EE365: Linear Quadratic Stochastic Control

Continuous state Markov decision process

Affine and quadratic functions

Linear quadratic Markov decision process

Continuous state Markov decision process

Continuous state Markov decision problem

- dynamics: $x_{t+1} = f_t(x_t, u_t, w_t)$
- $x_0, w_0, w_1, \ldots$ independent
- stage cost: $g_t(x_t, u_t, w_t)$
- state feedback policy: $u_t = \mu_t(x_t)$
- choose policy to minimize

$$J = \mathbf{E}\left(\sum_{t=0}^{T-1} g_t(x_t, u_t, w_t) + g_T(x_T)\right)$$

- we consider the case $\mathcal{X} = \mathbb{R}^n$, $\mathcal{U} = \mathbb{R}^m$

Continuous state Markov decision problem

- many (mostly mathematical) pathologies can occur in this case
  - but not in the special case we'll consider
- a basic issue: how do you even *represent* the functions f_t, g_t, and μ_t?
  - for n and m very small (say, 2 or 3) we can use *gridding*
  - we can give the coefficients in some (dense) basis of functions
  - most generally, we assume we have a method to compute function values, given the arguments
  - exponential growth that occurs in gridding called *curse of dimensionality*

Continuous state Markov decision problem: Dynamic programming

- set $v_T(x) = g_T(x)$
- for $t = T-1, \ldots, 0$

$$\mu_t(x) \in \underset{u}{\operatorname{argmin}}\, \mathbf{E}\left(g_t(x,u,w_t) + v_{t+1}(f_t(x,u,w_t))\right)$$
$$v_t(x) = \mathbf{E}\left(g_t(x,\mu_t(x),w_t) + v_{t+1}(f_t(x,\mu_t(x),w_t))\right)$$

- this gives value functions and optimal policy, *in principle only*
- but you can't in general represent, much less compute, v_t or μ_t

Continuous state Markov decision problem: Dynamic programming

for DP to be tractable, f_t and g_t need to have special form for which we can

- represent v_t, μ_t in some tractable way
- carry out expectation and minimization in DP recursion

one of the few situations where this holds: *linear quadratic problems*

- f_t is an affine function of x_t, u_t ('linear dynamical system')
- g_t are convex quadratic functions of x_t, u_t

Linear quadratic problems

for linear quadratic problems

- value functions $v_t^\star$ are quadratic
- hence representable by their coefficients
- we can carry out the expectation and the minimization in DP recursion explicitly *using linear algebra*
- optimal policy functions are affine: $\mu_t^\star(x) = K_t x + l_t$
- we can compute the coefficients K_t and l_t explicitly

in other words:
we can solve linear quadratic stochastic control problems in practice

Affine and quadratic functions

Affine functions

- $f : \mathbb{R}^p \to \mathbb{R}^q$ is affine if it has the form
$$f(x) = Ax + b$$
i.e., it is a linear function plus a constant
- a linear function is special case, with $b = 0$
- affine functions closed under sum, scalar multiplication, composition (with explicit formulas for coefficients in each case)

Quadratic function

- $f : \mathbb{R}^n \to \mathbb{R}$ is quadratic if it has the form

$$f(x) = \tfrac{1}{2}x^\mathsf{T}Px + q^\mathsf{T}x + \tfrac{1}{2}r$$

  with $P = P^\mathsf{T} \in \mathbb{R}^{n\times n}$ (the $\frac{1}{2}$ on r is for convenience)

- often write as quadratic form in $(x, 1)$:

$$f(x) = \frac{1}{2}\begin{bmatrix} x \\ 1 \end{bmatrix}^\mathsf{T} \begin{bmatrix} P & q \\ q^\mathsf{T} & r \end{bmatrix} \begin{bmatrix} x \\ 1 \end{bmatrix}$$

- special cases:
  - quadratic form: $q = 0$, $r = 0$
  - affine (linear) function: $P = 0$ $(P = 0$, $r = 0)$
  - constant: $P = 0$, $q = 0$
- uniqueness: $f(x) = \tilde{f}(x) \iff P = \tilde{P},\ q = \tilde{q},\ r = \tilde{r}$

Calculus of quadratic functions

- quadratic functions on $\mathbb{R}^n$ form a vector space of dimension

$$\frac{n(n+1)}{2} + n + 1$$

- *i.e.*, they are closed under addition, scalar multiplication

Composition of quadratic and affine functions

- suppose
  - $f(z) = \frac{1}{2}z^{\mathsf{T}}Pz + q^{\mathsf{T}}z + \frac{1}{2}r$ is quadratic function on $\mathbb{R}^m$
  - $g(x) = Ax + b$ is affine function from $\mathbb{R}^n$ into $\mathbb{R}^m$
- then composition $h(x) = (f \circ g)(x) = f(Ax + b)$ is quadratic
- write $h(x)$ as

$$\frac{1}{2}\begin{bmatrix} x \\ 1 \end{bmatrix}^{\mathsf{T}} \left(\begin{bmatrix} A & b \\ 0 & 1 \end{bmatrix}^{\mathsf{T}} \begin{bmatrix} P & q \\ q^{\mathsf{T}} & r \end{bmatrix} \begin{bmatrix} A & b \\ 0 & 1 \end{bmatrix} \right) \begin{bmatrix} x \\ 1 \end{bmatrix}$$

- so matrix multiplication gives us the coefficient matrix of h

Convexity and nonnegativity of a quadratic function

- f is convex (graph does not curve down) if and only if $P \geq 0$ (matrix inequality)
- f is strictly convex (graph curves up) if and only if $P > 0$ (matrix inequality)
- f is nonnegative (*i.e.*, $f(x) \geq 0$ for all x) if and only if

$$\begin{bmatrix} P & q \\ q^\mathsf{T} & r \end{bmatrix} \geq 0$$

- $f(x) > 0$ if and only if matrix inequality is strict
- nonnegative $\Rightarrow$ convex

Checking convexity and nonnegativity

- we can check convexity or nonnegativity in $O(n^3)$ operations by eigenvalue decomposition, Cholesky factorization, ...
- composition with affine function preserves convexity, nonnegativity:

$$f \text{ convex}, \; g \text{ affine} \implies f \circ g \text{ convex}$$

- linear combination of convex quadratics, with nonnegative coefficients, is convex quadratic
- if $f(x, w)$ is convex quadratic in x for each w (a random variable) then

$$g(x) = \mathop{\mathbf{E}}_{w} f(x, w)$$

  is convex quadratic (*i.e.*, convex quadratics closed under expectation)

Minimizing a quadratic

- if f is not convex, then $\min_x f(x) = -\infty$
- otherwise, x minimizes f if and only if $\nabla f(x) = Px + q = 0$
- for $q \notin \mathbf{range}(P)$, $\min_x f(x) = -\infty$
- for $P > 0$, unique minimizer is $x = -P^{-1}q$
- minimum value is
$$\min_x f(x) = -\frac{1}{2}q^\mathsf{T} P^{-1} q + \frac{1}{2}r$$
(a concave quadratic function of q)
- for case $P \geq 0$, $q \in \mathbf{range}(P)$, replace P^{-1} with $P^\dagger$

Partial minimization of a quadratic

- suppose f is a quadratic function of (x, u), convex in u
- then the partial minimization function

$$g(x) = \min_u f(x, u)$$

  is a quadratic function of x; if f is convex, so is g
- the minimizer $\mathop{\mathrm{argmin}}_u f(x, u)$ is an affine function of x
- minimizing a convex quadratic function over some variables yields a convex quadratic function of the remaining ones
- *i.e.*, convex quadratics closed under partial minimization

Partial minimization of a quadratic

- let's take
$$f(x,u) = \frac{1}{2} \begin{bmatrix} x \\ u \\ 1 \end{bmatrix}^{\mathsf{T}} \begin{bmatrix} P_{xx} & P_{xu} & q_x \\ P_{ux} & P_{uu} & q_u \\ q_x^{\mathsf{T}} & q_u^{\mathsf{T}} & r \end{bmatrix} \begin{bmatrix} x \\ u \\ 1 \end{bmatrix}$$
with $P_{uu} > 0$, $P_{ux} = P_{xu}^{\mathsf{T}}$
- minimizer of f over u satisfies
$$0 = \nabla_u f(x,u) = P_{uu}u + P_{ux}x + q_u$$
so $u = -P_{uu}^{-1}(P_{ux}x + q_u)$ is an affine function of x

Partial minimization of a quadratic

- substituting u into expression for f gives

$$g(x) = \frac{1}{2} \begin{bmatrix} x \\ 1 \end{bmatrix}^{\mathsf{T}} \begin{bmatrix} P_{xx} - P_{xu}P_{uu}^{-1}P_{ux} & q_x - P_{xu}P_{uu}^{-1}q_u \\ q_x^{\mathsf{T}} - q_u^{\mathsf{T}}P_{uu}^{-1}P_{ux} & r - q_u^{\mathsf{T}}P_{uu}^{-1}q_u \end{bmatrix} \begin{bmatrix} x \\ 1 \end{bmatrix}$$

- $P_{xx} - P_{xu}P_{uu}^{-1}P_{ux}$ is the Schur complement of P w.r.t. u
- $P_{xx} - P_{xu}P_{uu}^{-1}P_{ux} \geq 0$ if $P \geq 0$
- or simpler: g is composition of f with affine function $x \mapsto (x, u)$

$$\begin{bmatrix} x \\ u \end{bmatrix} = \begin{bmatrix} I \\ -P_{uu}^{-1}P_{ux} \end{bmatrix} x + \begin{bmatrix} 0 \\ -P_{uu}^{-1}q_u \end{bmatrix}$$

- we already know how to form composition quadratic (affine), and the result is convex

Summary

convex quadratics are closed under

- addition
- expectation
- pre-composition with an affine function
- partial minimization

in each case, we can explicitly compute the coefficients of the result using linear algebra

Linear quadratic Markov decision process

(Random) linear dynamical system

- dynamics $x_{t+1} = f_t(x_t, u_t, w_t) = A_t(w_t)x_t + B_t(w_t)u_t + c_t(w_t)$
- for each w_t, f_t is affine in (x_t, u_t)
- $x_0, w_0, w_1, \ldots$ are independent
- $A_t(w_t) \in \mathbb{R}^{n \times n}$ is dynamics matrix
- $B_t(w_t) \in \mathbb{R}^{n \times m}$ is input matrix
- $c_t(w_t) \in \mathbb{R}^n$ is offset

Linear quadratic stochastic control problem

- stage cost $g_t(x_t, u_t, w_t)$ is convex quadratic in (x_t, u_t) for each w_t
- choose policy $u_t = \mu_t(x_t)$ to minimize objective

$$J = \mathbf{E}\left(\sum_{t=0}^{T-1} g_t(x_t, u_t, w_t) + g_T(x_T)\right)$$

Dynamic programming

- set $v_T(x) = g_T(x)$
- for $t = T-1, \ldots, 0$,

$$\mu_t(x) \in \underset{u}{\operatorname{argmin}} \, \mathbf{E} \left(g_t(x, u, w_t) + v_{t+1}(f_t(x, u, w_t)) \right)$$
$$v_t(x) = \mathbf{E} \left(g_t(x, \mu_t(x), w_t) + v_{t+1}(f_t(x, \mu_t(x), w_t)) \right)$$

- all v_t are convex quadratic, and all μ_t are affine
- this gives value functions and optimal policy, *explicitly*

Dynamic programming

we show v_t are convex quadratic by (backward) induction

- suppose $v_T, \ldots, v_{t+1}$ are convex quadratic
- since f_t is affine in (x, u), $v_{t+1}(f_t(x, u, w_t))$ is convex quadratic
- so $g_t(x, u, w_t) + v_{t+1}(f_t(x, u, w_t))$ is convex quadratic
- and so is its expectation over w_t
- partial minimization over u leaves convex quadratic of x, which is $v_t(x)$
- argmin is affine function of x, so optimal policy is affine

Linear equality constraints

- can add (deterministic) linear equality constraints on x_t, u_t into g_t, g_T:

$$g_t(x,u,w) = g_t^{\mathrm{quad}}(x,u,w) + \begin{cases} 0 & \text{if } F_t x + G_t u = h_t \\ \infty & \text{otherwise} \end{cases}$$

- everything still works:
  - v_t is convex quadratic, possibly with equality constraints
  - μ_t is affine
- reason: minimizing a convex quadratic over some variables, subject to equality constraints, yields a convex quadratic in remaining variables