Optimization-Based Control

Richard M. Murray
Control and Dynamical Systems
California Institute of Technology

DRAFT v1.2a, 14 January 2008

Chapter 2
Optimal Control

This set of notes expands on Chapter 6 of *Feedback Systems* by Åström and Murray (ÅM08), which introduces the concepts of reachability and state feedback. We also expand on topics in Section 7.5 of ÅM08 in the area of feedforward compensation. Beginning with a review of optimization, we introduce the notion of Lagrange multipliers and provide a summary of the Pontryagin's maximum principle. Using these tools we derive the linear quadratic regulator for linear systems and describe its usage.

Prerequisites. Readers should be familiar with modeling of input/output control systems using differential equations, linearization of a system around an equilibrium point and state space control of linear systems, including reachability and eigenvalue assignment.

2.1 Review: Optimization

Consider first the problem of finding the maximum of a smooth function $F : \mathbb{R}^n \to \mathbb{R}$. That is, we wish to find a point $x^* \in \mathbb{R}^n$ such that $F(x^*) \geq F(x)$ for all $x \in \mathbb{R}^n$. A necessary condition for x^* to be a maximum is that the gradient of the function be zero at x^*,

$$\frac{\partial F}{\partial x}(x^*) = 0.$$

Figure 2.1 gives a graphical interpretation of this condition. Note that these are *not* sufficient conditions; the points x_1 and x_2 and x^* in the figure all

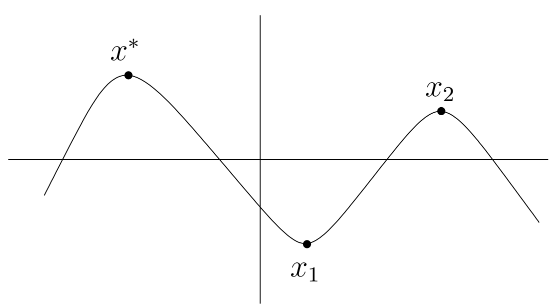

Figure 2.1: Optimization of functions. The maximum of a function occurs at a point where the gradient is zero.

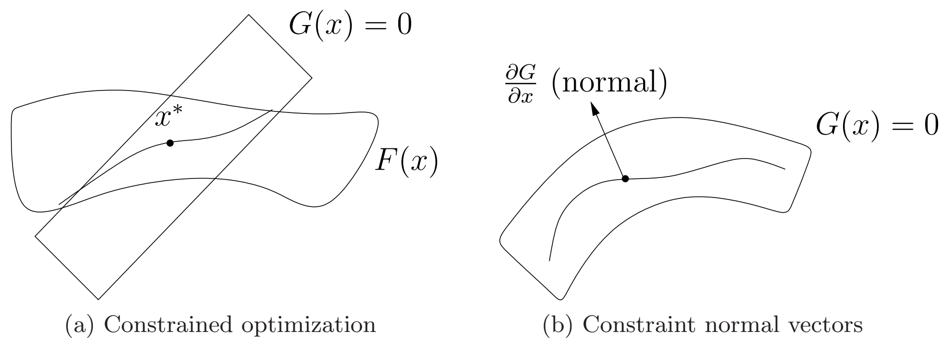

Figure 2.2: Optimization with constraints. (a) We seek a point x^* that maximizes $F(x)$ while lying on the surface $G(x) = 0$. (b) We can parameterize the constrained directions by computing the gradient of the constraint G.

satisfy the necessary condition but only one is the (global) maximum.

The situation is more complicated if constraints are present. Let $G_i : \mathbb{R}^n \to \mathbb{R}$, $i = 1, \dots, k$ be a set of smooth functions with $G_i(x) = 0$ representing the constraints. Suppose that we wish to find $x^* \in \mathbb{R}^n$ such that $G_i(x^*) = 0$ and $F(x^*) \geq F(x)$ for all $x \in \{x \in \mathbb{R}^n : G_i(x) = 0, i = 1, \dots, k\}$. This situation can be visualized as constraining the point to a surface (defined by the constraints) and searching for the maximum of the cost function along this surface, as illustrated in Figure 2.2.

A necessary condition for being at a maximum is that there are no directions tangent to the constraints that also increase the cost. The normal directions to the surface are spanned by $\partial G_i / \partial x$, as shown in Figure **??**. A necessary condition is that the gradient of F is spanned by vectors that are normal to the constraints, so that the only directions that increase the cost violate the constraints. We thus require that there exist scalars λ_i, $i = 1, \dots, k$ such that

$$\frac{\partial F}{\partial x}(x^*) + \sum_{i=1}^{k} \lambda_i \frac{\partial G_i}{\partial x}(x^*) = 0.$$

If we let $G = \begin{bmatrix} G_1 & G_2 & \dots & G_k \end{bmatrix}^T$, then we can write this condition as

$$\frac{\partial F}{\partial x} + \lambda^T \frac{\partial G}{\partial x} = 0$$

the term $\frac{\partial F}{\partial x}$ is the usual (gradient) optimality condition while the term $\frac{\partial G}{\partial x}$ is used to "cancel" the gradient in the directions normal to the constraint.

An alternative condition can be derived by modifying the cost function to incorporate the constraints. Defining $\widetilde{F} = F + \sum \lambda_i G_i$, the necessary condition becomes

$$\frac{\partial \widetilde{F}}{\partial x}(x^*) = 0.$$

The scalars λ_i are called *Lagrange multipliers*. Minimize $\widetilde{F}$ is equivalent to the optimization given by

$$\min_x \left(F(x) + \lambda^T G(x)\right).$$

The variables λ can be regarded as free variables, which implies that need to choose x such that $G(x) = 0$. Otherwise, we could choose λ to generate a large cost.

Example 2.1 Two free variables with a constraint
Consider the cost function given by

$$F(x) = F_0 - (x_1 - a)^2 - (x_2 - b)^2,$$

which has an unconstrained maximum at $x = (a, b)$. Suppose that we add a constraint $G(x) = 0$ given by

$$G(x) = x_1 - x_2.$$

With this constrain, we seek to optimize F subject to $x_1 = x_2$. Although in this case we could easily do this by simple substitution, we instead carry out the more general procedure using Lagrange multipliers.

The augmented cost function is given by

$$\tilde{F}(x) = F_0 - (x_1 - a)^2 - (x_2 - b)^2 + \lambda(x_1 - x_2),$$

where λ is the Lagrange multiplier for the constraint. Taking the derivative of F, we have

$$\frac{\partial F}{\partial x} = \begin{bmatrix} -2x_1 + 2a + \lambda & -2x_2 + 2b - \lambda \end{bmatrix}.$$

Setting each of these equations equal to zero, we have that at the maximum

$$x_1^* = a + \lambda/2, \qquad x_2^* = b - \lambda/2.$$

The remaining equation that we need is the constraint, which requires that $x_1^* = x_2^*$. Using these three equations, we see that $\lambda^* = b - a$ and we have

$$x_1^* = \frac{a+b}{2}, \qquad x_2^* = \frac{a+b}{2}.$$

To verify the geometric view described above, note that the gradients of F and G are given by

$$\frac{\partial F}{\partial x} = \begin{bmatrix} -2x_1 + 2a & -2x_2 + 2b \end{bmatrix}, \qquad \frac{\partial G}{\partial x} = \begin{bmatrix} 1 & -1 \end{bmatrix}.$$

At the optimal value of the (constrained) optimization, we have

$$\frac{\partial F}{\partial x} = \begin{bmatrix} a - b & b - a \end{bmatrix}, \qquad \frac{\partial G}{\partial x} = \begin{bmatrix} 1 & -1 \end{bmatrix}.$$

Although the derivative of F is not zero, it is pointed in a direction that is normal to the constraint, and hence we cannot decrease the cost while staying on the constraint surface. ∇

We have focused on finding the maximum of a function. We can switch back and forth between max and min by simply negating the cost function:

$$\max_x F(x) = \min_x \big(-F(x)\big)$$

We see that the conditions that we have derived are independent of the sign of F since they only depend on the gradient begin zero in approximate directions. Thus finding x^* that satisfies the conditions corresponds to finding an *extremum* for the function.

Very good software is available for solving optimization problems numerically of this sort. The NPSOL and SNOPT libraries are available in FORTRAN (and C). In MATLAB, the fmin function can be used to solve a constrained optimization problem.

2.2 Optimal Control of Systems

Consider now the *optimal control problem*:

$$\min_u \underbrace{\int_0^T L(x,u)\,dt}_{\text{integrated cost}} + \underbrace{V\big(x(T),u(T)\big)}_{\text{final cost}}$$

subject to the constraint

$$\dot{x} = f(x,u) \qquad x \in \mathbb{R}^n,\ u \in \mathbb{R}^m.$$

Abstractly, this is a constrained optimization problem where we seek a *feasible trajectory* $(x(t), u(t))$ that minimizes the cost function

$$J(x,u) = \int_0^T L(x,u)\,dt + V\big(x(T),u(T)\big).$$

More formally, this problem is equivalent to the "standard" problem of minimizing a cost function $J(x,u)$ where $(x,u) \in L_2[0,T]$ (the set of square integral functions) and $h(z) = \dot{x}(t) - f(x(t),u(t)) = 0$ models the dynamics.

There are many variations and special cases of the optimal control problem. We mention a few here:

Infinite Horizon. if we let $T = \infty$ and set $V = 0$, then we seek to optimize a cost function over all time. This is called the *infinite horizon* optimal control problem, versus the *finite horizon* problem with $T < \infty$.

Linear Quadratic. If the dynamical system is linear and the cost function is quadratic, we obtain the *linear quadratic* optimal control problem:

$$\dot{x} = Ax + Bu \qquad J = \int_0^T \left(x^T Q x + u^T R u\right)\, dt + x^T(T) P_1 x(T).$$

In this formulation, $Q \geq 0$ penalizes state error (assumes $x_d = 0$), $R > 0$ penalizes the input (*must* be positive definite) and $P_1 > 0$ penalizes terminal state.

Terminal Constraints. It is often convenient to ask that the final value of the trajectory, denoted x_f, be specified. We can do this by requiring that $x(T) = x_f$ or by using a more general form of constraint:

$$\psi_i(x(T)) = 0, \qquad i = 1, \ldots, q.$$

The fully constrained case is obtained by setting $q = n$ and defining $\psi_i(x(T)) = x_i(T) - x_{i,f}$.

Time Optimal. If we constrain the terminal condition to $x(T) = x_f$, let the terminal time T be free (so that we can optimize over it) and choose $L(x, u) = 1$, we can find the *time-optimal* trajectory between an initial and final condition. This problem is usually only well-posed if we additionally constrain the inputs u to be bounded.

A very general set of conditions are available for the optimal control problem that captures most of these special cases in a unifying framework. Consider a nonlinear system

$$\begin{aligned} &\dot{x} = f(x, u) && x = \mathbb{R}^n \\ &x(0) \text{ given} && u \in \Omega \subset \mathbb{R}^p \end{aligned}$$

where $f(x, u) = (f_1(x, u), \ldots f_n(x, u)) : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$. We wish to minimize a cost function J with terminal constraints:

$$J = \int_0^T L(x, u)\, dt + V(x(T)), \qquad \psi(x(T)) = 0.$$

The function $\psi : \mathbb{R}^n \to \mathbb{R}^q$ gives a set of q terminal constraints. Analogous to the case of optimizing a function subject to constraints, we construct the *Hamiltonian*:

$$H = L + \lambda^T f = L + \sum \lambda_i f_i.$$

A set of necessary conditions for a solution to be optimal was derived by Pontryagin [PBGM62].

Theorem 2.1 (Maximum Principle)**.** *If (x^*, u^*) is optimal, then there exists $\lambda^*(t) \in \mathbb{R}^n$ and $\nu^* \in \mathbb{R}^q$ such that*

$$\dot{x}_i = \frac{\partial H}{\partial \lambda_i} \qquad -\dot{\lambda}_i = \frac{\partial H}{\partial x_i} \qquad \begin{aligned} &x(0) \textit{ given}, \ \psi(x(T)) = 0 \\ &\lambda(T) = \frac{\partial V}{\partial x}(x(T)) + \nu^T \frac{\partial \psi}{\partial x} \end{aligned}$$

and

$$H(x^*(t), u^*(t), \lambda^*(t)) \leq H(x^*(t), u, \lambda^*(t)) \quad \textit{for all} \quad u \in \Omega$$

The form of the optimal solution is given by the solution of a differential equation with boundary conditions. If $u = \operatorname{argmin} H(x, u, \lambda)$ exists, we can use this to choose the control law u and solve for the resulting feasible trajectory that minimizes the cost. The boundary conditions are given by the n initial states $x(0)$, the q terminal constraints on the state $\psi(x(T)) = 0$ and the $n - q$ final values for the Lagrange multipliers

$$\lambda(T) = \frac{\partial V}{\partial x}(x(T)) + \nu^T \frac{\partial \psi}{\partial x}.$$

In this last equation, ν is a free variable and so there are n equations in $n+q$ free variables, leaving $n - q$ constraints on $\lambda(T)$. In total, we thus have $2n$ boundary values.

The maximum principle is a very general (and elegant) theorem. It allows the dynamics to be nonlinear and the input to be constrained to lie in a set Ω, allowing the possibility of bounded inputs. If $\Omega = \mathbb{R}^m$ (unconstrained input) and H is differentiable, then a necessary condition for the optimal input is

$$\frac{\partial H}{\partial u} = 0.$$

We note that even though we are *minimizing* the cost, this is still usually called the maximum principle (artifact of history).

Sketch of proof. We follow the proof given by Lewis and Syrmos [LS95], omitting some of the details required for a fully rigorous proof. We use the method of Lagrange multipliers, augmenting our cost function by the dynamical constraints and the terminal constraints:

$$\begin{aligned}\tilde{J}(x(\cdot), u(\cdot)) &= J(x, u) + \int_0^T \lambda^T(t)\big(\dot{x}(t) - f(x, u)\big)\, dt + \nu^T \psi(x(T), u(T)) \\ &= \int_0^T \big(L(x, u) + \lambda^T(t)\big(\dot{x}(t) - f(x, u)\big)\, dt \\ &\qquad + V(x(T), u(T)) + \nu^T \psi(x(T), u(T)).\end{aligned}$$

Note that λ is a function of time, with each $\lambda(t)$ corresponding to the instantaneous constraint imposed by the dynamics. The integral over the interval $[0, T]$ plays the role of the sum of the finite constraints in the regular optimization.

Making use of the definition of the Hamiltonian, the augmented cost becomes

$$\tilde{J}(x(\cdot), u(\cdot)) = \int_0^T \big(H(x, u) - \lambda^T(t)\dot{x}\big)\, dt + V(x(T), u(T)) + \nu^T \psi(x(T), u(T)).$$

We can now "linearize" the cost function around the optimal solution $x(t) = x^*(t) + \delta x(t)$, $u(t) = u^*(t) + \delta u(t)$. Using Leibnitz's rule, we have □

2.3 Examples

To illustrate the use of the maximum principle, we consider a number of analytical examples. Additional examples are given in the exercises.

Example 2.2 Scalar linear system
Consider the optimal control problem for the system

$$\dot{x} = ax + bu, \tag{2.1}$$

where $x = \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. We wish to find a trajectory $(x(t), u(t))$ that minimizes the cost function

$$J = \tfrac{1}{2} \int_{t_0}^{t_f} u^2(t)\, dt + \tfrac{1}{2} c x^2(t_f),$$

where the terminal time t_f is given and $c > 0$ is a constant. This cost function balances the final value of the state with the input required to get to that position.

To solve the problem, we define the various elements used in the maximum principle. Our integrated and terminal costs are given by

$$L = \tfrac{1}{2} u^2(t) \qquad V = \tfrac{1}{2} c x^2(t_f).$$

We write the Hamiltonian of this system and derive the following expressions:

$$\begin{aligned} H &= L + \lambda f = \tfrac{1}{2} u^2 + \lambda(ax + bu) \\ \dot{\lambda} &= -\frac{\partial H}{\partial x} = -a\lambda, \qquad \lambda(t_f) = \frac{\partial V}{\partial x} = c x(t_f). \end{aligned}$$

This is a final value problem for a linear differential equation and the solution can be shown to be

$$\lambda(t) = c x(t_f) e^{a(t_f - t)}$$

The optimal control is given by

$$\frac{\partial H}{\partial u} = u + b\lambda = 0 \quad \Rightarrow \quad u^*(t) = -b\lambda(t) = -bcx(t_f) e^{a(t_f - t)}.$$

Substituting this control into the dynamics given by equation (2.1) yields a first-order ODE in x:

$$\dot{x} = ax - b^2 c x(t_f) e^{a(t_f - t)}.$$

This can be solved explicitly as

$$x^*(t) = x(t_o) e^{a(t - t_o)} + \frac{b^2 c}{2a} x^*(t_f) \left[e^{a(t_f - t)} - e^{a(t + t_f - 2t_o)} \right].$$

Setting $t = t_f$ and solving for $x(t_f)$ gives

$$x^*(t_f) = \frac{2a\, e^{a(t_f - t_o)} x(t_o)}{2a - b^2 c \left(1 - e^{2a(t_f - t_o)}\right)}$$

and finally we can write

$$u^*(t) = -\frac{2abc\, e^{a(2t_f - t_o - t)} x(t_o)}{2a - b^2 c \left(1 - e^{2a(t_f - t_o)}\right)}$$
$$x^*(t) = x(t_o) e^{a(t - t_o)} + \frac{b^2 c\, e^{a(t_f - t_o)} x(t_o)}{2a - b^2 c \left(1 - e^{2a(t_f - t_o)}\right)} \left[e^{a(t_f - t)} - e^{a(t + t_f - 2t_o)} \right].$$

We can use the form of this expression to explore how our cost function affects the optimal trajectory. For example, we can ask what happens to the terminal state $x^*(t_f)$ and $c \to \infty$. Setting $t = t_f$ in equation (2.2) and taking the limit we find that

$$\lim_{c \to \infty} x^*(t_f) = 0.$$

∇

Example 2.3 Bang-bang control
The time-optimal control program for a linear system has a particularly simple solution. Consider a linear system with bounded input

$$\dot{x} = Ax + Bu, \qquad |u| \leq 1$$

and suppose we wish to minimize the time required to move from an initial state x_0 to a final state x_f. Without loss of generality we can take $x_f = 0$. We choose the cost functions and terminal constraints to satisfy

$$J = \int_0^T 1\, dt, \qquad \psi(x(T)) = x(T)$$

To find the optimal control, we form the Hamiltonian

$$H = 1 + \lambda^T (Ax + Bu) = 1 + (\lambda^T A)x + (\lambda^T B)u.$$

Now apply the conditions in the maximum principle:

$$\dot{x} = \frac{\partial H}{\partial \lambda} = Ax + Bu$$
$$-\dot{\lambda} = \frac{\partial H}{\partial x} = A^T \lambda$$
$$u = \text{arg min } H = -\text{sgn}(\lambda^T B)$$

The optimal solution always satisfies this equation (necessary condition) with $x(0) = x_0$ and $x(T) = 0$. It follows that the input is always $u = \pm 1 \implies$ "bang-bang". ∇

2.4 Linear Quadratic Regulators

The finite horizon, linear quadratic regulator (LQR) is given by

$$\dot{x} = Ax + Bu \qquad x \in \mathbb{R}^n, u \in \mathbb{R}^n, x_0 \text{ given}$$
$$\tilde{J} = \frac{1}{2}\int_0^T \left(x^T Q_x x + u^T Q_u u\right)\, dt + \frac{1}{2}x^T(T)P_1 x(T)$$

where $Q_x \geq 0$, $Q_u > 0$, $P_1 \geq 0$ are symmetric, positive (semi-) definite matrices. Note the factor of $\frac{1}{2}$ is left out, but we included it here to simplify the derivation. Gives same answer (with $\frac{1}{2}x$ cost).

Solve via maximum principle:

$$H = x^T Q_x x + u^T Q_u u + \lambda^T(Ax + Bu)$$
$$\dot{x} = \left(\frac{\partial H}{\partial \lambda}\right)^T = Ax + Bu \qquad x(0) = x_0$$
$$-\dot{\lambda} = \left(\frac{\partial H}{\partial x}\right)^T = Q_x x + A^T\lambda \qquad \lambda(T) = P_1 x(T)$$
$$0 = \frac{\partial H}{\partial u} = Q_u u + \lambda^T B \implies u = -Q_u^{-1}B^T\lambda.$$

This gives the optimal solution. Apply by solving *two point boundary value problem* (hard).

Alternative: guess the form of the solution, $\lambda(t) = P(t)x(t)$. Then

$$\dot{\lambda} = \dot{P}x + P\dot{x} = \dot{P}x + P(Ax - BQ_u^{-1}B^T P)x$$
$$-\dot{P}x - PAx + PBQ_u^{-1}BPx = Q_x x + A^T Px.$$

This equation is satisfied if we can find $P(t)$ such that

$$-\dot{P} = PA + A^T P - PBQ_u^{-1}B^T P + Q_x \qquad P(T) = P_1$$

Remarks:

1. This ODE is called *Riccati ODE*.
2. Can solve for $P(t)$ backwards in time and then apply
$$u(t) = -Q_u^{-1}B^T P(t)x.$$
This is a (time-varying) *feedback* control $\implies$ tells you how to move from *any* state to the origin.
3. Variation: set $T = \infty$ and eliminate terminal constraint:
$$J = \int_0^\infty (x^T Q_x x + u^T Q_u u)\, dt$$
$$u = -\underbrace{Q_u^{-1}B^T P}_{K}\, x \qquad \text{Can show } P \text{ is } \textit{constant}$$
$$0 = PA + A^T P - PBQ_u^{-1}B^T P + Q_x$$

This equation is called the *algebraic Riccati equation*.

4. In MATLAB, `K = lqr(A, B, Q_x, Q_u)`.

5. Require $Q_u > 0$ but $Q_x \geq 0$. Let $Q_x = H^T H$ (always possible) so that $L = \int_0^\infty x^T H^T H x + u^T Q_u u \, dt = \int_0^\infty \|Hx\|^2 + u^T Q_u u \, dt$. Require that (A, H) is *observable*. Intuition: if not, dynamics may not affect cost $\Longrightarrow$ ill-posed.

2.5 Choosing LQR weights

$$\dot{x} = Ax + Bu \qquad J = \int_0^\infty \overbrace{\left(x^T Q_x x + u^T Q_u u + x^T S u\right)}^{L(x,u)} dt,$$

where the S term is almost always left out.

Q: How should we choose Q_x and Q_u?

1. Simplest choice: $Q_x = I$, $Q_u = \rho I \implies L = \|x\|^2 + \rho\|u\|^2$. Vary ρ to get something that has good response.

2. Diagonal weights

$$Q_x = \begin{bmatrix} q_1 & & \\ & \ddots & \\ & & q_n \end{bmatrix} \qquad Q_u = \rho \begin{bmatrix} r_1 & & \\ & \ddots & \\ & & r_n \end{bmatrix}$$

   Choose each q_i to given equal effort for same "badness". E.g., $x_1 =$ distance in meters, $x_3 =$ angle in radians:

$$1 \text{ cm error OK} \implies q_1 = \left(\frac{1}{100}\right)^2 \qquad q_1 x_1^2 = 1 \text{ when } x_1 = 1 \text{ cm}$$

$$\frac{1}{60} \text{ rad error OK} \implies q_3 = (60)^2 \qquad q_3 x_3^2 = 1 \text{ when } x_3 = \frac{1}{60} \text{ rad}$$

   Similarly with r_i. Use ρ to adjust input/state balance.

3. Output weighting. Let $z = Hx$ be the output you want to keep small. Assume (A, H) observable. Use

$$Q_x = H^T H \qquad Q_u = \rho I \qquad \implies \text{trade off } \|z\|^2 \text{ vs } \rho\|u\|^2$$

4. Trial and error (on *weights*)

2.6 Further Reading

Exercises

2.1 (a) Let $G_1, G_2, \ldots, G_k$ be a set of row vectors on a $\mathbb{R}^n$. Let F be another row vector on $\mathbb{R}^n$ such that for every $x \in \mathbb{R}^n$ satisfying $G_i x = 0$, $i = 1, \ldots, k$, we have $Fx = 0$. Show that there are constants $\lambda_1, \lambda_2, \ldots, \lambda_k$ such that

$$F = \sum_{i=1}^{k} \lambda_k G_k.$$

(b) Let $x^* \in \mathbb{R}^n$ be an the extremal point (maximum or minimum) of a function f subject to the constraints $g_i(x) = 0$, $i = 1, \ldots, k$. Assuming that the gradients $\partial g_i(x^*)/\partial x$ are linearly independent, show that there are k scalers λ_k, $i = 1, \ldots, n$ such that the function

$$\tilde{f}(x) = f(x) + \sum_{i=1}^{n} \lambda_i g_i(x)$$

has an extremal point at x^*.

2.2 Consider the following control system

$$\begin{aligned} \dot{q} &= u \\ \dot{Y} &= qu^T - uq^T \end{aligned}$$

where $u \in \mathbb{R}^m$ and $Y \in reals^{m \times m}$ is a skew symmetric matrix.

(a) For the fixed end point problem, derive the form of the optimal controller minimizing the following integral

$$\frac{1}{2} \int_0^1 u^T u \, dt.$$

(b) For the boundary conditions $q(0) = q(1) = 0$, $Y(0) = 0$ and

$$Y(1) = \begin{bmatrix} 0 & -y_3 & y_2 \\ y_3 & 0 & -y_1 \\ -y_2 & y_1 & 0 \end{bmatrix}$$

for some $y \in \mathbb{R}^3$, give an explicit formula for the optimal inputs u.

(c) (Optional) Find the input u to steer the system from $(0, 0)$ to $(0, \tilde{Y}) \in \mathbb{R}^m \times \mathbb{R}^{m \times m}$ where $\tilde{Y}^T = -\tilde{Y}$.

(Hint: if you get stuck, there is a paper by Brockett on this problem.)

2.3 In this problem, you will use the maximum principle to show that the shortest path between two points is a straight line. We model the problem by constructing a control system

$$\dot{x} = u$$

where $x \in \mathbb{R}^2$ is the position in the plane and $u \in \mathbb{R}^2$ is the velocity vector along the curve. Suppose we wish to find a curve of minimal length connecting $x(0) = x_0$ and $x(1) = x_f$. To minimize the length, we minimize the integral of the velocity along the curve,

$$J = \int_0^1 \sqrt{\|\dot{x}\|}\, dt,$$

subject to to the initial and final state constraints. Use the maximum principle to show that the minimal length path is indeed a straight line at maximum velocity. (Hint: minimizing $\sqrt{\|\dot{x}\|}$ is the same as minimizing $\dot{x}^T \dot{x}$; this will simplify the algebra a bit.)

2.4 Consider the optimal control problem for the system

$$\dot{x} = -ax + bu$$

where $x = \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. (Note that this system is not quite the same as the one in Example **??**.) The cost function is given by

$$J = \tfrac{1}{2} \int_{t_0}^{t_f} u^2(t)\, dt + \tfrac{1}{2} c x^2(t_f),$$

where the terminal time t_f is given and c is a constant.

(a) Solve explicitly for the optimal control $u^*(t)$ and the corresponding state $x^*(t)$ in terms of t_0, t_f, $x(t_0)$ and t and describe what happens to the terminal state $x^*(t_f)$ as $c \to \infty$.

(b) Show that the system is differentially flat with appropriate choice of output(s) and compute the state and input as a function of the flat output(s).

(c) Using the polynomial basis $\{t^k,\ k = 0, \dots, M-1\}$ with an appropriate choice of M, solve for the (non-optimal) trajectory between $x(t_0)$ and $x(t_f)$. Your answer should specify the explicit input $u_d(t)$ and state $x_d(t)$ in terms of t_0, t_f, $x(t_0)$, $x(t_f)$ and t.

(d) Let $a = 1$ and $c = 1$. Use your solution to the optimal control problem and the flatness-based trajectory generation to find a trajectory between $x(0) = 0$ and $x(1) = 1$. Plot the state and input trajectories for each solution and compare the costs of the two approaches.

(e) (Optional) Suppose that we choose more than the minimal number of basis functions for the differentially flat output. Show how to use the additional degrees of freedom to minimize the cost of the flat trajectory and demonstrate that you can obtain a cost that is closer to the optimal.

2.5 Consider the optimal control problem for the system

$$\dot{x} = -ax^3 + bu$$

where $x = \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. The cost function is given by

$$J = \tfrac{1}{2} \int_{t_0}^{t_f} u^2(t)\, dt + \tfrac{1}{2} c x^2(t_f),$$

where the terminal time t_f is given and c is a constant.

(a) Derive a set of differential equations for the optimal control $u^*(t)$ and the corresponding state $x^*(t)$ in terms of t_0, t_f, $x(t_0)$ and t. Be sure to provide any initial or final conditions required for your equations to be solved.

(b) Show that the system is differentially flat with appropriate choice of output(s) and compute the state and input as a function of the flat output(s).

(c) Using the polynomial basis $\{t^k,\, k = 0, \dots, M-1\}$ with an appropriate choice of M, solve for the (non-optimal) trajectory between $x(t_0)$ and $x(t_f)$. Your answer should specify the explicit input $u_d(t)$ and state $x_d(t)$ in terms of t_0, t_f, $x(t_0)$, $x(t_f)$ and t.

(d) Increase M by one and show how to choose the free parameter to minimize the cost function.

2.6 Consider the problem of moving a two-wheeled mobile robot (eg, a Segway) from one position and orientation to another. The dynamics for the system is given by the nonlinear differential equation

$$\begin{aligned} \dot{x} &= \cos\theta\, v \\ \dot{y} &= \sin\theta\, v \\ \dot{\theta} &= \omega \end{aligned}$$

where (x, y) is the position of the rear wheels, θ is the angle of the robot with respect to the x axis, v is the forward velocity of the robot and ω is spinning rate. We wish to choose an input (v, ω) that minimizes the time that it takes to move between two configurations (x_0, y_0, θ_0) and (x_f, y_f, θ_f), subject to input constraints $|v| \le L$ and $|\omega| \le M$.

Use the maximum principle to show that any optimal trajectory consists of segments in which the robot is traveling at maximum velocity in either the forward or reverse direction, and going either straight, hard left ($\omega = -M$) or hard right ($\omega = +M$).

Note: one of the cases is a bit tricky and can't be completely proven with the tools we have learned so far. However, you should be able to show the other cases and verify that the tricky case is possible.

2.7 Consider a linear system with input u and output y and suppose we wish to minimize the quadratic cost function

$$J = \int_0^\infty \left(y^T y + \rho u^T u \right)\, dt.$$

Show that if the corresponding linear system is observable, then the closed loop system obtained by using the optimal feedback $u = -Kx$ is guaranteed to be stable.

2.8 Consider the control system transfer function

$$H(s) = \frac{s+b}{s(s+a)} \qquad a, b > 0$$

with state space representation

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ 0 & -a \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u$$
$$y = \begin{bmatrix} b & 1 \end{bmatrix} x$$

and performance criterion

$$V = \int_0^\infty (x_1^2 + u^2) dt.$$

(a) Let

$$P = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}$$

with $p_{12} = p_{21}$ and $P > 0$ (positive definite). Write the steady state Riccati equation as a system of four explicit equations in terms of the elements of P and the constants a and b.

(b) Find the gains for the optimal controller assuming the full state is available for feedback.

(c) Find the closed loop natural frequency and damping ratio.

2.9 The output $c(t)$ in a position-control system is governed by

$$J\ddot{c} = u,$$

where $u(t)$ is applied force.

(a) Write down a state space realization (find A and B).

(b) Use the matrix Riccati equation to find the feedback control law minimizing

$$\int_0^\infty (c^2 + q^2 u^2) dt.$$

(c) Show that the optimal control system has damping ratio $\frac{1}{\sqrt{2}}$.

(d) What is the corresponding optimal value of natural frequency?

(See AM05, Sec 4.4 if you don't remember how damping ratio (or factor) and natural frequency are defined.)

2.10 Consider the optimal control problem for the system

$$\dot{x} = ax + bu \qquad J = \tfrac{1}{2}\int_{t_0}^{t_f} u^2(t)\,dt + \tfrac{1}{2}cx^2(t_f),$$

where $x \in \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. We take the terminal time t_f as given and let $c > 0$ be a constant that balances the final value of the state with the input required to get to that position. The optimal is derived in the lecture notes for week 6 and is shown to be

$$\begin{aligned} u^*(t) &= -\frac{2abc\,e^{a(2t_f-t_o-t)}x(t_o)}{2a - b^2c\left(1 - e^{2a(t_f-t_o)}\right)} \\ x^*(t) &= x(t_o)e^{a(t-t_o)} + \frac{b^2c\,e^{a(t_f-t_o)}x(t_o)}{2a - b^2c\left(1 - e^{2a(t_f-t_o)}\right)}\left[e^{a(t_f-t)} - e^{a(t+t_f-2t_o)}\right]. \end{aligned} \tag{2.2}$$

Now consider the infinite horizon cost

$$J = \tfrac{1}{2}\int_{t_0}^{\infty} u^2(t)\,dt$$

with $x(t)$ at $t = \infty$ constrained to be zero.

(a) Solve for $u^*(t) = -bPx^*(t)$ where P is the positive solution corresponding to the algebraic Riccati equation. Note that this gives an explicit feedback law ($u = -bPx$).

(b) Plot the state solution of the finite time optimal controller for the following parameter values

$$\begin{gathered} a = 2 \qquad b = 0.5 \qquad x(t_0) = 4 \\ c = 0.1,\ 10 \qquad t_f = 0.5,\ 1,\ 10 \end{gathered}$$

(This should give you a total of 6 curves.) Compare these to the infinite time optimal control solution. Which finite time solution is closest to the infinite time solution? Why?

2.11 In this problem we will explore the effect of constraints on control of the linear unstable system given by

$$\begin{aligned} \dot{x}_1 &= 0.8x_1 - 0.5x_2 + 0.5u \\ \dot{x}_2 &= x_1 + 0.5u \end{aligned}$$

subject to the constraint that $|u| \leq a$ where a is a postive constant.

(a) Ignore the constraint ($a = \infty$) and design an LQR controller to stabilize the system. Plot the response of the closed system from the initial condition given by $x = (1, 0)$.

(b) Use SIMULINK or `ode45` to simulate the the system for some finite value of a with an initial condition $x(0) = (1, 0)$. Numerically (trial and error) determine the smallest value of a for which the system goes unstable.

(c) Let $a_{\min}(\rho)$ be the smallest value of a for which the system is unstable from $x(0) = (\rho, 0)$. Plot $a_{\min}(\rho)$ for $\rho = 1$, 4, 16, 64, 256.

(d) *Optional*: Given $a > 0$, design and implement a receding horizon control law for this system. Show that this controller has larger region of attraction than the controller designed in part (b). (Hint: solve the finite horizon LQ problem analytically, using the bang-bang example as a guide to handle the input constraint.)

2.12 Consider the lateral control problem for an autonomous ground vehicle from Example 1.1. We assume that we are given a reference trajectory $r = (x_d, y_d)$ corresponding to the desired trajectory of the vehicle. For simplicity, we will assume that we wish to follow a straight line in the x direction at a constant velocity $v_d > 0$ and hence we focus on the y and θ dynamics:

$$\dot{y} = \sin\theta\, v_d$$
$$\dot{\theta} = \frac{1}{\ell}\tan\phi\, v_d.$$

We let $v_d = 10$ m/s and $\ell = 2$ m.

(a) Design an LQR controller that stabilizes the position y to the origin. Plot the step and frequency response for your controller and determine the overshoot, rise time, bandwidth and phase margin for your design. (Hint: for the frequency domain specifications, break the loop just before the process dynamics and use the resulting SISO loop transfer function.)

(b) Suppose now that $y_d(t)$ is not identically zero, but is instead given by $y_d(t) = r(t)$. Modify your control law so that you track $r(t)$ and demonstrate the performance of your controller on a "slalom course" given by a sinusoidal trajectory with magnitude 1 meter and frequency 1 Hz.