ME 7247: Advanced Control Systems **Fall 2022–23**

Lecture 04: Multi-objective optimization

Tuesday, September 20, 2022

Lecturer: Laurent Lessard *Scribe: Paola Kefallinos*

Often a problem asks us to optimize more than one characteristic of a system. However there are usually trade-offs in doing so; that is, one can optimize a certain trait at the expense of another. For these kinds of problems, there isn't a solution: it's a matter of choice. Today we investigate how to optimize a problem with multiple objectives.

1 Review

Let's review what we've covered in the last couple of lectures considering the equation $Ax = b$ with $A \in \mathbb{R}^{m \times n}$.

Least Squares

- Typically, A is a tall matrix (more equations than variables)
- We would like to find an approximate solution $A\hat{x} \approx b$ as there is typically no x satisfying $Ax = b$ (called the *estimation* setup)
- In optimization notation, the least-squares (LS) problem is written as

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \|Ax - b\|^2 \tag{1}$$

- The set of solutions of Eq. (1) is precisely the same as the set of solutions to the *normal equations* (2) below.

$$A^{\mathsf{T}} A\hat{x} = A^{\mathsf{T}} b \tag{2}$$

- A solution to the normal equations always exists. That solution is unique if and only if $\text{null}(A) = \{0\}$, i.e. if the columns of A are linearly independent.

Least Norm

- Typically, A is a wide matrix (more variables than equations)
- There are typically infinitely many x satisfying $Ax = b$, so we want to find the "best" x among all solutions (called the *control* setup)
- In optimization notation, the least-norm (LN) problem is written as

$$\begin{aligned} \underset{x \in \mathbb{R}^n}{\text{minimize}} \quad & \|x\|^2 \\ \text{such that} \quad & Ax = b \end{aligned} \tag{3}$$

- The set of solutions of (3) is precisely the same as the set of solutions $\hat{x}$ to the system of equations

$$AA^\mathsf{T}w = b \quad \text{and} \quad \hat{x} = A^\mathsf{T}w \tag{4}$$

- A solution to Eq. (4) exists if and only if $b \in \text{range}(A)$, i.e. if $Ax = b$ has at least one solution. If a solution exists, it is always unique. Note: there may be many w that solve (4), but they all lead to the same $\hat{x}$.

2 Defining the Multi-Objective Optimization Problem

Introduction to the cost notation Consider a hybrid version of LS and LN, where we are trying to make both $\|Ax - b\|^2$ and $\|x\|^2$ small at the same time. First, we write these as two separate costs ($J_1 : \mathbb{R}^n \to \mathbb{R}$ and $J_2 : \mathbb{R}^n \to \mathbb{R}$). In this case, we have

$$J_1(x) = \|Ax - b\|^2 \tag{5a}$$

$$J_2(x) = \|x\|^2 \tag{5b}$$

Where Equation (5a) above represents the Least Squares problem and Equation (5b) represents the Least norm problem. One can see that a value of x that minimizes one function does not minimize the other. Figure 1 below depicts a 2D plot of $J_1(x)$ and $J_2(x)$ for all the possible x given that A and b are fixed.

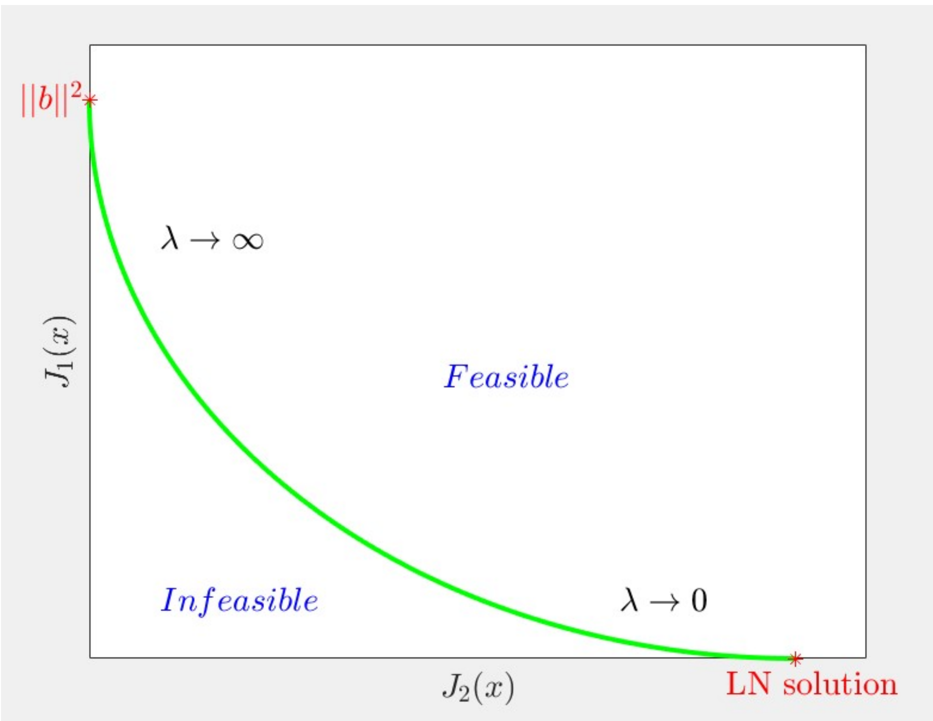

Figure 1: Pareto-optimal front

The green line is the optimal solution, and is called the "Pareto-optimal front". The areas to the right and left of the line are "feasible" and "infeasible", respectively. A natural question may arise from this plot: which point on the line is the best? This question has no answer! Any design on this curve is an optimal solution to the multi-objective problem, and points on the Pareto-optimal front are not comparable unless we assign weights on the objective functions to prioritize them.

A single cost function We can write a single expression that combines these with the help of a weighting parameter $\lambda > 0$ that weights the costs of each accordingly. Eq. (6) below shows that increasing the value of the weight parameter makes the cost function more sensitive to $J_2(x)$, and naturally the converse is true.

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad J_1(x) + \lambda J_2(x) \tag{6}$$

This is actually a least squares problem! Let's do some manipulation to prove it.

Rearrange into Least Squares problem

1. Substitute Eqs. (5a) and (5b) into Eq. (6).

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \|Ax - b\|^2 + \|x\|^2 \tag{7}$$

2. Recall block matrix with norm relationship

$$\|x_1\|^2 + \|x_2\|^2 = \left\| \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right\|^2 \tag{8}$$

3. Rewrite (7) using (8)

$$\|Ax - b\|^2 + \|x\|^2 = \min_x \left\| \begin{bmatrix} Ax - b \\ \sqrt{\lambda} x \end{bmatrix} \right\|^2 \tag{9a}$$

$$= \min_x \left\| \begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix} x - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\|^2 \tag{9b}$$

4. This is a LS problem, so its solution set is the same as that of the normal equations

$$\begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix}^\mathsf{T} \begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix} x = \begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix}^\mathsf{T} \begin{bmatrix} b \\ 0 \end{bmatrix} \tag{10a}$$

$$\begin{bmatrix} A^\mathsf{T} & \sqrt{\lambda} I \end{bmatrix} \begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix} = \begin{bmatrix} A^\mathsf{T} & \sqrt{\lambda} I \end{bmatrix} \begin{bmatrix} b \\ 0 \end{bmatrix} \tag{10b}$$

$$A^\mathsf{T} A + \sqrt{\lambda} I (\sqrt{\lambda} I) = A^\mathsf{T} b + \sqrt{\lambda} I (0) \tag{10c}$$

$$(A^\mathsf{T} A + \lambda I) x = A^\mathsf{T} b \tag{10d}$$

5. The matrix $A^\mathsf{T} A + \lambda I$ is invertible for any $\lambda > 0$ and *any* A.

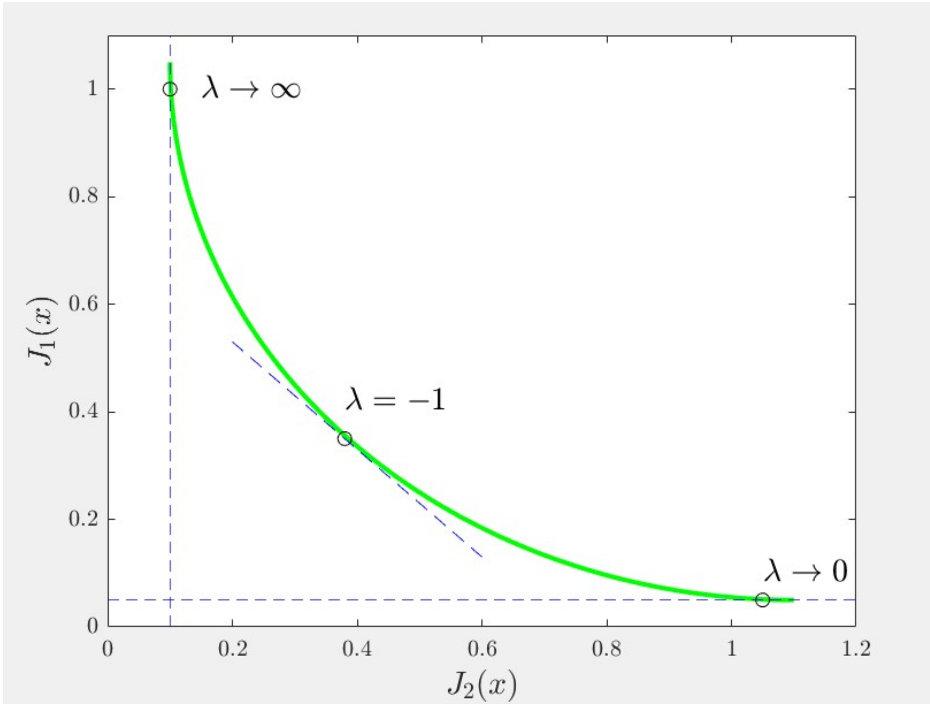

Figure 2: Slope of Pareto-optimal Curve

Geometric interpretation: The trade-off parameter λ is the negative slope of the Pareto-optimal curve. So as we vary $\lambda = 0 \to \infty$, we start on the bottom-right with a slope of 0 and end on the top-left with a slope of $-\infty$.

3 Special Cases

When $\lambda \to 0$ in the multi-objective optimization problem, we can recover either the LS or the LN norm solution, depending on the assumptions we make.

LS problem In the LS setting, when $\lambda \to 0$ and A has full column rank, then $A^\mathsf{T}A$ is invertible. Thus, we can take the limit by just setting $\lambda = 0$ and solving (10d).

$$\begin{aligned} \hat{x} &= \lim_{\lambda \to 0} (A^\mathsf{T}A + \lambda I)^{-1} A^\mathsf{T} b \\ &= (A^\mathsf{T}A)^{-1} A^\mathsf{T} b \end{aligned} \tag{11}$$

This is the same solution we found when solving the LS problem in Lecture 2.

LN problem In the LN setting, when $\lambda \to 0$ and A has full row rank, AA^T is invertible. However, we can't simply set $\lambda = 0$ as we did for LS because $A^\mathsf{T}A$ is *not* invertible. We can nevertheless evaluate the limit using the push-through identity (presented in the next section).

$$\begin{aligned}\hat{x} &= \lim_{\lambda\to 0}(A^\mathsf{T}A+\lambda I)^{-1}A^\mathsf{T}b \\ &= \lim_{\lambda\to 0} A^\mathsf{T}(AA^\mathsf{T}+\lambda I)^{-1}b \\ &= A^\mathsf{T}(AA^\mathsf{T})^{-1}b\end{aligned} \tag{12}$$

This is the same solution we found when solving the LN problem in Lecture 3.

4 The Push-Through Identity

Push-Through Identity: If $A \in \mathbb{R}^{m\times n}$ and $B \in \mathbb{R}^{n\times m}$, then

$$A(BA+\lambda I)^{-1} = (AB+\lambda I)^{-1}A \tag{13}$$

Moreover, $AB+\lambda I$ is invertible if and only if $BA+\lambda I$ is invertible.

Proof:

1. Factor out A from the right and left of $ABA+\lambda A$.

$$\begin{aligned}A(BA+\lambda I_n) &= ABA+\lambda A \\ &= (AB+\lambda I_m)A\end{aligned} \tag{14}$$

2. Multiply both sides by $(BA+\lambda I_n)^{-1}$ on the right to isolate A.

$$A = (AB+\lambda I_m)A(BA+\lambda I_n)^{-1} \tag{15}$$

3. Finally, multiply both sides by $(AB+\lambda I_m)^{-1}$ on the left.

$$(AB+\lambda I_m)^{-1}A = A(BA+\lambda I_n)^{-1} \tag{16}$$

Regarding invertibility, suppose $AB+\lambda I$ is *not* invertible. Then there must exist a nonzero element in the nullspace. So there is some $v \neq 0$ such that $ABv+\lambda v = 0$. Multiply both sides by B on the left and obtain $0 = BABv+\lambda Bv = (BA+\lambda I)Bv$. Therefore $Bv \in \text{null}(BA+\lambda I)$. We can't have $Bv = 0$, because then $ABv+\lambda v = \lambda v = 0$, which contradicts the fact that $\lambda > 0$ and $v \neq 0$. Therefore, we have identified a nonzero element of the nullspace of $BA+\lambda I$, which means that $BA+\lambda I$ is not invertible. Applying the same argument starting with $BA+\lambda I$, we conclude that $AB+\lambda I$ is invertible if and only if $BA+\lambda I$ is invertible.

This is called the "Push-Through" identity because in Eq. (13), matrix A is pushed from the left side of $(BA+\lambda I_n)^{-1}$ to the other. Notice in Equation 16, the left hand side requires the inverse of an $m\times m$ matrix while the right hand side requires an inverse of an $n\times n$ matrix. When $m \gg n$ (m much larger than n), this property becomes computationally helpful as we need only compute the inverse of the smaller matrix.

Figure 3: Depiction of mass transfer

5 Revisiting the Mass Transfer Example

We want to move a mass (initially at rest) a distance close to 1 unit in 10 seconds by applying a force every second (time is discretized into 1 second units). First, we start by defining the following variables:

- x_t = position at time t
- v_t = velocity at time t
- f_t = force applied at time t

We will assume that initial conditions are $x_0 = 0$ and $v_0 = 0$ and the dynamics of the system can described by the following equations

1. $v_{t+1} = v_t + f_t$
2. $x_{t+1} = x_t + v_t$

Goals

1. Make $(x_{10} - 1)^2$ small (get as close to final position as we can)
2. Make $f_0^2 + f_1^2 + ... + f_9^2$ small (use as little fuel as possible)

Solution Process .
First, write all goals in terms of f

$$\begin{aligned} x_{10} &= v_0 + v_1 + ... + v9 + x_0 \\ v_1 &= v_0 + f_0 \\ v_2 &= v_0 + f_0 + f_1 \\ v_3 &= v_0 + f_0 + f_1 + f_2 \\ &\vdots \\ v_{10} &= f_0 + f_1 + f_2 + \cdots + f_9 \end{aligned} \tag{17}$$

Substitute the value for velocity (v_10) into the expression for position(x_10) to yield

$$
\begin{aligned}
x_{10} &= 9f_0 + 8f_1 + 7f_2 + \cdots + f_8 \\
&= \begin{bmatrix} 9 & 8 & 7 & \cdots & 2 & 1 \end{bmatrix} \begin{bmatrix} f_0 \\ f_1 \\ \vdots \\ f_8 \end{bmatrix} \\
&= a^\mathsf{T} f
\end{aligned}
\tag{18}
$$

Then write the costs in optimization notation

$$
\begin{aligned}
&\textbf{Goal 1:}\ J_1(f) = \|a^\mathsf{T} f - 1\|^2 \\
&\textbf{Goal 2:}\ J_2(f) = \|f\|^2
\end{aligned}
\tag{19}
$$

and then combine them to form a single cost function of the form

$$
\min_f \|a^\mathsf{T} f - 1\|^2 + \|f\|^2 \tag{20}
$$

and then we solve for $\hat{f}$ using the push through identity

$$
\begin{aligned}
\hat{f} &= (aa^\mathsf{T} + \lambda I)^{-1} a \\
&= a(a^\mathsf{T} a + \lambda I)^{-1} \\
&= \frac{1}{\|a\|^2 + \lambda} \cdot a
\end{aligned}
\tag{21}
$$

As λ increases (heavily weighting J_2 i.e the cost of fuel) the model would decide not to move and pay the price for not reaching the destination. Conversely, as λ decreases (heavily weighting J_1 i.e distance) the model doesn't care how much fuel is used as long as it ends up in the right place. Fig. 4 below depicts a 2D plot of the trade-off curve (Pareto-optimal front).

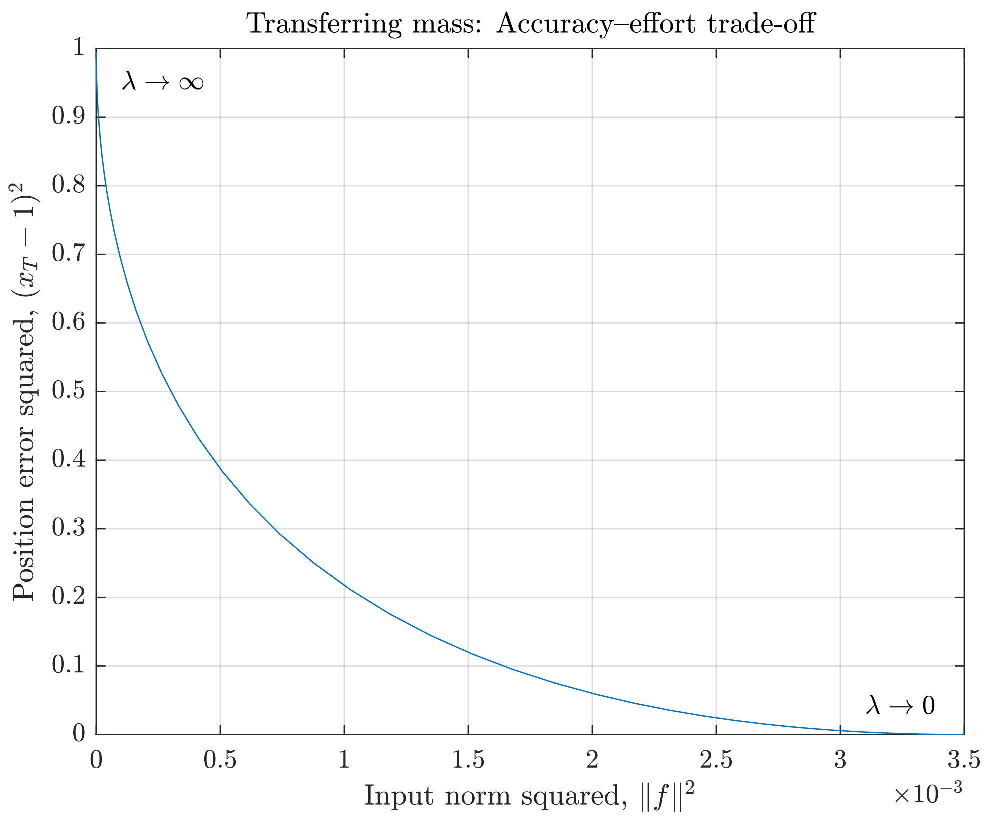

Figure 4: Trade-off curve between the squared position error and the squared norm of the force applied.