Complexity Bounds of Iterative Linearization Algorithms for Discrete-Time Nonlinear Control

Speaker: V. Roulet
Co-authors: S. Srinivasa, M. Fazel, Z. Harchaoui
Soon on ArXiv, paper and code available on request
Send me an email!

INFORMS Optimization Society Conference
03/15/2022

Nonlinear Control Problems

Continuous-Time Control problem

- System driven by dynamics $\dot{x}(t) = \bar{f}(x(t), u(t))$
- Minimize cost $\bar{h}(x(t), t)$ over $t \in [0, T]$ for $x(0)$ fixed

Discrete-Time Control Problem

- Discretize dynamics as $x_{t+1} = f(x_t, u_t)$
- Minimize costs $h_t(x_t)$ over $t \in \{0, \dots, \tau\}$ for x_0 fixed

Algorithms Principle

Current controls $u_0, \dots, u_{\tau-1}$ with trajectory $x_0, \dots, x_\tau$

1. Linearize dynamics f around x_t, u_t
2. Take quadratic approx. of the costs h_t around x_t
3. Solve resulting LQ pb
4. Repeat from 1.

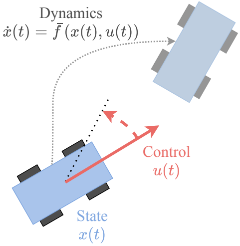

Dynamics of a car

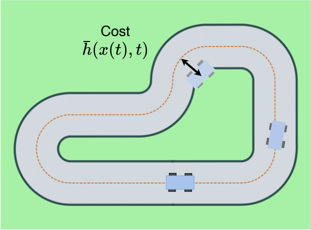

Tracking objective

Autonomous Car Racing

Simple model of a car

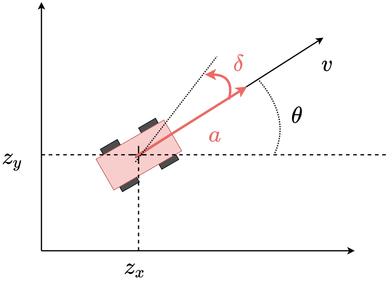

$$x = (z_x, z_y, \theta, v), \qquad u = (\delta, a)$$
$$\dot{z}_x = v\cos\theta \qquad \dot{\theta} = v\tan(\delta)$$
$$\dot{z}_y = v\sin\theta \qquad \dot{v} = a$$

Algo. converges *fast* to *optimal trajectory*

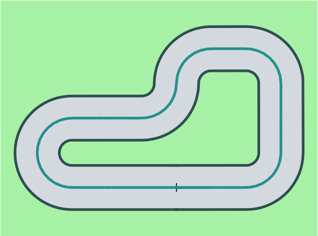

Optimized trajectory horizon $\tau = 100$

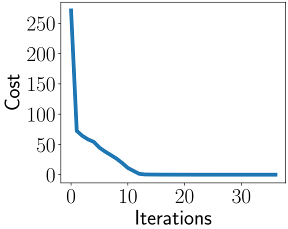

Convergence of the algorithm

Autonomous Car Racing

Bicycle model of a car (Liniger et al. 2015)

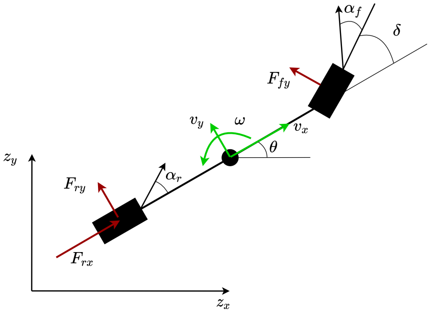

Models tire forces (highly non-linear)

Unclear whether the algorithm succeeded...

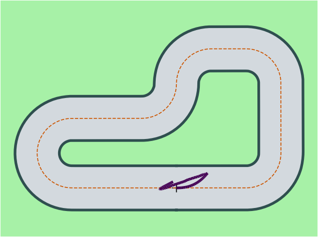

Optimized trajectory horizon $\tau = 100$

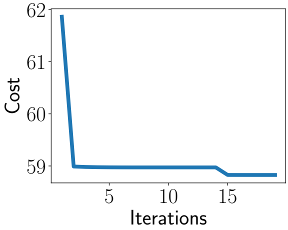

Convergence of the algorithm

Objectives and Outline

Questions

1. What are sufficient conditions to ensure global convergence?
2. What are the worst-case complexity bounds of these algorithms?

Related work

- Sufficient optimality conditions in continuous-time (Mangasarian 1966)
  → Translatable in discrete-time, requires convexity of implicitly defined functions
- Local convergence of Differential Dynamic Programming
  (Polak 2011, Murray & Yakowitz 1984, Liao & Shoemaker 1991)
- Local convergence of generalized Gauss-Newton
  e.g. (Yamashita & Fukushima 2001, Diehl & Messerer 2019)
- Global convergence of regularized Gauss-Newton a.k.a. Levenberg-Marquardt
  e.g. (Bergou et al. 2020)

Plan

Discrete-Time Nonlinear Control Problems

Continuous-Time Control problem

$$\min_{x(t),u(t)} \quad \int_0^T \bar{h}(x(t),t)$$
$$\text{s.t.} \quad \dot{x}(t) = \bar{f}(x(t),u(t)),\; x(0)=\bar{x}_0$$

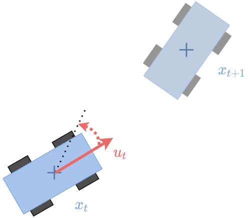

Euler discretization

Discrete-Time Control Problem

$$\min_{\substack{x_0;\dots;x_\tau \\ u_0;\dots;u_{\tau-1}}} \quad \sum_{t=1}^{\tau} h_t(x_t)$$
$$\text{s.t.} \quad x_{t+1} = f(x_t,u_t),\; x_0=\bar{x}_0$$

Discretization schemes:

$$\text{Euler:} \quad f(x_t,u_t) = x_t + \Delta \bar{f}(x_t,u_t)$$
$$\text{Multi-step:} \quad f(x_t,u_t) = x_{t+1}$$
$$\text{s.t.} \quad x_{t+(s+1)/k} = x_s + \Delta \bar{f}(x_{t+s/k},u_{t+s/k})$$
$$\dim(u_t) = k \dim(u(t))$$

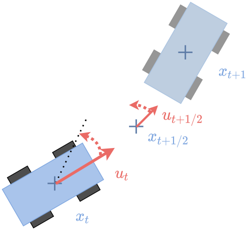

2-step discretization

Nonlinear Control Algorithms for Discrete-Time Control Problems

Forward Given a sequence of controls $u_0, \ldots, u_{\tau-1}$

a. Compute associated trajectory $x_{t+1} = f(x_t, u_t)$

b. Record linear expansions $\ell_f^{x_t, u_t}$ of the dynamics f around x_t, u_t

c. Record quadratic expansions $q_{h_t}^{x_t}$ of the costs around x_t

Backward Compute optimal policies π_t for the regularized linear-quadratic control problem

$$\min_{\substack{y_0, \ldots, y_\tau \\ v_0, \ldots, v_{\tau-1}}} \quad \sum_{t=1}^{\tau} q_{h_t}^{x_t}(y_t) + \frac{\nu}{2} \sum_{t=0}^{\tau-1} \|v_t\|_2^2$$

$$\text{s.t} \quad y_{t+1} = \ell_f^{x_t, u_t}(y_t, v_t), \quad y_0 = 0$$

by back-propagating the cost-to-go functions, starting from $c_\tau = q_{h_\tau}$,

$$c_t(y_t) = q_{h_t}(y_t) + \min_{v_t} \left\{ \frac{\nu}{2} \|v_t\|_2^2 + c_{t+1}(\ell_f^{x_t, u_t}(y_t, v_t)) \right\}$$

Roll-out Update the iterates as $u_t^{\text{next}} = u_t + v_t$

where v_t are computed by rolling-out the policies along either

- the linearized dynamics → Iterative Linear Quadratic Regulator (ILQR) (Li & Todorov 2007)

$$v_t = \pi_t(y_t) \quad y_{t+1} = \ell_f^{x_t, u_t}(y_t, v_t)$$

- the original dynamics → Iterative Differential Dynamic Programming (IDDP) (Tassa et al. 2012)

$$v_t = \pi_t(y_t) \quad y_{t+1} = f(x_t + y_t, u_t + v_t) - f(x_t, u_t)$$

ILQR Computational Scheme

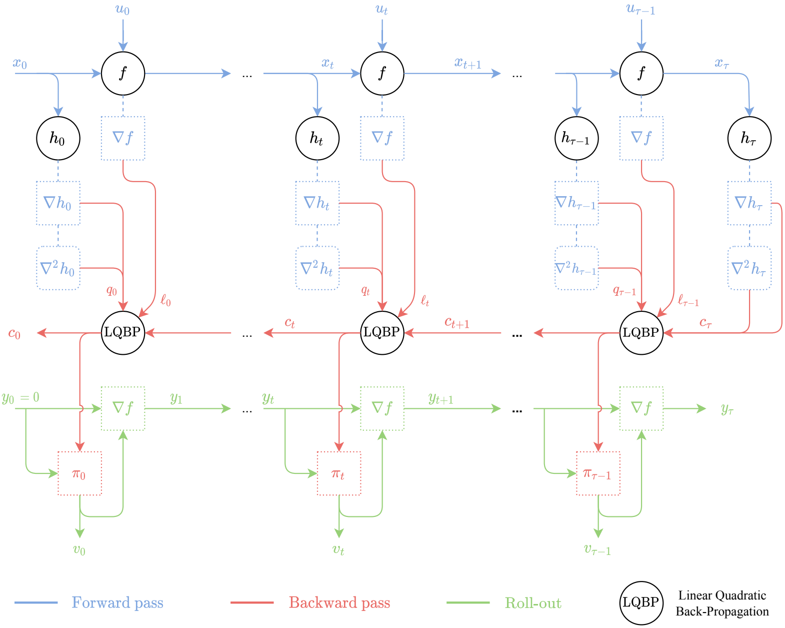

Plan

Nonlinear Control Problems and Algorithms

A Sufficient Condition for Global Convergence

Convergence Analysis of Iterative Linear-Quadratic Approximations

Objective Decomposition

Control of τ steps of f for $\boldsymbol{u} = (u_0; \dots; u_{\tau-1})$

$$f^{[\tau]}(x_0, \boldsymbol{u}) = (x_1; \dots; x_\tau)$$
$$\text{s.t.} \quad x_{t+1} = f(x_t, u_t)$$

Total cost for $\boldsymbol{x} = (x_1, \dots, x_\tau)$ $h(\boldsymbol{x}) = \sum_{t=1}^{\tau} h_t(x_t)$

Composite objective

$$\mathcal{J}(\boldsymbol{u}) = h(f^{[\tau]}(\bar{x}_0, \boldsymbol{u})) = \sum_{t=1}^{\tau} h_t(x_t)$$
$$\text{s.t.} \quad x_{t+1} = f(x_t, u_t), \quad x_0 = \bar{x}_0$$

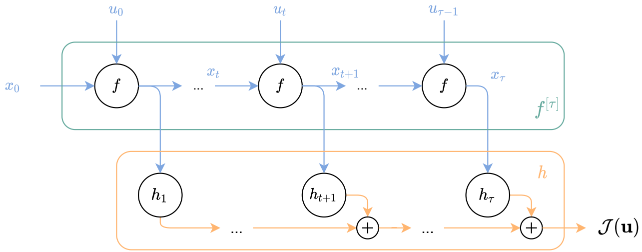

Sufficient Condition for Global Convergence

Idea:

- Prove sufficient condition for global conv. of 1st order methods, such as, for $c > 0$,

$$\|\nabla \mathcal{J}(\boldsymbol{u})\|_2^2 \geq c(\mathcal{J}(\boldsymbol{u}) - \mathcal{J}^*)$$

Gradient dominated objective $\mathcal{J}$

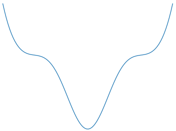

Non-convex gradient dominated function

Derivation:

- Here consider that the total cost h is e.g. μ-strongly convex s.t.

$$\|\nabla h(\boldsymbol{x})\|_2^2 \geq \mu(h(\boldsymbol{x}) - h^*)$$

- We have $\mathcal{J}(\boldsymbol{u}) = h(f^{[\tau]}(\bar{x}_0, \boldsymbol{u}))$ so $\|\nabla \mathcal{J}(\boldsymbol{u})\|_2^2 = \|\nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u}) \nabla h(\boldsymbol{x})\|_2^2$
- So if $f^{[\tau]}(\bar{x}_0, \boldsymbol{u})$ satisfies

$$\forall \boldsymbol{u} \quad \sigma_{\min}(\nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})) \geq \sigma > 0$$

where $\sigma_{\min}(A) = \inf_{\|z\|>0} \|Az\|_2 / \|z\|_2$ is the minimal singular value of A then

$$\|\nabla \mathcal{J}(\boldsymbol{u})\|_2^2 \geq \sigma^2 \|\nabla h(\boldsymbol{x})\|_2^2 \geq \sigma^2 \mu(h(\boldsymbol{x}) - h^*) = \sigma^2 \mu(\mathcal{J}(\boldsymbol{u}) - \mathcal{J}^*)$$

Interpretation of a Sufficient Condition for Global Convergence

Interpretation

$$\begin{aligned}\sigma_{\min}(\nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})) > 0 &\iff \boldsymbol{\lambda} \mapsto \nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})\boldsymbol{\lambda} \text{ is injective} \\ &\iff \boldsymbol{v} \mapsto \nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})^\top \boldsymbol{v} \text{ is surjective}\end{aligned}$$

Here $\boldsymbol{y} = \nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})^\top \boldsymbol{v}$ is the linearization of the trajectories given as

$$y_{t+1} = \nabla_{x_t} f(x_t, u_t)^\top y_t + \nabla_{u_t} f(x_t, u_t)^\top v_t, \quad y_0 = 0$$

So $\sigma_{\min}(\nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})) > 0$ if the linearization of the trajectories are *surjective*

How to verify this condition from f only?

Characterization of a Sufficient Condition for Global Convergence

Lemma

If the linearization, $v \to \nabla_u f(x,u)^\top v$, of l_f-Lip. cont. dynamics f is surjective,

$$\forall x, u, \quad \sigma_{\min}(\nabla_u f(x,u)) \geq \sigma_f > 0, \tag{Surj}$$

then the linearization of the trajectories, $\boldsymbol{v} \to \nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})^\top \boldsymbol{v}$, is surjective,

$$\forall \boldsymbol{u} \quad \sigma_{\min}(\nabla_{\boldsymbol{u}} f^{[\tau]}(\bar{x}_0, \boldsymbol{u})) \geq \frac{\sigma_f}{1 + l_f} > 0,$$

Take-away: Simply need to check that the dynamic have surj. linearizations

Problem:

- Usually less control variables than state variables $\dim(u(t)) < \dim(x(t))$...
  So $\sigma_{\min}(\nabla_u f(x(t), u(t)) > 0$ impossible

$\to$ Use multistep schemes s.t. $\dim(u_t) = k \dim(x_t)$

Intuition for a Sufficient Condition for Global Convergence

Pendulum dynamics

$$ml^2\ddot{\theta}(t) = -mlg\sin\theta(t) - \mu\dot{\theta}(t) + u(t)$$

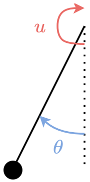

One step Euler scheme
$f(x_t, u_t) = x_{t+1}$ for $x_t = (\theta_t, \omega_t)$ with $\omega = \dot{\theta}$

$$\begin{aligned}\theta_{t+1} &= \theta_t + \Delta\omega_t \\ \omega_{t+1} &= \omega_t - \Delta(g/l\sin\theta_t - \mu\omega_t) + \Delta u_t\end{aligned}$$

Linearization surjective? ✗

Two steps Euler scheme $f(x_t, u_t) = x_{t+1}$ with $u_t = (v_t, v_{t+1/2})$

$$\begin{aligned}\theta_{t+1/2} &= \theta_t + \Delta\omega_t & \theta_{t+1} &= \theta_t + ... + \Delta^2 v_t \\ \omega_{t+1/2} &= \omega_t - \Delta(g/l\sin\theta_t - \mu\omega_t) + \Delta v_t & \omega_{t+1} &= \omega_t + \ldots + \Delta v_{t+1/2}\end{aligned}$$

Linearization surjective w.r.t. $u_t = (v_t, v_{t+1/2})$? ✓

Overall Analysis

Multistep scheme

$$f(x_t, u_t) = x_{t+1}$$
$$x_{t+(s+1)/k} = x_s + \Delta \bar{f}(x_{t+s/k}, u_{t+s/k})$$

summarized as

$$f(x_t, u_t) = \phi^{\{k\}}(x_t, u_t)$$

Control in k steps of a dynamic ϕ
For $\boldsymbol{v} = (v_0, \ldots, v_{k-1})$,

$$\phi^{\{k\}}(y_0, \boldsymbol{v}) = y_k$$
$$\text{s.t.} \quad y_{s+1} = \phi(y_s, v_s)$$

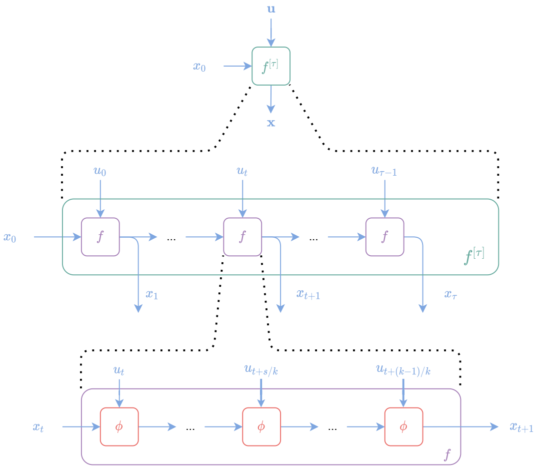

Zooming in the dynamical structure

Sufficient condition for global convergence can be verified by analyzing whether ϕ can be *linearized by static feedback*

Plan

Assumptions

Problem

$$\min_{\boldsymbol{u}} \left\{ \mathcal{J}(\boldsymbol{u}) = h(g(\boldsymbol{u})) \right\}, \text{ where } g(\boldsymbol{u}) = f^{[\tau]}(\bar{x}_0, \boldsymbol{u}), \quad h(\boldsymbol{x}) = \sum_{t=1}^{\tau} h_t(x_t)$$

Algorithm

$$\boldsymbol{u}^{(k+1)} = \boldsymbol{u}^{(k)} + \mathrm{LQR}_{\nu_k}(\mathcal{J})(\boldsymbol{u}^{(k)}) \qquad \text{(ILQR)}$$

where $\mathrm{LQR}_{\nu_k}(\mathcal{J})(\boldsymbol{u}^{(k)})$ is the oracles returning a direction computed by dynamic programming with a regularization ν_k.

Assumptions

- costs h_t
  - μ_h-strongly convex $\rightarrow$ same for total cost h
  - L_h-smooth $\rightarrow$ same for total cost h
  - M_h-smooth Hessian $\rightarrow$ same for total cost h
- dynamic f
  - l_f-Lip. continuous, L_f smooth $\rightarrow$ g is l_g-Lip.continous, L_g-smooth

$$l_g \leq l_f S, \qquad L_g \leq L_f S (l_f S + 1)^2 \qquad \text{where } S = \sum_{t=0}^{\tau-1} l_f^t$$

  - $\sigma_{\min}(\nabla_u f(x, u)) \geq \sigma_f > 0 \qquad \rightarrow \sigma_{\min}(\nabla g(\boldsymbol{u})) \geq \sigma_g = \sigma_f / (1 + l_f) > 0$

Convergence Analysis Viewpoint

ILQR as a generalized Gauss-Newton (Sideris & Bobrow 2005)

- Overall ILQR minimizes a quadratic approx. of h on top of a linear approx. of g
- So it can be summarized as

$$\begin{aligned} \mathrm{LQR}_\nu(\mathcal{J})(\boldsymbol{u}) &= \arg\min_{\boldsymbol{v}} q_h^{g(\boldsymbol{u})}(\ell_g^{\boldsymbol{u}}(\boldsymbol{v})) + \frac{\boldsymbol{\nu}}{2}\|\boldsymbol{v}\|_2^2 \\ &= -(\nabla g(\boldsymbol{u})\nabla^2 h(g(\boldsymbol{u}))\nabla g(\boldsymbol{u})^\top + \boldsymbol{\nu}\,\mathrm{I})^{-1}\nabla g(\boldsymbol{u})\nabla h(g(\boldsymbol{u})) \end{aligned}$$

  which is a *regularized generalized Gauss-Newton method*
- The *regularization* helps to interpolate between Grad Desc and Gauss Newton

Contributions

- Analysis of regularized generalized Gauss-Newton method given h strongly convex, g with surjective Jacobians $\nabla g(\boldsymbol{u})^\top$
- Similar approach as (Nesterov 2007) for modified Gauss-Newton a.k.a. prox-linear
  Global conv. of prox-linear under error bound cond. $\rightarrow$ (Drusvyatskiy & Lewis 2018)

Convergence Analysis

Global convergence idea

- Select regularization ν to ensure sufficient decrease, i.e., for $\boldsymbol{v} = \mathrm{LQR}_\nu(\boldsymbol{u})$, $\boldsymbol{x} = g(\boldsymbol{u})$, $G = \nabla g(\boldsymbol{u})$, $H = \nabla^2 h(\boldsymbol{x})$,

$$\begin{aligned}\mathcal{J}(\boldsymbol{u}+\boldsymbol{v}) &\leq \mathcal{J}(\boldsymbol{u}) + q_h^{\boldsymbol{x}} \circ \ell_g^{\boldsymbol{u}}(\boldsymbol{v}) + \frac{\nu}{2}\|\boldsymbol{v}\|_2^2 \qquad \text{(Suff Dec)} \\ &= \mathcal{J}(\boldsymbol{u}) - \frac{1}{2}\nabla h(\boldsymbol{x})^\top G^\top (GHG^\top + \nu \mathrm{I})^{-1} G \nabla h(\boldsymbol{x})\end{aligned}$$

- Given that $\sigma_{\min}(G) \geq \sigma_g$, we have $\lambda_{\min}(G^\top G) > \sigma_g^2$, so

$$\mathcal{J}(\boldsymbol{u}+\boldsymbol{v}) - \mathcal{J}(\boldsymbol{u}) \leq -\frac{\sigma_g^2}{l_g^2 L_h + \nu}\|\nabla h(\boldsymbol{x})\|_2^2 \leq -\frac{\sigma_g^2 \mu_h}{l_g^2 L_h + \nu}(\mathcal{J}(\boldsymbol{u}) - \mathcal{J}^*)$$

→ linear convergence ensured for constant ν satisfying (Suff Dec)

- Condition (Suff Dec) is ensured for $\nu(\boldsymbol{u}) = c(\|\nabla h(\boldsymbol{x})\|_2)$ with c increasing
→ decreasing regularizations can be taken to get better rates

Convergence Analysis

Local convergence idea

- By standard linear algebra, for $\boldsymbol{x} = g(\boldsymbol{u})$, $G = \nabla g(\boldsymbol{u})$, $H = \nabla^2 h(\boldsymbol{x})$,

$$
\begin{aligned}
\mathrm{LQR}_\nu(\mathcal{J})(\boldsymbol{u}) &= -(GHG^\top + \nu \mathrm{I})^{-1} G \nabla h(\boldsymbol{x}) \\
&= -G(HG^\top G + \nu \mathrm{I})^{-1} \nabla h(\boldsymbol{x}) && \text{(Push-Forward Identity)} \\
&= -G(G^\top G)^{-1}(H + \nu (G^\top G)^{-1})^{-1} \nabla h(\boldsymbol{x}) && (G^\top G \text{ invertible})
\end{aligned}
$$

- So denoting $\boldsymbol{x}^{\text{next}} = g(\boldsymbol{u} + \boldsymbol{v})$ for $\boldsymbol{v} = \mathrm{LQR}_\nu(\mathcal{J})(\boldsymbol{u})$,

$$
\boldsymbol{x}^{\text{next}} \approx g(\boldsymbol{u}) + \nabla g(\boldsymbol{u})^\top \boldsymbol{v} = \boldsymbol{x} - (\nabla^2 h(\boldsymbol{x}) + \nu (\nabla g(\boldsymbol{u})^\top \nabla g(\boldsymbol{u}))^{-1})^{-1} \nabla h(\boldsymbol{x}).
$$

$\rightarrow$ Approximate Newton method on the trajectories $\boldsymbol{x}$ for $\nu \ll 1$

$\rightarrow$ Quadratic local convergence can be ensured for decreasing regularizations ν

Complexity Bound for ILQR

Theorem

Consider strongly convex, smooth, Hessian-smooth costs h_t and Lip. cont., smooth dynamics f with surjective linearizations, the ILQR algorithm equipped with $\nu(\boldsymbol{u}) = \bar{\nu}\|\nabla h(g(\boldsymbol{u}))\|_2$ for $\bar{\nu}$ large enough converges to accuracy ε in at most

$$\underbrace{4\theta_g(\sqrt{\delta_0} - \sqrt{\delta})}_{\text{1st phase}} + \underbrace{2\rho_h \ln\left(\frac{\delta_0}{\delta}\right) + 2\alpha \ln\left(\frac{\theta_g\sqrt{\delta_0} + \rho_g}{\theta_g\sqrt{\delta} + \rho_g}\right)}_{\text{2nd phase}} + \underbrace{O(\ln\ln(\varepsilon))}_{\text{3rd phase}}$$

iterations, each having a comput. complexity $O(\tau(\dim(x) + \dim(u))^3)$, where

- $\delta_0 = \mathcal{J}(\boldsymbol{u}^{(0)}) - \mathcal{J}^*$ *is the initial gap*
- $\delta = 1/(32\rho_h(\theta_h(1 + \sqrt{\rho_h}{\rho_g}^3/3) + \sqrt{\rho_h}\theta_g(1 + \rho_g\rho_h))^2)$ *is the gap of quadratic conv.*
- $\rho_h = L_h/\mu_h$ *is the cond. nb of the costs*
- $\rho_g = l_g/\sigma_g$ *is the cond. nb of the linearized traj.*
- $\theta_h = M_h/\mu_h^{3/2}$ *is the param. of self-concordance of the costs*
- $\theta_g = L_g/(\sigma_g^2\sqrt{\mu_h})$ *acts a self-concordance param. for the linear-quadratic decomp.*
- $\alpha = 4{\rho_g}^2(2{\rho_g}^2\theta_h/(3\theta_g) + \rho_h)$ *is another cond. nb*

Conclusion

Outcomes

- Identified a simple sufficient condition for global convergence
- Provided detailed complexity bounds for ILQR and IDDP

Long-term Objectives

- Identify the impact of
  - discretization stepsize Δ
  - discretization method
- Inform optimal window size for Model Predictive Control

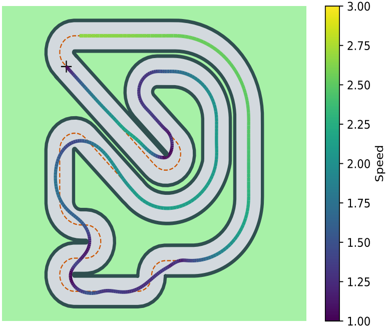

Optimized traj. with MPC
& contouring objective

Thank you for your attention!

Static Feedback Linearization

Definition (Static Feedback Linearization for $\dim(y_t) = d, \dim(v_t) = 1$)

A dynamical system $y_{t+1} = \phi(y_t, v_t)$ is linearizable by static feedback
if there exists some diffeomorphism a and $b(y, \cdot)$ s.t.
the reparameterized sytem $z_t = a(y_t)$, $w_t = b(y_t, v_t)$ is linear of the form

$$z_{t+1}^{(i)} = z_t^{(i+1)} \text{ for all } i \in \{1, \dots, d-1\}, \quad z_{t+1}^{(d)} = w_t,$$

Examples

- System driven by its acceleration, with $|\partial_{v_t} \psi(y_t, v_t)| > 0$

$$y_{t+1}^{(1)} = y_t^{(1)} + \Delta y_t^{(2)}, \quad y_{t+1}^{(2)} = y_t^{(2)} + \Delta \psi(y_t, v_t)$$

- System driven by its d^{th} derivative
- More generally, (Aranda-Bricaire et al. 1996) essentially showed that local feedback linearization $\iff$ reachability of any state by ϕ
  Proof is constructive and might be quantified

Multistep Schemes and Static Feedback Linearization

Idea

- If $z_{t+1}^{(i)} = z_t^{(i+1)}$ for all $i \in \{1, \dots, d-1\}$, $z_{t+1}^{(d)} = w_t$, then

$$
\begin{array}{ccccc}
z^{(1)} & z^{(2)} & z^{(3)} & & w_0 \\
\vdots & \vdots & \vdots & & \vdots \\
\vdots & \vdots & z^{(d)} & \dots & \vdots \\
\vdots & z^{(d)} & w_0 & & \vdots \\
z^{(d)} & w_0 & w_1 & & w_{d-1} \\
 & & & & \\
t=0 & t=1 & t=2 & \dots & t=d
\end{array}
$$

$\rightarrow$ By considering d steps $z_d = w_{0:d-1}$,
$\rightarrow$ so the control in d steps of the reparameterized system is the identity
$\rightarrow$ so it clearly has surjective linearizations

- This property is kept under the diffeomorphisms a, b

Multistep Schemes and Static Feedback Linearization

Theorem (for $\dim(y_t) = d, \dim(v_t) = 1$)

If the system defined by $y_{t+1} = \phi(y_t, v_t)$ *is linearizable by static feedback with transformations* a *and* b *that are Lipschitz-continuous and such that*

$$\forall y \ \sigma_{\min}(\nabla a(y)) \geq \sigma_a > 0, \quad \inf_{y,v} \sigma_{\min}(\nabla_v b(y, v)) \geq \sigma_b > 0,$$

then the control in $k \geq d$ *steps of the dynamic* ϕ *satisfies,*

$$\inf_{y_0, \boldsymbol{v}} \sigma_{\min}(\nabla_{\boldsymbol{v}} \phi^{\{k\}}(y_0, \boldsymbol{v})) \geq \frac{\sigma_b}{l_a} \frac{1}{1 + (d-1) l_b / \sigma_a} > 0.$$

Take-away:

- Having access to the exact diffeomorphisms a, b may be intractable
- But showing their existence may be possible
  and global convergence guarantees follow

Extended Analysis for ILQR

Theorem

Given smooth convex costs h_t s.t., for some $\mu > 0$ and $r \in [1/2, 1)$,

$$\|\nabla h(x)\|_2 \geq \mu_h^r (h(x) - h^*)^r$$

and Lip. cont., smooth dynamics f with surjective linearizations, the ILQR algorithm converges globally with a complexity

$$O(\varepsilon^{2r-1}/(2r-1) + \delta_0^{1-r}/(1-r)), \quad \textit{i.e.} \quad O(\ln(\varepsilon) + \sqrt{\delta_0}) \textit{ if } r = 1/2$$

Theorem

Given convex, smooth, Hessian-smooth, self-concordant cost h and Lip. cont., smooth dynamics f with surjective linearizations, the ILQR algorithm converges locally with a quadratic rate

- Precise rates given in the paper in terms of the cond nb defined before

Complexity bound for IDDP

Idea
Analyze IDDP as an approximate ILQR similar as (Murray & Yakowitz 1984) for local conv.

Lemma

Given strongly convex, smooth, Hessian-smooth costs h_t,
Lip. cont., smooth dynamics f with surj. linearizations, there exists $\eta > 0$ s.t.

$$\forall \boldsymbol{u}, \nu \quad \| \mathrm{DDP}_\nu(\mathcal{J})(\boldsymbol{u}) - \mathrm{LQR}_\nu(\mathcal{J})(\boldsymbol{u})\|_2 \leq \eta \| \mathrm{LQR}_\nu(\mathcal{J})(\boldsymbol{u})\|_2^2$$

Theorem

Consider strongly convex, smooth, Hessian-smooth costs h_t
and Lip. cont., smooth dynamics f with surjective linearizations,
the IDDP algo. equipped with appropriate regularization
converges globally with a local quadratic rate.

Numerical Illustrations

Simple Car with Tracking Cost

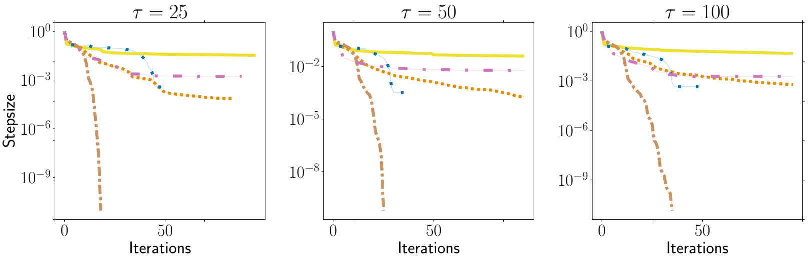

Realistic Car with Contouring Cost

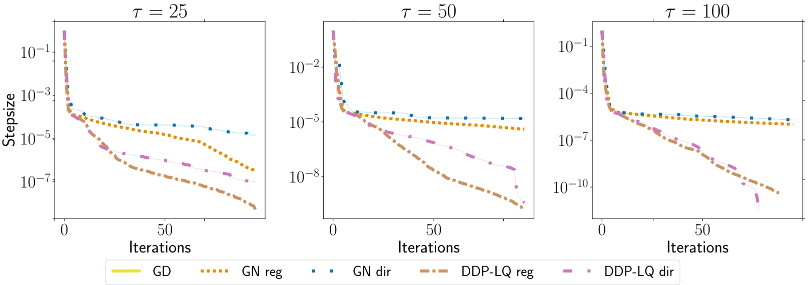

Aranda-Bricaire, E., Kotta, Ü. & Moog, C. (1996), 'Linearization of discrete-time systems', *SIAM Journal on Control and Optimization* **34**(6), 1999–2023.

Bergou, E. H., Diouane, Y. & Kungurtsev, V. (2020), 'Convergence and complexity analysis of a Levenberg-Marquardt algorithm for inverse problems', *Journal of Optimization Theory and Applications* **185**(3), 927–944.

Diehl, M. & Messerer, F. (2019), Local convergence of generalized Gauss-Newton and sequential convex programming, *in* '2019 IEEE 58th Conference on Decision and Control (CDC)', pp. 3942–3947.

Drusvyatskiy, D. & Lewis, A. S. (2018), 'Error bounds, quadratic growth, and linear convergence of proximal methods', *Mathematics of Operations Research* **43**(3), 919–948.

Grizzle, J. (1986), Feedback linearization of discrete-time systems, *in* 'Analysis and optimization of systems', Springer, pp. 273–281.

Isidori, A. (1995), *Nonlinear Control Systems*, 3rd edn, Springer-Verlag.

Jakubczyk, B. (1987), 'Feedback linearization of discrete-time systems', *Systems & Control Letters* **9**(5), 411–416.

Jakubczyk, B. & Sontag, E. D. (1990), 'Controllability of nonlinear discrete-time systems: A Lie-algebraic approach', *SIAM Journal on Control and Optimization* **28**(1), 1–33.

Li, W. & Todorov, E. (2007), 'Iterative linearization methods for approximately optimal control and estimation of non-linear stochastic system', *International Journal of Control* **80**(9), 1439–1453.

Liao, L.-Z. & Shoemaker, C. (1991), 'Convergence in unconstrained discrete-time differential dynamic programming', *IEEE Transactions on Automatic Control* **36**(6), 692–706.

Liniger, A., Domahidi, A. & Morari, M. (2015), 'Optimization-based autonomous racing of 1: 43 scale rc cars', *Optimal Control Applications and Methods* **36**(5), 628–647.

Mangasarian, O. L. (1966), 'Sufficient conditions for the optimal control of nonlinear systems', *SIAM Journal on Control* **4**(1), 139–152.

Murray, D. & Yakowitz, S. (1984), 'Differential dynamic programming and Newton's method for discrete optimal control problems', *Journal of Optimization Theory and Applications* **43**(3), 395–414.

Nesterov, Y. (2007), 'Modified Gauss-Newton scheme with worst case guarantees for global performance', *Optimisation methods and software* **22**(3), 469–483.

Polak, E. (2011), 'On the role of optimality functions in numerical optimal control', *Annual Reviews in Control* **35**(2), 247–253.

Sideris, A. & Bobrow, J. E. (2005), An efficient sequential linear quadratic algorithm for solving nonlinear optimal control problems, *in* 'Proceedings of the 2005 American Control Conference', pp. 2275–2280.

Sontag, E. D. (2013), *Mathematical control theory: deterministic finite dimensional systems*, Vol. 6, Springer Science & Business Media.

Tassa, Y., Erez, T. & Todorov, E. (2012), Synthesis and stabilization of complex behaviors through online trajectory optimization, *in* '2012 IEEE/RSJ International Conference on Intelligent Robots and Systems', pp. 4906–4913.

Yamashita, N. & Fukushima, M. (2001), On the rate of convergence of the Levenberg-Marquardt method, *in* 'Topics in numerical analysis', Springer, pp. 239–249.