

Nonlinear Control Algorithms as Globally Convergent Generalized Gauss-Newton Methods in a Differentiable Programming Framework

Vincent Roulet

Co-authors: Siddhartha Srinivasa, Maryam Fazel, Zaid Harchaoui

paper: <https://arxiv.org/abs/2204.02322>

code: <https://github.com/vroulet/ilqc>


04/13/2022



Optimization in a Differentiable Programming Framework

Differentiable programming framework

- To solve $\min_u F(u)$, needs oracle as $\nabla F(u)$
- 1. Record gradients of elementary computations
→ Needs differentiable programming framework
- 2. Use chain-rule along graph of computations
→ Back-propagate gradients




Generic graph of computations

Today's problem

- Simple dynamical structure $x_{t+1} = f(x_t, u_t)$
- Canonical example: **nonlinear control**

Why?

- Algorithms used are not just a gradient descent
- Surprising empirical performance
- May serve as a starting point to extend differentiable programming methods



Graph of computations in nonlinear control

Nonlinear Control Problems

Continuous Time Control problem

- System driven by dynamics $\dot{x}(t) = f(x(t), u(t))$
- Minimize cost $h(x(t), t)$ 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_{τ}

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




Dynamics of a car



Tracking objective

Autonomous Car Racing

Simple model of a car




$$x = (z_x, z_y, \theta, v), \quad u = (\delta, a)$$


$$\dot{z}_x = v \cos \theta \quad \dot{\theta} = v \tan(\delta)$$

$$\dot{z}_y = v \sin \theta \quad \dot{v} = a$$

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

Questions

1. What are sufficient conditions to ensure global convergence?
2. How can we understand these algorithms from an optimization viewpoint?
3. 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 or generalized Gauss-Newton
(Polak 2011, Murray & Yakowitz 1984, Liao & Shoemaker 1991, Yamashita & Fukushima 2001, Diehl & Messerer 2019)
- **(Unregularized)** Gauss-Newton, Newton methods for nonlinear control
(Sideris & Bobrow 2005, Dunn & Bertsekas 1989, Wright 1990)

Outline

Iterative Linear Quadratic Optimization Algorithms for Nonlinear Control

A Sufficient Condition for Global Convergence

Convergence Analysis of ILQR and IDDP

Discrete Time Control Problems

Continuous Time Control problem

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

Discrete Time Control Problem


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

Discretization schemes: (time-step Δ)


Euler: $f(x_t, u_t) = x_t + \Delta f(x_t, u_t)$

Multi-step: $f(x_t, u_t) = x_{t+1}$

s.t. $x_{t+(s+1)/k} = x_s + \Delta f(x_{t+s/k}, u_{t+s/k})$
 $\dim(u_t) = k \dim(u(t))$ for k steps



Euler discretization



2-step discretization

Nonlinear Control Algorithms for Discrete Time Control Problems

- Forward** Given a sequence of controls $u_0, \dots, u_{\tau-1}$
- Compute associated trajectory $x_{t+1} = f(x_t, u_t)$
 - Record linear expansions $\ell_f^{x_t, u_t}$ of the dynamics f around x_t, u_t
 - Record quadratic expansions $q_{h_t}^{x_t}$ of the costs around x_t

Backward Solve the associated regularized linear-quadratic control problem

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

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

by computing recursively the cost-to-go from y_t at time t , from $c_\tau = q_{h_\tau}$,

$$c_t : y_t \mapsto \underbrace{q_{h_t}^{x_t}(y_t)}_{\text{current cost}} + \underbrace{\min_{v_t} \left\{ \frac{\nu}{2} \|v_t\|_2^2 + c_{t+1}(\ell_f^{x_t, u_t}(y_t, v_t)) \right\}}_{\substack{\text{optimal move at time } t \\ \text{associated policy } \pi_t : y_t \mapsto v_t^*}} \quad \left(\begin{array}{l} \text{lin. quad. problem} \\ \rightarrow \text{closed form sol.} \end{array} \right)$$

Roll-out Update the iterates as $u_t^{\text{next}} = u_t + v_t$
where v_t are computed by rolling-out the policies π_t along either

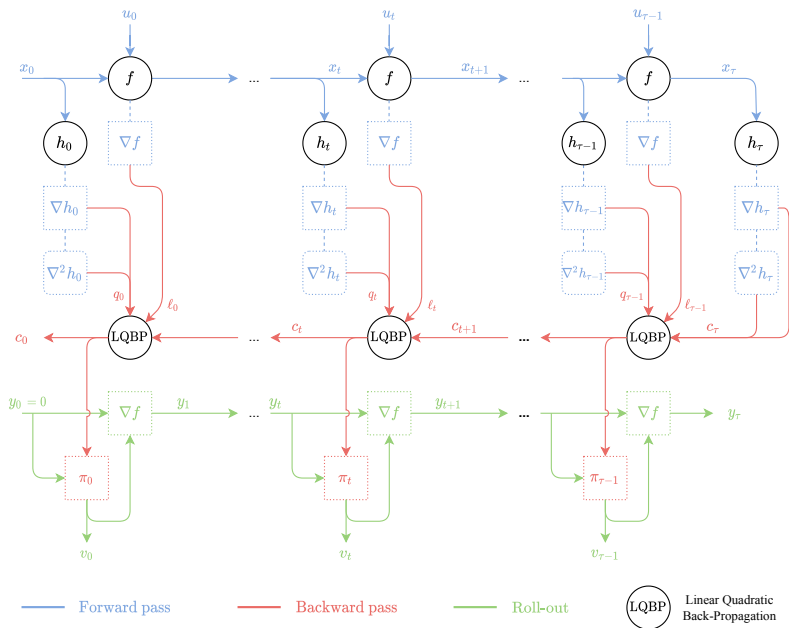
- the linearized dynamics \rightarrow **Iterative Linear Quadratic Regulator (ILQR)** (Li & Todorov 2007)

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

- the original dynamics \rightarrow **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



Outline

Iterative Linear Quadratic Optimization Algorithms for Nonlinear Control

A Sufficient Condition for Global Convergence

Convergence Analysis of ILQR and IDDP

Objective Decomposition

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

$$f^{[\tau]}(x_0, \mathbf{u}) = (x_1; \dots; x_\tau)$$

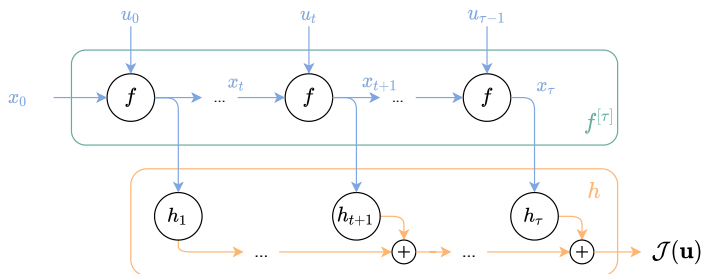
$$\text{s.t. } x_{t+1} = f(x_t, u_t)$$

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

Composite objective for $x_0 = \bar{x}_0$

$$\mathcal{J}(\mathbf{u}) = h(f^{[\tau]}(x_0, \mathbf{u})) = \sum_{t=1}^{\tau} h_t(x_t)$$

$$\text{s.t. } x_{t+1} = f(x_t, u_t)$$




A Sufficient Condition for Global Convergence

Idea:

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

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

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.

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

- We have $\mathcal{J}(\mathbf{u}) = h(f^{[\tau]}(x_0, \mathbf{u}))$ so $\|\nabla \mathcal{J}(\mathbf{u})\|_2^2 = \|\nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u}) \nabla h(\mathbf{x})\|_2^2$
- So if $f^{[\tau]}(x_0, \mathbf{u})$ satisfies

$$\forall \mathbf{u} \quad \underline{\sigma}(\nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u})) := \inf_{\boldsymbol{\lambda}} \frac{\|\nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u}) \boldsymbol{\lambda}\|_2}{\|\boldsymbol{\lambda}\|_2} \geq \sigma > 0$$

then

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

Interpretation of a Sufficient Condition for Global Convergence

Interpretation

$$\underline{\sigma}(\nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u})) > 0$$

\iff Reverse mode of auto-diff $\boldsymbol{\lambda} \mapsto \nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u}) \boldsymbol{\lambda}$ is injective

\iff Forward mode of auto-diff $\mathbf{v} \mapsto \nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u})^\top \mathbf{v}$ is surjective

Here $\mathbf{y} = \nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u})^\top \mathbf{v}$ is the linearization of the trajectories given as

$$\mathbf{y}_{t+1} = \nabla_{x_t} f(x_t, u_t)^\top \mathbf{y}_t + \nabla_{u_t} f(x_t, u_t)^\top \mathbf{v}_t, \quad \mathbf{y}_0 = 0$$

So $\underline{\sigma}(\nabla_{\mathbf{u}} f^{[\tau]}(x_0, \mathbf{u})) > 0$ 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 (R. et al. (2022))

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

$$\forall x, u, \quad \underline{\sigma}(\nabla_u f(x, u)) \geq \sigma_f > 0,$$

then the linearization of the trajectories, $v \rightarrow \nabla_u f^{[\tau]}(x_0, u)^\top v$, is surjective,

$$\forall x_0, u, \quad \underline{\sigma}(\nabla_u f^{[\tau]}(x_0, u)) \geq \frac{\sigma_f}{1 + l_f} > 0,$$

→ Simply need to check that the dynamic has 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
- Use multistep schemes s.t. $\dim(u_t) = k \dim(u(t))$

Intuition for a Sufficient Condition for Global Convergence

Pendulum dynamics


$$m\ddot{\theta}(t) = -mg \sin \theta(t) - \mu\dot{\theta}(t) + u(t)$$

One step Euler scheme

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

$$\text{angle } \theta_{t+1} = \theta_t + \Delta\omega_t$$

$$\text{angle speed } \omega_{t+1} = \omega_t - \Delta(g \sin \theta_t - \mu\omega_t) + \Delta u_t$$



Linearization surjective? **X**

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

$$\theta_{t+1/2} = \theta_t + \Delta\omega_t$$

$$\theta_{t+1} = \theta_t + \dots + \Delta^2 v_t$$

$$\omega_{t+1/2} = \omega_t - \Delta(g \sin \theta_t - \mu\omega_t) + \Delta v_t$$

$$\omega_{t+1} = \omega_t + \dots + \Delta v_{t+1/2}$$

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

Overall Analysis

Multistep scheme

$$f(x_t, u_t) = x_{t+1}$$

$$\begin{aligned} \text{s.t. } x_{t+(s+1)/k} &= x_{t+s/k} + \Delta f(x_{t+s/k}, u_{t+s/k}) \\ &:= \phi(x_{t+s/k}, u_{t+s/k}) \end{aligned}$$


→ study dynamical struct. of f itself

Control of a dynamic ϕ in k steps

for $\mathbf{v} = (v_0; \dots; v_{k-1})$,

$$\phi^{\{k\}}(y_0, \mathbf{v}) = y_k$$

$$\text{s.t. } y_{s+1} = \phi(y_s, v_s)$$



Zooming in the dynamical structure

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

Linearization Scheme General Idea

Definition (Simplified see e.g. (Isidori 1995, Sontag 2013))

A dynamical system $y_{s+1} = \phi(y_s, v_s)$ is linearizable by static feedback if there exists some diffeomorphisms a and $b(y, \cdot)$ such that the reparam. system $z_s = a(y_s)$, $w_s = b(y_s, v_s)$ is linear, i.e., $z_{s+1} = Az_s + Bw_s$

Simple Example

- System driven by its d^{th} derivative (like acceleration in the pendulum example)

$$y_{t+1}^{(i)} = y_t^{(i)} + \Delta y_t^{(i+1)} \text{ for } i \in \{1, \dots, d-1\}, \quad y_{t+1}^{(d)} = y_t^{(d)} + \Delta \psi(y_t, v_t)$$

s.t. $|\psi(y_t, v_t)| \neq 0$ for all y_t, v_t , with Δ the time step

Theorem (R. et al. (2022) simplified¹)

If a d -dimensional system defined by $y_{t+1} = \phi(y_t, v_t)$ is linearizable by static feedback then its control in d steps $\phi^{\{d\}}(y, \mathbf{v})$ has surjective linearizations. Hence a control problem with dynamic $f = \phi^{\{d\}}$ and strongly convex costs h satisfy a gradient dominating property.

The problem could be solved by gradient descent
But the algorithms are not a gradient descent!

¹Quantitative results available

Outline

Iterative Linear Quadratic Optimization Algorithms for Nonlinear Control

A Sufficient Condition for Global Convergence

Convergence Analysis of ILQR and IDDP

Setup

Problem

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

Algorithm

$$\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \text{LQR}_{\nu_k}(\mathcal{J})(\mathbf{u}^{(k)}) \quad (\text{ILQR})$$

where $\text{LQR}_{\nu_k}(\mathcal{J})(\mathbf{u}^{(k)})$ is the oracle returning a direction computed by dynamic programming with a regularization ν_k

Assumptions

- costs h_t : μ_h -strongly convex, L_h -smooth, M_h -smooth Hessian
→ same for overall cost h
- dynamic f : l_f -Lip. continuous, L_f smooth with $\underline{\sigma}(\nabla_u f(x, u)) \geq \sigma_f > 0$
→ mapping g : l_g -Lip.continuous, L_g -smooth with $\underline{\sigma}(\nabla g(\mathbf{u})) \geq \sigma_g > 0$
with l_g, L_g, σ_g estimable from l_f, L_f, σ_f

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

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

which is a *regularized generalized Gauss-Newton method*

Convergence proof idea

1. For large enough regularization, $\text{LQR}_{\nu}(\mathcal{J})(\mathbf{u}) \approx -\nu^{-1} \nabla g(\mathbf{u}) \nabla h(g(\mathbf{u}))$
→ linear global convergence possible as for a gradient descent
2. Denoting $\mathbf{x}^{\text{next}} = g(\mathbf{u} + \mathbf{v})$ for $\mathbf{v} = \text{LQR}_{\nu}(\mathcal{J})(\mathbf{u})$, with simple linear algebra,

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

- so for small enough regularization $\mathbf{x}^{\text{next}} \approx \mathbf{x} - \nabla^2 h(\mathbf{x})^{-1} \nabla h(\mathbf{x})$
→ local quadratic convergence possible as for a Newton method

3. Can show that a regularization $\nu \propto \|\nabla h(\mathbf{x})\|_2$ ensures both!

Complexity Bound for ILQR

Theorem (R. et al. (2022)¹)

Under the aforementioned assumptions, the ILQR algorithm equipped with $\nu(\mathbf{u}) = \bar{\nu} \|\nabla h(g(\mathbf{u}))\|_2$ for $\bar{\nu}$ large enough converges to accuracy ε in

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

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

- $\delta_0 = \mathcal{J}(\mathbf{u}^{(0)}) - \mathcal{J}^*$ 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 **condition number** of the costs
- $\rho_g = l_g/\sigma_g$ is the **condition number** 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 as **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

¹Extensions to self-concordant costs or gradient dominated costs available

Complexity Bound for IDDP

Idea

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

Lemma (R. et al. (2022))

Under the aforementioned assumptions, denoting $\text{DDP}_\nu(\mathcal{J})(\mathbf{u})$, $\text{LQR}_\nu(\mathcal{J})(\mathbf{u})$ the oracles returned by IDDP and ILQR resp., there exists $\eta > 0$ s.t.

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

Theorem (R. et al. (2022))

Under the aforementioned assumptions, the IDDP algorithm equipped with appropriate regularization converges globally with a local quadratic rate.

Code Example from Toolbox ILQC

```
import torch
from envs.car import Car
from envs.backward import lin_quad_backward, quad_backward
from envs.rollout import roll_out_lin

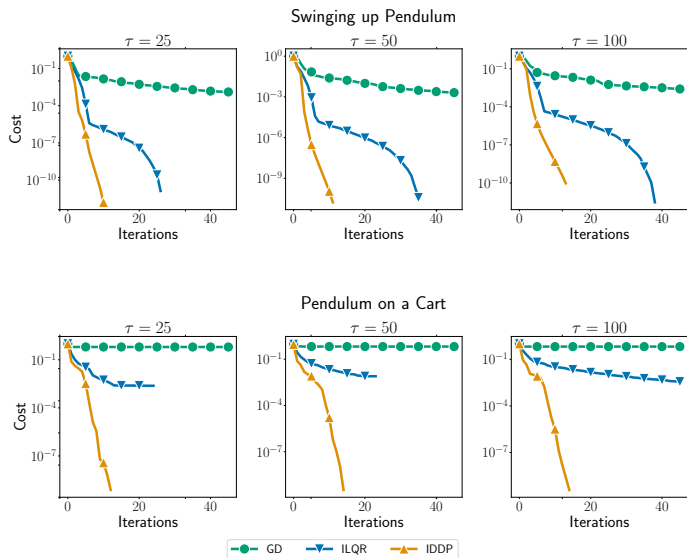
# Define control problem and candidate control variables
env = Car(model='simple', discretization='euler', cost='exact',
horizon=50, dt=0.02)
ctrls = torch.randn(env.horizon, env.dim_ctrl, requires_grad=True)

# ILQR/Gauss-Newton step
traj, costs = env.forward(ctrls, approx='linquad')
policies = lin_quad_backward(traj, costs, reg_ctrl=1.)[0]
gauss_newton_dir = roll_out_lin(traj, policies)
gauss_newton_step = ctrls + gauss_newton_dir

# IDDP step
iddp_dir = roll_out_exact(traj, policies)
iddp_step = ctrls + iddp_dir

# Newton and DDP with quad. approx. also available
```


Numerical Illustrations

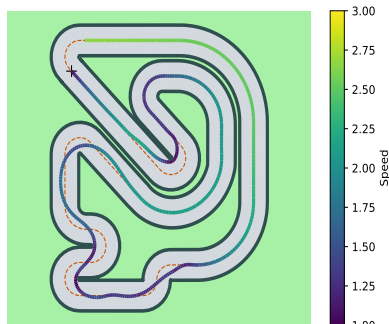


IDDP exploits differentiable programming but is not a classical GN method
Can we derive similar algorithms that exploit the problem structure?

Conclusion

Summary

- Conv. guarantees for canonical noncvx pb
→ analyze problem at elementary scale
as done in a diff. prog. implementation
- Complexity bounds for ILQR and IDDP
→ quad. convergence at low iteration cost
by using a diff. prog. implementation
- Generalized back-propagation as in IDDP
→ consider alternate sol. for oracle subpbs,
use similar graph of computations



Model Predictive Control
& contouring objective

Thank you for your attention!

- 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.
- Dunn, J. & Bertsekas, D. (1989), 'Efficient dynamic programming implementations of Newton's method for unconstrained optimal control problems', *Journal of Optimization Theory and Applications* **63**(1), 23–38.
- Isidori, A. (1995), *Nonlinear Control Systems*, 3rd edn, Springer-Verlag.
- 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. (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.
- Polak, E. (2011), 'On the role of optimality functions in numerical optimal control', *Annual Reviews in Control* **35**(2), 247–253.
- R., V., Srinivasa, S., Fazel, M. & Harchaoui, Z. (2022), 'Complexity bounds of iterative linear quadratic optimization algorithms for discrete time nonlinear control', *arXiv preprint arXiv:2204.02322*.
- Sideris, A. & Bobrow, J. (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. (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.
- Wright, S. (1990), 'Solution of discrete-time optimal control problems on parallel computers', *Parallel Computing* **16**(2-3), 221–237.
- Yamashita, N. & Fukushima, M. (2001), On the rate of convergence of the Levenberg-Marquardt method, *in* 'Topics in numerical analysis', Springer, pp. 239–249.