


02465: Introduction to reinforcement learning and control

Linear-quadratic problems in control

Tue Herlau

DTU Compute, Technical University of Denmark (DTU)



DTU Compute
Department of Applied Mathematics and Computer Science

Lecture Schedule

Dynamical programming

- 1 The finite-horizon decision problem
31 January
- 2 Dynamical Programming
7 February
- 3 DP reformulations and introduction to Control
14 February
- 4 Discretization and PID control
21 February
- 5 Direct methods and control by optimization
28 February
- 6 **Linear-quadratic problems in control**
7 March
- 7 Linearization and iterative LQR
14 March

Reinforcement learning

- 8 Exploration and Bandits
21 March
- 9 Policy and value iteration
4 April
- 10 Monte-carlo methods and TD learning
11 April
- 11 Model-Free Control with tabular and linear methods
18 April
- 12 Eligibility traces and value-function approximations
25 April
- 13 Q-learning and deep-Q learning
2 May

Control

- 4 Discretization and PID control
21 February
- 5 Direct methods and control by optimization
28 February
- 6 **Linear-quadratic problems in control**
7 March
- 7 Linearization and iterative LQR
14 March

Syllabus: <https://02465material.pages.compute.dtu.dk/02465public>
Help improve lecture by giving feedback on DTU learn

Reading material:

- [Her24, Chapter 16]

Learning Objectives

- Linear-quadratic regulator (LQR)
- Derivation of the LQR from DP
- Applications and variations

Recap Practicals

- Project evaluations will be ready in about a week
- Programming evaluations see <https://02465material.pages.compute.dtu.dk/02465public/projects/project1.html>
- Part 2:
 - Less programming
 - A bit more emphasis on linear algebra; don't be afraid to write short answers if they are correct.
 - Be inspired by existing examples

Recap Useful linear algebra



- A matrix A is **positive semi-definite** if it is symmetric and $x^T A x \geq 0$ for all x
 - This means A behaves like a positive number: $ax^2 \geq 0$.
- if A is a symmetric matrix then:

$$\frac{1}{2} x^T A x + b^T x = \frac{1}{2} (x + A^{-1}b)^T A (x + A^{-1}b) - \frac{1}{2} b^T A^{-1}b$$

- This allows us to quickly find minimum

Recap Dynamical programming algorithm

The Dynamical Programming algorithm

For every initial state x_0 , the optimal cost $J^*(x_0)$ is equal to $J_0(x_0)$, and optimal policy π^* is $\pi^* = \{\mu_0, \dots, \mu_{N-1}\}$, computed by the following algorithm, which proceeds backward in time from $k = N$ to $k = 0$ and for each $x_k \in S_k$ computes

$$J_N(x_N) = g_N(x_N) \quad (1)$$

$$J_k(x_k) = \min_{u_k \in \mathcal{A}_k(x_k)} \mathbb{E} \{g_k(x_k, u_k, w_k) + J_{k+1}(f_k(x_k, u_k, w_k))\} \quad (2)$$

$$\mu_k(x_k) = u_k^* \quad (u_k^* \text{ is the } u_k \text{ which minimizes the above expression}). \quad (3)$$

Recap
Assumptions today



- For $k = 0, 1, \dots, N - 1$

$$x_{k+1} = f_k(x_k, u_k, w_k) = A_k x_k + B_k u_k,$$

$$g_k(x_k, u_k, w_k) = \frac{1}{2} x_k^T Q_k x_k + \frac{1}{2} u_k^T R_k u_k,$$

$$g_N(x_k) = \frac{1}{2} x_k^T Q_N x_k$$

- Note:** This is not the most general case, but will illustrate the main ideas

Linear Quadratic Regulator
Apply dynamical programming!



- Define $V_N \equiv Q_N$ and initialize:

$$J_N^*(x_N) = \frac{1}{2} x_N^T Q_N x_N = \frac{1}{2} x_N^T V_N x_N$$

- DP iteration (start at $k = N - 1$)

$$J_k(x_k) = \min_{u_k} \mathbb{E}_{w_k} \{g_k(x_k, u_k, w_k) + J_{k+1}(f_k(x_k, u_k, w_k))\}$$

- Remember to store optimal u_k^* as $\pi_k(x_k) = u_k^*$

Linear Quadratic Regulator
LQR, simplified form



DP solution gives the controller:

- $V_N = Q_N$
- $L_k = -(R_k + B_k^T V_{k+1} B_k)^{-1} (B_k^T V_{k+1} A_k)$
- $V_k = Q_k + L_k^T R_k L_k + (A_k + B_k L_k)^T V_{k+1} (A_k + B_k L_k)$
- $u_k^* = L_k x_k$
- $J_k^*(x_k) = \frac{1}{2} x_k^T V_k x_k$

Linear Quadratic Regulator
Double Integrator Example



- True dynamics

$$\dot{x}(t) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t) \quad (4)$$

- Euler discretization** using $\Delta = 1$ System evolves according to:

$$x_{k+1} = \underbrace{\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}}_{=A} x_k + \underbrace{\begin{bmatrix} 0 \\ 1 \end{bmatrix}}_{=B} u_k$$

- Cost function:

$$J(x_0) = \sum_{k=0}^N \frac{1}{2\rho} x_{k,1}^2 + \sum_{k=0}^{N-1} \frac{1}{2} u_k^2$$

- Can be put into standard form using matrices/start position:

$$Q_k = Q_N = \begin{bmatrix} \frac{1}{\rho} & 0 \\ 0 & 0 \end{bmatrix} \quad R = 1$$


Linear Quadratic Regulator
Exponential integrator



- Apply discrete LQR
- Simulate starting in $x_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ using policy

$$\pi_k(x_k) = L_k x_k$$

- What about the true system $\dot{x}(t) = f(x, u)$?



Linear Quadratic Regulator
Quiz: LQR



Consider a (generic) LQR problem of the form:

$$x_{k+1} = A x_k + B u_k \quad (5)$$

$$\text{cost} = \sum_{k=0}^{N-1} \frac{1}{2} x_k^T Q x_k + \frac{1}{2} R_0 u_k^T u_k \quad (6)$$

Where $R_0 > 0$ is a constant. After LQR, the controller selects actions using $u_k = L_k x_k$. What do you think typically happens with the matrix L_k when $R_0 \rightarrow \infty$ (**very big** R_0)


- The entries in L_k becomes very small, negative numbers
- The entries in L_k becomes very big, positive numbers
- It is not possible to say anything about the typical case
- The entries in L_k gets closer to zero
- Don't know.

Linear Quadratic Regulator Double integrator example



• **Blue:** LQR using Euler $\mathbf{x}_{k+1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \mathbf{u}_k$

• **Red:** LQR using Exponential $\mathbf{x}_{k+1} = e^{A\Delta} \mathbf{x}_k + A^{-1} (e^{A\Delta} - I) B \mathbf{u}_k$



- LQR is optimal in discrete problem
- Discrete controller can be bad in real problem (always check!)
- Always use EI for linear dynamics

Linear Quadratic Regulator Example: The locomotive



Steer locomotive (starting at $x = -1$) to goal ($x^* = 0$)

$$\ddot{x}(t) = \frac{1}{m} u(t) \quad (7)$$

Can be re-written as:

$$\dot{\mathbf{x}} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \mathbf{x} + \begin{bmatrix} 0 \\ \frac{1}{m} \end{bmatrix} u \quad (8)$$


Discretized to $\mathbf{x}_{k+1} = A \mathbf{x}_k + B \mathbf{u}_k$.

Linear Quadratic Regulator Locomotive: PID and LQR



$$e_k = x^* - x_k$$

$$u_k = e_k K_p + K_d \frac{e_k - e_{k-1}}{\Delta}$$



- Alternatively: Use a cost function $\sum_k \mathbf{x}_k^T Q \mathbf{x}_k + \mathbf{u}_k^T \mathbf{u}_k$ and use LQR!

📄 lecture_04_pid_d.py 📄 lecture_06_lqr_locomotive.py

Linear Quadratic Regulator Planning on an infinite horizon



Recall LQR has the form:

1. $V_N = Q_N$
2. $L_k = -(R_k + B_k^T V_{k+1} B_k)^{-1} (B_k^T V_{k+1} A_k)$
3. $V_k = Q_k + L_k^T R_k L_k + (A_k + B_k L_k)^T V_{k+1} (A_k + B_k L_k)$
4. $\mathbf{u}_k^* = L_k \mathbf{x}_k$
5. $J_k^*(\mathbf{x}_k) = \frac{1}{2} \mathbf{x}_k^T V_k \mathbf{x}_k$

- What happens if we repeat step 2 and 3 *many* times?
- The method will converge: $L_k \rightarrow L$
 - Select actions $\mathbf{u}_k = L \mathbf{x}_k$ ("plan until convergence")
- If you think about it, this corresponds to planning on $N \rightarrow \infty$ horizon.
- **This is quite popular in control theory; what we will do in RL.**

Linear Quadratic Regulator Observations



- The cost term $\frac{1}{2} \mathbf{x}^T Q \mathbf{x} + \frac{1}{2} \mathbf{u}^T R \mathbf{u}$ is **smallest** when $\mathbf{x} = \mathbf{u} = \mathbf{0}$
- Implies that LQR will control system to state $\mathbf{x} = \mathbf{u} = \mathbf{0}$
- Suppose we want to drive system towards $\mathbf{x}_g, \mathbf{u}_g$?
 - Use $c(\mathbf{x}, \mathbf{u}) = \frac{1}{2} (\mathbf{x} - \mathbf{x}_g)^T Q (\mathbf{x} - \mathbf{x}_g) + \frac{1}{2} (\mathbf{u} - \mathbf{u}_g)^T R (\mathbf{u} - \mathbf{u}_g)$
- more generally assume

$$c_k(\mathbf{x}_k, \mathbf{u}_k) = \frac{1}{2} \mathbf{x}_k^T Q_k \mathbf{x}_k + \frac{1}{2} \mathbf{u}_k^T R_k \mathbf{u}_k + \mathbf{u}_k^T H_k \mathbf{x}_k + \mathbf{q}_k^T \mathbf{x}_k + \mathbf{r}_k^T \mathbf{u}_k + q_k \quad (9)$$

$$c_N(\mathbf{x}_k) = \frac{1}{2} \mathbf{x}_k^T Q_N \mathbf{x}_k + \mathbf{q}_N^T \mathbf{x}_k + q_N \quad (10)$$

and dynamics

$$\mathbf{x}_{k+1} = A_k \mathbf{x}_k + B_k \mathbf{u}_k + \mathbf{d}_k$$

Linear Quadratic Regulator General discrete LQR algorithm

How to start living in luxury and never work again!




1. $V_N = Q_N; \mathbf{v}_N = \mathbf{q}_N; \mathbf{v}_N = \mathbf{q}_N$
2.
$$L_k = -S_{uu,k}^{-1} S_{ux,k} \quad S_{u,k} = \mathbf{r}_k + B_k^T \mathbf{v}_{k+1} + B_k^T V_{k+1} \mathbf{d}_k$$

$$l_k = -S_{uu,k}^{-1} S_{ux,k} \quad S_{uu,k} = R_k + B_k^T V_{k+1} B_k$$

$$S_{ux,k} = H_k + B_k^T V_{k+1} A_k$$
3.
$$V_k = Q_k + A_k^T V_{k+1} A_k - L_k^T S_{uu,k} L_k$$

$$\mathbf{v}_k = \mathbf{q}_k + A_k^T (\mathbf{v}_{k+1} + V_{k+1} \mathbf{d}_k) + S_{ux,k}^T l_k$$

$$\mathbf{v}_k = \mathbf{v}_{k+1} + \mathbf{q}_k + \mathbf{d}_k^T \mathbf{v}_{k+1} + \frac{1}{2} \mathbf{d}_k^T V_{k+1} \mathbf{d}_k + \frac{1}{2} l_k^T S_{ux,k}$$
4. $\mathbf{u}_k^* = l_k + L_k \mathbf{x}_k$ Doctors hate this one weird trick!
5. $J_k(\mathbf{x}_k) = \frac{1}{2} \mathbf{x}_k^T V_k \mathbf{x}_k + \mathbf{v}_k^T \mathbf{x}_k + v_k$ $V_k \leftarrow \frac{1}{2} (V_k^T + V_k)$



$$\begin{bmatrix} \dot{u} \\ \dot{w} \\ \dot{q} \\ \dot{\theta} \end{bmatrix} = \underbrace{\begin{bmatrix} -0.003 & 0.039 & 0 & -0.322 \\ -0.065 & -0.319 & 7.74 & 0 \\ 0.02 & -0.101 & -0.429 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} u - u_w \\ w - w_w \\ q \\ \theta \end{bmatrix}}_x + \underbrace{\begin{bmatrix} 0.01 & 1 \\ -0.18 & -0.04 \\ -1.16 & 0.598 \\ 0 & 0 \end{bmatrix}}_B \underbrace{\begin{bmatrix} c \\ t \end{bmatrix}}_u$$

$$\begin{bmatrix} y_1(t) \\ y_2(t) \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 7.74 \end{bmatrix}}_P \begin{bmatrix} u(t) - u_w(t) \\ w(t) - w_w(t) \\ q(t) \\ \theta(t) \end{bmatrix}$$

- y_1 and y_2 corresponds to the airspeed and climb rate.
- **Start:** $x = 0$ (steady flight)

• **Want:** Airspeed of 10: $y^* = \begin{bmatrix} 10 \\ 0 \end{bmatrix}$

- Write dynamics as $\dot{x} = Ax + Bu$
- Introduce cost function:


$$\int_0^{t_F} \left(\frac{1}{2} (y - y^*)^\top (y - y^*) + \frac{1}{2} u^\top u \right) dt$$

- Discretize dynamics using Exponential Integration to get $x_{k+1} = \bar{A}x_k + \bar{B}u_k$
- Discretize cost to get one of the form

$$\sum_{k=0}^{\infty} \frac{1}{2} x_k^\top Q x_k + q x_k + q_0 + \frac{1}{2} u_k^\top R u_k$$

- Apply LQR!

- Control law $u_k = Lx_k$



Left: airspeed and climb rate. **Right:** Elevator and throttle
Why does the output adjust quickly but fail to get entirely to the goal y^* ?

- Something bad happened to the dynamics with the exponential integration
 - The explanation has to do with planning on a finite horizon
 - The explanation is that R in $u_k^\top R u_k$ should be bigger
- d. Don't know.

21 DTU Compute

- Consider the case where there is additive Gaussian noise:

$$x_{k+1} = A_k x_k + B_k u_k + \omega_k$$

- We can still solve the problem, and (amazingly!) the noise has **no influence** on the control law

$$u_k = L_k x_k$$

- LQR is robust to noise

- Stability/controllability of LQR?
 - **Important subject which we ignore**
- What if matrices A_k, B_k are random?
 - **This too can be solved[Ber05, Chapter 4]**
- What about partial observation?
 - **i.e. assume we observe $o_k = D_k x_k$ [Ber05, Chapter 4]**
- What about constraints? What if we know $u_L \leq u_k \leq u_B$?
- Euler integration is often not ideal.
 - **Alternatives including error analysis**

- D.P. Bertsekas. *Dynamic Programming and Optimal Control*. Number v. 1 in Athena Scientific optimization and computation series. Athena Scientific, 2005.
- Tue Herlau. *Sequential decision making*. (Freely available online), 2024.