# 6.231 DYNAMIC PROGRAMMING 

## LECTURE 5

## LECTURE OUTLINE

- Review of approximate PI based on projected Bellman equations
- Issues of policy improvement
- Exploration enhancement in policy evaluation
- Oscillations in approximate PI
- Aggregation - An alternative to the projected equation/Galerkin approach
- Examples of aggregation
- Simulation-based aggregation
- Relation between aggregation and projected equations


## REVIEW

## DISCOUNTED MDP

- System: Controlled Markov chain with states $i=1, \ldots, n$ and finite set of controls $u \in U(i)$
- Transition probabilities: $p_{i j}(u)$

- Cost of a policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$ starting at state $i$ :

$$
J_{\pi}(i)=\lim _{N \rightarrow \infty} E\left\{\sum_{k=0}^{N} \alpha^{k} g\left(i_{k}, \mu_{k}\left(i_{k}\right), i_{k+1}\right) \mid i=i_{0}\right\}
$$

with $\alpha \in[0,1)$

- Shorthand notation for DP mappings

$$
\begin{aligned}
& (T J)(i)=\min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)(g(i, u, j)+\alpha J(j)), \quad i=1, \ldots, n, \\
& \left(T_{\mu} J\right)(i)=\sum_{j=1}^{n} p_{i j}(\mu(i))(g(i, \mu(i), j)+\alpha J(j)), \quad i=1, \ldots, n
\end{aligned}
$$

## APPROXIMATE PI



- Evaluation of typical policy $\mu$ : Linear cost function approximation

$$
\tilde{J}_{\mu}(r)=\Phi r
$$

where $\Phi$ is full rank $n \times s$ matrix with columns the basis functions, and $i$ th row denoted $\phi(i)^{\prime}$.

- Policy "improvement" to generate $\bar{\mu}$ :

$$
\bar{\mu}(i)=\arg \min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha \phi(j)^{\prime} r\right)
$$

## EVALUATION BY PROJECTED EQUATIONS

- Approximate policy evaluation by solving

$$
\Phi r=\Pi T_{\mu}(\Phi r)
$$

$\Pi$ : weighted Euclidean projection; special nature of the steady-state distribution weighting.

- Implementation by simulation (single long trajectory using current policy - important to make $\Pi T_{\mu}$ a contraction). LSTD, LSPE methods.
- Multistep option: Solve $\Phi r=\Pi T_{\mu}^{(\lambda)}(\Phi r)$ with

$$
T_{\mu}^{(\lambda)}=(1-\lambda) \sum_{\ell=0}^{\infty} \lambda^{\ell} T_{\mu}^{\ell+1}, \quad 0 \leq \lambda<1
$$

- As $\lambda \uparrow 1, \Pi T_{\mu}^{(\lambda)}$ becomes a contraction for any projection norm (allows changes in $\Pi$ )
- Bias-variance tradeoff



## ISSUES OF POLICY IMPROVEMENT

## EXPLORATION

- 1st major issue: exploration. To evaluate $\mu$, we need to generate cost samples using $\mu$
- This biases the simulation by underrepresenting states that are unlikely to occur under $\mu$.
- As a result, the cost-to-go estimates of these underrepresented states may be highly inaccurate, and seriously impact the "improved policy" $\bar{\mu}$.
- This is known as inadequate exploration - a particularly acute difficulty when the randomness embodied in the transition probabilities is "relatively small" (e.g., a deterministic system).
- To deal with this we must change the sampling mechanism and modify the simulation formulas.
- Solve

$$
\Phi r=\bar{\Pi} T_{\mu}(\Phi r)
$$

where $\bar{\Pi}$ is projection with respect to an explorationenhanced norm [uses a weight distribution $\zeta=$ $\left.\left(\zeta_{1}, \ldots, \zeta_{n}\right)\right]$.

- $\zeta$ is more "balanced" than $\xi$ the steady-state distribution of the Markov chain of $\mu$.
- This also addresses any lack of ergodicity of $\mu$.


## EXPLORATION MECHANISMS

- One possibility: Use multiple short simulation trajectories instead of single long trajectory starting from a rich mixture of states. This is known as geometric sampling, or free-form sampling.
- By properly choosing the starting states, we enhance exploration
- The simulation formulas for $\operatorname{LSTD}(\lambda)$ and $\operatorname{LSPE}(\lambda)$ have to be modified to yield the solution of $\Phi r=\bar{\Pi} T_{\mu}^{(\lambda)}(\Phi r)$ (see the DP text)
- Another possibility: Use a modified policy to generate a single long trajectory. This is called an off-policy approach.
- Modify the transition probabilities of $\mu$ to enhance exploration
- Again the simulation formulas for $\operatorname{LSTD}(\lambda)$ and $\operatorname{LSPE}(\lambda)$ have to be modified to yield the solution of $\Phi r=\bar{\Pi} T_{\mu}^{(\lambda)}(\Phi r)$ (use of importance sampling; see the DP text)
- With larger values of $\lambda>0$ the contraction property of $\bar{\Pi} T_{\mu}^{(\lambda)}$ is maintained.
- LSTD may be used without $\bar{\Pi} T_{\mu}^{(\lambda)}$ being a contraction ... LSPE and TD require a contraction.


## POLICY ITERATION ISSUES: OSCILLATIONS

- 2nd major issue: oscillation of policies
- Analysis using the greedy partition of the space of weights $r$ : $R_{\mu}$ is the set of parameter vectors $r$ for which $\mu$ is greedy with respect to $\tilde{J}(\cdot ; r)=\Phi r$

$$
R_{\mu}=\left\{r \mid T_{\mu}(\Phi r)=T(\Phi r)\right\} \quad \forall \mu
$$

If we use $r$ in $R_{\mu}$ the next "improved" policy is $\mu$


- If policy evaluation is exact, there is a finite number of possible vectors $r_{\mu}$, (one per $\mu$ )
- The algorithm ends up repeating some cycle of policies $\mu^{k}, \mu^{k+1}, \ldots, \mu^{k+m}$ with

$$
r_{\mu^{k}} \in R_{\mu^{k+1}}, r_{\mu^{k+1}} \in R_{\mu^{k+2}}, \ldots, r_{\mu^{k+m}} \in R_{\mu^{k}}
$$

- Many different cycles are possible


## MORE ON OSCILLATIONS/CHATTERING

- In the case of optimistic policy iteration a different picture holds (policy evaluation does not produce exactly $r_{\mu}$ )

- Oscillations of weight vector $r$ are less violent, but the "limit" point is meaningless!
- Fundamentally, oscillations are due to the lack of monotonicity of the projection operator, i.e., $J \leq J^{\prime}$ does not imply $\Pi J \leq \Pi J^{\prime}$.
- If approximate PI uses an evaluation of the form

$$
\Phi r=\left(W T_{\mu}\right)(\Phi r)
$$

with $W$ : monotone and $W T_{\mu}$ : contraction, the policies converge (to a possibly nonoptimal limit).

- These conditions hold when aggregation is used


## AGGREGATION

## PROBLEM APPROXIMATION - AGGREGATION

- Another major idea in ADP is to approximate $J^{*}$ or $J_{\mu}$ with the cost-to-go functions of a simpler problem.
- Aggregation is a systematic approach for problem approximation. Main elements:
- Introduce a few "aggregate" states, viewed as the states of an "aggregate" system
- Define transition probabilities and costs of the aggregate system, by relating original system states with aggregate states
- Solve (exactly or approximately) the "aggregate" problem by any kind of VI or PI method (including simulation-based methods)
- If $\hat{R}(y)$ is the optimal cost of aggregate state $y$, we use the approximation

$$
J^{*}(j) \approx \sum \phi_{j y} \hat{R}(y), \quad \forall j
$$

where $\phi_{j y}$ are the aggregation probabilities, encoding the "degree of membership of $j$ in the aggregate state $y "$

- This is a linear architecture: $\phi_{j y}$ are the features of state $j$


## HARD AGGREGATION EXAMPLE

- Group the original system states into subsets, and view each subset as an aggregate state
- Aggregation probs.: $\phi_{j y}=1$ if $j$ belongs to aggregate state $y$ (piecewise constant approx).

$$
\begin{aligned}
& \Phi=\left(\begin{array}{llll}
1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0 \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right)
\end{aligned}
$$

- What should be the "aggregate" transition probs. out of $x$ ?
- Select $i \in x$ and use the transition probs. of $i$. But which $i$ should I use?
- The simplest possibility is to assume that all states $i$ in $x$ are equally likely.
- A generalization is to randomize, i.e., use "disaggregation probabilities" $d_{x i}$ : Roughly, the "degree to which $i$ is representative of $x$."


## AGGREGATION/DISAGGREGATION PROBS



- Define the aggregate system transition probabilities via two (somewhat arbitrary) choices.
- For each original system state $j$ and aggregate state $y$, the aggregation probability $\phi_{j y}$
- Roughly, the "degree of membership of $j$ in the aggregate state $y$."
- In hard aggregation, $\phi_{j y}=1$ if state $j$ belongs to aggregate state/subset $y$.
- For each aggregate state $x$ and original system state $i$, the disaggregation probability $d_{x i}$
- Roughly, the "degree to which $i$ is representative of $x$."
- Aggregation scheme is defined by the two matrices $D$ and $\Phi$. The rows of $D$ and $\Phi$ must be probability distributions.


## AGGREGATE SYSTEM DESCRIPTION



- The transition probability from aggregate state $x$ to aggregate state $y$ under control $u$
$\hat{p}_{x y}(u)=\sum_{i=1}^{n} d_{x i} \sum_{j=1}^{n} p_{i j}(u) \phi_{j y}, \quad$ or $\hat{P}(u)=D P(u) \Phi$
where the rows of $D$ and $\Phi$ are the disaggregation and aggregation probs.
- The expected transition cost is

$$
\hat{g}(x, u)=\sum_{i=1}^{n} d_{x i} \sum_{j=1}^{n} p_{i j}(u) g(i, u, j), \quad \text { or } \hat{g}=D P(u) g
$$

## AGGREGATE BELLMAN'S EQUATION



- The optimal cost function of the aggregate problem, denoted $\hat{R}$, is
$\hat{R}(x)=\min _{u \in U}\left[\hat{g}(x, u)+\alpha \sum_{y} \hat{p}_{x y}(u) \hat{R}(y)\right], \quad \forall x$ Bellman's equation for the aggregate problem.
- The optimal cost function $J^{*}$ of the original problem is approximated by $\tilde{J}$ given by

$$
\tilde{J}(j)=\sum_{y} \phi_{j y} \hat{R}(y), \quad \forall j
$$

## EXAMPLE I: HARD AGGREGATION

- Group the original system states into subsets, and view each subset as an aggregate state
- Aggregation probs.: $\phi_{j y}=1$ if $j$ belongs to aggregate state $y$.

- Disaggregation probs.: There are many possibilities, e.g., all states $i$ within aggregate state $x$ have equal prob. $d_{x i}$.
- If optimal cost vector $J^{*}$ is piecewise constant over the aggregate states/subsets, hard aggregation is exact. Suggests grouping states with "roughly equal" cost into aggregates.
- A variant: Soft aggregation (provides "soft boundaries" between aggregate states).


## EXAMPLE II: FEATURE-BASED AGGREGATION

- Important question: How do we group states together?
- If we know good features, it makes sense to group together states that have "similar features"

- A general approach for passing from a featurebased state representation to a hard aggregationbased architecture
- Essentially discretize the features and generate a corresponding piecewise constant approximation to the optimal cost function
- Aggregation-based architecture is more powerful (it is nonlinear in the features)
- ... but may require many more aggregate states to reach the same level of performance as the corresponding linear feature-based architecture


## EXAMPLE III: REP. STATES/COARSE GRID

- Choose a collection of "representative" original system states, and associate each one of them with an aggregate state

- Disaggregation probabilities are $d_{x i}=1$ if $i$ is equal to representative state $x$.
- Aggregation probabilities associate original system states with convex combinations of representative states

$$
j \sim \sum_{y \in \mathcal{A}} \phi_{j y} y
$$

- Well-suited for Euclidean space discretization
- Extends nicely to continuous state space, including belief space of POMDP


## EXAMPLE IV: REPRESENTATIVE FEATURES

- Here the aggregate states are nonempty subsets of original system states. Common case: Each $S_{x}$ is a group of states with "similar features"

- Restrictions:
- The aggregate states/subsets are disjoint.
- The disaggregation probabilities satisfy $d_{x i}>$ 0 if and only if $i \in x$.
- The aggregation probabilities satisfy $\phi_{j y}=1$ for all $j \in y$.
- Hard aggregation is a special case: $\cup_{x} S_{x}=$ $\{1, \ldots, n\}$
- Aggregation with representative states is a special case: $S_{x}$ consists of just one state


## APPROXIMATE PI BY AGGREGATION



- Consider approximate PI for the original problem, with policy evaluation done by aggregation.
- Evaluation of policy $\mu: \tilde{J}=\Phi R$, where $R=$ $D T_{\mu}(\Phi R)$ ( $R$ is the vector of costs of aggregate states for $\mu$ ). Can be done by simulation.
- Looks like projected equation $\Phi R=\Pi T_{\mu}(\Phi R)$ (but with $\Phi D$ in place of $\Pi$ ).
- Advantage: It has no problem with oscillations.
- Disadvantage: The rows of $D$ and $\Phi$ must be probability distributions.


## ADDITIONAL ISSUES OF AGGREGATION

## ALTERNATIVE POLICY ITERATION

- The preceding PI method uses policies that assign a control to each aggregate state.
- An alternative is to use PI for the combined system, involving the Bellman equations:

$$
\begin{gathered}
R^{*}(x)=\sum_{i=1}^{n} d_{x i} \tilde{J}_{0}(i), \quad \forall x \\
\tilde{J}_{0}(i)=\min _{u \in U(i)} \sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha \tilde{J}_{1}(j)\right), \quad i=1, \ldots, n \\
\tilde{J}_{1}(j)=\sum_{y \in \mathcal{A}} \phi_{j y} R^{*}(y), \quad j=1, \ldots, n
\end{gathered}
$$

Original


- Simulation-based PI and VI are still possible.


## RELATION OF AGGREGATION/PROJECTION

- Compare aggregation and projected equations

$$
\Phi R=\Phi D T(\Phi R), \quad \Phi r=\Pi T(\Phi r)
$$

- If $\Phi D$ is a projection (with respect to some weighted Euclidean norm), then the methodology of projected equations applies to aggregation
- Hard aggregation case: $\Phi D$ can be verified to be projection with respect to weights $\xi_{i}$ proportional to the disaggregation probabilities $d_{x i}$
- Aggregation with representative features case: $\Phi D$ can be verified to be a semi-norm projection with respect to weights $\xi_{i}$ proportional to $d_{x i}$
- A (weighted) Euclidean semi-norm is defined by $\|J\|_{\xi}=\sqrt{\sum_{i=1}^{n} \xi_{i}(J(i))^{2}}$, where $\xi=\left(\xi_{1}, \ldots, \xi_{n}\right)$, with $\xi_{i} \geq 0$.
- If $\Phi^{\prime} \Xi \Phi$ is invertible, the entire theory and algorithms of projected equations generalizes to semi-norm projected equations [including multistep methods such as LSTD/LSPE/TD $(\lambda)]$.
- Reference: Yu and Bertsekas, "Weighted Bellman Equations and their Applications in Approximate Dynamic Programming," MIT Report, 2012.


## DISTRIBUTED AGGREGATION I

- We consider decomposition/distributed solution of large-scale discounted DP problems by hard aggregation.
- Partition the original system states into subsets $S_{1}, \ldots, S_{m}$.
- Distributed VI Scheme: Each subset $S_{\ell}$
- Maintains detailed/exact local costs
$J(i)$ for every original system state $i \in S_{\ell}$ using aggregate costs of other subsets
- Maintains an aggregate cost $R(\ell)=\sum_{i \in S_{\ell}} d_{\ell i} J(i)$
- Sends $R(\ell)$ to other aggregate states
- $J(i)$ and $R(\ell)$ are updated by VI according to

$$
J_{k+1}(i)=\min _{u \in U(i)} H_{\ell}\left(i, u, J_{k}, R_{k}\right), \quad \forall i \in S_{\ell}
$$

with $R_{k}$ being the vector of $R(\ell)$ at time $k$, and
$H_{\ell}(i, u, J, R)=\sum_{j=1}^{n} p_{i j}(u) g(i, u, j)+\alpha \sum_{j \in S_{\ell}} p_{i j}(u) J(j)$

$$
+\alpha \sum_{j \in S_{\ell^{\prime}}, \ell^{\prime} \neq \ell} p_{i j}(u) R\left(\ell^{\prime}\right)
$$

## DISTRIBUTED AGGREGATION II

- Can show that this iteration involves a supnorm contraction mapping of modulus $\alpha$, so it converges to the unique solution of the system of equations in $(J, R)$

$$
\begin{aligned}
J(i)=\min _{u \in U(i)} H_{\ell}(i, u, J, R), & R(\ell)=\sum_{i \in S_{\ell}} d_{\ell i} J(i), \\
& \forall i \in S_{\ell}, \ell=1, \ldots, m .
\end{aligned}
$$

- This follows from the fact that $\left\{d_{\ell i} \mid i=\right.$ $1, \ldots, n\}$ is a probability distribution.
- View these equations as a set of Bellman equations for an "aggregate" DP problem. The difference is that the mapping $H$ involves $J(j)$ rather than $R(x(j))$ for $j \in S_{\ell}$.
- In an asynchronous version of the method, the aggregate costs $R(\ell)$ may be outdated to account for communication "delays" between aggregate states.
- Convergence can be shown using the general theory of asynchronous distributed computation, briefly described in the 2nd lecture (see the text).

MIT OpenCourseWare
http://ocw.mit.edu

### 6.231 Dynamic Programming and Stochastic Control

Fall 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

