

Convexity Exploiting Newton-Type Optimization for Learning and Control

Moritz Diehl

Systems Control and Optimization Laboratory

Department of Microsystems Engineering and Department of Mathematics
University of Freiburg

joint work with Florian Messerer and Katrin Baumgärtner

IPAM Workshop on
Intersections between Control, Learning and Optimization
February 24, 2020

Overview

- Embedded Optimization
 - Universal Approximation Theorem for Convex Optimization
 - Model Predictive Control (and two Applications)
- Convexity Exploiting Newton-Type Optimization
 - Sequential Convex Programming (SCP)
 - Generalized Gauss-Newton (GGN)
 - Sequential Convex Quadratic Programming (SCQP)
 - Local Convergence Analysis and Desirable Divergence
- Zero-Order Optimization-based Iterative Learning Control
 - Tutorial Example
 - Bounding the Loss of Optimality
 - Local Convergence Analysis

Complex Sensor Actuator Systems

SENSORS


- GPS
- acceleration
- radar
- vision
- ...



How to connect ?

ACTUATORS

- flight surfaces
- steering wheel
- motor speeds
- torques
- ...



Complex Sensor Actuator Systems

SENSORS

- GPS
- acceleration
- radar
- vision
- ...



How to connect ?
Linear Filters ?

ACTUATORS

- flight surfaces
- steering wheel
- motor speeds
- torques
- ...



Complex Sensor Actuator Systems

SENSORS

- GPS
- acceleration
- radar
- vision
- ...



How to connect ?
Linear Filters ?
Deep Neural Networks ?

ACTUATORS

- flight surfaces
- steering wheel
- motor speeds
- torques
- ...



Complex Sensor Actuator Systems

SENSORS


- GPS
- acceleration
- radar
- vision
- ...

EMBEDDED OPTIMIZATION




ACTUATORS

- flight surfaces
- steering wheel
- motor speeds
- torques
- ...



Complex Sensor Actuator Systems



Solve, in real-time and repeatedly, an optimization problem that depends on the incoming stream of input data, to generate a stream of output data.

Embedded Optimization: a computationally expensive map

Example: Parametric Quadratic Programming (pQP)

$$\boxed{x} \rightarrow \boxed{u = \arg \min_u \begin{bmatrix} x \\ u \end{bmatrix}^T \begin{bmatrix} Q & S \\ S^T & R \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} + \begin{bmatrix} x \\ u \end{bmatrix}^T \begin{bmatrix} g \\ h \end{bmatrix} \text{ s.t. } Au + Bx \leq b} \rightarrow \boxed{u}$$

Note: pQP problem is jointly convex in (x, u) . This fact can be exploited algorithmically, as e.g. done in the solver qpOASES [Ferreau et al., 2008-2014]

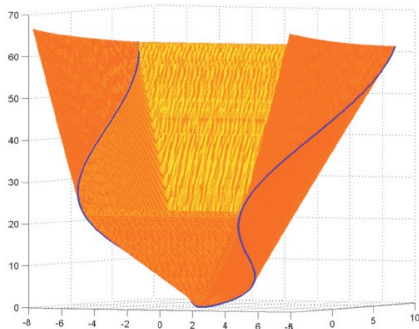
Example: Parametric Quadratic Programming (pQP)

$$\boxed{x} \rightarrow \boxed{u = \arg \min_u \begin{bmatrix} x \\ u \end{bmatrix}^T \begin{bmatrix} Q & S \\ S^T & R \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} + \begin{bmatrix} x \\ u \end{bmatrix}^T \begin{bmatrix} g \\ h \end{bmatrix} \text{ s.t. } Au + Bx \leq b} \rightarrow \boxed{u}$$

Note: pQP problem is jointly convex in (x, u) . This fact can be exploited algorithmically, as e.g. done in the solver qpOASES [Ferreau et al., 2008-2014]

Universal approximation theorem of convex optimization:

Every continuous map $\mu : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_u}$ can be obtained as minimizer of a *parametric convex program*:




$$\mu(x) = \arg \min_u g(x, u) \text{ s.t. } (x, u) \in \Gamma$$

where g and Γ are jointly convex in (x, u)

[Baes, D., Necoara, IEEE TAC, 2008]

Model Predictive Control (MPC)

Always look a bit into the future




Example: driver predicts and optimizes, and therefore slows down before a curve

MPC from RL/DP perspective:

“Multistep Lookahead for Deterministic Systems with Constraints”

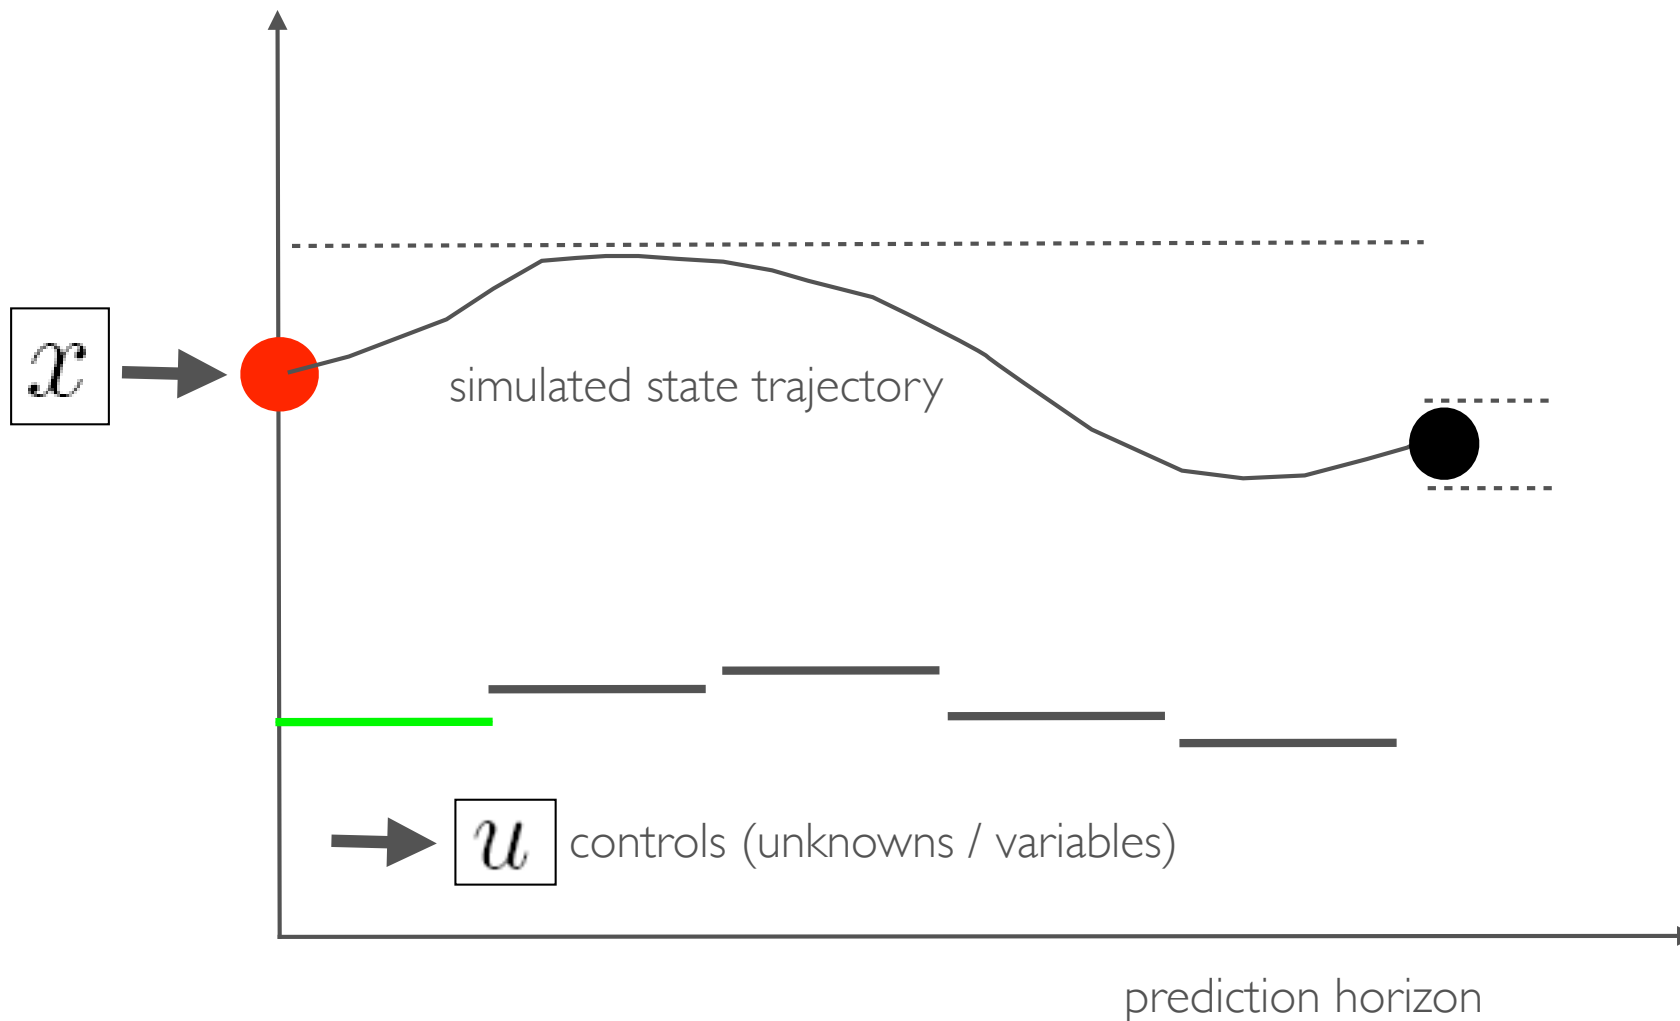
Optimal Control Problem in MPC

For given system state \mathbf{x} , which controls \mathbf{u} lead to the best objective value without violation of constraints?




Optimal Control Problem in MPC

For given system state \mathbf{x} , which controls \mathbf{u} lead to the best objective value without violation of constraints ?



Model Predictive Control of RC Race Cars (in Freiburg)



Minimize least squares distance to centerline, respect constraints. Use nonlinear embedded optimization software *acados* coupled to ROS, sample at 100 Hz.

[Kloeser et al., submitted]

Time-Optimal Point-To-Point Motions [PhD Vandebroek 2011]



Fast oscillating systems (cranes, plotters, wafer steppers, ...)


Control aims:

- reach end point as fast as possible
- do not violate constraints
- no residual vibrations

Idea: formulate as embedded optimization problem
in form of Model Predictive Control (MPC)




Time Optimal MPC of a Crane




Time Optimal MPC of a Crane

Univ. Leuven [Vandenbrouck, Swevers, D.]



Optimal Solutions in qpOASES Varying in Time



Optimal Control Problem in Model Predictive Control

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && \sum_{i=0}^N \varphi_i(F_i(s_i, u_i)) + \varphi_N(F_N(s_N)) \\ & \text{subject to} && s_0 = x, \\ & && s_{i+1} = S_i(s_i, u_i), \quad i = 0, \dots, N-1, \\ & && H_i(s_i, u_i) \in \Omega_i, \quad i = 0, \dots, N-1, \\ & && H_N(s_N) \in \Omega_N \end{aligned}$$

- ▶ variables $w = (s, u)$ with $s = (s_0, \dots, s_N)$ and $u = (u_0, \dots, u_{N-1})$
- ▶ convexities in φ_i (e.g. quadratic) and Ω_i (e.g. polyhedral, ellipsoidal)
- ▶ nonlinearities in dynamic system S_i and constraint functions F_i, H_i
- ▶ often: S_i result of time integration (direct multiple shooting)

Overview

- Embedded Optimization
 - Universal Approximation Theorem for Convex Optimization
 - Model Predictive Control (and two Applications)
- • Convexity Exploiting Newton-Type Optimization
 - Sequential Convex Programming (SCP)
 - Generalized Gauss-Newton (GGN)
 - Sequential Convex Quadratic Programming (SCQP)
 - Local Convergence Analysis and Desirable Divergence
- Zero-Order Optimization-based Iterative Learning Control
 - Tutorial Example
 - Bounding the Loss of Optimality
 - Local Convergence Analysis

Nonlinear optimization with convex substructure

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && \phi_0(F_0(w)) \\ & \text{subject to} && F_i(w) \in \Omega_i \quad i = 1, \dots, m, \\ & && G(w) = 0 \end{aligned}$$

Assumptions:

- ▶ twice continuously differentiable functions $G : \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_g}$ and $F_i : \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_{F_i}}$ for $i = 0, 1, \dots, m$.
- ▶ outer function $\phi_0 : \mathbb{R}^{n_{F_0}} \rightarrow \mathbb{R}$ convex.
- ▶ sets $\Omega_i \subset \mathbb{R}^{n_{F_i}}$ convex for $i = 1, \dots, m$,
(possibly $z \in \Omega_i \Leftrightarrow \phi_i(z) \leq 0$ with smooth convex ϕ_i)

Idea:

exploit convex substructure via *iterative convex approximations*.

Sequential Convex Programming (SCP)

- ▶ linearize $F_i^{\text{lin}}(w; \bar{w}) := F_i(\bar{w}) + J_i(\bar{w})(w - \bar{w})$ with $J_i(\bar{w}) := \frac{\partial F_i}{\partial w}(\bar{w})$
- ▶ formulate convex subproblems:

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && \phi_0(F_0^{\text{lin}}(w; \bar{w})) \\ & \text{subject to} && F_i^{\text{lin}}(w; \bar{w}) \in \Omega_i, \quad i = 1, \dots, m, \\ & && G^{\text{lin}}(w; \bar{w}) = 0 \end{aligned}$$

- ▶ start at w_0 with $k = 0$
- ▶ solve convex subproblem at $\bar{w} = w_k$ to obtain next iterate w_{k+1}

Historical notes:

The SCP idea with LP subproblems was originally called "Method of Approximation Programming" in [Griffith & Stewart, 1961]. SCP with more general convex sets (matrix cones) was proposed in [Fares, Apkarian, Noll 2002].

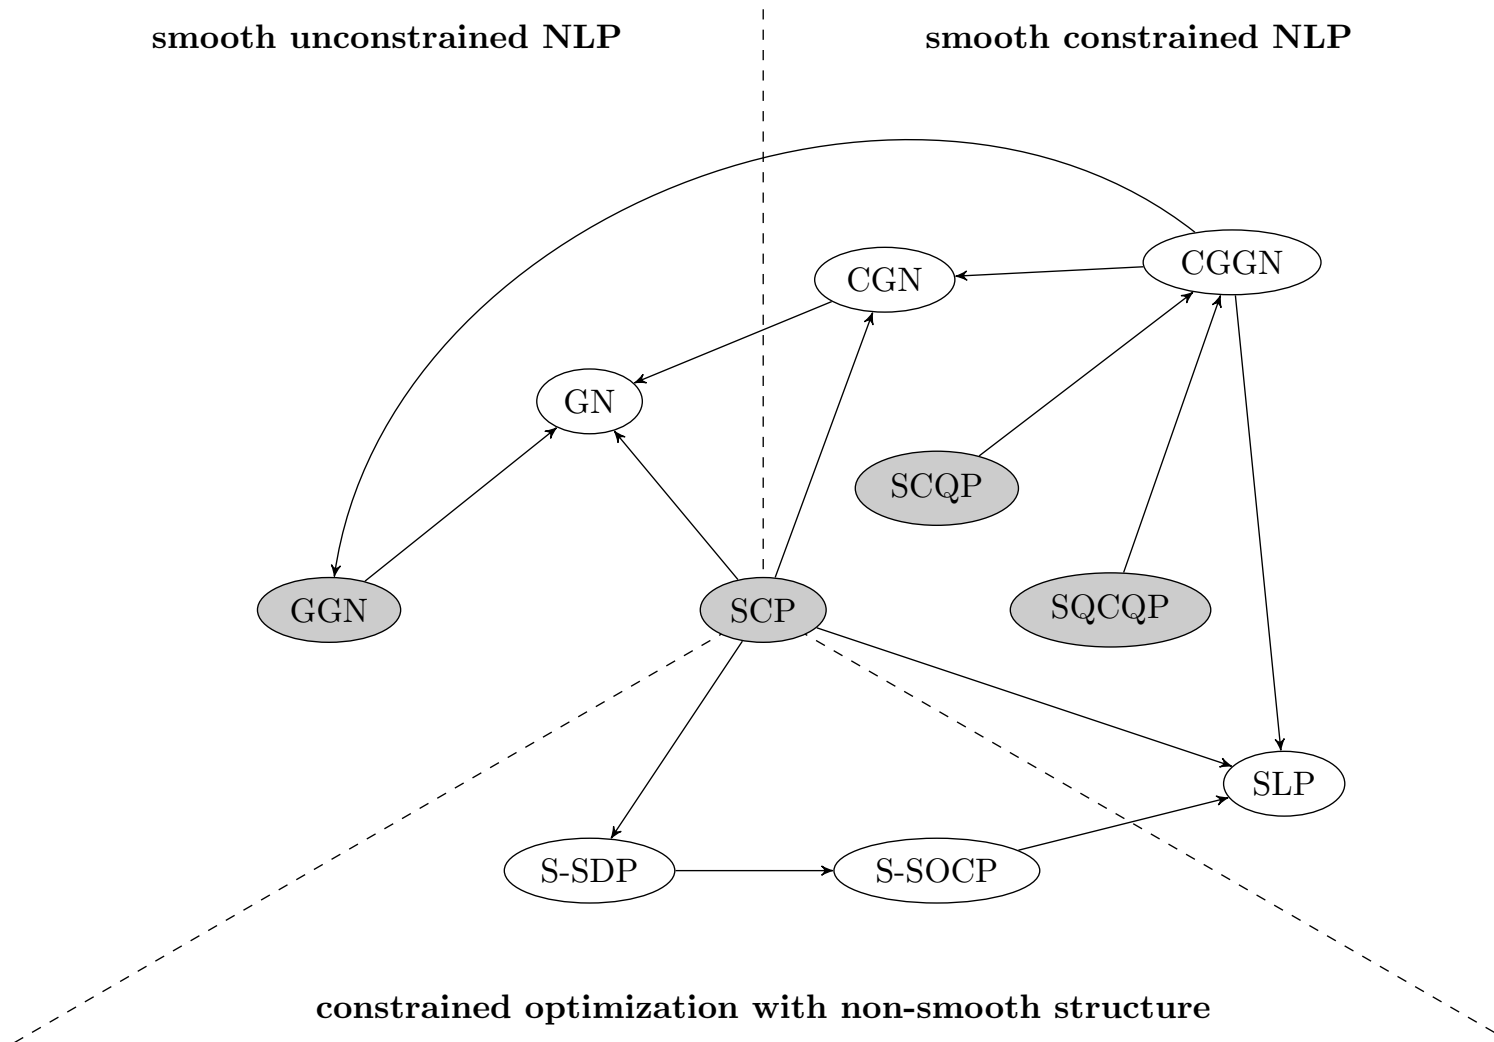
Why is this class of problems and algorithms interesting ?

- ▶ many optimization problems have "convex-over-nonlinear" structure
- ▶ standard NLP solvers cannot address all non-smooth convex constraints
- ▶ there exist many mature and efficient convex optimization solvers

Some application areas:

- ▶ nonlinear least squares for estimation and tracking
[Gauss 1809; Bock 1983; Li and Biegler 1989; Sideris and Bobrow 2004]
- ▶ nonlinear matrix inequalities for reduced order controller design
[Fares, Noll, Apkarian 2002; Tran-Dinh et al. 2012]
- ▶ ellipsoidal terminal regions in nonlinear model predictive control
[Chen and Allgöwer 1998; Verschueren 2016]
- ▶ robustified inequalities in nonlinear optimization
[Nagy and Braatz 2003; D., Bock, Kostina 2006]
- ▶ tube-following optimal control problems [Van Duijkeren 2019]
- ▶ non-smooth composite minimization [Apkarian et al. 2008; Lewis and Wright 2016]
- ▶ deep neural network training with convex loss functions
[Schraudolph 2002; Martens 2016]

Class Picture of Iterative Convex Approximation Methods



[Messerer and D., in preparation]

Simplest case: smooth unconstrained problems

Unconstrained minimization of "convex over nonlinear" function

$$\underset{w \in \mathbb{R}^n}{\text{minimize}} \quad \underbrace{\phi(F(w))}_{=: f(w)}$$

Assumptions:

- Inner function $F : \mathbb{R}^n \rightarrow \mathbb{R}^N$ of class C^2
- Outer function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$ of class C^2 and convex

Simplest case: smooth unconstrained problems

Unconstrained minimization of "convex over nonlinear" function

$$\underset{w \in \mathbb{R}^n}{\text{minimize}} \quad \underbrace{\phi(F(w))}_{=: f(w)}$$

Assumptions:

- Inner function $F : \mathbb{R}^n \rightarrow \mathbb{R}^N$ of class C^2
- Outer function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$ of class C^2 and convex

SCP subproblem becomes



$$\underset{w \in \mathbb{R}^n}{\text{minimize}} \quad \underbrace{\phi\left(F^{\text{lin}}(w; \bar{w})\right)}_{=: f_{\text{SCP}}(w; \bar{w})} \quad (1)$$

Tutorial Example: Pseudo Huber Loss Minimization


Experiments conducted by Florian Messerer

$$\underset{w \in \mathbb{R}}{\text{minimize}} \quad \underbrace{\sum_{i=1}^n \phi_{\delta}(y_i - m(x_i + w))}_{=: f(w)}$$

Aim: fit $n=3$ measurements y_i to a model $m(w + x_i)$ with $m(x) = \frac{3}{4}x + \sin(x)$ using the pseudo Huber loss $\phi_{\delta}(x) := \sqrt{\delta^2 + x^2}$.



Cost function and SCP approximation



SCP for Least Squares = Gauss-Newton

With quadratic $\phi(z) = \frac{1}{2}\|z\|_2^2 = \frac{1}{2}z^\top z$, SCP subproblems become

$$\underset{w \in \mathbb{R}^n}{\text{minimize}} \quad \frac{1}{2} \|F(w_k) + J(w_k)(w - w_k)\|_2^2 \quad (2)$$

If $\text{rank}(J) = n$ this is uniquely solvable, giving

$$w_{k+1} = w_k - \underbrace{\left(J(w_k)^\top J(w_k) \right)^{-1}}_{=: B_{\text{GN}}(w_k)} \underbrace{J(w_k)^\top F(w_k)}_{= \nabla f(w_k)}$$

SCP applied to LS = Newton method with "Gauss-Newton Hessian"

$$B_{\text{GN}}(w) \approx \nabla^2 f(w)$$

Generalized Gauss-Newton (GGN) [Schraudolph 2002]

For general convex $\phi(\cdot)$ we have for $f(w) = \phi(F(w))$

$$\nabla^2 f(w) = \underbrace{J(w)^\top \nabla^2 \phi(F(w)) J(w)}_{=: B_{\text{GGN}}(w)} + \underbrace{\sum_{j=1}^N \nabla^2 F_j(w) \nabla_{z_j} \phi(F(w))}_{=: E_{\text{GGN}}(w)}$$

"GGN Hessian" "Error matrix"

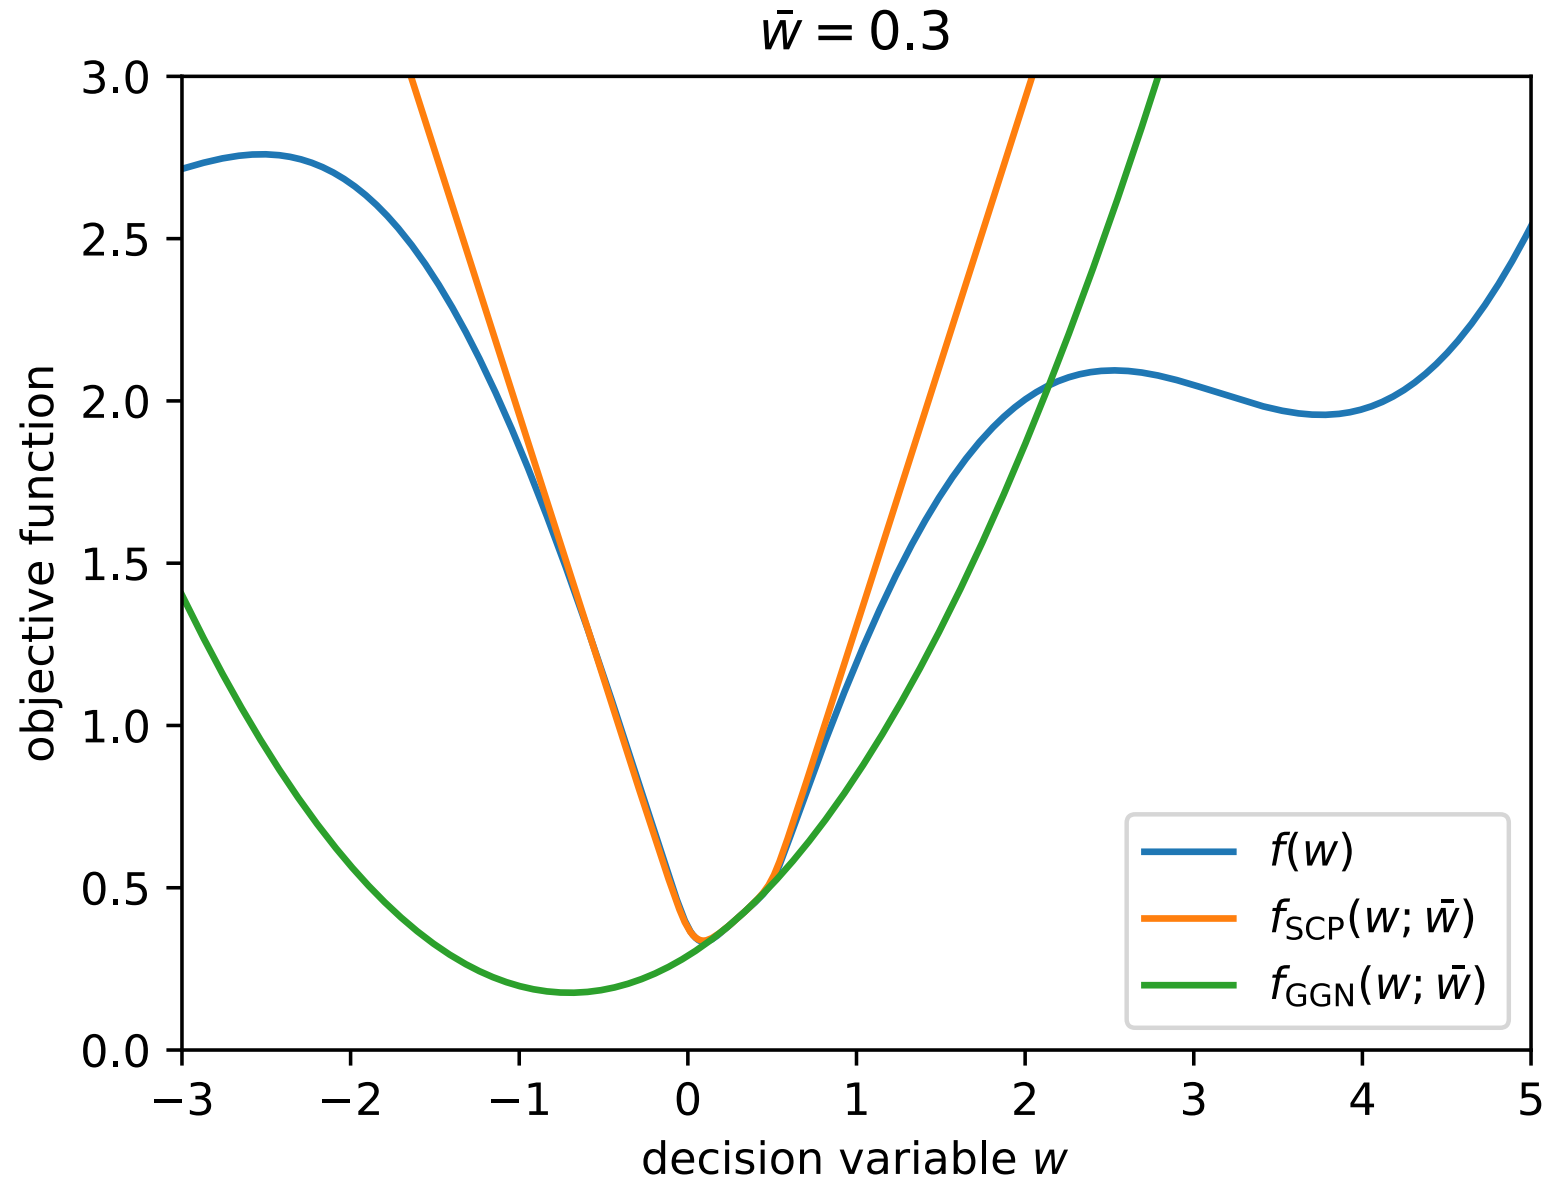
Generalized Gauss-Newton (GGN) method iterates according to

$$w_{k+1} = w_k - B_{\text{GGN}}(w)^{-1} \nabla f(w_k)$$

Note: GGN solves convex quadratic subproblems

$$\min_{w \in \mathbb{R}^n} \underbrace{f(w_k) + \nabla f(w_k)^\top (w - w_k) + \frac{1}{2} (w - w_k)^\top B_{\text{GGN}}(w_k) (w - w_k)}_{=: f_{\text{GGN}}(w; w_k)}$$

Tutorial Example: SCP and GGN Approximation



General smooth NLP formulation with constraints

Now regard an NLP with smooth convex $\phi_0, \phi_1, \dots, \phi_m$

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && \underbrace{\phi_0(F_0(w))}_{=: f_0(w)} \\ & \text{subject to} && \underbrace{\phi_i(F_i(w))}_{=: f_i(w)} \leq 0, \quad i = 1, \dots, m, \\ & && G(w) = 0 \end{aligned}$$

SCP subproblem becomes

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && \phi_0(F_0^{\text{lin}}(w; \bar{w})) \\ & \text{subject to} && \phi_i(F_i^{\text{lin}}(w; \bar{w})) \leq 0, \quad i = 1, \dots, m, \\ & && G^{\text{lin}}(w; \bar{w}) = 0 \end{aligned}$$

(SCP algorithm is expensive, but multiplier-free and affine-invariant)

Constrained Gauss-Newton [Bock 1983]

Use $B_{\text{CGN}}(w) := J_0(w)^\top \nabla^2 \phi_0(F_0(w)) J_0(w)$ and solve convex quadratic program (QP)

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && f_0^{\text{lin}}(w; \bar{w}) + \frac{1}{2}(w - \bar{w})^\top B_{\text{CGN}}(\bar{w})(w - \bar{w}) \\ & \text{subject to} && f_i^{\text{lin}}(w; \bar{w}) \leq 0, \quad i = 1, \dots, m, \\ & && G^{\text{lin}}(w; \bar{w}) = 0 \end{aligned}$$

- ▶ like SCP, the method is multiplier free and affine invariant
- ▶ QPs are potentially cheaper to solve
- ▶ but CGN diverges on some problems where SCP converges

Remark: for least-squares objectives, this method is due to [Bock 1983]. In many papers, Bock's method is called "the Generalized Gauss-Newton (GGN) method". To avoid a notation clash with Schraudolph and the computer science literature, we prefer to call Bock's method "the Constrained Gauss-Newton (CGN) method".

Sequential Convex Quadratic Programming (SCQP)

[Verschuere et al 2016]

$$B_{\text{SCQP}}(w, \mu) := J_0(w)^\top \nabla^2 \phi_0(F_0(w)) J_0(w) + \sum_{i=1}^m \mu_i J_i(w)^\top \nabla^2 \phi_i(F_i(w)) J_i(w)$$

$$\begin{aligned} & \underset{w \in \mathbb{R}^{n_w}}{\text{minimize}} && f_0^{\text{lin}}(w; \bar{w}) + \frac{1}{2} (w - \bar{w})^\top B_{\text{SCQP}}(\bar{w}, \bar{\mu}) (w - \bar{w}) \\ & \text{subject to} && f_i^{\text{lin}}(w; \bar{w}) \leq 0, \quad i = 1, \dots, m, \quad | \quad \mu^+, \\ & && G^{\text{lin}}(w; \bar{w}) = 0 \end{aligned}$$

- ▶ obtain pair (w_{k+1}, μ_{k+1}) from solution at $(\bar{w}, \bar{\mu}) = (w_k, \mu_k)$
- ▶ "optimizer state" contains both, \bar{w} and inequality multipliers $\bar{\mu}$
- ▶ again, only a QP needs to be solved in each iteration
- ▶ again, affine invariant
- ▶ $B_{\text{SCQP}}(w, \mu) \succeq B_{\text{CGN}}(w)$ (more likely to converge than CGN)
- ▶ for unconstrained problems, SCQP becomes GGN
- ▶ in fact, SCQP has same contraction rate as SCP [Messerer & D., ECC 2020]

Identical local convergence of SCP and SCQP

Theorem 1 [Messerer and Diehl, ECC 2020]

Regard KKT point $z^* := (w^*, \mu^*, \lambda^*)$ with LICQ and strict complementarity. Denote the reduced Hessian by $\tilde{\Lambda}_*$, the reduced SCQP Hessian by \tilde{B}_* (*) and assume that $\tilde{B}_* \succ 0$. Then

- ▶ z^* is a fixed point for both the SCP and SCQP iterations
- ▶ both methods are well-defined in a neighborhood of z^*
- ▶ their linear contraction rates are equal and given by the smallest $\alpha \in \mathbb{R}$ that satisfies the linear matrix inequality

$$-\alpha\tilde{B}_* \preceq \tilde{\Lambda}_* - \tilde{B}_* \preceq \alpha\tilde{B}_* \quad (3)$$

(*) $\tilde{\Lambda}_* := Z^\top \nabla^2 \mathcal{L}(w^*, \mu^*, \lambda^*) Z$ and $\tilde{B}_* := Z^\top B_{\text{SCQP}}(w^*, \mu^*) Z$ with Z a fixed nullspace basis of the Jacobian of active constraints

Identical local convergence of SCP and SCQP/GGN

Theorem 1 [Messerer and Diehl, ECC 2020]

Regard KKT point $z^* := (w^*, \mu^*, \lambda^*)$ with LICQ and strict complementarity. Denote the reduced Hessian by $\tilde{\Lambda}_*$, the reduced SCQP Hessian by \tilde{B}_* (*) and assume that $\tilde{B}_* \succ 0$. Then

- ▶ z^* is a fixed point for both the SCP and SCQP iterations
- ▶ both methods are well-defined in a neighborhood of z^*
- ▶ their linear contraction rates are equal and given by the smallest $\alpha \in \mathbb{R}$ that satisfies the linear matrix inequality

$$-\alpha \tilde{B}_* \preceq \tilde{\Lambda}_* - \tilde{B}_* \preceq \alpha \tilde{B}_* \quad (3)$$

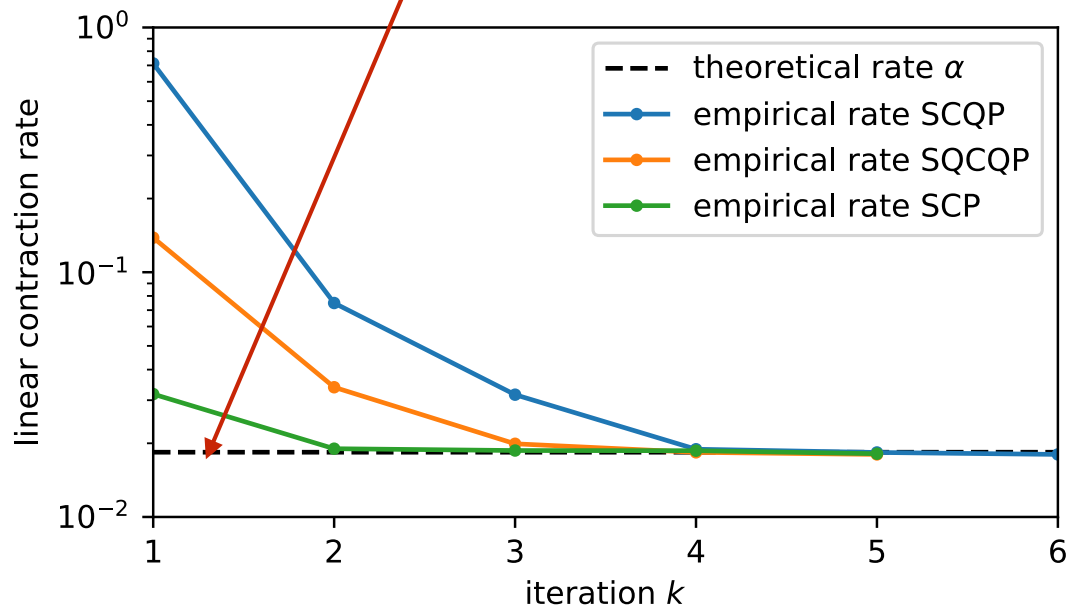

(*) $\tilde{\Lambda}_* := Z^\top \nabla^2 \mathcal{L}(w^*, \mu^*, \lambda^*) Z$ and $\tilde{B}_* := Z^\top B_{\text{SCQP}}(w^*, \mu^*) Z$ with Z a fixed nullspace basis of the Jacobian of active constraints

Corollary

Necessary condition for local convergence of both methods is $\tilde{B}_* \succeq \frac{1}{2} \tilde{\Lambda}_* \succeq 0$


Proof of corollary: Set $\alpha = 1$ in (3).

Tutorial Example: Objective and Local Contraction Rate




Desirable Divergence and Mirror Problem [cf. Bock 1987]

SCP and GGN do not converge to every local minimum. This can help to avoid "bad" local minima, as discussed next.



Regard maximum likelihood estimation problem $\min_w \phi(M(w) - y)$ with nonlinear model $M : \mathbb{R}^n \rightarrow \mathbb{R}^N$ and measurements $y \in \mathbb{R}^N$. Assume penalty ϕ is symmetric with $\phi(-z) = \phi(z)$ as is the case for symmetric error distributions. At a solution w^* , we can generate "mirror measurements" $y_{\text{mr}} := 2M(w^*) - y$ obtained by reflecting the residuals. From a statistical point of view, y_{mr} should be as likely as y .


SCP divergence \Leftrightarrow minimum unstable under mirroring



Theorem [Messerer and D., 2019/2020] generalizing [Bock 1987]

Regard a local minimizer w^* of $\phi(M(w) - y)$ that satisfies SOOSC. If the necessary SCP convergence condition $\tilde{B}_* \succeq \frac{1}{2}\tilde{\Lambda}_*$ does not hold, then w^* is a stationary point of the mirror problem but **not** a local minimizer.

SCP divergence \Leftrightarrow minimum unstable under mirroring

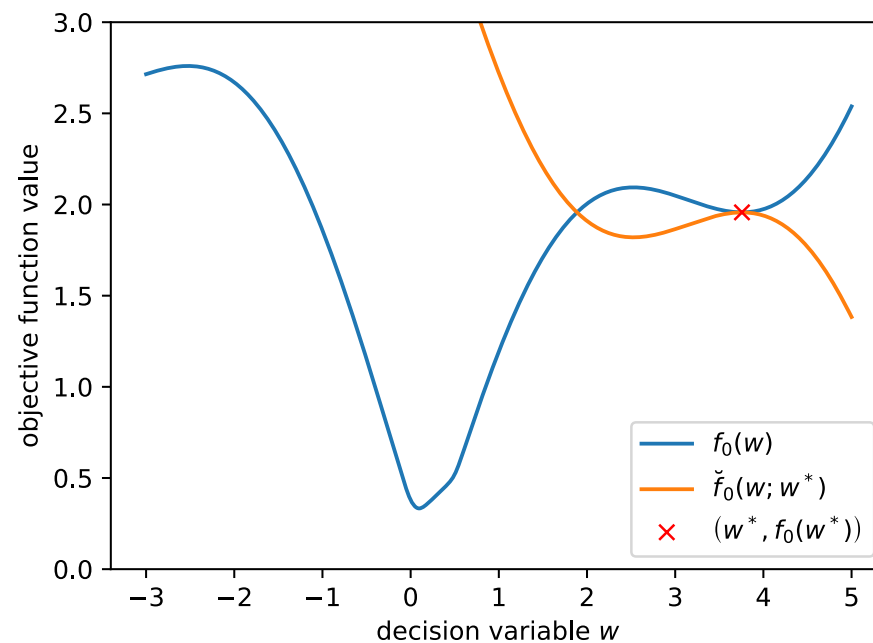
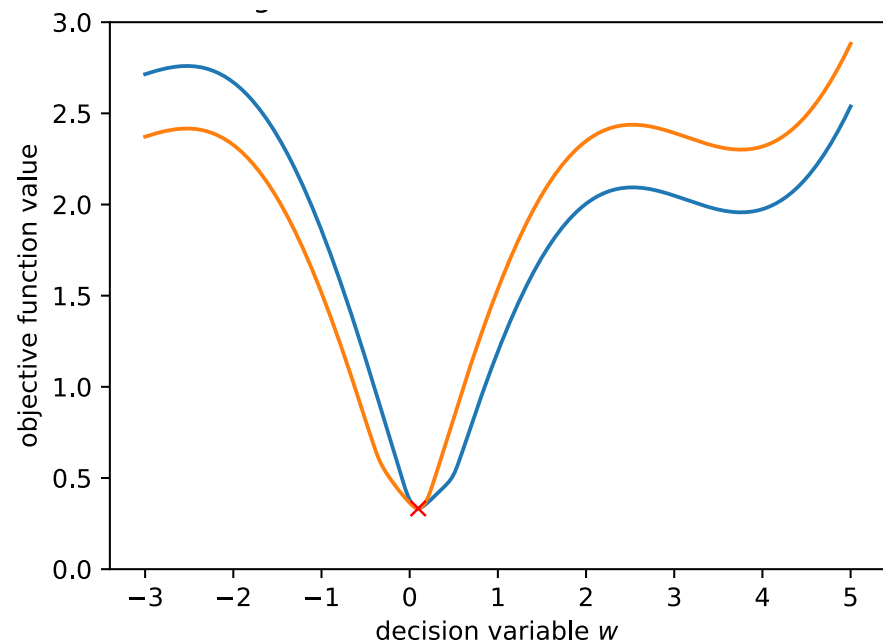

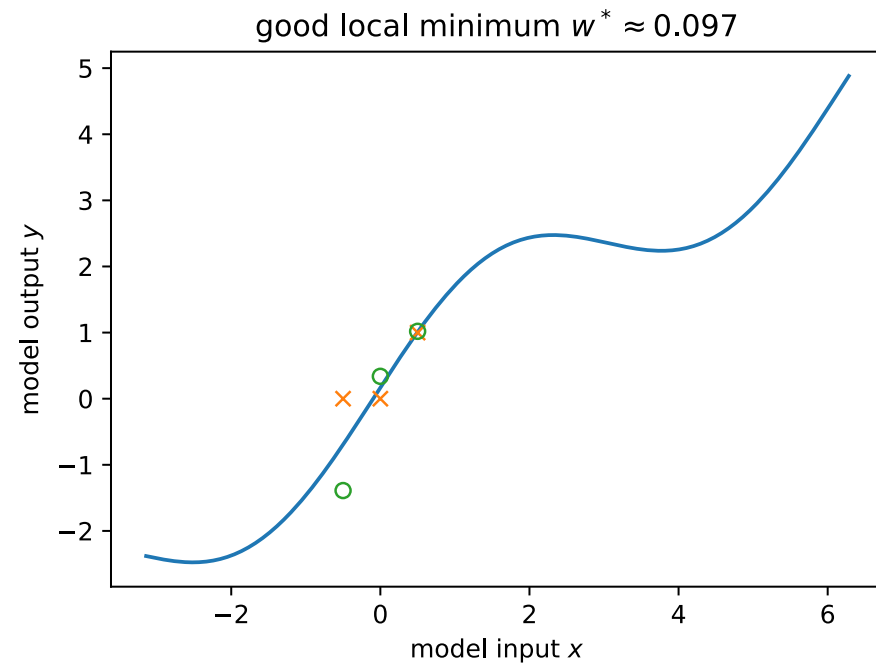


Theorem [Messerer and D., 2019/2020] generalizing [Bock 1987]

Regard a local minimizer w^* of $\phi(M(w) - y)$ that satisfies SOOSC. If the necessary SCP convergence condition $\tilde{B}_* \succeq \frac{1}{2}\tilde{\Lambda}_*$ does not hold, then w^* is a stationary point of the mirror problem but **not** a local minimizer.

Sketch of proof (unconstrained): use $M(w^) - y_{\text{mr}} = y - M(w^*)$ to show that $\nabla f_{\text{mr}}(w^*) = J(w^*)^\top (y - M(w^*)) = 0$ and $\nabla^2 f_{\text{mr}}(w^*) = B_{\text{GGN}}(w^*) - E_{\text{GGN}}(w^*) = 2B_{\text{GGN}}(w^*) - \nabla^2 f(w^*) \not\preceq 0$

Tutorial Example and Mirror Problems at Different Local Minima



Overview

- Embedded Optimization
 - Universal Approximation Theorem for Convex Optimization
 - Model Predictive Control (and two Applications)
- Convexity Exploiting Newton-Type Optimization
 - Sequential Convex Programming (SCP)
 - Generalized Gauss-Newton (GGN)
 - Sequential Convex Quadratic Programming (SCQP)
 - Local Convergence Analysis and Desirable Divergence
- ➔ • Zero-Order Optimization-based Iterative Learning Control
 - Tutorial Example
 - Bounding the Loss of Optimality
 - Local Convergence Analysis

Two Ingredients of Newton-Type Optimization

The convexity exploiting algorithms presented so far need two ingredients:

1. a good nonlinear model and its linearisation, and
2. convex substructure in objective and constraints

Two Ingredients of Newton-Type Optimization

The convexity exploiting algorithms presented so far need two ingredients:

1. a good nonlinear model and its linearisation, and
2. convex substructure in objective and constraints

Which of the two is more important for success in data-driven optimization?

Two Ingredients of Newton-Type Optimization

The convexity exploiting algorithms presented so far need two ingredients:

1. a good nonlinear model and its linearisation, and
2. **convex substructure** in objective and constraints

Which of the two is more important for success in data-driven optimization?

Iterative Learning Control for Lemon-Ball Throwing




Iterative Learning of Ball Throwing with Minimal Energy


Experiments conducted by Katrin Baumgärtner

- Model $F_M(u)$ maps initial velocity $u \in \mathbb{R}^2$ to landing position $y \in \mathbb{R}$
- Aim: throw ball further than $y \geq 10$ with minimal initial velocity
- Experiments with “real plant” give pairs (u_k, y_k) [shorter distance than predicted]
- We can use (u_k, y_k) to correct the model, and iteratively obtain u_{k+1} by solving the following optimization problem:


$$\begin{array}{ll} \text{minimize} & \|u\|_2^2 \\ u \in \mathbb{R}^2, y \in \mathbb{R} & \\ \text{subject to} & y = \underbrace{F_M(u) - F_M(u_k) + y_k}_{=: \tilde{F}_M(u; u_k, y_k)}, \\ & y \geq 10 \end{array}$$



Iterations of Algorithm and Reduced Problem Visualization



$$\begin{aligned} & \text{minimize} && \|u\|_2^2 \\ & u \in \mathbb{R}^2 \\ & \text{subject to} && \tilde{F}_M(u; u_k, y_k) \geq 10 \end{aligned}$$



Zero Order Optimization-based Iterative Learning Control (ZOO-ILC)

Aim: optimization with unknown input-output system $y = F_R(u)$ ("reality"):

$$\begin{aligned} & \underset{u, y}{\text{minimize}} && \phi(u, y) \\ & \text{subject to} && F_R(u) - y = 0, \\ & && H(u, y) \leq 0 \end{aligned} \tag{4}$$

Zero Order Optimization-based Iterative Learning Control (ZOO-ILC)

Aim: optimization with unknown input-output system $y = F_R(u)$ ("reality"):

$$\begin{aligned} & \underset{u, y}{\text{minimize}} && \phi(u, y) \\ & \text{subject to} && F_R(u) - y = 0, \\ & && H(u, y) \leq 0 \end{aligned} \tag{4}$$

ZOO-ILC idea [cf. Schöllig, Volkaert, Zeilinger]: use trial input u_k with output y_k and a model F_M to obtain new trial input u_{k+1} from solution of

$$\begin{aligned} & \underset{u, y}{\text{minimize}} && \phi(u, y) \\ & \text{subject to} && F_M(u) - y = F_M(u_k) - y_k, \\ & && H(u, y) \leq 0 \end{aligned} \tag{5}$$

Questions: Does this method converge? What is its loss of optimality?

Feasibility and Loss of Optimality of ZOO-ILC

Theorem 2 [Baumgärtner et al., in preparation]

For any fixed point (\bar{u}, \bar{y}) of the ZOO-ILC algorithm with multipliers $(\bar{\lambda}, \bar{\mu})$ holds under mild conditions:

- ▶ (\bar{u}, \bar{y}) is feasible for the real problem
- ▶ the loss of optimality compared to a real solution (u_R, y_R) is bounded by:

$$\phi(\bar{u}, \bar{y}) - \phi(u_R, y_R) \leq \bar{\lambda}^\top (J_M(\bar{u}) - J_R(\bar{u})) (u_R - \bar{u})$$

Here, the Lagrangian of the model problem is given by

$$\mathcal{L}(u, y, \lambda, \mu) = \phi(u, y) + \lambda^\top (F_M(u) - y - b_k) + \mu^\top H(u, y)$$

and $J_M(u)$ and $J_R(u)$ are the Jacobians of $F_M(u)$ and $F_R(u)$.

Special cases where ZOO-ILC delivers a lossless solution

$$\phi(\bar{u}, \bar{y}) - \phi(u_R, y_R) \leq \bar{\lambda}^\top (J_M(\bar{u}) - J_R(\bar{u})) (u_R - \bar{u})$$

ZOO-ILC delivers lossless solution in the following three cases:

1. Tracking ILC with zero residual (standard ILC):

$$\bar{\lambda} = 0$$

2. Model and real Jacobian coincide at solution (rarely the case):

$$J_M(\bar{u}) - J_R(\bar{u}) = 0$$

3. Constrained problems where solution u_R is in vertex of the reduced feasible set:

$$u_R - \bar{u} = 0$$

(if the Jacobian error is small enough, LICQ and strict complementarity hold)

Special cases where ZOO-ILC delivers a lossless solution

$$\phi(\bar{u}, \bar{y}) - \phi(u_R, y_R) \leq \bar{\lambda}^\top (J_M(\bar{u}) - J_R(\bar{u})) (u_R - \bar{u})$$

ZOO-ILC delivers lossless solution in the following three cases:

1. Tracking ILC with zero residual (standard ILC):

$$\bar{\lambda} = 0$$

2. Model and real Jacobian coincide at solution (rarely the case):

$$J_M(\bar{u}) - J_R(\bar{u}) = 0$$


3. Constrained problems where solution u_R is in vertex of the reduced feasible set:

$$u_R - \bar{u} = 0$$

(if the Jacobian error is small enough, LICQ and strict complementarity hold)

Solutions for L_2 - and L_∞ -norm minimisation

$$\begin{array}{ll} \text{minimize} & \|u\|_2^2 \\ u \in \mathbb{R}^2 & \\ \text{subject to} & \tilde{F}_M(u; u_k, y_k) \geq 10 \end{array}$$




suboptimality: $0.874 \leq 1.377$ (bound)


Solutions for L_2 - and L_∞ -norm minimisation

$$\begin{aligned} &\text{minimize} && \|u\|_2^2 \\ &u \in \mathbb{R}^2 \\ &\text{subject to} && \tilde{F}_M(u; u_k, y_k) \geq 10 \end{aligned}$$

$$\begin{aligned} &\text{minimize} && \|u\|_\infty \\ &u \in \mathbb{R}^2 \\ &\text{subject to} && \tilde{F}_M(u; u_k, y_k) \geq 10 \end{aligned}$$



suboptimality: $0.874 \leq 1.377$ (bound)



solution in vertex, no loss of optimality



Time-Optimal Motion of an oscillator (L1-tracking)

Real plant: $T^2\ddot{y} + 2Td\dot{y} + y + \beta y^3 = K_R u$
 with $T = 1, d = 0.5, \beta = 2, K_R = 0.9$



Model: $T^2\ddot{y} + 2Td\dot{y} + y = K_M u$
 with $K_R = 1$

$$\begin{aligned} & \text{minimize}_{y(\cdot), u(\cdot)} \int_0^{T_H} |y(t) - y_{\text{ref}}| + \alpha u(t)^2 dt \\ & \text{subject to } y(t) = F_M(t; u) + y_k(t) - F_M(t; u_k), \\ & |u(t)| \leq 1, \quad t \in [0, T_H] \end{aligned}$$





$$\alpha = 10^{-4}$$









Time-Optimal Motion of an oscillator (L1-tracking)



Time-Optimal Motion of an oscillator (L1-tracking)



Time-Optimal Motion of an oscillator (L1-tracking)



When does the ZOO-ILC method converge?

Theorem 3 (Convergence of ZOO-ILC) [Baumgärtner et al., in preparation]

Regard a fixed point $\bar{z} = (\bar{u}, \bar{y}, \bar{\lambda}, \bar{\mu}_{\mathcal{A}})$ of ZOO-ILC and assume it satisfies LICQ, SOSC and strict complementarity in the model problem. Then the local contraction rate is given by the spectral radius $\rho(A)$ of the matrix

$$A := \begin{bmatrix} \mathbb{I}_{n_u} & 0 & 0 & 0 \end{bmatrix} \left(\frac{\partial R}{\partial z}(\bar{z}; \bar{u}, \bar{y}) \right)^{-1} \begin{bmatrix} 0 \\ 0 \\ J_{\mathbf{M}}(\bar{u}) - J_{\mathbf{R}}(\bar{u}) \\ 0 \end{bmatrix}$$

The ZOO-ILC method converges if $\rho(A) < 1$ and diverges if $\rho(A) > 1$.

Here, $\mu_{\mathcal{A}}$ are the active constraint multipliers and $R(z; u', y')$ is defined by

$$R(z; u', y') := \begin{bmatrix} \nabla_u \mathcal{L}_{\mathbf{M}}(u, y, \lambda, \mu_{\mathcal{A}}; u', y') \\ \nabla_y \mathcal{L}_{\mathbf{M}}(u, y, \lambda, \mu_{\mathcal{A}}; u', y') \\ F_{\mathbf{M}}(u) - y + y' - F_{\mathbf{M}}(u') \\ H_{\mathcal{A}}(u, y) \end{bmatrix}$$

where the Lagrangian of the model problem is given by

$$\mathcal{L}_{\mathbf{M}}(u, y, \lambda, \mu_{\mathcal{A}}; u', y') = \phi(u, y) + \lambda^{\top} (F_{\mathbf{M}}(u) - y + y' - F_{\mathbf{M}}(u')) + \mu_{\mathcal{A}}^{\top} H_{\mathcal{A}}(u, y)$$

and $J_{\mathbf{M}}(u)$ and $J_{\mathbf{R}}(u)$ are the Jacobians of $F_{\mathbf{M}}(u)$ and $F_{\mathbf{R}}(u)$.

When does the ZOO-ILC method converge?

$$\begin{bmatrix} 0 \\ 0 \\ J_{\text{M}}(\bar{u}) - J_{\text{R}}(\bar{u}) \\ 0 \end{bmatrix}$$

When does the ZOO-ILC method converge?

$$\begin{bmatrix} 0 \\ 0 \\ J_{\text{M}}(\bar{u}) - J_{\text{R}}(\bar{u}) \\ 0 \end{bmatrix}$$

Contraction rate grows with distance between model and real Jacobian.

Summary

- Embedded optimization and Model Predictive Control are powerful tools with proven success in control engineering practice.
- Exploiting convex structures in nonlinear problems is key for reliable and fast algorithms.
- Sequential Convex Programming (SCP) and its variants converge linearly. They avoid minimizers where the nonlinearity dominates the convex substructure (which can be desirable in estimation)
- Zero-Order Optimization allows us to design theoretically solid Iterative Learning Control algorithms. They can recover an optimal solution if the convex substructure dominates the model errors.