EE/AA547-W22 2022.02.17

Topic: NLP and Intro to DP

Lecturer: L.J. Ratliff

1 Review

Some great references:

- Bertsekas Nonlinear Programming
- Bertsekas Dynamic Programming and Optimal Control Vols I, II
- Bertsekas Neurodynamic Programming
- Kirk Optimal Control

1.1 Non-Linear Programming Overview

The basics of nonlinear programming you have probably seen in your calculus sequence even if you did not know it. The idea is that for an unconstrained problem, we can simply look at derivatives (first and second order) to determine necessary and sufficient conditions for (global/local) "optimality"—i.e., a critical point that is either a (global/local) minimum or maximum.

For example consider the sufficiently smooth cost function $F : \mathbb{R}^n \to \mathbb{R}$, and suppose we want to find a minimum of this cost. Consider the following cartoon graphic:

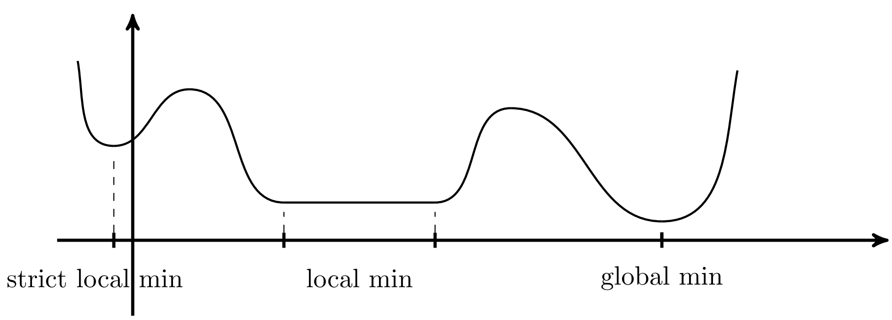

Definition 1. A vector $x^\star$ is a local minimum of F if it is no worse than its neighbors; i.e., if there exists an $\epsilon > 0$ such that

$$F(x^\star) \le F(x) \quad \forall\ x :\ \|x - x^\star\| \le \epsilon$$

It is a global minimum if this inequality holds for all x. And, it is strict if the inequality is strict for all $x \neq x^\star$.

Necessary Conditions:

- $DF(x^\star) = 0$
- $D^2F(x^\star) \succeq 0$

Sufficient Conditions:

- $DF(x^\star) = 0$

- $D^2 F(x^\star) \succ 0$

These sufficient conditions in particular are equivalent to the following: there exists $\gamma > 0$ and $\epsilon > 0$ such that

$$F(x) \geq F(x^\star) + \frac{\gamma}{2}\|x - x^\star\|^2, \quad \forall\, x \text{ with } \|x - x^\star\| < \epsilon$$

1.1.1 Constrained NLP

The more interesting and more apropos NLP setting is the constrained setting. This will be relevant for us since we typically are seeking an optimal controller for an objective of the form $F(x,u)$ where $f(x,u) = 0$ is the constrained defined by the dynamics—e.g., $f(x,u) = \dot{x} - (Ax + Bu)$. And in particular we seek to minimize costs $J(u)$ were the state constraints implicit specify x in terms of u.

Consider the optimization problem

$$\begin{aligned} \min_u \quad & F(x,u) \\ \text{s.t.} \quad & f(x,u) = 0 \end{aligned}$$

There are several types of "constraint satisfaction/qualification" conditions specifying sufficient conditions for such problems. We will take the approach of writing out the Lagrangian and looking at the sufficient conditions of the corresponding unconstrained problem. Indeed, let $\lambda \in \mathbb{R}^n$ (where there are n constraints) be the Lagrange multiplier.

$$D_x f(x,u) = 0 \implies$$

Then the Lagrangian is

$$L(x,u,\lambda) = F(x,u) + \lambda^\top f(x,u)$$

Since $f(x,u) = 0$ and we are optimizing over u here we can use the elimination method to remove one of the variables. The best way to see this is with a simple example.

Problem 1. (Elimination Method.)Consider the optimization problem

$$\begin{aligned} \min_u \quad & F(x,u) \\ \text{s.t.} \quad & Gx + Hu = b \end{aligned}$$

where G is invertible. The matrix $[G\ H]$ is an $n \times (n+m)$ matrix where $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$. Suppose it has linearly independent rows and $b \in \mathbb{R}^n$ is given. Here $G \in \mathbb{R}^{n\times n}$ and $H \in \mathbb{R}^{n\times m}$. Write out first order optimality conditions.
Solution. Since G is invertible, we can eliminate x using the constraints:

$$x = G^{-1}(b - Hu)$$

Hence the problem becomes

$$\begin{aligned} \min_u \quad & \tilde{F}(u) := F(G^{-1}(b - Hu), u) \\ \text{s.t.} \quad & u \in \mathbb{R}^m \end{aligned}$$

If $(x^\star, u^\star)$ where $x^\star := G^{-1}(b - Hu^\star)$ is a local minimum of the constrained problem then $u^\star$ is a local minimum of the reduced cost function $\tilde{F}$, so that

$$0 = \nabla \tilde{F}(u^\star) = -H^\top (G^\top)^{-1} \nabla_x F(x^\star, u^\star) + \nabla_u F(x^\star, u^\star)$$

Let

$$\lambda^\star := -(G^\top)^{-1} \nabla_x F(x^\star, u^\star)$$

so that

$$0 = \nabla \tilde{F}(u^\star) = -H^\top \lambda^\star + \nabla_u F(x^\star, u^\star)$$

and

$$\lambda^\star := -(G^\top)^{-1} \nabla_x F(x^\star, u^\star) \iff \nabla_x F(x^\star, u^\star) + G^\top \lambda^\star = 0$$

Putting these together we have that

$$\nabla F(x^\star, u^\star) + \begin{bmatrix} G & H \end{bmatrix}^\top \lambda^\star = 0$$

which is precisely the Lagrange multiplier condition of the original problem—i.e.,

$$L(x, u, \lambda) = F(x, u) + \lambda^\top (Gx + Hu - b) \implies \begin{bmatrix} \nabla_x F(x^\star, u^\star) & \nabla_u F(x^\star, u^\star) \end{bmatrix} + (\lambda^\star)^\top \begin{bmatrix} G & H \end{bmatrix} = 0$$

The problems we will consider are convex (quadratic) and are subject to linear constraints. Hence, checking the second order conditions won't be strictly necessary. But let's see them for posterity.

We know from unconstrained optimization that for the problem in terms of $\tilde{F}$ that it is necessary that

$$z^\top \nabla^2 \tilde{F}(u^\star) z \geq 0 \qquad u \in \mathbb{R}^m$$

Using our conditions for $\tilde{F}$ we have that

$$\nabla^2 \tilde{F}(u) = \nabla \left(-H^\top (G^\top)^{-1} \nabla_x F(G^{-1}(b - Hu), u) + \nabla_u F(G^{-1}(b - Hu), u) \right)$$

Plugging in $x^\star$ and partitioning the Hessian $\nabla^2 F(x^\star, u^\star)$ we have

$$\nabla^2 F(x^\star, u^\star) = \begin{bmatrix} \nabla_x^2 F(x^\star, u^\star) & \nabla_{xu} F(x^\star, u^\star) \\ \nabla_{ux} F(x^\star, u^\star) & \nabla_u^2 F(x^\star, u^\star) \end{bmatrix}$$

Hence we have that

$$\begin{aligned} \nabla^2 \tilde{F}(u^\star) &= \nabla \left(-H^\top (G^\top)^{-1} \nabla_x F(G^{-1}(b - Hu^\star), u^\star) + \nabla_u F(G^{-1}(b - Hu^\star), u^\star) \right) \\ &= H^\top (G^\top)^{-1} \nabla_x^2 F(x^\star, u^\star) G^{-1} H - H^\top (G^\top)^{-1} \nabla_{xu} F(x^\star, u^\star) - \nabla_{ux} F(x^\star, u^\star) G^{-1} H + \nabla_u^2 F(x^\star, u^\star) \end{aligned}$$

Now since $\nabla^2 F(x^\star, u^\star)$ is necessarily positive semidefinite and the constraints are linear (i.e., $\nabla^2 f_i(x^\star, u^\star) = 0$) we have that

$$0 \leq z^\top \nabla^2 \tilde{F}(x^\star, u^\star) z = y^\top \nabla^2 F(x^\star, u^\star) y, \qquad y = \begin{bmatrix} -G^{-1} H z \\ z \end{bmatrix}$$

The point is that the Lagrange multiplier conditions for the constrained problem are nothing more than the zero gradient and positive semidefinite Hessian conditions for the unconstrained problem.

Optimality Conditions. The general optimality conditions we care about are the following:

- optimal Lagrange multiplier:

$$\nabla_x L(x, u, \lambda) = \nabla_x F(x, u) + \lambda^\top \nabla_x f(x, u) = 0 \implies (\lambda^\star)^\top = -\nabla_x F(x, u) (\nabla_x f(x, u))^{-1}$$

- Constraint satisfaction:

$$f(x, u) = 0$$

- Optimality of u:

$$\nabla_u L(x, u, \lambda) = \nabla_u F(x, u) + \lambda^\top \nabla_u f(x, u) = 0$$

1.2 DT LQR Problem

Consider

$$x(t+1) = Ax(t) + Bu(t),\ x(0) = x_0$$

and

$$J(u) = x_N^\top Q_f x_N + \sum_{k=0}^{N-1} \left(x_k^\top Q x_k + u_k^\top R u_k\right)$$

where $u = (u_0, \ldots, u_{N-1})$ and $Q = Q^\top \succeq 0$, $Q_f = Q_f^\top \succeq 0$, $R = R^\top \succ 0$ are the given state cost, final state cost, and input cost matrices.

- N is the time horizon
- first term measures state deviation
- second term measures input size or actuator authority
- last term measures final state deviation
- Q, R set relative weights of state deviation and input usage
- $R \succ 0$ means any (non-zero) input adds cost to J

LQR Problem: find $u^\star$ that minimizes $J(u)$.

Problem 2. (LQR as an NLP.) Write out the LQR problem as an NLP. In particular a least squares problem.

Solution. Unwrapping the dynamics we have that

$$\begin{aligned}
x_0 &= I \cdot x_0 \\
x_1 &= Ax_0 + Bu_0 \\
x_2 &= A(x_0 + Bu_0) + Bu_1 = Ax_0 + \begin{bmatrix} AB & B \end{bmatrix} \begin{bmatrix} u_0 \\ u_1 \end{bmatrix} \\
&\vdots \\
x_N &= A^N x_0 + \begin{bmatrix} A^{N-1}B & A^{N-2}B & \cdots & B \end{bmatrix} \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{N-1} \end{bmatrix}
\end{aligned}$$

Stacking this up we have

$$\underbrace{\begin{bmatrix} x_0 \\ \vdots \\ x_N \end{bmatrix}}_{:=X} = \underbrace{\begin{bmatrix} 0 & \cdots & \cdots & \cdots \\ B & 0 & \cdots & \cdots \\ AB & B & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1}B & A^{N-2}B & \cdots & B \end{bmatrix}}_{=:H} \underbrace{\begin{bmatrix} u_0 \\ \vdots \\ u_{N-1} \end{bmatrix}}_{=:u} + \underbrace{\begin{bmatrix} I \\ A \\ \vdots \\ A^N \end{bmatrix}}_{=:G} x_0$$

Now our constraint is going to be

$$X = Hu + Gx_0, \qquad G \in \mathbb{R}^{Nn \times m},\ H \in \mathbb{R}^{Nn \times Nm}$$

Now for the cost we have that

$$\begin{aligned}
J(u) &= x_N^\top Q_f x_N + \sum_{k=0}^{N-1} \left(x_k^\top Q x_k + u_k^\top R u_k\right) \\
&= X^\top \text{blkdiag}(Q, \cdots, Q, Q_f) X + u^\top \text{blkdiag}(R, \ldots, R) u \\
&= \left\| \text{blkdiag}(Q^{1/2}, \cdots, Q^{1/2}, Q_f^{1/2}) X \right\|^2 + \left\| \text{blkdiag}(R^{1/2}, \ldots, R^{1/2}) u \right\|^2
\end{aligned}$$

where the last equality is possible since Q and R are positive (semi)definite symmetric matrices. So the least squares problem is

$$\begin{aligned} \min_u \quad & \left\| \text{blkdiag}(Q^{1/2}, \cdots, Q^{1/2}, Q_f^{1/2}) X \right\|^2 + \left\| \text{blkdiag}(R^{1/2}, \ldots, R^{1/2}) u \right\|^2 \\ \text{s.t.} \quad & X = Hu + Gx_0 \end{aligned}$$

And this is equivalent to

$$\min_u \quad \left\| \text{blkdiag}(Q^{1/2}, \cdots, Q^{1/2}, Q_f^{1/2})(Hu + Gx_0) \right\|^2 + \left\| \text{blkdiag}(R^{1/2}, \ldots, R^{1/2}) u \right\|^2$$

which is an unconstrained NLP.