let $\tilde{\pi}$ be any policy and define $V^{\tilde{\pi}}$ via

Then,

Proof: First note that for the optimal policy π^* , $V^{\pi^*} = T V^{\pi^*}$. We have

Using the value difference identity and that T is a contraction, we calculate

where the inequality follows because T is a contraction, the sum of positive linear operators, is a positive linear operator itself and hence is also monotone. Plugging the inequality obtained into $V^{\tilde{\pi}} - T V^{\tilde{\pi}}$ gives

Taking the maximum norm of both sides and using the triangle inequality and that T is a contraction gives the desired result.

Approximate Policy Iteration

Notice that the progress lemma makes no assumptions about the origin of the errors. This motivates considering a generic version of **approximate policy iteration** where for $k \geq 0$ in the k th update set, the new policy π_{k+1} is approximately greedy with respect to J_{π_k} in that sense that

The progress lemma implies that the resulting sequence of policies will have value functions that converge to a neighborhood of J^* where the size of the neighborhood is governed by the magnitude of the error terms ϵ_k .

Theorem (Approximate Policy Iteration): Let ϵ_k , δ_k be such that $\epsilon_k \leq \delta_k$ holds for all k . Then, for any $\epsilon > 0$,

Proof: Left as an exercise.

Consider now a version of approximate policy iteration where the sequence of policies π_k is defined as follows:

That is, for each $k \geq 0$, π_{k+1} is greedy with respect to J_{π_k} .

Corollary (Approximate Policy Iteration with Approximate Action-value Functions): The sequence defined in (1) is such that

Proof: To simplify the notation consider policies π_k and functions J_{π_k} over the state-action space such that $J_{\pi_k} = T_{\pi_k} J_{\pi_k}$ and $J_{\pi_k} = T_{\pi_k} J_{\pi_k}$. We have

where we used that Φ is linear, monotone, and that T is monotone, and both are nonexpansions in the maximum norm.

Hence, if π is defined by $\pi = T(\Phi \pi)$ then π is optimal and the result follows from the previous theorem.

Global planning with least-squares policy iteration

Putting things together gives the following planning method:

- 1 Given the feature map Φ , find λ and β as in the Kiefer-Wolfowitz theorem
- 2 Let π_0 be an arbitrary policy
- 3 For $k = 0, 1, 2, \dots$ do
- 4 Roll out with policy π_k for n steps to get the targets $\{r_t + \gamma V_{t+1}^{\pi_k}\}$ where $V_{t+1}^{\pi_k} = \Phi^T \pi_k$ and $r_t + \gamma V_{t+1}^{\pi_k} = \Phi^T \pi_k + r_t$
- 5 Solve the weighted least-squares problem given by Eq. (10.10) to get π_{k+1}
- 6 Return π_k

We call this method least-squares policy iteration (LSPI) for obvious reasons. Note that this is a **global planning method**: The method makes no use of an input state and the parameter vector returned can be used to get the policy π (as in the method above).

Theorem (LSPI performance): Fix an arbitrary full rank feature-map Φ and let $\lambda > 0$. Assume that B2 holds. Then, for any $\beta > 0$, with probability at least $1 - \epsilon$, the policy π which is greedy with respect to $\Phi^T \pi + r$ is β -suboptimal with

$$\frac{2\lambda \beta \|\Phi\|_{\infty}^2}{\lambda \beta \|\Phi\|_{\infty}^2 - \epsilon}$$

In particular, for any $\epsilon > 0$, choosing $\beta = \frac{\epsilon}{2\lambda \|\Phi\|_{\infty}^2}$ so that

$$\lambda \beta \|\Phi\|_{\infty}^2 - \epsilon > 0$$

policy is ϵ -optimal with

$$\frac{1}{1 - \gamma} \epsilon$$

while the total computation cost is $\frac{1}{1 - \gamma} \frac{1}{\epsilon^2}$.

Thus, with a polynomial cost, LSPI with the specific configuration at the cost of polynomial computation cost, but importantly, with a cost that is independent of the size of the state space, can result in a good policy as long as ϵ , the worst-case error of approximating action-value functions of policies using the features provided, is sufficiently small.

Proof: Note that B2 and that Φ is full rank implies that for any memoryless policy π there exists a parameter vector w such that $\Phi w = v_\pi$ (cf. Part 2 of Question 3 of Assignment 2). Hence, we can use the “LSPE extrapolation error bound” (cf. (1)). By this result, a union bound and of course by B2, we get that for any π , with probability at least $1 - \delta$, for any n ,

$$\| \hat{v}_\pi - v_\pi \| \leq \frac{1}{\sqrt{n}} \left(\frac{1}{1 - \gamma} \sqrt{\log \frac{1}{\delta}} + \frac{1}{1 - \gamma} \sqrt{\log \frac{1}{\delta}} \right)$$

where we also used that $\| \Phi^\top \Phi \| \leq \frac{1}{1 - \gamma^2}$. Call the quantity on the right-hand side in the above inequality ϵ .

Take the event when the above inequalities hold and for now assume this event holds. By the previous theorem, $\hat{\pi}$ is ϵ -optimal with

$$\frac{1}{1 - \gamma} \epsilon$$

To obtain the second part of the result, we split ϵ into two equal parts: $\epsilon/2$ is set to force the iteration error to be at most $\epsilon/2$, while $\epsilon/2$ and $\epsilon/2$ are chosen to force the policy evaluation error to be at most $\epsilon/2$. Here, to choose $\epsilon/2$ and $\epsilon/2$, ϵ is again split into two equal parts. The details of this calculation are left to the reader.

Notes

Approximate Dynamic Programming (ADP)

Value iteration and policy iteration are specific instances of dynamic programming methods. In general, dynamic programming refers to methods that use value functions to calculate good policies. In **approximate dynamic programming** the methods are modified by introducing “errors” when calculating the values. The idea is that the origin of the errors does not matter (e.g., whether they come due to imperfect function approximation, linear, or nonlinear, or due to the sampling): The analysis is done in a general form. While here we met approximate policy iteration, one can also use the same ideas as shown here to study an approximate version of value iteration. A homework in problem set 2 asks you to study this method, which is usually called **approximate value iteration**. In an earlier homework you were asked to study how linear programming can also be used to compute optimal value functions. Adding approximations we then get **approximate linear programming**.

What function approximation technique to use?

We note in passing that fans of neural networks should like that the general, ADP-style results, like the theorem in the middle of this lecture, can be also applied to the case when neural networks are used as the function approximation technique. However, one main lesson of the lecture is that to control extrapolation errors, one should be quite careful in how the training data is chosen. For linear prediction and least-squares fitting, optimal design gives a complete answer, but the analog questions are completely open in the case of nonlinear function approximation, such as neural networks. There is also a sizable literature that connects nonparametric techniques (an analysis friendly relative of neural networks) to ADP methods.

Concentrability coefficients and all that jazz

The idea of introducing approximate calculations has been introduced at the same time people got interested in Markov Decision Processes in the 1960s. Hence, the literature is quite enormous. However, the approach taken here which asks for error bounds where the algorithmic (not approximation-) error is uniformly controlled regardless of the MDP is quite recent and where the term that involves the approximation error is also uniformly bounded (for a fixed dimension and discount factor).

Earlier literature often presented bounds where the magnification factor of the approximation and the algorithmic error involved terms which depended on the MDP. Often these came in the form of “concentrability coefficients” (and yours truly was quite busy with working on these results a while ago). The main conclusion of this earlier analysis is that more stochasticity in the transitions means less control, less concentrability, which is advantageous for the ADP algorithms. While this makes sense and this indicates that these earlier results are complementary to the results presented here, the issue is that these

results are quite pessimistic for example when the MDP is deterministic (as in this case the concentrability coefficients can be as large as the size of the state space).

While here we emphasized the importance of using a good design to control the extrapolation errors, in these earlier results, no optimal design was used. The upshot is that this saves the effort of coming up with a good design, but the obvious downside is that the extrapolation error may become uncontrolled. In the batch setting (which we will come back to later), of course, there is no way to control the sample collection, and this is in fact the setting where this earlier analysis was done.

The strength of hints

A critical assumption in the analysis of API was that the approximation error is controlled uniformly for all policies. This feels limiting. Yet, there are some interesting sufficient conditions when this assumption is clearly satisfied. In general, these require that the transition dynamics and the reward are both “compressible”. For example, if the MDP is such that \mathcal{P} , the immediate reward as a function of the state–action pairs satisfies $\|r\|_{\infty} \leq \gamma$ and the transition matrix, \mathcal{P} , satisfies $\|\mathcal{P}\|_{\infty} \leq \gamma$ with some matrix γ , then for any policy π , \mathcal{P}^{π} has a range which is a subset of \mathcal{R}^{γ} . Since π^* is the fixed–point of \mathcal{P}^{π^*} , i.e., $\mathcal{P}^{\pi^*} \pi^* = \pi^*$, it follows that π^* is also necessarily in the range space of \mathcal{P}^{π^*} . As such, \mathcal{P}^{π^*} and \mathcal{P}^{π} . MDPs that satisfy the above two constraints are called **linear in** (or sometimes, just “linear MDPs”). Exact linearity can be relaxed: If \mathcal{P}^{π} and \mathcal{P}^{π^*} , then for any policy π , \mathcal{P}^{π} with \mathcal{P}^{π^*} — \mathcal{P}^{π} . Nevertheless, later we will investigate whether this assumption can be relaxed.

The tightness of the bounds

It is not known whether the bound presented in the final result is tight. In fact, the dependence of ϵ on the γ is almost certainly not tight; in similar scenarios it has been shown in the past that replacing Hoeffding’s inequality with Bernstein’s inequality allows the reduction of this factor. It is more interesting whether the amplification factor of the approximation error, $\frac{1}{1-\gamma}$, is best possible. In the next lecture we will show that the $\frac{1}{1-\gamma}$ approximation error amplification factor cannot be removed while keeping the runtime under control. In a later lecture, we will show that the dependence on γ cannot be improved either – at least for this algorithm. However, we will see that if the main concern is the amplification of the approximation error, while keeping the runtime polynomial (perhaps with a higher order though) then under B2 better algorithms exist.

The cost of optimal experimental design

The careful reader would not miss that to run the proposed method one needs to find the set and the weighting function . The first observation here is that it is not crucial to find the best possible pair. The Kiefer-Wolfowitz theorem showed that with this best possible choice, . However, if one finds a pair such that , the price of this is that wherever appears in the final performance bound, a submultiplicative factor of will also need to be introduced. This should be acceptable. In relation to this note that by relaxing this optimality requirement, the cardinality of can be reduced. For example, by introducing the factor of as suggested above allows one to reduce the cardinality to ; which may actually be a good tradeoff as this can save much on the runtime.

However, the question still remains of who computes these (approximately) optimal designs and at what cost. While this calculation only needs to be done once and is independent of the MDP (just depends on the feature map), the value of these methods remains unclear because of this compute cost. General methods to compute approximately optimal designs needed here are known, but their runtime for our case will be proportional to the number of state-action pairs. In the very rare cases when simulating transitions is very costly but the number of state-action pairs is not too high, this may be a viable option. However, these cases are rare. For special choices of the feature-map, optimal designs may be known. However, this reduces the general applicability of the method presented here. Thus, a major question is whether the optimal experimental design can be avoided. What is known is that for linear prediction with least-squares, clearly, they cannot be avoided. One suspects that this is true more generally.

Can optimal designs be avoided while keeping the results essentially unchanged? Of particular interest would be if the feature-map would also be only “locally explored” as the planner interacts with the simulator. Altogether, one suspects that two factors contributed here for the appearance of optimal experimental design: One factor is that the planner is global: It comes up with a parameter vector that leads to a policy that can be used regardless of the state. The other (perhaps) factor is that the approach was based on simple “patching up” a dynamic programming algorithm with a function approximator. While this is a common approach, controlling the extrapolation errors in this approach is critical and is likely only possible with something like an optimal experimental design. As we shall see soon, there are indeed approaches that avoid the optimal experimental design step and which are based on online planning and they also deviate from the ADP approach.

Policy evaluation alternatives

The policy evaluation method presented here feels unsophisticated. It uses simple Monte-Carlo rollouts, with truncation, averaging and least-squares regression. The reinforcement learning literature offers many alternatives, such as the “temporal difference” learning

type methods that are based on solving the fixed point equation $T_{\pi}^* v = v$. One can indeed try to use this equation to avoid the crude Monte-Carlo approach presented here, in the hope of reducing the variance (which is currently rather crudely upper bounded using the $\frac{1}{\sqrt{n}}$ term in the Hoeffding bound). Rewriting the fixed point as $v = T_{\pi} v$, and then plugging in $v + \epsilon$, we see that the trouble is that to control the extrapolation errors, the optimal design must likely depend on the policy to be evaluated (because of the appearance of ϵ).

Alternative error control: Bellman residuals

Let v and v^* be so that

Here, $v - v^*$ is called the “Bellman residual” of v . The policy evaluation alternatives above aim at controlling these residuals. The reader is invited to derive the analogue of the “approximate policy iteration” error bound in [\(1\)](#) for this scenario.

The role of ϵ in the Kiefer-Wolfowitz result

One may wonder about how critical is the presence of ϵ in the results presented. For this, we can say that it is not critical. Unweighted least-squares does not perform much worse.

Least-squares error bound

The error bound presented for least-squares does not use the full power of randomness. When part of the errors ϵ_i with ϵ_i are random, some helpful averaging effects can appear, which we ignored for now, but which could be used in a more refined analysis.

Optimal experimental design – a field on its own

Optimal experimental design is a subfield of statistics. The design considered here is just one possibility. In fact, this design which is called G-optimal design (G stands, uninspiringly, for the word “general”). The Kiefer-Wolfowitz theorem actually also states that this is equivalent to the D-optimal designs.

Lack of convergence

The results presented show convergence to a ball around the optimal target. Some people think this is a major concern. While having a convergent method may look more appealing, as long as one controls the size of the ball, I will not be too concerned.

Approximate value iteration (AVI)

Similarly to what is done here, one can introduce an approximate version of value-iteration. This is the subject of Question 3 of [homework 2](#). While the conditions are different, the qualitative behavior of AVI is similar to that of approximate policy iteration.

In particular, as for approximate policy iteration, there are two steps to this proof: One is to show that the residuals can be controlled and the second is that if they are controlled then the policy that is greedy with respect to (say) is ϵ -optimal with controlled by ϵ . For this second part, we have the following bound:

where ϵ . The procedure that uses least-squares fitting to get the iterates is known under various names, such as **least-squares value iteration** (LSVI), **fitted Q-iteration** (FQI), **least-squares Q iteration** (LSQI). This proliferation of abbreviations and names is unfortunate, but there is not much that can be done at this stage. To add insult to injury, when neural networks are used to represent the iterates and an incremental stochastic gradient descent algorithm is used for “fitting” the weights of these networks by resampling old data from a “replay buffer”, the resulting procedure is coined “Deep Q-Networks” (training), or DQN for short.

Bounds on the parameter vector

The Kiefer-Wolfowitz theorem implies the following:

Proposition: Let \mathcal{X} and \mathcal{Y} be such that \mathcal{X} and \mathcal{Y} and \mathcal{Y} . Then, there exist a matrix \mathbf{A} such that for

there exists ϵ such that the following hold:

- 1 $\mathbf{A}^{-1} \mathbf{A} \mathbf{A}^{-1} = \mathbf{A}^{-1}$;
- 2 $\mathbf{A}^{-1} \mathbf{A} \mathbf{A}^{-1} = \mathbf{A}^{-1}$;
- 3 $\mathbf{A}^{-1} \mathbf{A} \mathbf{A}^{-1} = \mathbf{A}^{-1}$.

Proof: Let π^* be the γ -optimal design whose existence is guaranteed by the Kiefer-Wolfowitz theorem. Let Σ be the underlying moment matrix. Then, by the definition of π^* ,

Define μ and Σ . The first property is clearly satisfied. As to the second property,

Finally, for the third property,

finishing the proof.

Thus, if one has access to the full feature-map then knowing that a function realized is bounded, one may as well assume that the feature map is bounded and the parameter vector is bounded just by \bar{M} .

Regularized least-squares

The linear least-squares predictor given by a feature-map ϕ and data \mathcal{D} predicts a response at x via $\hat{y}(x)$ where

with

Here, by abusing notation for the sake of minimizing clutter, we use Σ , μ . The problem is that Σ may not be invertible (i.e., Σ^{-1} may not be defined as written above). “By continuity”, it is nearly equally problematic when Σ is ill-conditioned (i.e., its minimum eigenvalue is “much smaller” than its maximum eigenvalue). In fact, this leads to poor “generalization”. One remedy, often used, is to modify Σ by shifting it with a small constant multiple of the identity matrix:

Here, λ is a tuning parameter, whose value is often chosen based on cross-validation or with a similar process. The modification guarantees that $A + \lambda I$ is invertible and it overall improves the quality of predictions, especially when A is tuned base on data.

Above, the choice of the identity matrix, while is common in the literature, is completely arbitrary. In particular, invertibility will be guaranteed if A is replaced with any other positive definite matrix B . In fact, the matrix one should use here should be one that makes $\|B^{-1}\|$ small (while, say, keeping the minimum eigenvalue of B at constant). That this is the choice that makes sense can be argued for by noting that with

the vector defined in (1) is the minimizer of

and thus, the extra penalty has the least impact for the choice of λ that makes the norm of the smallest. If we only know that λ^* , by our previous note, a good choice is λ^* , where λ^* where λ^* is a λ -optimal design. Indeed, with this choice, λ^* . Note also that if we apply the feature-standardization transformation of the previous note, we have

showing that the choice of using the identity matrix is justified when the features are standardized as in the proposition of the previous note.

References

We will only scratch the surface now; expect more references to be added later.

The bulk of this lecture is based on

- Tor Lattimore, Csaba Szepesvári, and Gellért Weisz. 2020. “Learning with Good Feature Representations in Bandits and in RL with a Generative Model.” ICML and

arXiv:1911.07676,

who introduced the idea of using ϵ -optimal designs for controlling the extrapolation errors. A very early reference on error bounds in “approximate dynamic programming” is the following:

- Whitt, Ward. 1979. “Approximations of Dynamic Programs, II.” *Mathematics of Operations Research* 4 (2): 179–85.

The analysis of the generic form of approximate policy iteration is a refinement of Proposition 6.2 from the book of Bertsekas and Tsitsiklis:

- Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, Massachusetts, 1996.

However, there are some differences between the “API” theorem presented here and Proposition 6.2. In particular, the theorem presented here appears to capture all sources of errors in a general way, while Proposition 6.2 is concerned with value function approximation errors and errors introduced in the “greedification step”. The form adopted here appears, for example, in Theorem 1 of a technical report of Scherrer, who also gives earlier references:

- Scherrer, Bruno. 2013. “On the Performance Bounds of Some Policy Search Dynamic Programming Algorithms.” [arxiv](#).

The earliest of these references is perhaps

- Munos, R. 2003. “Error Bounds for Approximate Policy Iteration.” *ICML*.

Least-squares policy iteration appears in

- Lagoudakis, M. G. and Parr, R. Least-squares policy iteration. *The Journal of Machine Learning Research*, 4:1107–1149, 2003.

The particular form presented in this work though uses value function approximation based on minimizing the Bellman residuals (using the so-called LSTD method).

Two books that advocate the ADP approach:

- Powell, Warren B. 2011. *Approximate Dynamic Programming. Solving the Curses of Dimensionality*. Hoboken, NJ, USA: John Wiley & Sons, Inc.
- Lewis, Frank L., and Derong Liu. 2013. *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control*. Hoboken, NJ, USA: John Wiley & Sons, Inc.

And a chapter:

- Bertsekas, Dimitri P. 2009. “Chapter 6: Approximate Dynamic Programming,” January, 1–118.

A paper that is concerned with API and least-squares methods, but uses concentrability is:

Antos, Andras, Csaba Szepesvári, and Rémi Munos. 2007. “Learning near-Optimal Policies with Bellman-Residual Minimization Based Fitted Policy Iteration and a Single Sample Path.” *Machine Learning* 71 (1): 89–129.

Optimal experimental design has a large literature. A nice book concerned with computation is this:

- M. J. Todd. *Minimum-volume ellipsoids: Theory and algorithms*. SIAM, 2016.

The Kiefer-Wolfowitz theorem is from:

- J. Kiefer and J. Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12(5):363–365, 1960.

More on computation here:

- E. Hazan, Z. Karnin, and R. Meka. Volumetric spanners: an efficient exploration basis for learning. *Journal of Machine Learning Research*, 17(119):1–34, 2016
- M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization, volume 2*. Springer Science & Business Media, 2012.

The latter book is a very good general starting point for convex optimization.

That the features are standardized as shown in the notes is assumed (and discussed), e.g., in

- Wang, Ruosong, Dean P. Foster, and Sham M. Kakade. 2020. “What Are the Statistical Limits of Offline RL with Linear Function Approximation?” *arXiv [cs.LG]*. [arXiv](#)

which we will meet later.

