ORF522 – Linear and Nonlinear Optimization

20. Sequential Convex Programming

Bartolomeo Stellato — Fall 2021

Ed Forum

$$\frac{1}{k^2}$$

$$\nabla f(x^k) - f(x^*) \leq$$

- Nesterov's theorem declares the existence of a fuction f, and gives its lower bound for first order methods; but how does it give lower bounds for all convex L-smooth functions?

- The part of the lecture that I struggled with most was the relationship between/difference between Nesterov momentum and accelerated proximal gradient methods, since it seemed that the weights achieve very similar results.

Today's lecture
[Chapter 4 and 17, NO][ee364b]

Convex algorithms to solve nonconvex optimization problems

- Sequential convex programming

- Trust region methods

- Building convex approximations

- Regularized trust region methods

- Difference of convex programming

Methods for nonconvex optimization

Convex optimization algorithms: global and typically **fast**

Nonconvex optimization algorithms: must give up one, global or fast

- **Local methods: fast** but **not global**   ⟶   **heuristics**
  Need not find a global (or even feasible) solution.
  They cannot certify global optimality because
  KKT conditions are not sufficient.

- **Global methods: global** but often **slow**
  They find a global solution and certify it.

Sequential Convex Programming

Sequential convex programming (SCP)

Local optimization method that leverages convex optimization

Subproblems are convex $\longrightarrow$ we can solve them efficiently

It is a **heuristic**

• It **can fail** to find an optimal (or even feasible point)

• Results **depend on the starting point.**
  We can run the algorithm from many initial points and take the best result.

It often works very well
it finds a feasible point with good objective value (often optimal!)

# Gradient descent as SCP

Problem

minimize $f(x)$

Iterates

$$x^{k+1} = x^k - t_k \nabla f(x^k)$$

Quadratic approximation, replace $\nabla^2 f(x^k)$ with $\dfrac{1}{t_k} I$

$$x^{k+1} = \underset{y}{\text{argmin}} \; f(x^k) + \nabla f(x^k)^T (y - x^k) + \frac{1}{2t_k} \|y - x^k\|_2^2$$

strongly convex problem

$f(x^{k+1})$ —

— $f(x^k)$

x^{k+1} x^k

The problem

minimize $f(x)$

subject to $g_i(x) \leq 0, \quad i = 1, \ldots, m$ with $x \in \mathbf{R}^n$

                $h_i(x) = 0, \quad i = 1, \ldots, p$

- f and g_i can be nonconvex

- h_i can be nonaffine

Trust region methods

Main idea

$$\begin{aligned}
\text{minimize} \quad & f(x) \\
\text{subject to} \quad & g_i(x) \le 0, \quad i = 1, \dots, m \\
& h_i(x) = 0, \quad i = 1, \dots, p
\end{aligned}$$

iterate x^k
trust region $\mathcal{T}^k$

approximate convex problem

$$\begin{aligned}
\text{minimize} \quad & \hat{f}(x) \\
\text{subject to} \quad & \hat{g}_i(x) \le 0, \quad i = 1, \dots, m \\
& \hat{h}_i(x) = 0, \quad i = 1, \dots, p \\
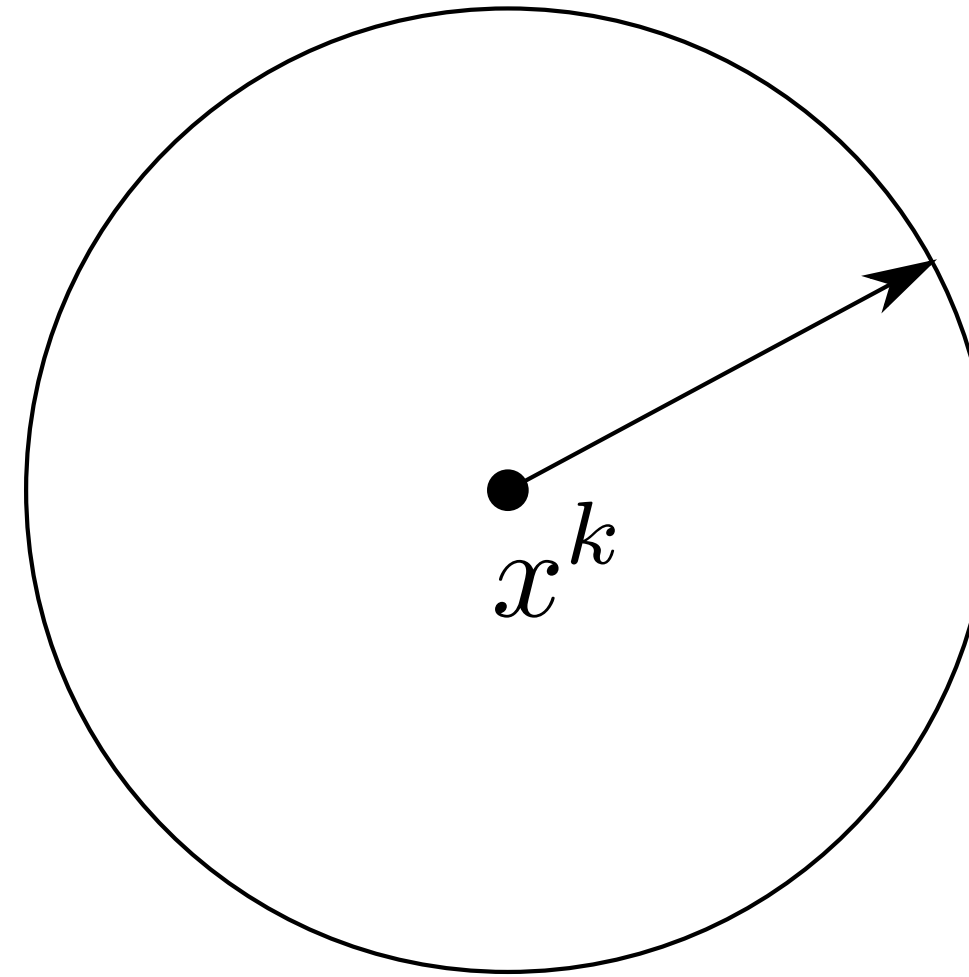& x \in \mathcal{T}^k
\end{aligned}$$

solve to get x^{k+1}

- $\hat{f}$ ($\hat{g}_i$) is a convex approximation of f (g_i) over $\mathcal{T}^k$
- $\hat{h}_i$ is an affine approximation of h_i over $\mathcal{T}^k$

The trust region

$$\mathcal{T}^k = \{x \mid \|x - x^k\| \leq \rho\}$$

Ball $\quad \mathcal{T}^k = \{x \mid \|x - x^k\|_2 \leq \rho\}$

Box $\quad \mathcal{T}^k = \{x \mid |x_i - x_i^k| \leq \rho_i\}$

Note: if f, g_i h_i are convex or affine in x_i, then we can take $\rho_i = \infty$

$$\mathcal{L}(x, \mu) = f(x) + \mu^{\gamma}\left(\|x - x^k\|_2 - \rho\right)$$

Proximal operator interpretation

proximal problem

minimize $\quad f(x) + \dfrac{1}{2\lambda}\|x - x^k\|_2^2$

trust region problem

minimize $\quad f(x)$

subject to $\quad \|x - x^k\|_2 \le \rho$

optimality conditions

$$0 \in \partial f(x^{\mathrm{pr}}) + \frac{1}{\lambda}(x^{\mathrm{pr}} - x^k)$$

$$\boxed{\lambda = \rho/\mu}$$

optimality conditions

$$0 \in \partial f(x^{\mathrm{tr}}) + \mu\frac{x^{\mathrm{tr}} - x^k}{\|x^{\mathrm{tr}} - x^k\|_2}$$

$$\|x^{\mathrm{tr}} - x^k\|_2 = \rho$$

Note

- Minimum outside tr: $\|x^{\mathrm{tr}} - x^k\| = \rho$
- $\partial\|z\|_2 = \nabla(z^T z)^{1/2} = z/\|z\|$ (if $z \ne 0$)

12

Building convex approximations

Convex Taylor expansions

$$(v)_+ = \max(0, v)$$

Given nonconvex function f

First order
$$\hat{f}(x) = f(x^k) + \nabla f(x^k)^T (x - x^k)$$

Second order
$$\hat{f}(x) = f(x^k) + \nabla f(x^k)^T (x - x^k) + (1/2)(x - x^k)^T P_+ (x - x^k)$$

where $P_+ = \Pi_{\mathbf{S}_+}(\nabla^2 f(x)) = U(\mathbf{diag}(\lambda))_+ U^T$ **positive semidefinite cone projection**

Local approximation
it does not depend on trust-region radius ρ

Quasi-linearization

Very easy and cheap method for affine approximation

$$\text{write } h \text{ as } h(x) = A(x)x + b(x)$$

$$\downarrow$$

$$\text{use } \hat{h}(x) = A(x^k)x + b(x^k)$$

Example $\quad f(x) = (1/2)x^T P x + q^T x = ((1/2)P x + q)^T x + r$

$\quad \hat{h}(x)$

Quasi-linear: $\quad \hat{h} = ((1/2)P x^k + q)^T x + r$

$\quad \hat{h}(x)$

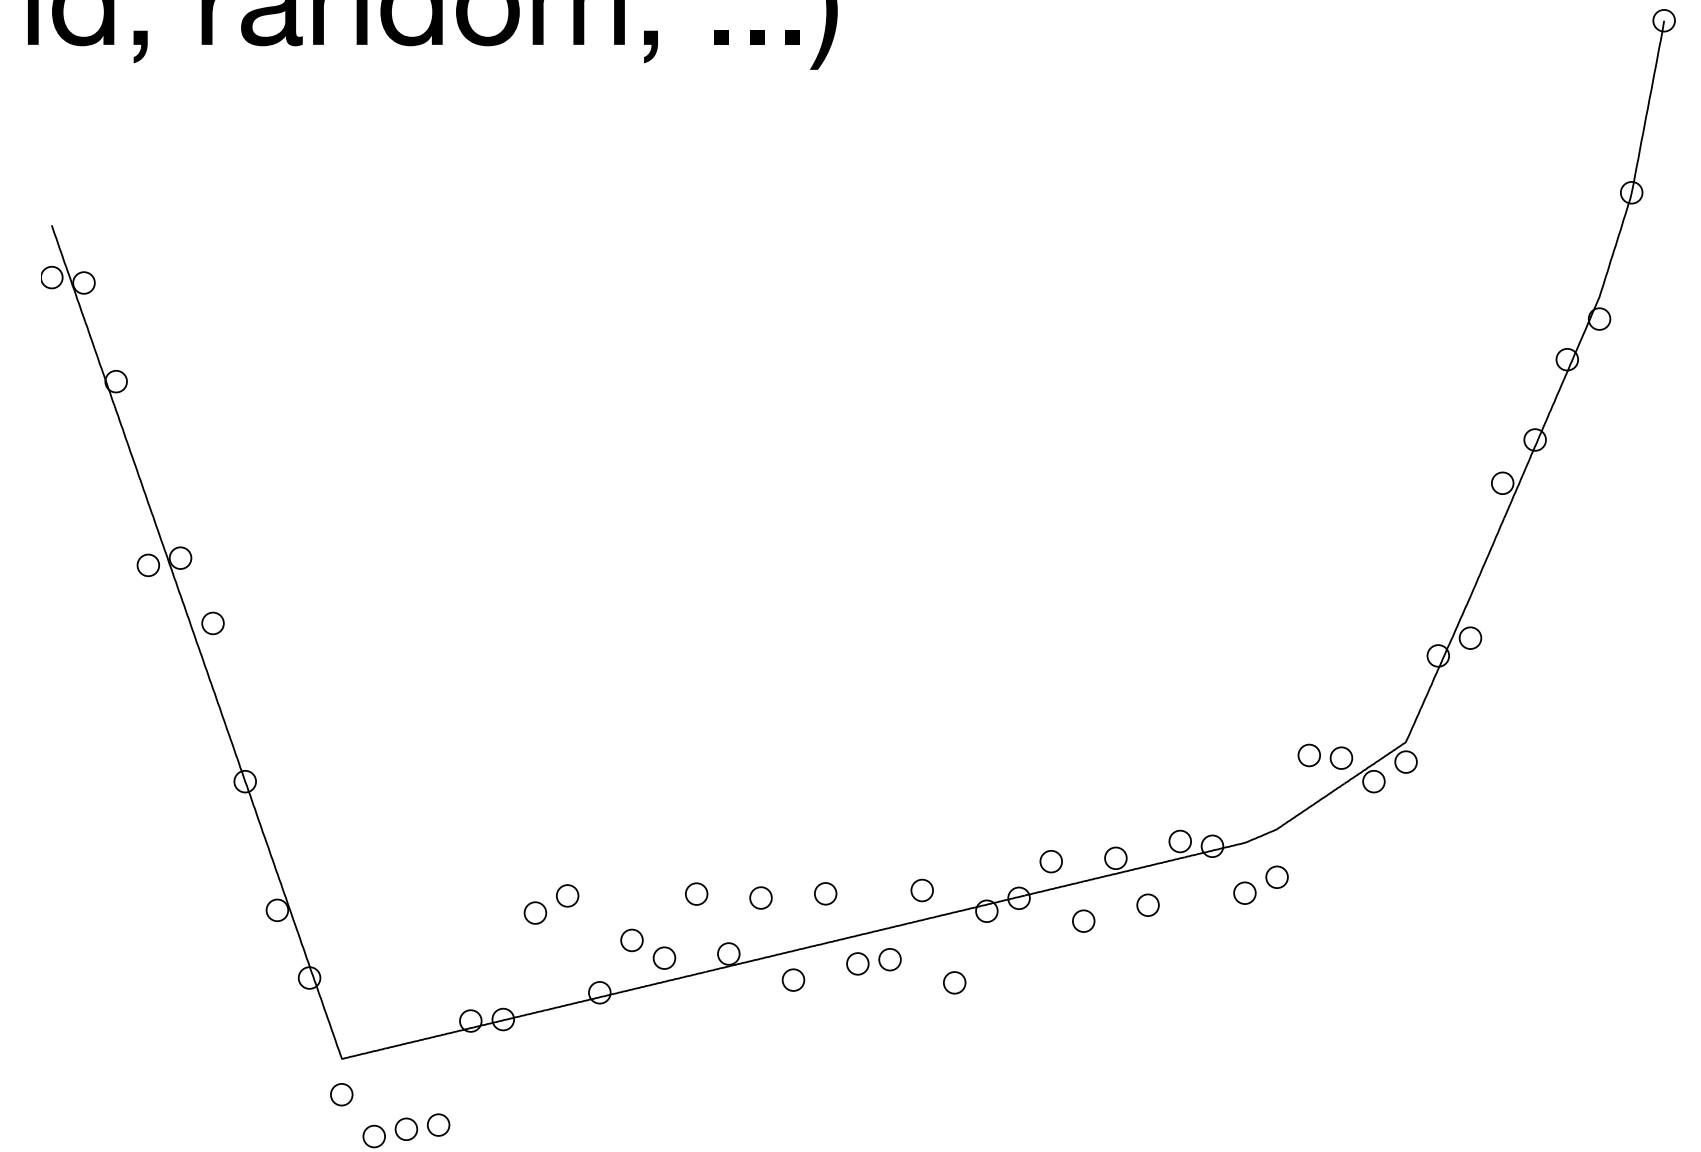Taylor: $\quad \hat{h} = h(x^k) + (P x^k + q)^T(x - x^k)$

Local approximation
it does not depend on trust-region radius ρ

Particle methods

Idea

- Choose points $z_1, \ldots, z_K \in \mathcal{T}^k$ (e.g., verticles, grid, random, ...)
- Evaluate function $y_i = f(z_i)$
- Fit data (z_i, y_i) with convex functions
  (convex optimization)

Advantages

- Nondifferentiable functions
- **regional models**: they depend on current x^k and radii ρ_i

Particle methods

$$\hat{f}(x) = \max_i \{\hat{y}_i + g_i^T(x - z_i)\}$$

Fit piecewise linear functions to data

$\hat{y}_i$ act as function values $\hat{f}(z_i)$

g_i act as subgradients $\partial \hat{f}(z_i)$

Fitting problem

minimize $\quad \sum_{i=1}^{K}(\hat{y}_i - y_i)^2$

subject to $\quad \hat{y}_j \geq \hat{y}_i + g_i^T(z_j - z_i), \quad i, j = 1, \ldots, K$ **convexity**

$\quad\quad\quad\quad \hat{y}_i \leq y_i, \quad i = 1, \ldots, K$ **lower bound**

$f(x) = x^4 - 2x^3 + 0.3x$

5 random

12 random

uniform grid

Particle methods

Fit quadratic functions to data

$$\hat{f}(x) = (1/2)(x - x^k)^T P (x - x^k) + q^T (x - x^k) + r$$

Fitting problem

$$\text{minimize} \quad \sum_{i=1}^{K} ((1/2)(z_i - x^k)^T P (z_i - x^k) + q^T (z_i - z^k) + r - y_i)^2$$

$$\text{subject to} \quad P \succeq 0$$

Remarks

- No necessarily upper/lower bound

- We can add other objectives, convex constraints and norm penalties

- Can be more sample efficient than piecewise linear

- Need to solve a **convex problem for every function at every SCP iteration**

Trust region example

Example: nonconvex quadratic program

$$\text{minimize} \quad f(x) = (1/2)x^T P x + q^T x$$

$$\text{subject to} \quad \|x\|_\infty \leq 1$$

P is symmetric but not positive semidefinite

Taylor approximation

$$\hat{f}(x) = f(x^k) + (Px^k + q)^T(x - x^k) + (1/2)(x - x^k)^T P_+(x - x^k)$$

Example: nonconvex quadratic program

Lower bound via convex duality

minimize $\quad f(x) = (1/2)x^T P x + q^T x$

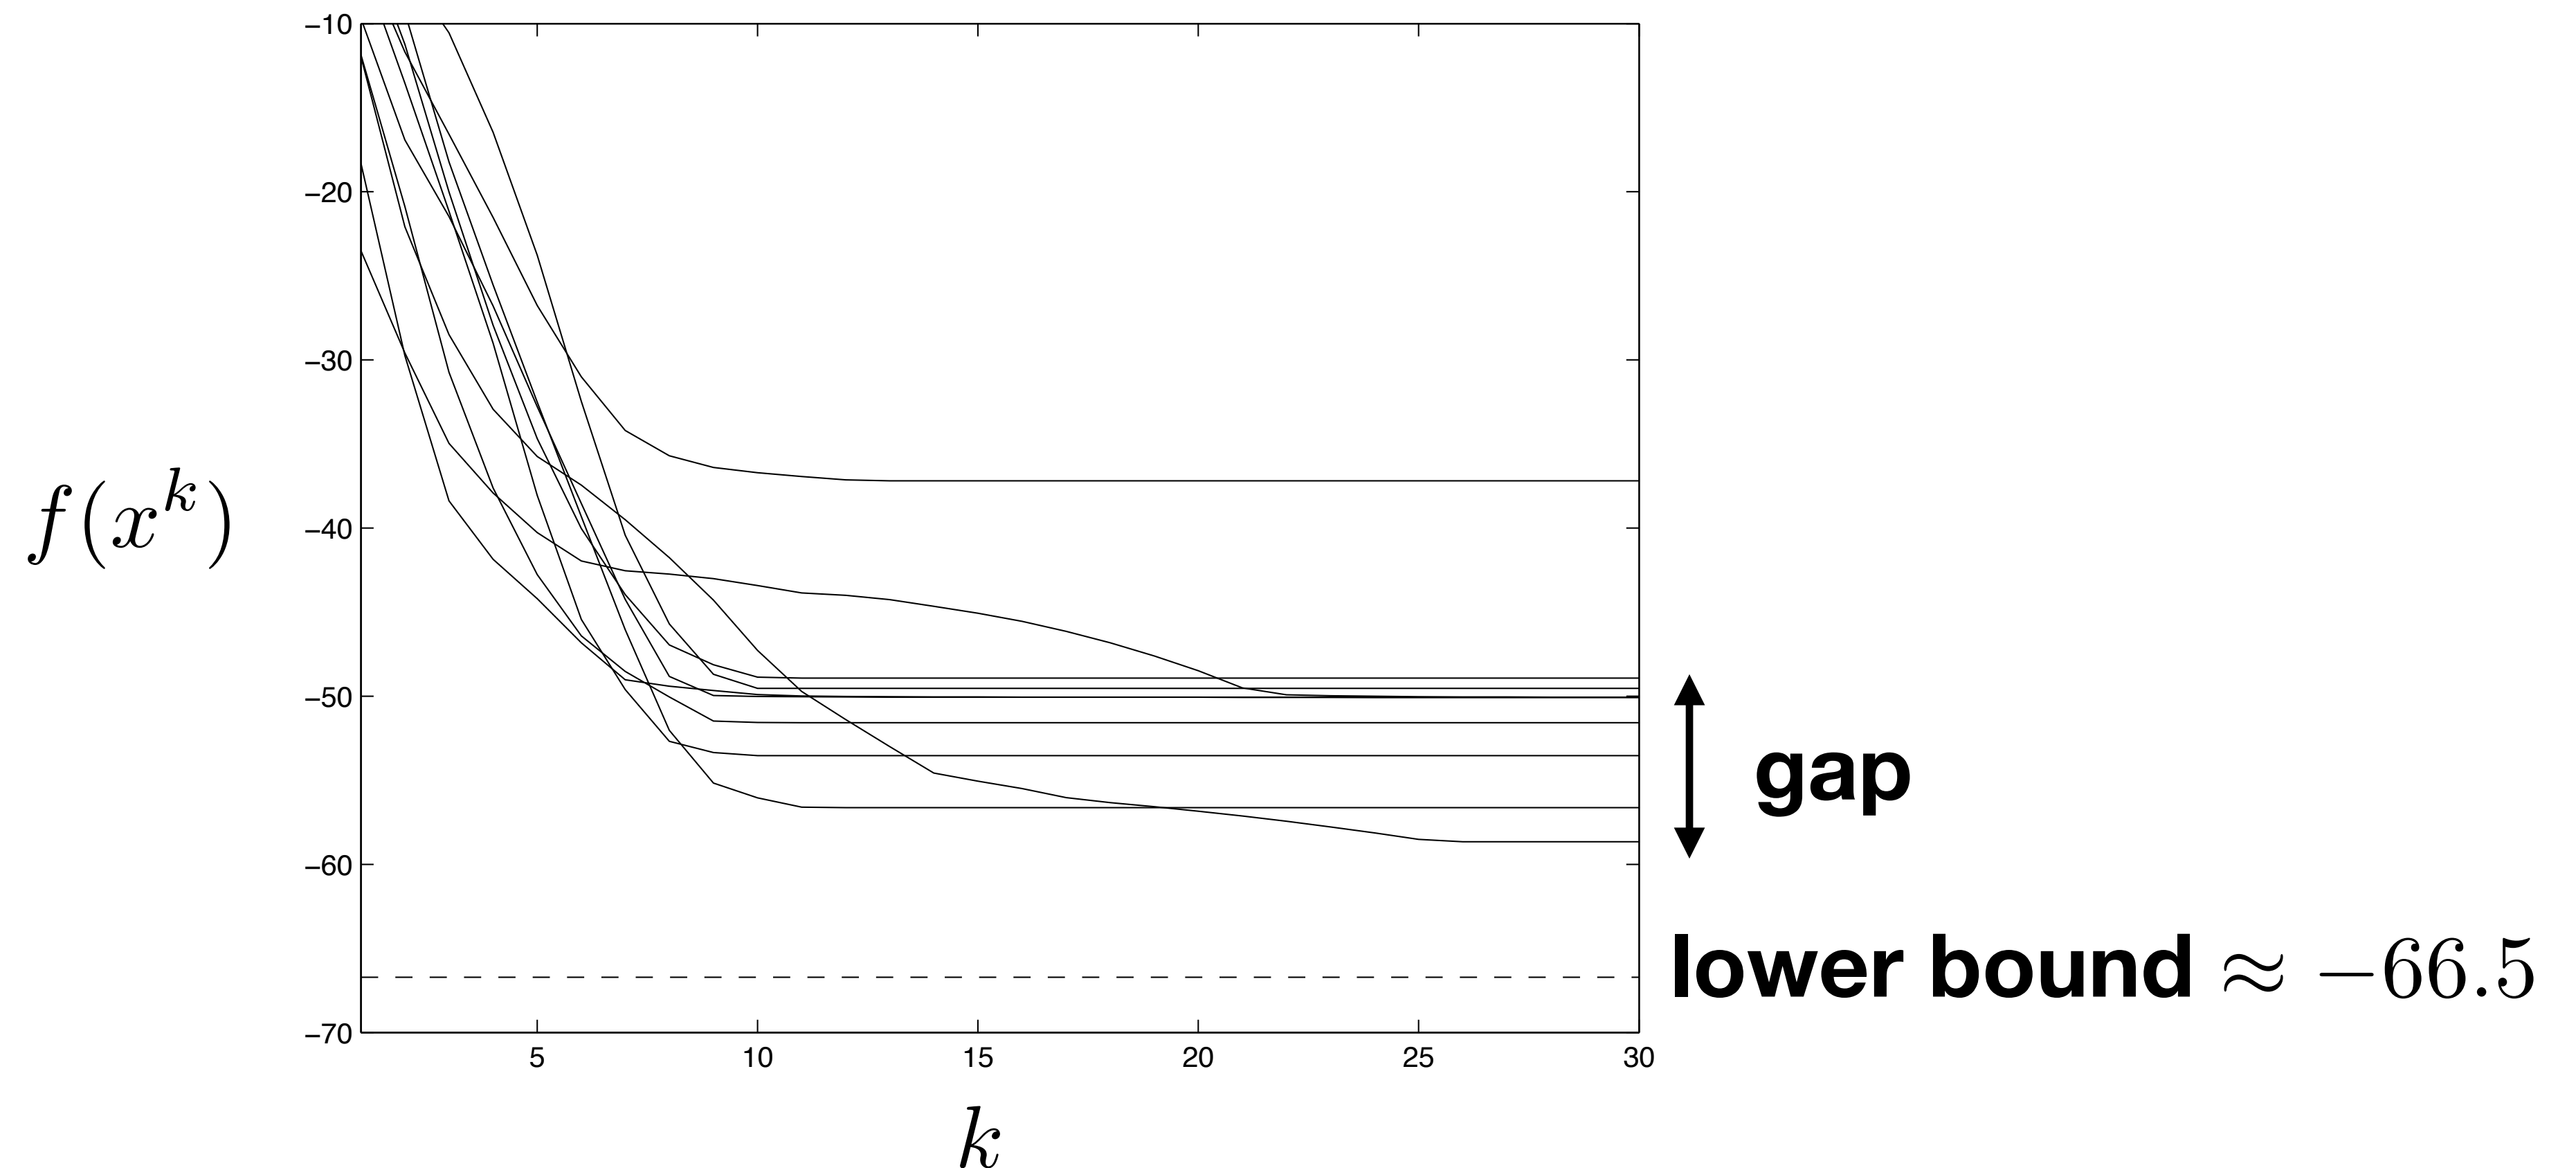subject to $\quad \|x\|_\infty \leq 1$ $\longleftrightarrow$ $x_i^2 - 1 \leq 0$

Lagrangian

$$L(x, \lambda) \quad = (1/2)x^T P x + q^T x + \sum_{i=1}^{n} \lambda_i(x_i^2 - 1)$$
$$= (1/2)x^T(P + 2\mathbf{diag}(\lambda))x + q^T x - \mathbf{1}^T \lambda$$

Dual problem (always convex)

minimize $\quad -(1/2)q^T(P + 2\mathbf{diag}(\lambda))^{-1}q - \mathbf{1}^T \lambda \qquad g(\lambda)$

$\lambda \geq 0$

Example: nonconvex quadratic program

SCP with $\rho = 0.2$ with 10 different random $x_0 \in \mathbf{R}^n$

$f(x^k)$ vs k, with **gap** and **lower bound** ≈ -66.5

Regularized trust region methods

Issues with vanilla sequential convex programming

minimize $\quad f(x)$

subject to $\quad g_i(x) \le 0, \quad i = 1, \dots, m$

$\qquad\qquad h_i(x) = 0, \quad i = 1, \dots, p$

$\longrightarrow$

minimize $\quad \hat{f}(x)$

subject to $\quad \hat{g}_i(x) \le 0, \quad i = 1, \dots, m$

$\qquad\qquad \hat{h}_i(x) = 0, \quad i = 1, \dots, p$

$\qquad\qquad x \in \mathcal{T}^k$

Infeasibility

Approximate problem can be infeasible (e.g. too small ρ)

Evaluate progress
when x^k infeasible

- Objective: $f(x^k)$
- Inequality violations: $g_i(x^k)_+$
- Equality violations: $|h_i(x^k)|$

Controlling trust region size

- ρ **too large**
  poor approximations $\to$ bad x^{k+1}
- ρ **too small**
  good approximations $\to$ slow progress

Exact penalty formulation

Solve unconstrained problem instead of the original problem

minimize $\quad \phi(x) = f(x) + \lambda \left(\sum_{i=1}^{m} (g_i(x))_+ + \sum_{i=1}^{p} |h_i(x)| \right), \qquad \lambda > 0$

For λ large enough $\longrightarrow$ $x^\star = \operatorname{argmin} \phi(x)$ solves the original problem

($\lambda > \|y^\star\|_\infty$ where $y^\star$ is the dual variable satisfying the KKT conditions)

SCP solves the convex approximation (always feasible)

$$\hat{\phi}(x) = \hat{f}(x) + \lambda \left(\sum_{i=1}^{m} (\hat{g}_i(x))_+ + \sum_{i=1}^{p} |\hat{h}_i(x)| \right)$$

If λ not large enough, we have **sparse violations**

Trust region update

Idea judge progress in ϕ using $\hat{x} = \operatorname{argmin} \hat{\phi}(x)$

Exact decrease

$$\delta = \phi(x^k) - \phi(\hat{x})$$

Approximate decrease

$$\hat{\delta} = \phi(x^k) - \hat{\phi}(\hat{x})$$

Updates

$\delta \geq \alpha\hat{\delta} \longrightarrow$
- accept: $x^{k+1} = \hat{x}$
- increase region $\rho = \beta^{\mathrm{acc}}\rho$

$\delta < \alpha\hat{\delta} \longrightarrow$
- reject: $x^{k+1} = x^k$
- decrease region $\rho = \beta^{\mathrm{rej}}\rho$

Parameters

tolerance α (e.g., $= 0.1$)
accept multiplier $\beta^{\mathrm{acc}} \geq 1$ (e.g., $= 1.1$)
reject multiplier $\beta^{\mathrm{rej}} \in (0, 1)$ (e.g., 0.5)

Interpretation

If actual decrease δ is more than α fraction of predicted decrease $\hat{\delta}$
then increase trust region size (longer steps). Otherwise decrease it.

Regularized trust region example

Nonlinear optimal control
Robotic arm

2-dimensional system

no gravity (horizontal)

controlled torques τ_1, τ_2

l_2, m_2

τ_2

θ_2

τ_1

l_1, m_1

θ_1

Nonlinear optimal control

The problem

minimize $\quad J = \displaystyle\int_0^T \|\tau(t)\|_2^2 \mathrm{d}t \qquad$ **minimum torque**

subject to $\quad \theta(0) = \theta_{\text{init}}, \quad \theta(T) = \theta_{\text{final}} \qquad$ **position**

$\quad\quad\quad\quad\quad \dot\theta(0) = 0, \quad \dot\theta(T) = 0 \qquad$ **velocity**

$$\|\tau(t)\|_\infty \leq \tau_{\max}, \quad 0 \leq t \leq T$$

Dynamics

$$M(\theta)\ddot\theta + W(\theta,\dot\theta)\dot\theta = \tau$$

$\downarrow$

Not convex!
(Hard to optimize)

Note: cheap to simulate

$$M(\theta) = \begin{bmatrix} (m_1 + m_2)l_1^2 & m_2 l_1 l_2 (s_1 s_2 + c_1 c_2) \\ m_2 l_1 l_2 (s_1 s_2 + c_1 c_2) & m_2 l_2^2 \end{bmatrix}$$

$$W(\theta,\dot\theta) = \begin{bmatrix} 0 & m_2 l_1 l_2 (s_1 c_2 - c_1 s_2)\dot\theta_2 \\ m_2 l_1 l_2 (s_1 c_2 - c_1 s_2)\dot\theta_1 & 0 \end{bmatrix}$$

where $s_i = \sin(\theta_i)$ and $c_i = \cos(\theta_i)$

Nonlinear optimal control

Discretization

Discretize with **time intervals** $h = T/N$

Objective
$$J = \int_0^T \|\tau(t)\|_2^2 \mathrm{d}t \approx h \sum_{i=1}^N \|\tau_i\|_2^2, \quad \text{with} \quad \tau_i = \tau(ih)$$

Dynamics: approximate derivatives

$$M(\theta)\ddot{\theta} + W(\theta, \dot{\theta})\dot{\theta} = \tau$$

$$\dot{\theta}(ih) \approx \frac{\theta_{i+1} - \theta_{i-1}}{2h}$$

$$\ddot{\theta}(ih) \approx \frac{\theta_{i+1} - 2\theta_i + \theta_{i-1}}{h^2}$$

zero initial velocities

$$\theta_0 = \theta_1 = \theta_{\text{init}}$$
$$\theta_N = \theta_{N+1} = \theta_{\text{final}}$$

nonlinear equality constraints

$$M(\theta_i)\frac{\theta_{i+1} - 2\theta_i + \theta_{i-1}}{h^2} + W\left(\theta_i, \frac{\theta_{i+1} - \theta_{i-1}}{2h}\right)\frac{\theta_{i+1} - \theta_{i-1}}{2h} = \tau_i$$

30

Nonlinear optimal control

Convexification

$$\text{minimize} \quad h \sum_{i=1}^{N} \|\tau_i\|_2^2$$

$$\text{subject to} \quad \theta_0 = \theta_1 = \theta_{\text{init}}, \quad \theta_N = \theta_{N+1} = \theta_{\text{final}}$$

$$\|\tau_i\|_\infty \leq \tau_{\max}$$

$$M(\theta_i)\frac{\theta_{i+1} - 2\theta_i + \theta_{i-1}}{h^2} + W\left(\theta_i, \frac{\theta_{i+1} - \theta_{i-1}}{2h}\right)\frac{\theta_{i+1} - \theta_{i-1}}{2h} = \tau_i$$

Quasi-linearization of the dynamics around previous x^k

$$M(\theta_i^k)\frac{\theta_{i+1} - 2\theta_i + \theta_{i-1}}{h^2} + W\left(\theta_i^k, \frac{\theta_{i+1}^k - \theta_{i-1}^k}{2h}\right)\frac{\theta_{i+1} - \theta_{i-1}}{2h} = \tau_i$$

Remarks

- trust region only on θ_i (cost and constraints convex in τ_i)
- initialize with straight line: $\theta_i = \frac{i-1}{N-1}(\theta_{\text{final}} - \theta_{\text{init}}), \quad i = 1, \ldots, N$

31

Nonlinear optimal control

Example

System

- $m_1 = 1$, $m_2 = 5$, $l_1 = l_2 = 1$
- $N = 40$, $T = 10$
- $\theta_{\text{init}} = (0, -2.9), \quad \theta_{\text{final}} = (3, 2.9)$
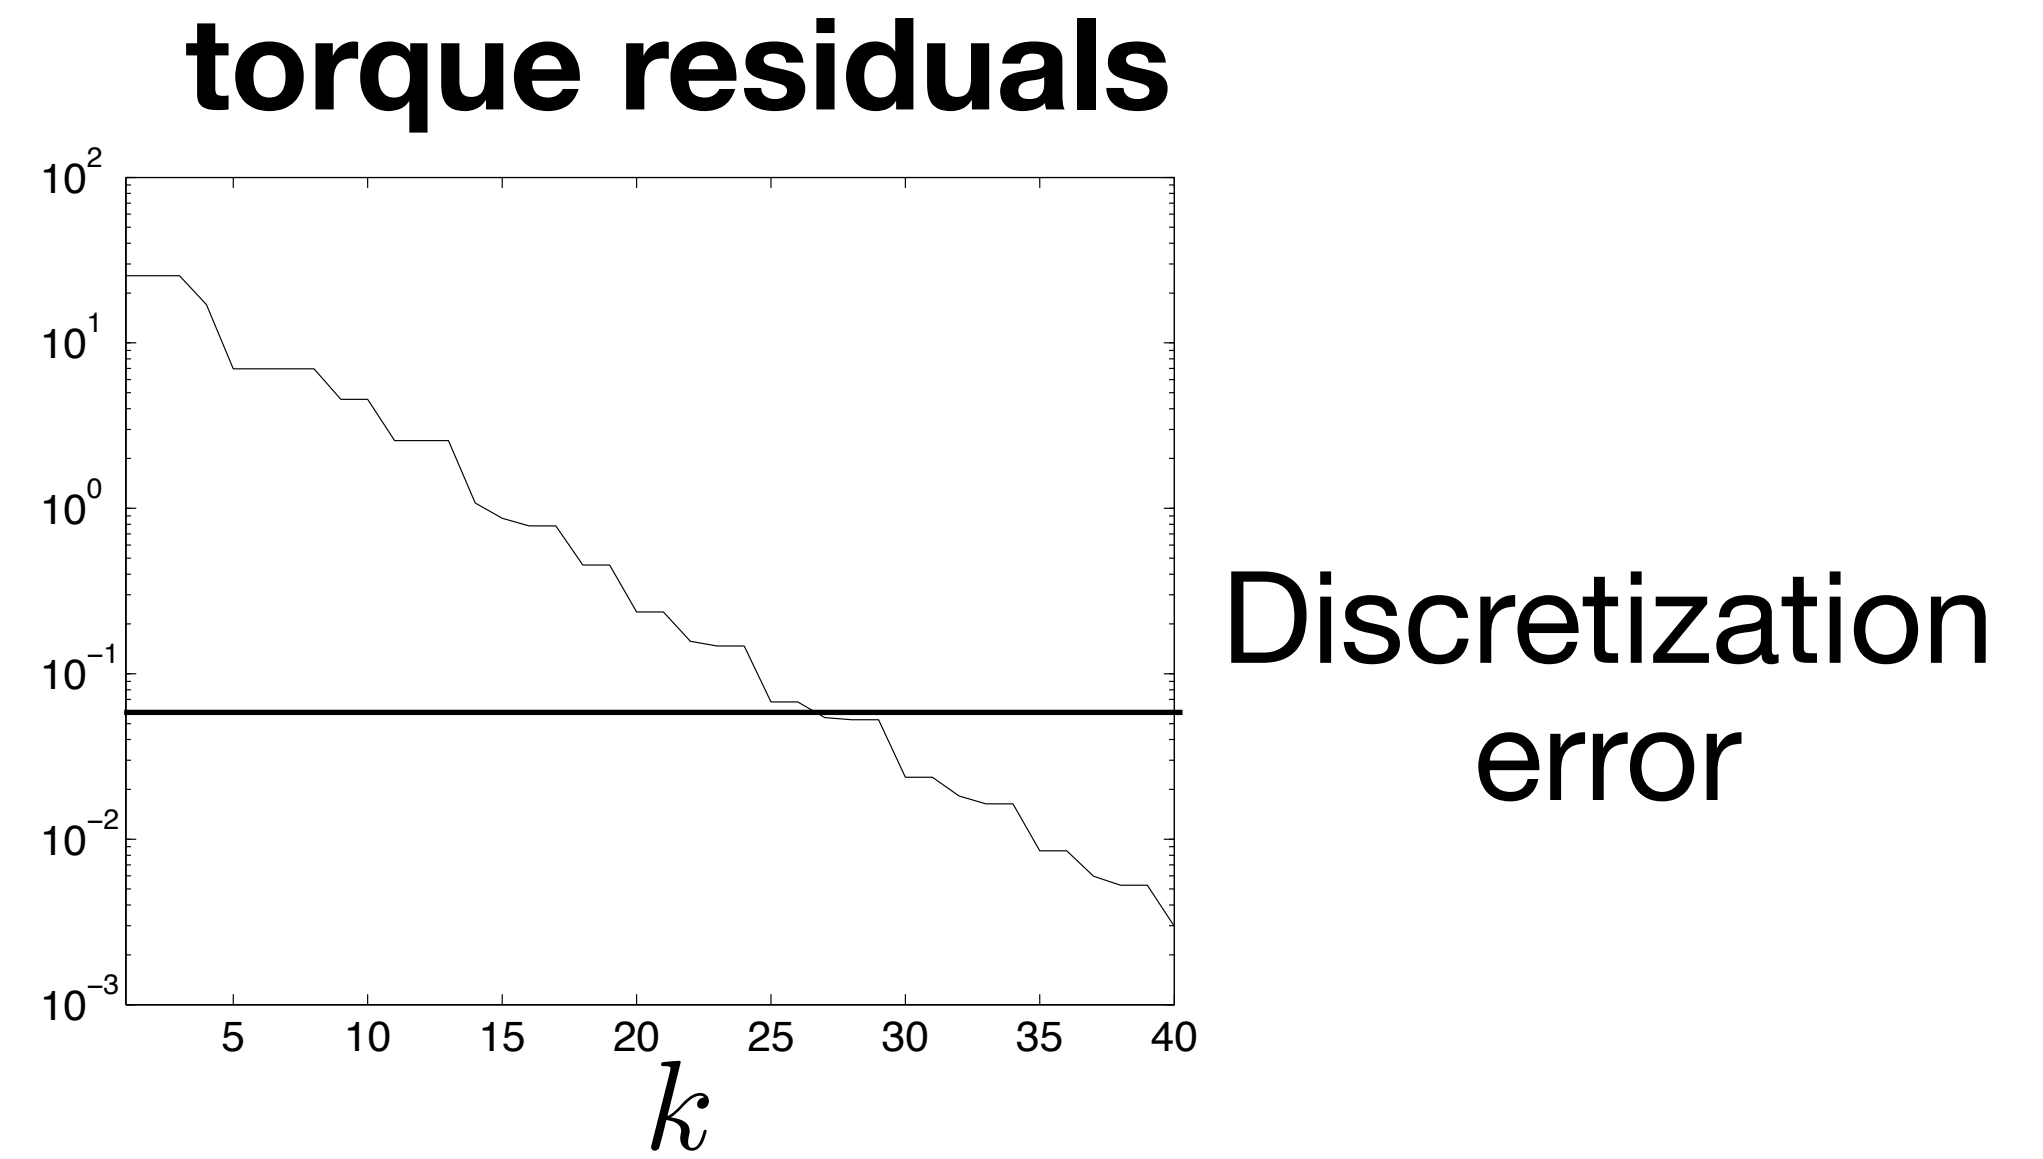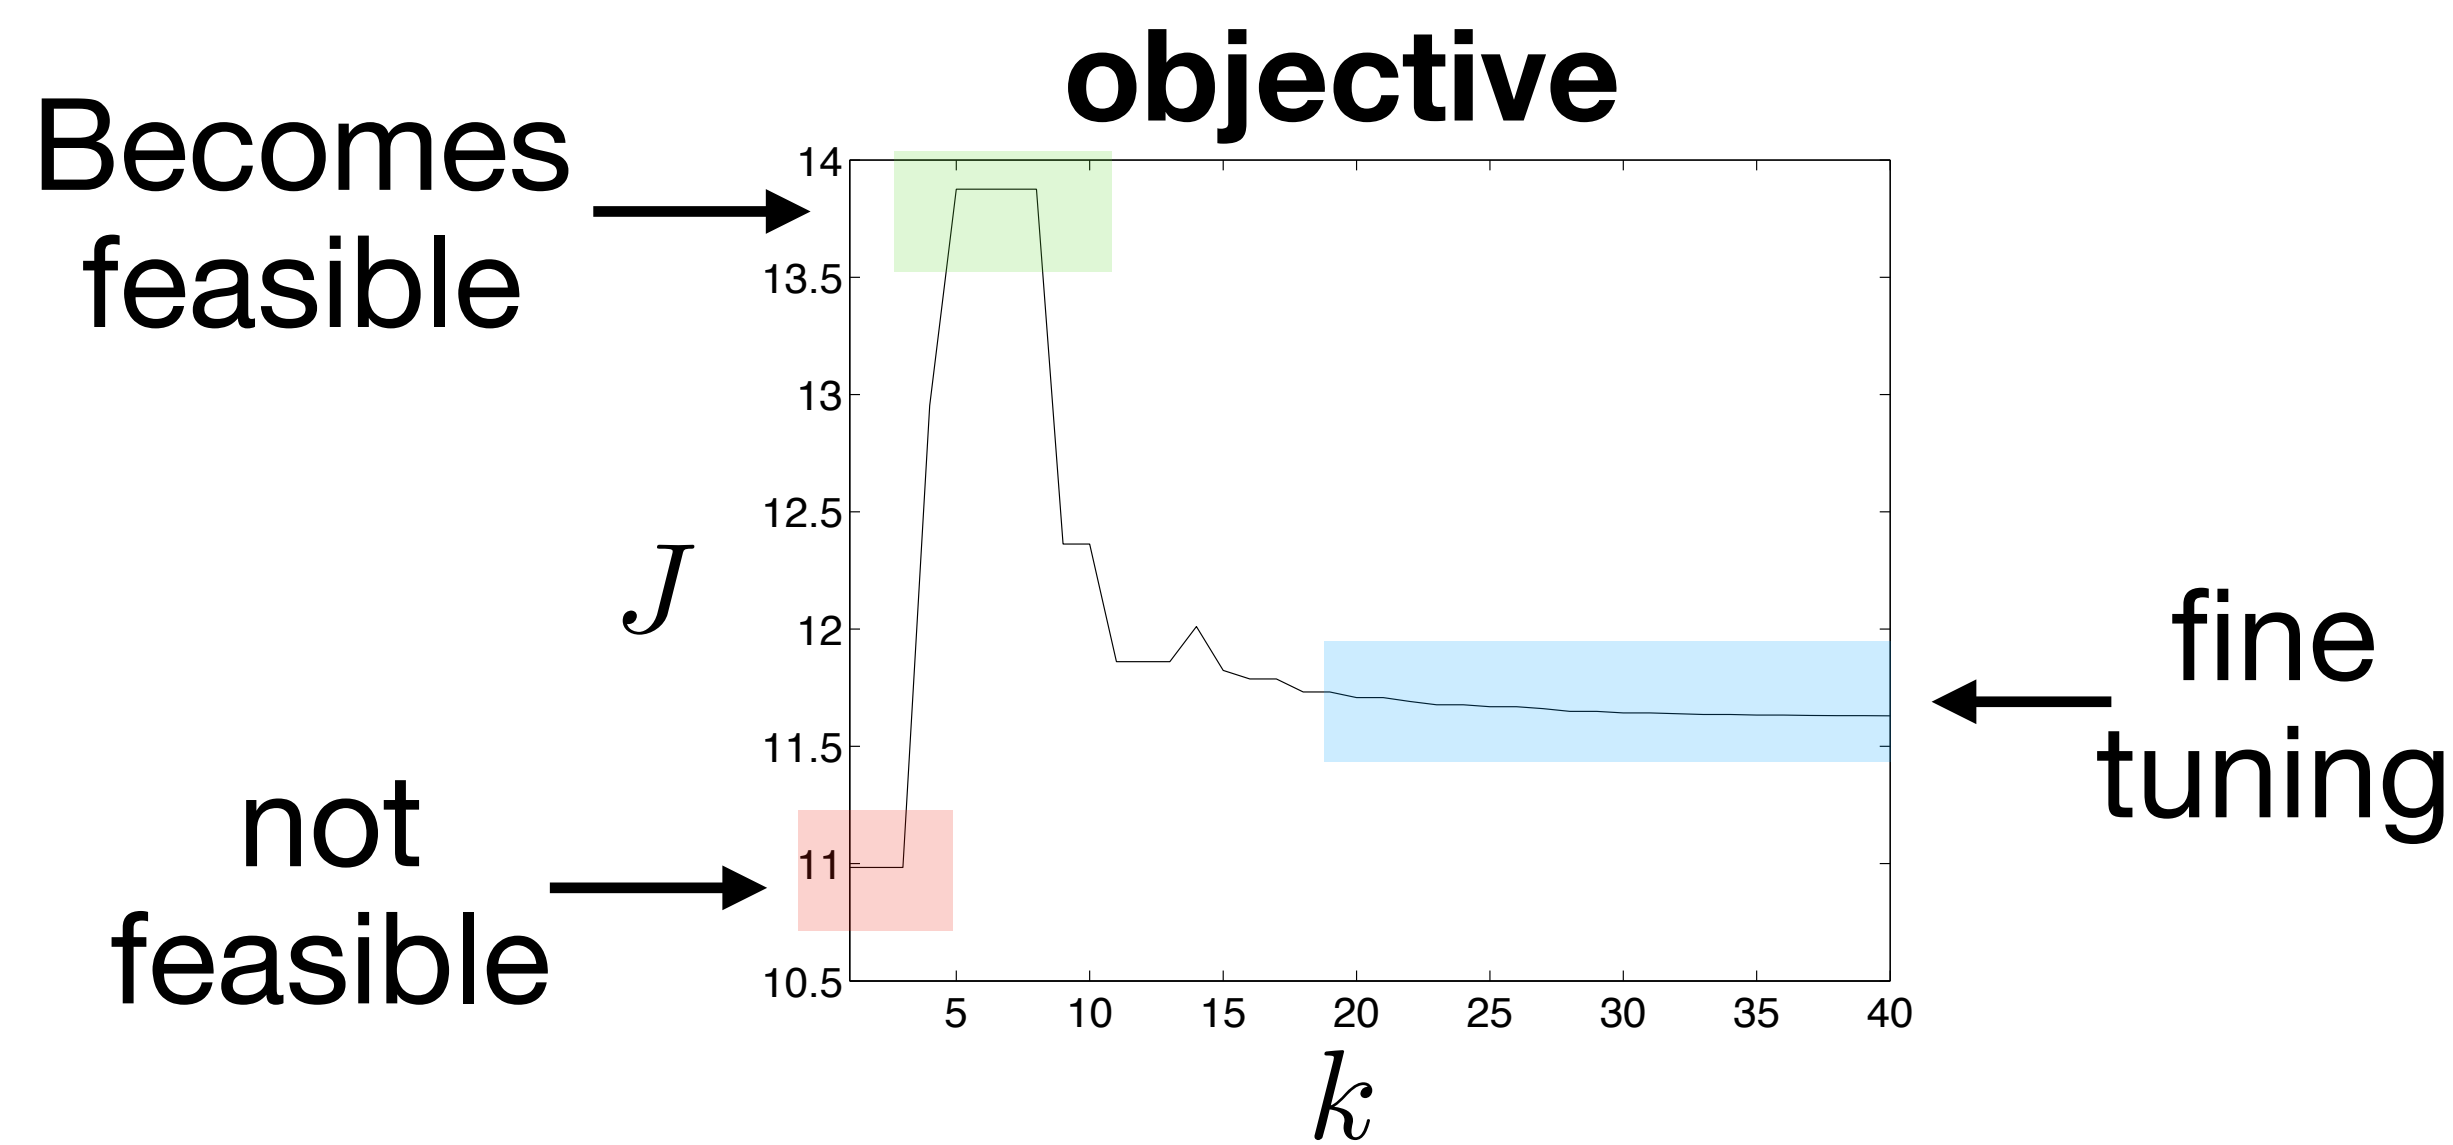- $\tau_{\max} = 1.1$

Algorithm

- $\lambda = 2$
- $\alpha = 0.1$, $\beta^{\text{acc}} = 1.1$, $\beta^{\text{rej}} = 0.5$
- $\rho_1 = 90°$ (very large)

Progress

**Note:** does not go to 0

Nonlinear optimal control

objective

Becomes feasible

not feasible

fine tuning

J

k

torque residuals

Discretization error

k

decrease in ϕ

$\hat{\delta}$: (dashed)

δ : (solid)

k

trust region size

ρ^k

k

Nonlinear optimal control

Trajectories

Difference of convex programming

Difference of convex programming

minimize $\quad f_0(x) - g_0(x)$

subject to $\quad f_i(x) - g_i(x) \leq 0, \quad i = 1, \ldots, m$

Difference of convex functions

where f_i and g_i are convex

Very powerful
it can represent any twice differentiable function

Hard
nonconvex problem unless g_i are affine

[On Functions Representable As A Difference Of Convex Functions, Hartman]

Difference of convex programming

Convexification

$$-g(x^k) - \nabla g(x^k)^T (x - x^k)$$

x^k

$-g(x)$

**Convexify** $f(x) - g(x)$

$$f(x) - \hat{g}(x) = f(x) - g(x^k) - \nabla g(x^k)^T (x - x^k)$$

$$f(x) - g(x) \leq f(x) - \hat{g}(x)$$

Remarks

• True objective better than convexified objective

• True feasible set contains convexified feasible set $\longrightarrow$ **No trust region needed**

Difference of convex programming
Iterations

Convex-concave procedure

1. Convexify: form $\hat{g}_i(x) = g_i(x^k) + \nabla g_i(x^k)^T(x - x^k)$ for $i = 0, \dots, m$

2. Solve to obtain x^{k+1}

$$\begin{aligned} \text{minimize} \quad & f_0(x) - \hat{g}_0(x) \\ \text{subject to} \quad & f_i(x) - \hat{g}_i(x) \leq 0 \end{aligned}$$

Remarks

It always converges to a stationary point (it might be a maximum)

[Variations and extension of the convex–concave procedure, Lipp, Boyd]

Path planning example

Find shortest path connecting a and b in $\mathbf{R}^d$

Avoid circles centered at c_j with radius r_j with $j = 1, \ldots, m$

$$\begin{aligned}
\text{minimize} \quad & L \\
\text{subject to} \quad & x_0 = a, \quad x_n = b \\
& \|x_i - x_{i-1}\|_2 \leq L/n, \quad i = 1, \ldots, n \\
& \|x_i - c_j\|_2 \geq r_j, \quad i = 1, \ldots, n, \quad j = 1, \ldots, m
\end{aligned}$$

path lengths ——

obstacle constraints
(not convex) ——

Path planning example

minimize L

subject to $x_0 = a, \quad x_n = b$

$\|x_i - x_{i-1}\|_2 \leq L/n, \quad i = 1, \ldots, n$

$\|x_i - c_j\|_2 \geq r_j, \quad i = 1, \ldots, n, \quad j = 1, \ldots, m$

Dimension: $d = 2$ Steps: $n = 50$

It converges in 26 iterations (convex problems)

[Disciplined Convex-Concave Programming, Shen, Diamond, Gu, Boyd]

Sequential convex programming

Today, we learned to:

- **Familiarize** with concepts of sequential convex programming

- **Develop** trust region algorithms

- **Build** convex approximations of nonlinear/nonsmooth functions

- **Develop** regularized trust region methods to account for infeasibility

- **Recognize** difference-of-convex programs and **apply** convex-concave procedure

Next lecture

- Branch and bound algorithms