Learning for Control
A Practitioner's View

Manfred Morari

Mahyar Fazlyab, Alex Robey, Hamed Hassani, George J. Pappas

oCPS Fall School, Eindhoven

October 31, 2019

A Practitioner's Perspective

- Chemical Process Control
  - Shell, BP, Exxon, DuPont, ICI PLC
- Building Climate/Energy Control (HVAC)
  - Siemens, Carrier
- Automotive Systems
  - Ford, Daimler-Chrysler
- Aircraft Systems
  - United Technologies
- Power Electronics, Electrical Power Systems
  - ABB

MPC Workshop 1998

Nonlinear Model Predictive Control Workshop
Frank Allgöwer, Alex Zheng
Ascona, 1998

Dominated by Process Control

MPC Workshop 2008

INTERNATIONAL WORKSHOP ON ASSESSMENT AND FUTURE DIRECTIONS OF NONLINEAR MODEL PREDICTIVE CONTROL

September 5-9, 2008
Pavia, Italy

University of Pavia
CeRS - IUSS Pavia

Lalo Magni, Davide Raimondo,
Frank Allgöwer

Process Control has almost disappeared

Applications in automotive, power electronics,…

Applications in Automotive

Developments in Predictive and Optimization-Based Control of Automotive Powertrain Systems

Ilya Kolmanovsky
Ford Research and Advanced Engineering

Research & Advanced Engineering
Ford Motor Company

ETH, November 2008

- Model Predictive Control of engine idle speed
- Preview control of boosted gasoline engines
- Optimal and predictive control of Hybrid Electric Vehicles

Applications in Power Electronics

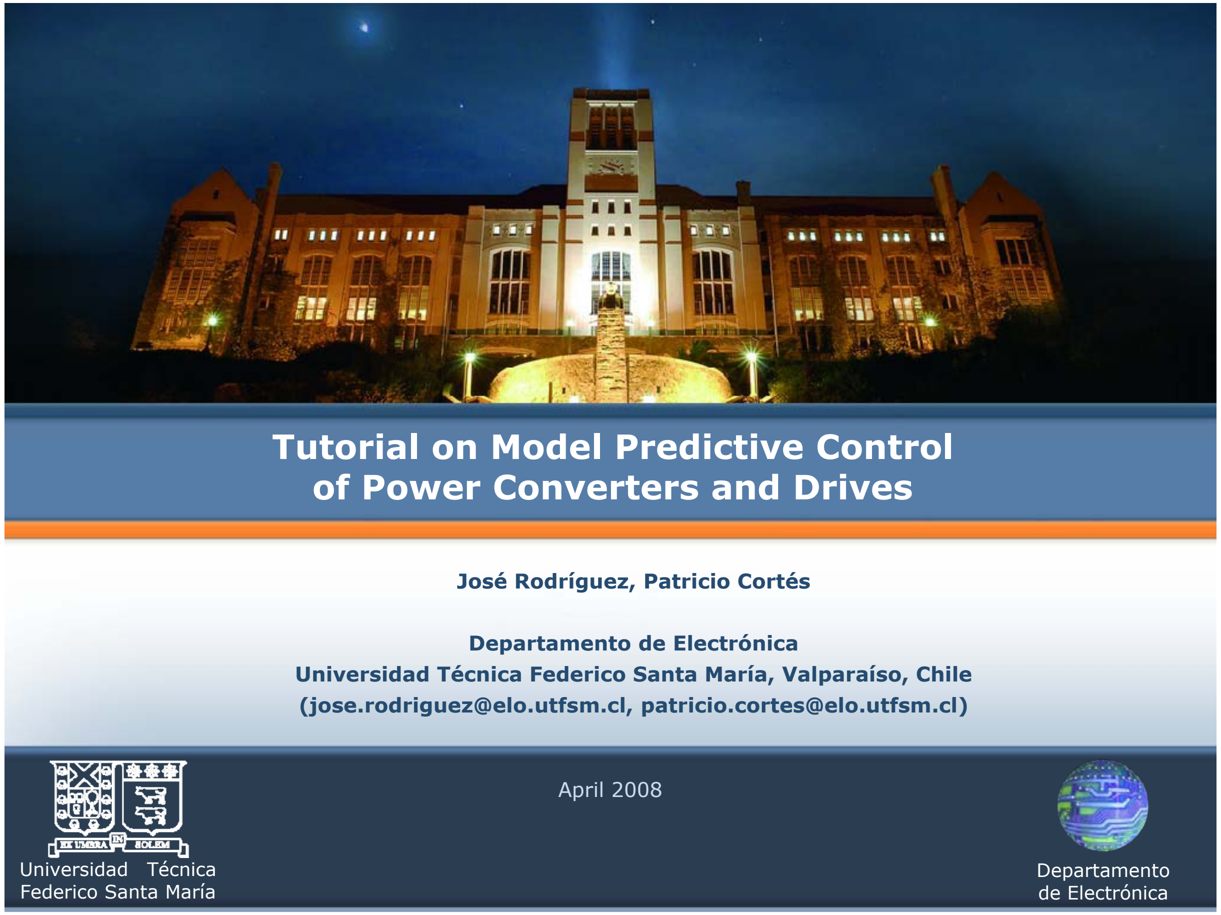

Speedup of software for MIPs

Calculations

Improvement in MIP Software from 1988-2017

- Algorithms: 147650x
- Machines: 17120x

  http://preshing.com/20120208/a-look-back-at-single-threaded-cpu-performance/
- NET: (Algorithm × Machine): 2,527,768,000x

What Does This "Mean"?

- A "typical" MILP that would have taken 124 years to solve in 1988 will solve in 1 second now.
- This is amazing, but your mileage may vary

Computation / Software

Formal specification
- YALMIP
- HYSDEL
- Linear + Hybrid models

Control law
- Explicit MPC
- Fixed-complexity solutions

Software synthesis
- Real-time workshop
- Bounded-time solvers
- Verifiable code generation

Verified controller

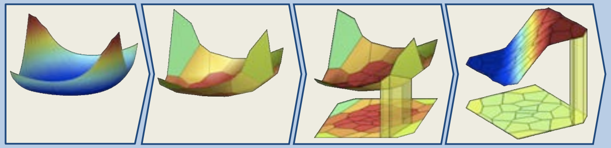

Multi-Parametric Toolbox (MPT)

- (Non)-Convex Polytopic Manipulation
- Multi-Parametric Programming
- Control of PWA and LTI systems

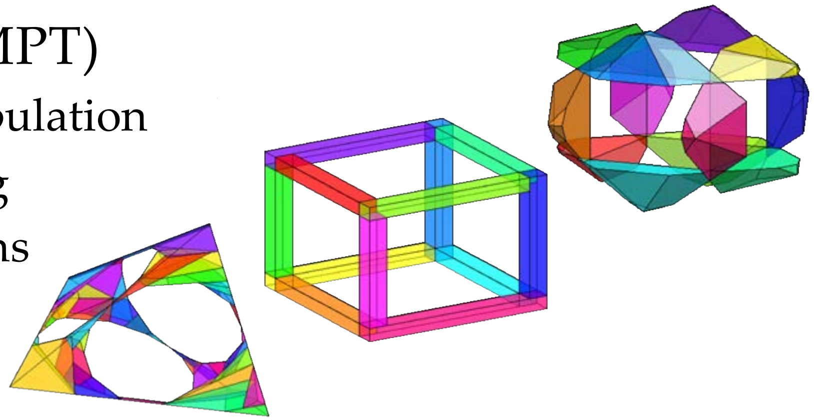

MPT 3.0

Methods and Tools for Embedded MPC under leadership of S. Richter and A. Domahidi

Offline

Certified Learning of
MPC Control Laws

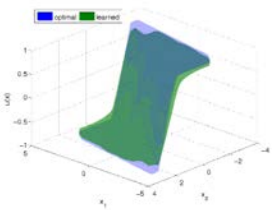

Learning Binary
Decision Rules

Goal:
Enable advanced
MPC formulations
on embedded
platforms

Online

Code Generation
for First-order Methods

FiOrdOs

Code Generation
for Multistage QCQPs

FORCES

Embedded SOCP
Programming

ECOS

FORCES PRO user map

UBS Global Finalist 2015
The UBS Future of Finance Challenge

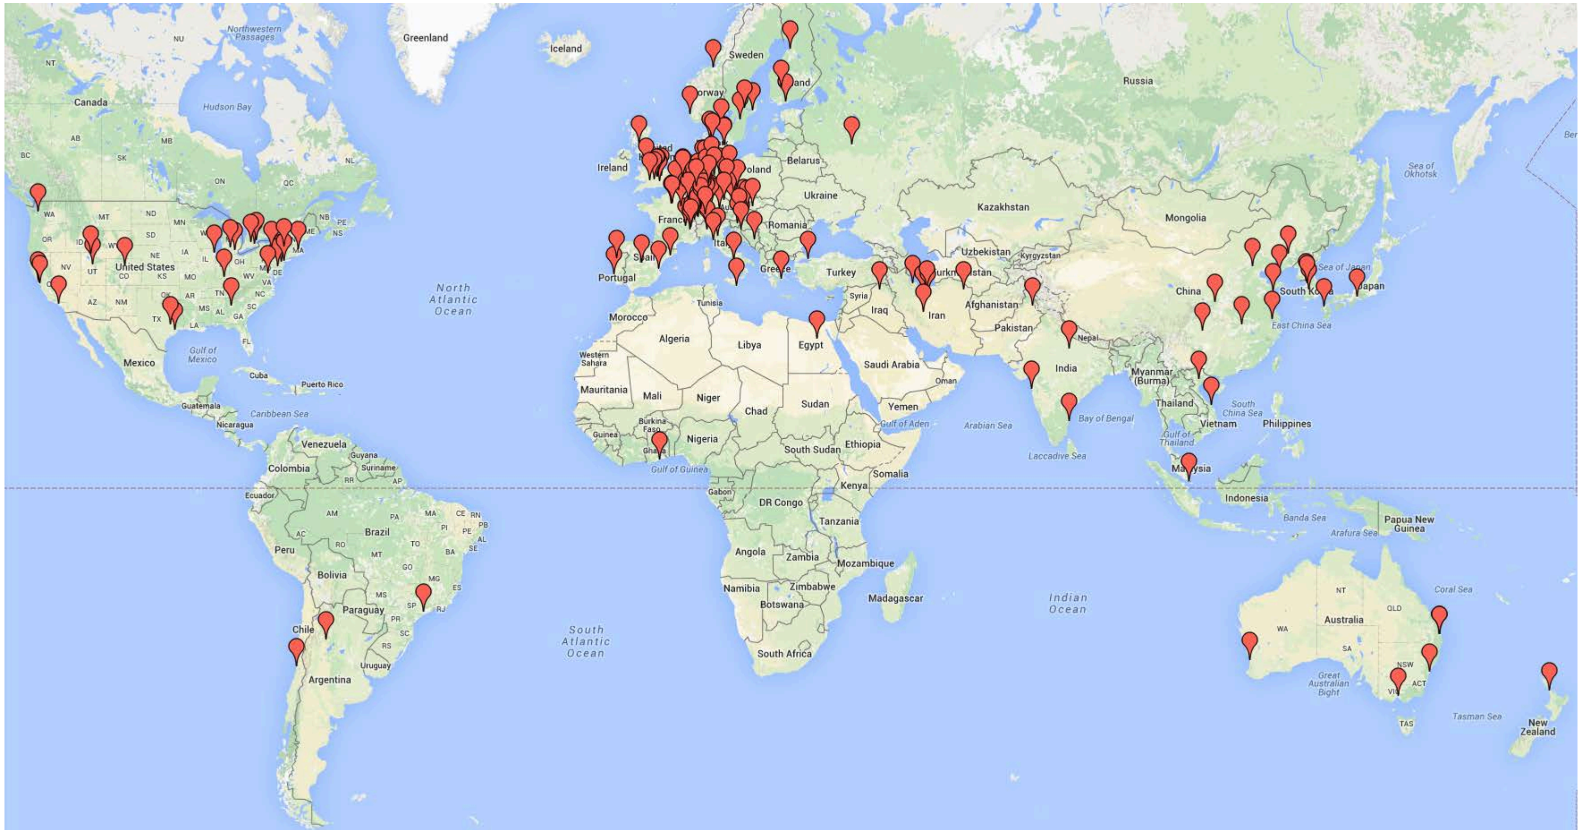

More than 150 users world wide & across industries in 35 countries

embotech

Spinoff ETHzürich

Embedded Online Optimization for Model Predictive Control at Megahertz Rates

Juan L. Jerez, *Student Member, IEEE*, Paul J. Goulart, Stefan Richter,
George A. Constantinides, *Senior Member, IEEE*, Eric C. Kerrigan, *Member, IEEE*, and
Manfred Morari, *Fellow, IEEE*

***Abstract*—Faster, cheaper, and more power efficient optimization solvers than those currently possible using general-purpose techniques are required for extending the use of model predictive control (MPC) to resource-constrained embedded platforms. We propose several custom computational architectures for different first-order optimization methods that can handle linear-quadratic MPC problems with input, input-rate, and soft state constraints. We provide analysis ensuring the reliable operation of the resulting controller under reduced precision fixed-point arithmetic. Implementation of the proposed architectures in FPGAs shows that satisfactory control performance at a sample rate beyond 1 MHz is achievable even on low-end devices, opening up new possibilities for the application of MPC on embedded systems.**

***Index Terms*—Embedded systems, optimization algorithms, predictive control of linear systems.**

ABB introduces ACS6080 drive for high performance motor control

Press release | Zurich, Switzerland | 2019-03-07

with MP^{3}C

Model Predictive Pulse Pattern Control

Model Predictive Control in Power Electronics

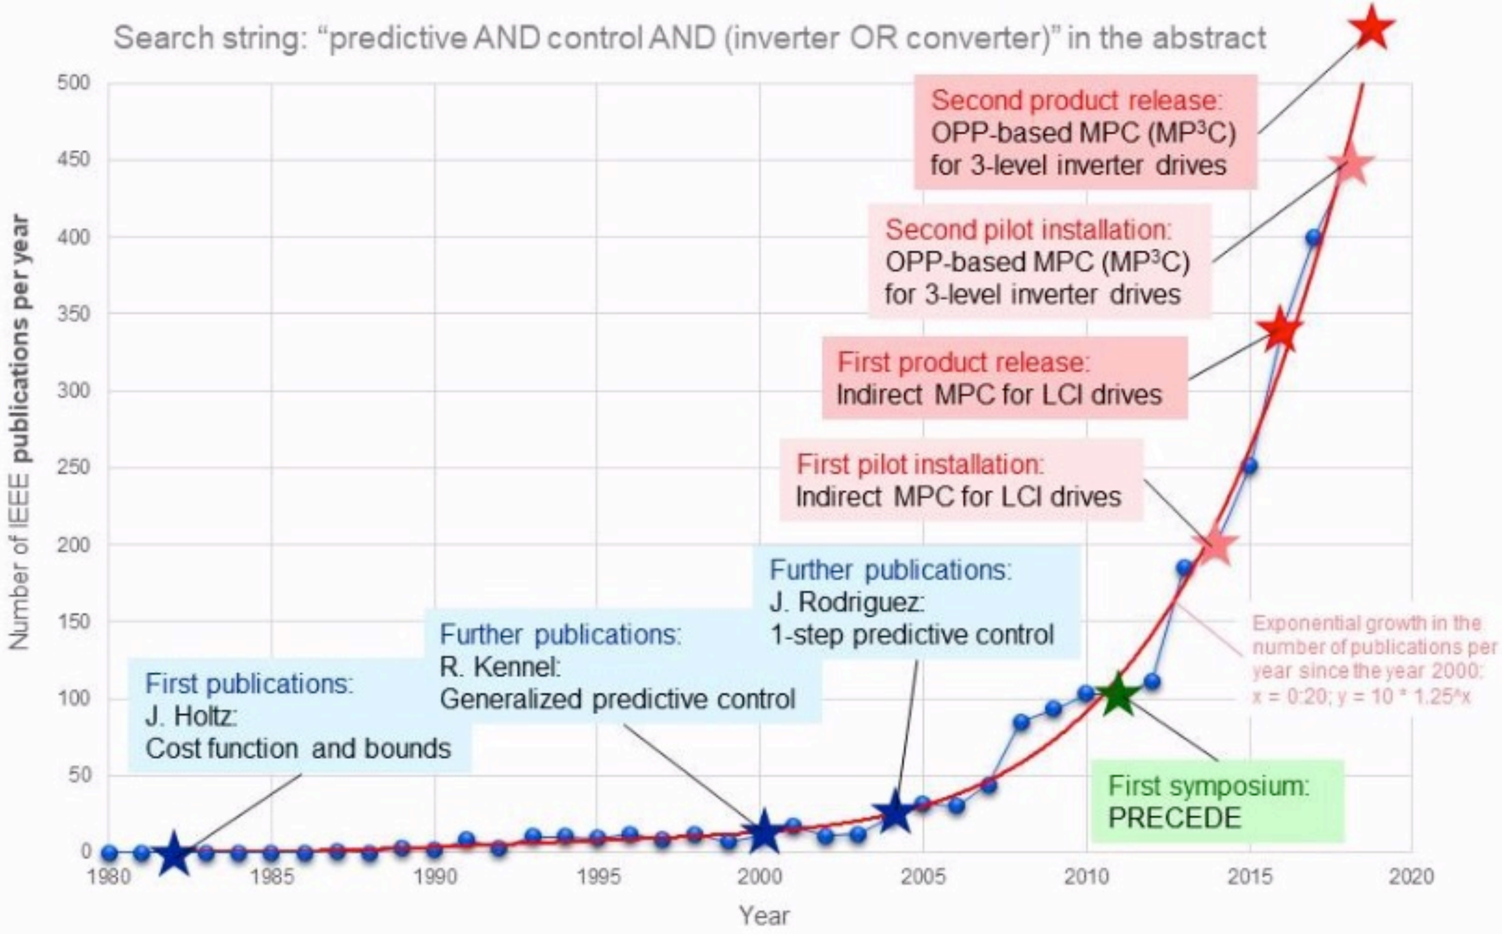

Tobias Geyer
July 2019

The number of annual publications **doubles every three years**

ABB

MPC Outlook

- Robust MPC
- Stochastic MPC
- Hierarchical / decentralized MPC
- MPC with “economic” objective function
- Output feedback MPC

- MPC for nonlinear systems
- Switched / hybrid systems
- Adaptive / Learning MPC

Embracing the Machine Learning and Artificial Intelligence contributions

Learning for Dynamics and Control (L4DC)

May 30 & 31, 2019 at the Ray and Maria Stata Center
Massachusetts Institute of Technology, Cambridge, MA

Over the next decade, the biggest generator of data is expected to be devices which sense and control the physical world.

This explosion of real-time data that is emerging from the physical world requires a rapprochement of areas such as machine learning, control theory, and optimization. While control theory has been firmly rooted in tradition of model-based design, the availability and scale of data (both temporal and spatial) will require rethinking of the foundations of our discipline. From a machine learning perspective, one of the main challenges going forward is to go beyond pattern recognition and address problems in data driven control and optimization of dynamical processes.

A Practitioner's Perspective

Manfred Morari

Mahyar Fazlyab, Alex Robey, Hamed Hassani, George J. Pappas

L4DC - Learning for Dynamics & Control, MIT

May 30, 2019

Penn
UNIVERSITY of PENNSYLVANIA

- **Looking back**

- Looking forward

- Some research directions

Idea: Get rid of “Modeling” …
…in Model-Based-Design

- Kalman (1958): Design of a self-optimizing control system. Trans. ASME
- Bellman (1961): Adaptive Control Processes
- Åström & Wittenmark (1973): On Self-Tuning Regulators. Automatica
- Landau (1974): A survey of model reference adaptive techniques, Automatica
- Narendra & Valavani (1976): Stable adaptive observers and controllers. Proc IEEE
- Åström, Borisson, Ljung, Wittenmark (1977): Theory and applications of self-tuning regulators. Automatica
- …..

ASEA Novatune introduced in 1982…

…and mostly abandoned by 1995

"Even if Novatune in many cases provides very good control, the experience is that the effort it takes, to make it work that well, is discouraging. ***It is worth the effort in some cases, but not as a general tool. What is needed is a tool that is much easier to use.*** You shouldn't be required to set any parameters, except to state what kind of result you desire."

Per Erik Maden (1995) Experiences with Adaptive Control since 1982. CDC Proc.

Why did Adaptive Control "fail"?
--- It was not appropriate

Easy

- no specs
- model simple

- general solution

Why did Adaptive Control “fail”?
--- It was not appropriate

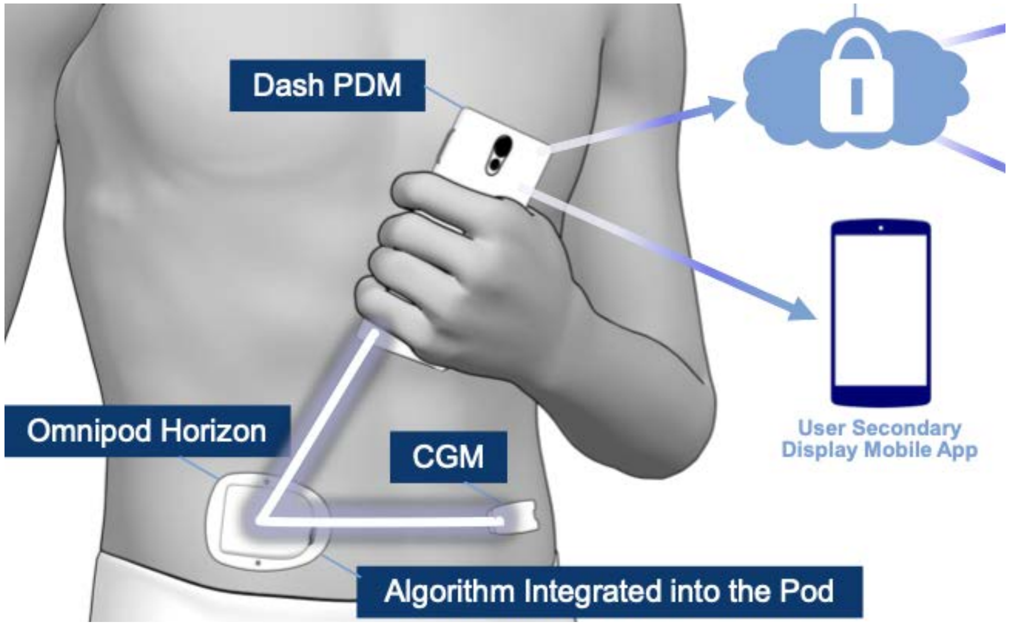

Courtesy: Insulet Corporation

Tough

- tight specs
- model complex
- specific solution

Why did Adaptive Control “fail”?
--- It was not appropriate

Anything works

- no specs
- model simple

- general solution

Nothing works

- tight specs
- model complex

- specific solution

Why did Adaptive Control "fail"?
--- It was not appropriate

Anything works

- no specs
- model simple

- general solution

Learning Control

Nothing works

- tight specs
- model complex

- specific solution

Why did Adaptive Control "fail"?
--- It was too complicated

- PID is **optimal**
  for most simple linear dynamic processes and performance specs
  - Low order dynamics, stable + integrator
- PID is **"practically optimal"** for many more dynamic processes
  - Approximated by low order dynamics, e.g. first order + dead time

Books: Morari & Zafiriou (1989), Åström & Hägglund (1995),
Skogestad & Postlethwaite (2005),

Why did Adaptive Control “fail”?
--- tuning all the time not needed

- Åström & Hägglund (2000). Supervision of adaptive control algorithms
- PID Autotuner: Tune on demand only

Commercial Autotuners

- One-button autotuning
- Three settings: fast, slow, delay dominated
- Automatic generation of gain schedules
- Adaptation of feedback gains
- Adaptation of feedforward gain
- Many versions
  - Single loop controllers
  - DCS systems
- Robust
- Excellent industrial experience
- Large numbers

Thanks: Karl Aström

Explore and Exploit - Tuning

- Essentially no prior knowledge needed
- Automatic generation of test signals specifically to estimate 2-4 parameters affecting the tuning
- Original ideas by Ziegler & Nichols from the 1940s extended by Åström & Hägglund

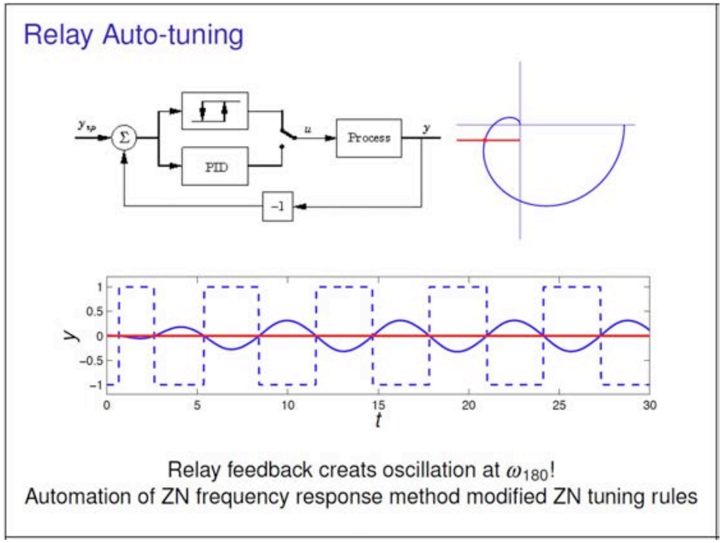

- Looking back

- **Looking forward**

- Some research directions

Learning Controllers

- Model-based vs. model-free
  - If you do not have a model, how can you verify the performance of the closed-loop control system?
  - If you do have a model, why would you use a model-free learning method?
- Policy learning based on reward function
  - Curse of dimensionality
  - Specification guarantees via definition of reward function

Design ≠ Optimization
Design ≈ Constraint Satisfaction

Propositional Logic Control Specifications for Refrigeration Cycle

Manipulated Inputs $u_{i,min} \leq u_i \leq u_{i,max}$ i=1,2,3

Controlled / Monitored Outputs $y_i = y_{i,set}$ i=1,2,3 / $z_i = z_{i,set}$ i=1,..4

Prioritized Objectives

$z_1 > z_{1,min}$	$z_3 < z_{3,max}$	$z_3 < z_{3,max}$	$y_3 = y_{3,set}$
$z_2 < z_{2,max}$	$z_4 < z_{4,max}$	$z_4 < z_{4,max}$	$y_2 = y_{2,set}$
$y_1 = y_{1,set}$	$z_4 > z_{4,min}$	$z_2 > z_{2,min}$	
	$y_2 = y_{2,set}$	$y_3 = y_{3,set}$	

Specification guarantees via definition of reward function?

GARY MARCUS BUSINESS 08.14.19 09:00 AM

WIRED

DEEPMIND'S LOSSES AND THE FUTURE OF ARTIFICIAL INTELLIGENCE

"…In some ways, deep reinforcement learning is a kind of turbocharged memorization; systems that use it are capable of awesome feats, but they have only a shallow understanding of what they are doing. As a consequence, current systems lack flexibility, and thus are unable to compensate if the world changes, sometimes even in tiny ways."

IAS
2019-10-16

Energy-Based Approaches To Representation Learning

Yann LeCun
New York University
Facebook AI Research
http://yann.lecun.com

facebook
Artificial Intelligence Research

Proposed Learning Controllers

- Learn Discrepancy Model
- Design robust MPC

MPC provides

guarantee of closed loop stability and specifications **by design**

Note: simpler alternatives may be preferable

with more transparent architectures

- Looking back
- Looking forward
- **Some research directions**

Some Research Directions

- **MPC Approximation via Neural Networks**
- Robustness Analysis of Learning Enabled Components
- Gaussian-Process based Model Predictive Control

MPC Theory : Properties

1) Recursive feasibility: Input and state constraints are satisfied

2) Stability of the closed-loop system

MPC = Nonlinear control synthesis
with feasibility / stability guarantees by design !!!

Assumption: real-time trajectory optimization problem
is solved to ε-optimality

Neural Network MPC Controllers with Guarantees

Steven Chen[1], Tianyu Wang[2], Nikolay Atanasov[2], Vijay Kumar[1], Manfred Morari[1]

Objective: Learn a **recursively feasible** (*RF*) and **asymptotically stable** (*AS*) MPC control law for large systems

Approach:

Offline (*Learning*):

- Generate a labeled dataset using correlated sampling
- Fit a neural network to optimal primal variables $\boldsymbol{z}$ (trajectory)

Online (*Guarantees*):

- Initialize a primal active set method with the neural network
- Terminate after achieving primal feasibility and suboptimality criteria
- These criteria guarantee *RF* and *AS*

Combining deep learning with traditional optimization provides *guarantees by construction* and *scales to large systems*

Chen et al. Learning Neural Network Model Predictive Control Law with Guarantees. arXiv: 1910.10835

1. University of Pennsylvania 2. UC San Diego

Generating a dataset in high dimensions

We need to sample x ∈ $\boldsymbol{X_0}$ to train and evaluate the neural network

Challenge: Computing $\boldsymbol{X_0}$ for large systems is computationally intractable

$\boldsymbol{X_0}$ is defined by a ***membership oracle***

Naïve Approach: Rejection sampling will not scale (0.4% points feasible)

Our Approach: Utilize ideas in ***geometric random walks*** and ***quasi Monte Carlo***

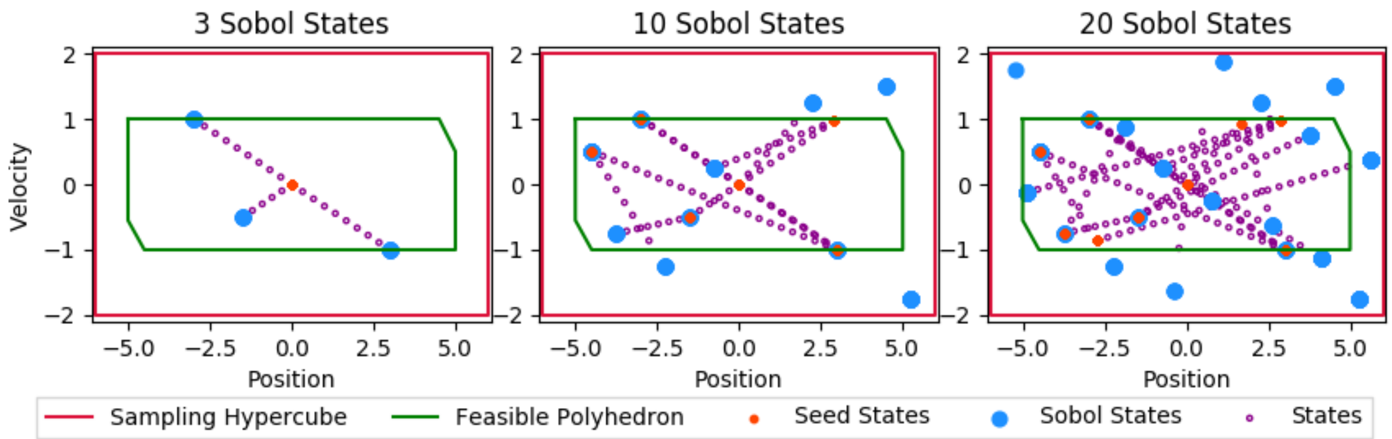

Our approach efficiently queries states $\boldsymbol{x}$ inside $\boldsymbol{X_0}$

Generating a dataset in high dimensions

We need to sample x ∈ $\boldsymbol{X_0}$ to train and evaluate the neural network

Challenge: Computing $\boldsymbol{X_0}$ for large systems is computationally intractable

$\boldsymbol{X_0}$ is defined by a ***membership oracle***

Naïve Approach: Rejection sampling will not scale (0.4% points feasible)

Our Approach: Utilize ideas in ***geometric random walks*** and ***quasi Monte Carlo***

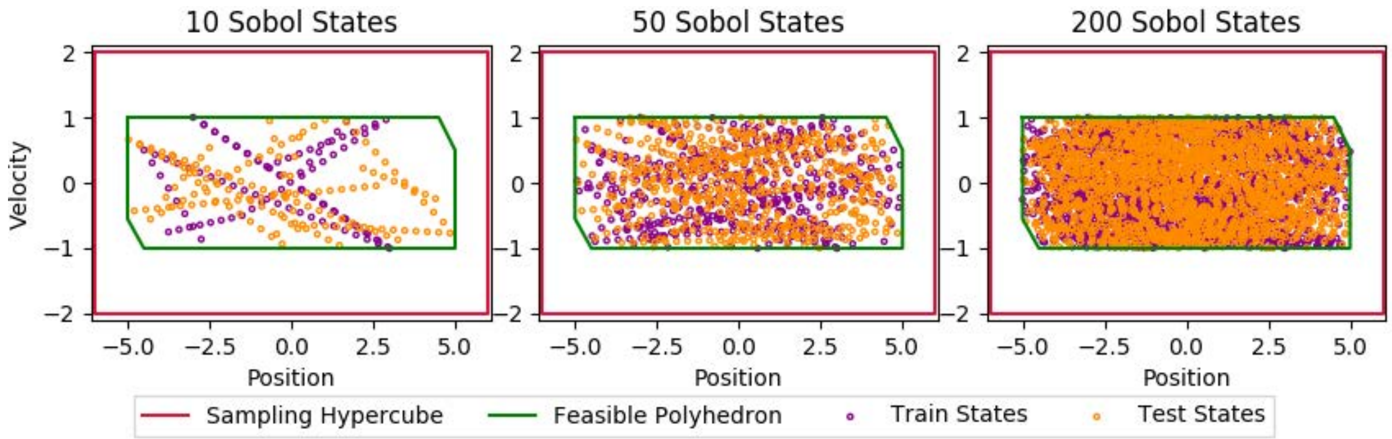

It is designed to fill up the volume of $\boldsymbol{X_0}$

Example: Oscillating Masses

18 oscillating masses [1]
State dim: 36
Action dim: 9
Horizon: 50

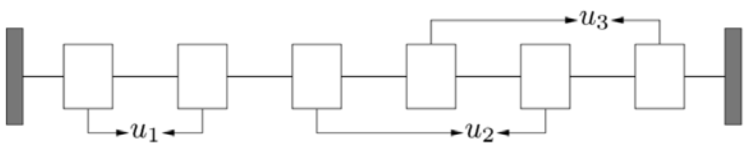

6 mass version

Training Parameter	Value
Training Set Size	2,500,000
Testing Set Size	250,000
# Training Epochs	200 Epochs (~40 hour)
Neural Network Depth	7 layers
Neural Network Hidden Width	128-512
# Neural Network Parameters	1,668,554

[1] Y. Wang and S. Boyd. Fast Model Predictive Control Using Online Optimization. IEEE Transactions on Control Systems Technology. Vol 18, No. 2, 2010.

Example: Oscillating Masses (preliminary results)

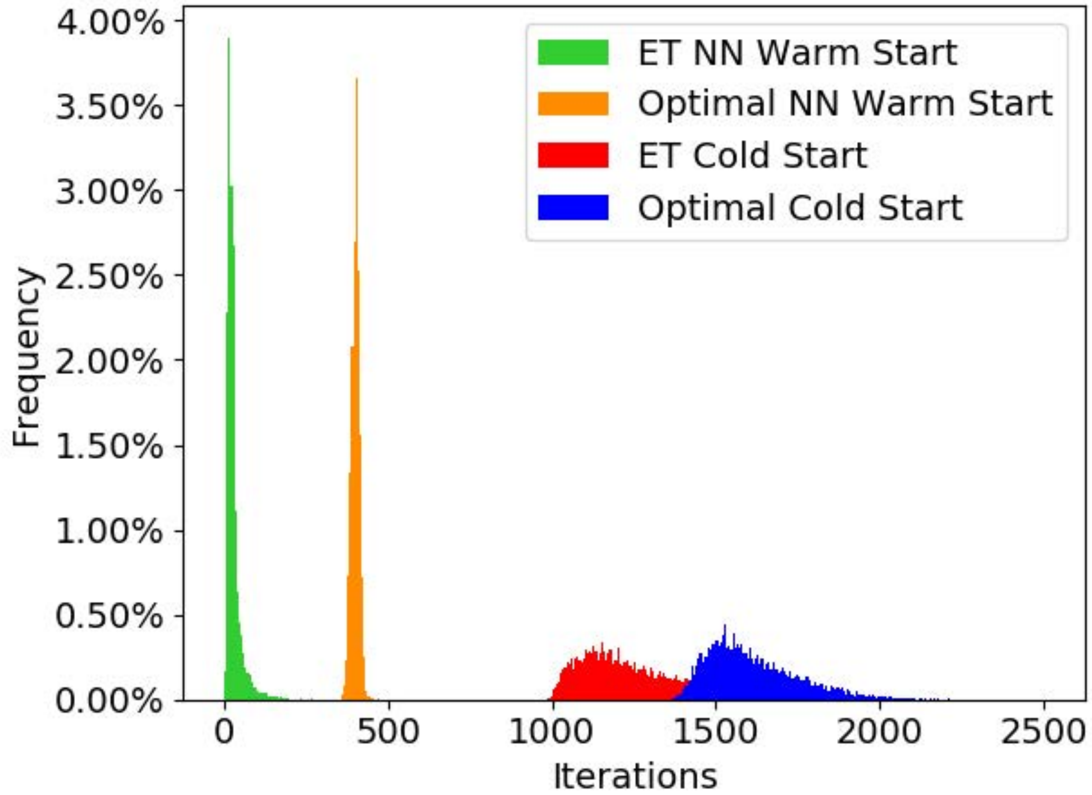

Our method reduces the number of required iterations by 98% and computation time by 90%

Some Research Directions

- MPC Approximation via Neural Networks
- **Robustness Analysis of Learning Enabled Components**
- Gaussian-Process based Model Predictive Control

Robustness Analysis of Learning Enabled Components (LECs)

1. ***Safety Verification*** and Robustness Analysis of Neural Networks via Quadratic Constraints and Semidefinite Programming

   M. Fazlyab, M. Morari, G.J. Pappas arXiv:1903.01287

2. Efficient and Accurate Estimation of ***Lipschitz Constants*** for Deep Neural Networks

   M. Fazlyab, A. Robey, H. Hassani , M. Morari, G.J. Pappas

   NeurIPS – Spotlight (2019) arXiv:1906.04893

Safety Verification of Neural Networks

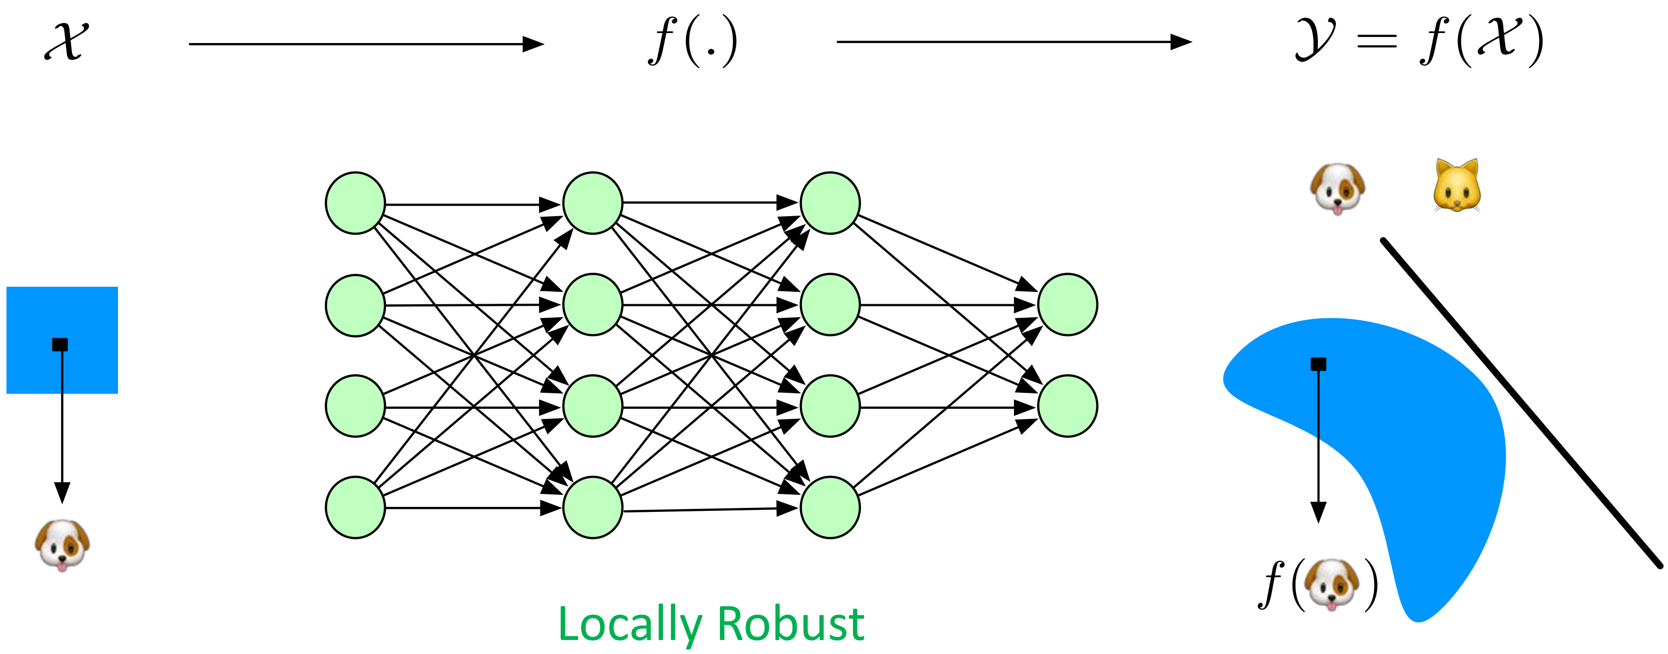

Safety Verification of Neural Networks

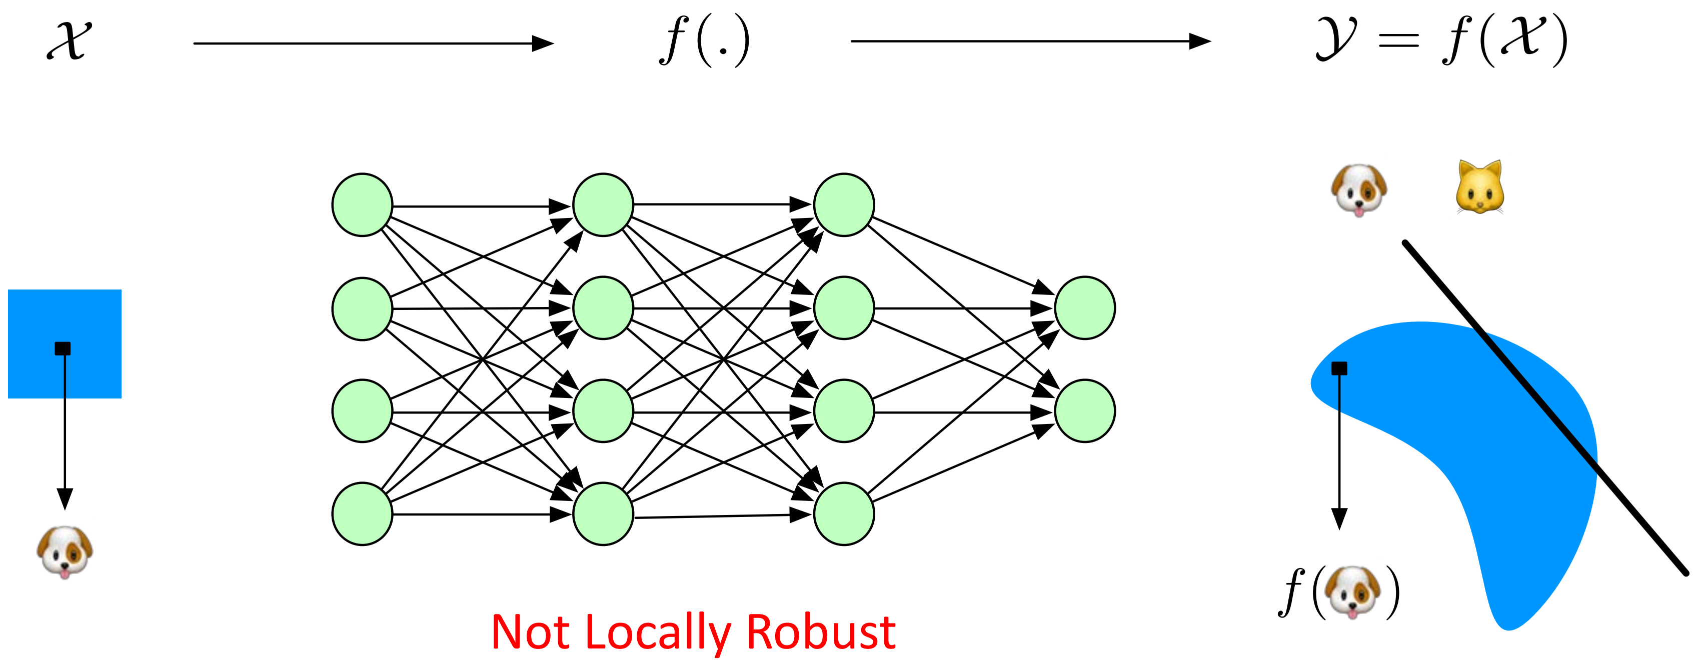

Safety Verification of Neural Networks

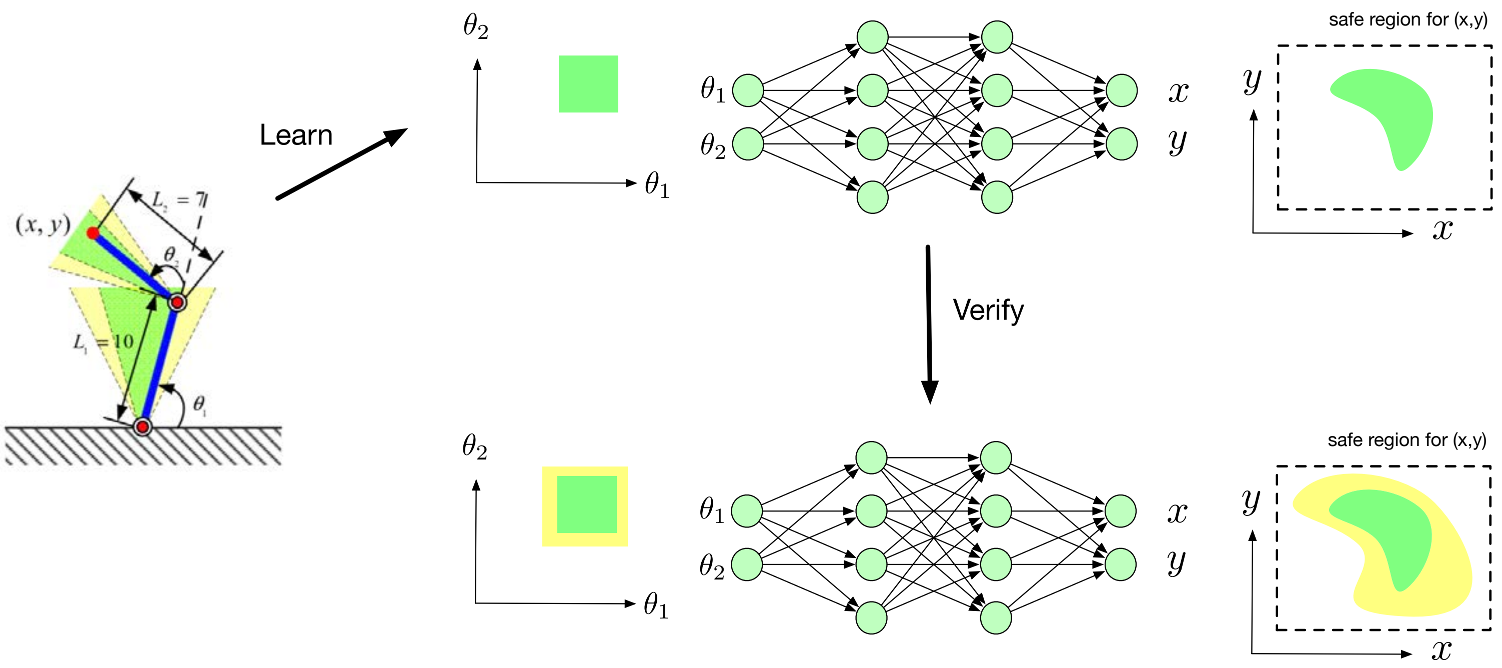

Learning forward kinematics of robotic arms [Xiang et al., 2018]

Safety Verification of Neural Networks

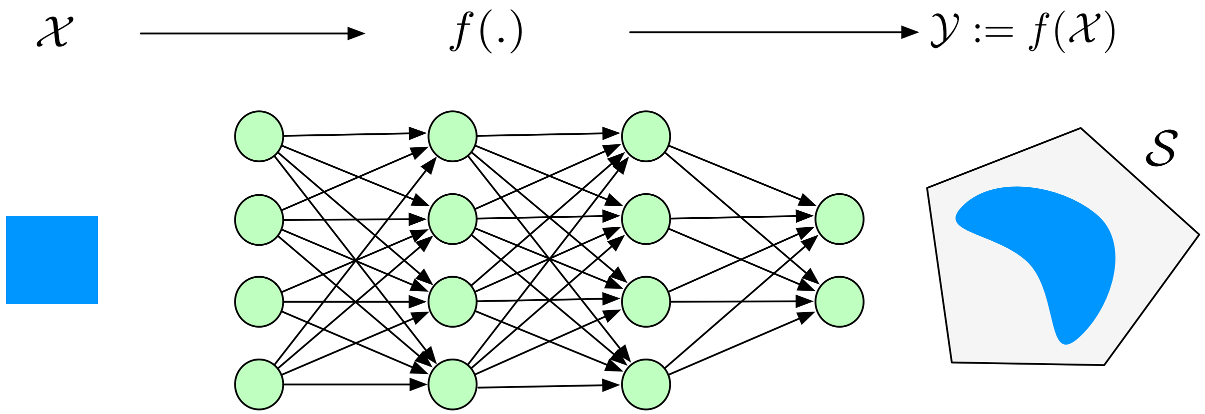

Guarantee $Y \subseteq S$

Safety Verification of Neural Networks

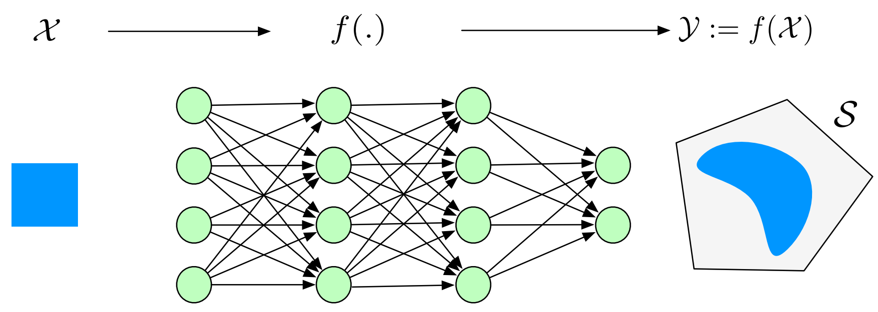

Problem is NP complete

Exact (complete) verifiers
Inexact (incomplete) verifiers
Survey: Liu, Arnon, Lazarus, Barrett, Kochenderfer arXiv:1903.06758

Our work: convex relaxation by adapting tools from robust control

Cf. Raghunathan, Steinhardt, Liang, NeurIPS 2018 [RSL]

$\Rightarrow$Allows direct analysis of NN in closed loop context

DARPA Project: Assured Autonomy Unmanned Underwater Vehicle (UUV)

- Sonar data
- NN to locate pipeline on sea floor
- Steering control loop

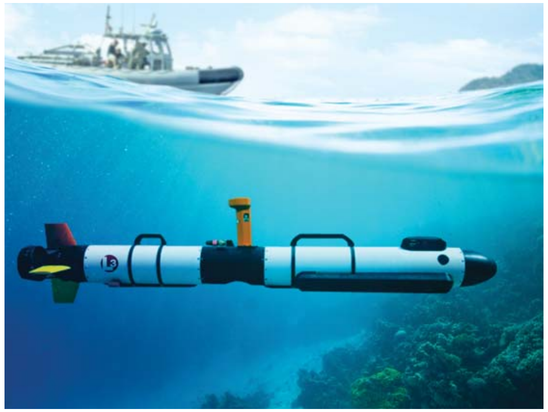

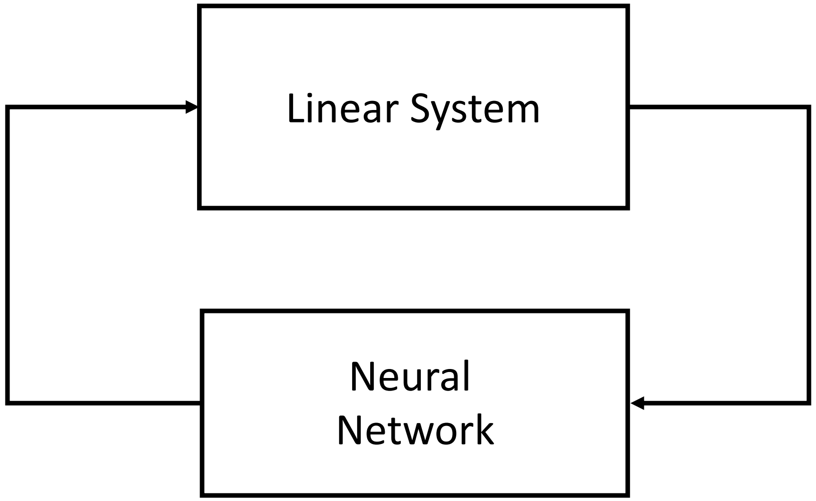

General Interconnection of Linear System and Quadratically-Constrained Nonlinearity

Kothare, Morari, Automatica (1999)

Big Picture of Our Result

- **Key Idea:** replace the activation functions by the "quadratic" constraints they impose on their input-output pairs
  $\hookrightarrow$ Quadratically Constrained Linear Network (**QCLN**)

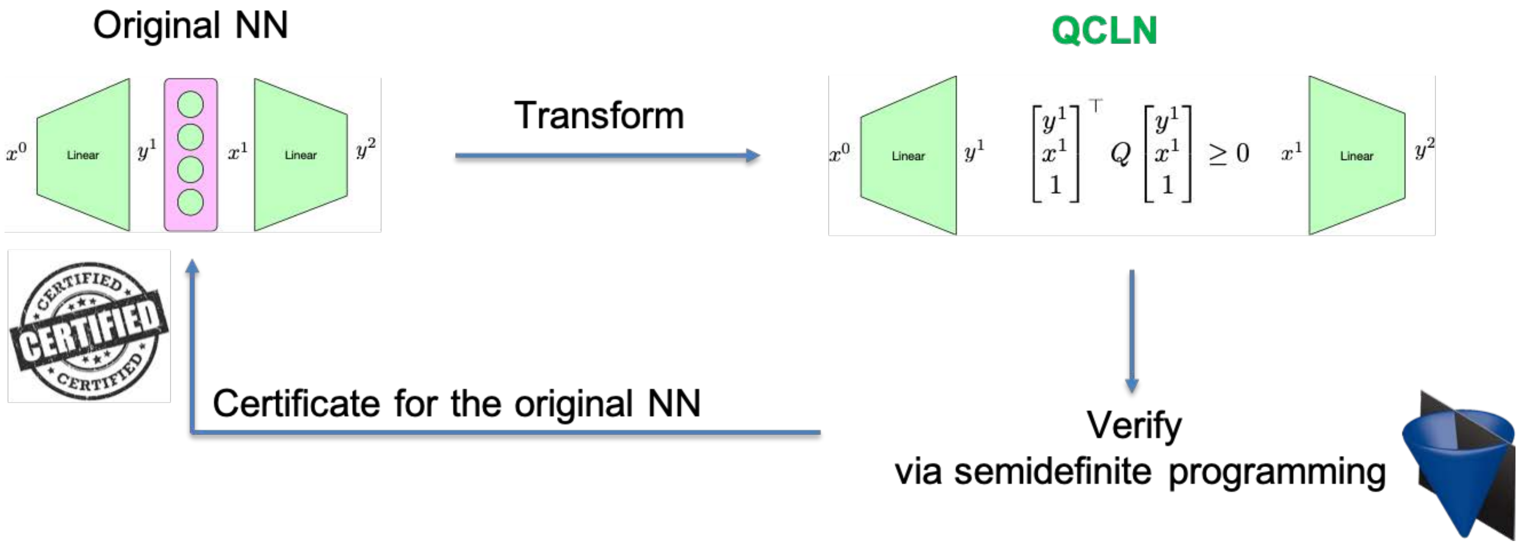

- **Key Insight:** any property (safety, robustness, Lipschitz continuity, etc.) that we can prove for QCLN will hold for the original NN as well

ReLU Function

- Precisely described by 3 constraints

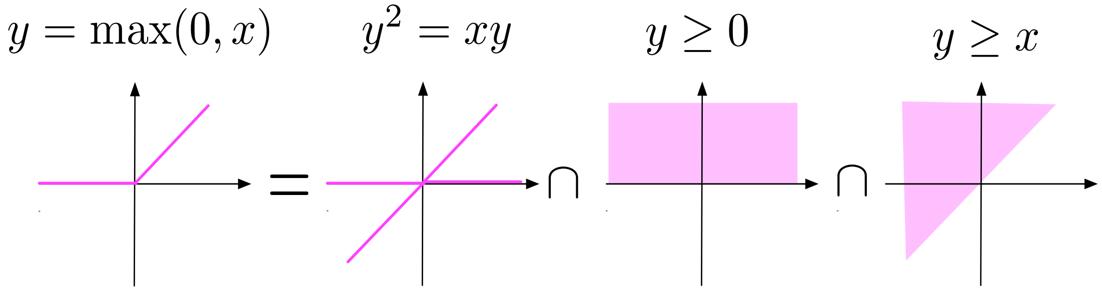

- **Relaxation:** for any $(\lambda, \nu, \eta) \in \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}_+$

$$\lambda(y^2 - xy) + \nu(y - x) + \eta y \geq 0$$

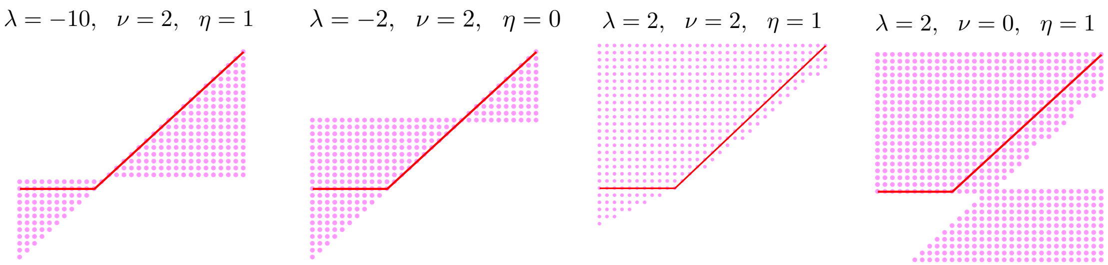

Quadratic Constraints Possible for Other Activation Functions

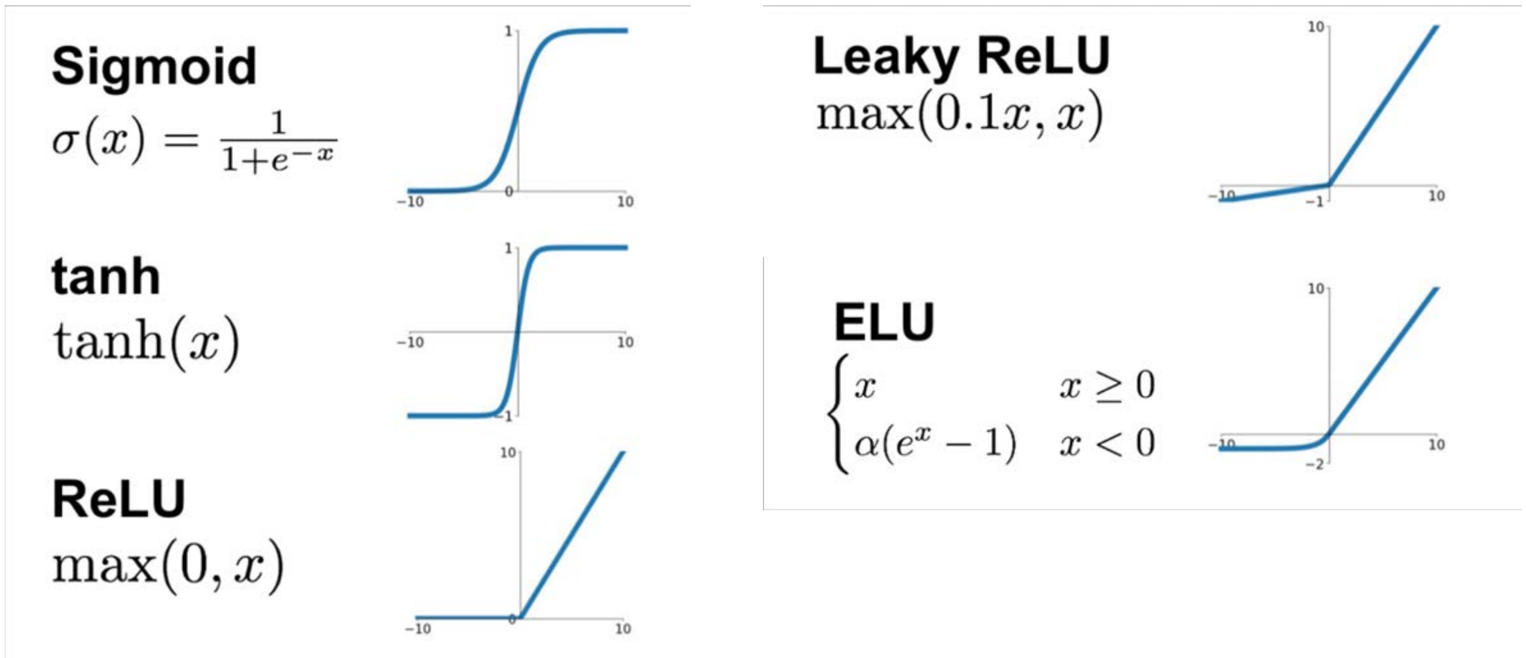

Numerical Experiments

- CVX and Mosek in MATLAB
- 4-core CPU with 16GB of RAM
- Comparison with

  [RSL] Raghunathan, Steinhardt, Liang, NeurIPS 2018

Effect of Number of Layers

- Network: one-layer with architecture (2-100-...-100-2) with $\ell = 1, 2, 4, 6, 8, 10$ layers.
- ReLU activation function
- Input set: ℓ_∞-norm: $\mathcal{X} = \{x \colon \|x - x^\star\|_\infty \le \epsilon\}$ with $x^\star = (0.5, 0.5)$ and $\epsilon = 0.1$.

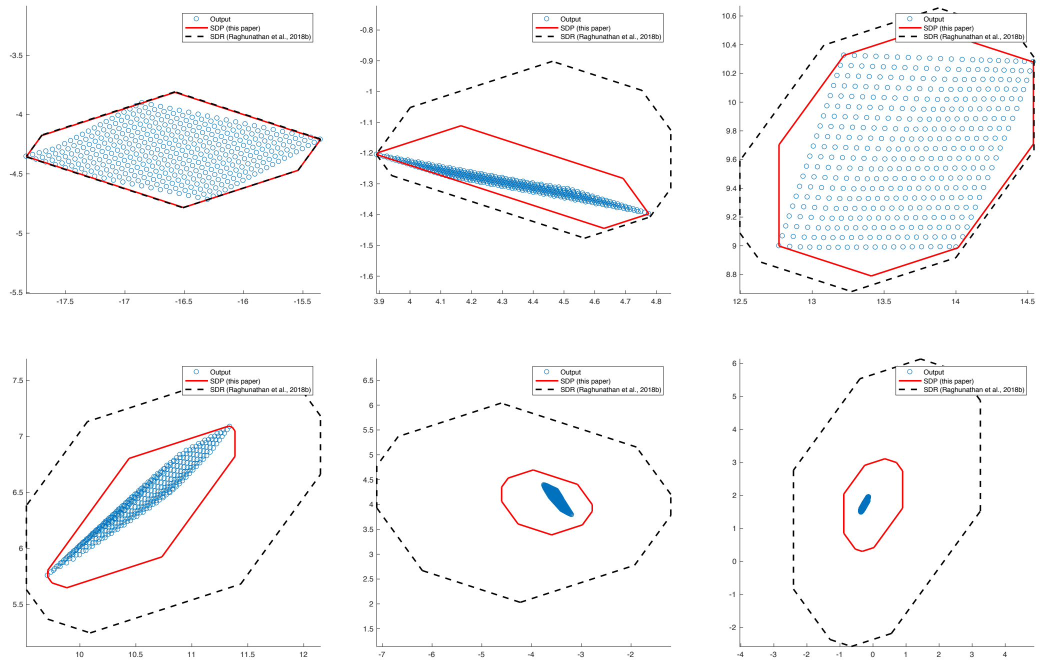

Solve Time

- Comparison between the solve times[2] of the SDP, and the SDR[3] of [Raghunathan et al., 2018b] for a varying number of neurons. The SDR approach runs out of memory (OOM) for networks larger than 1600 neurons.

Number of neurons	Solve time	
	SDP (this paper)	SDR
200	3.2	2.7
400	11.3	20.4
800	78.6	149.1
1200	311.2	799.1
1600	1072.6	OOM
2000	1249.7	OOM
3000	3126.5	OOM

[2] CVX overhead included

[3] Semidefinite Relaxation

Robustness Analysis of Learning Enabled Components (LECs)

Penn
UNIVERSITY of PENNSYLVANIA

1. ***Safety Verification*** and Robustness Analysis of Neural Networks via Quadratic Constraints and Semidefinite Programming

   M. Fazlyab, M. Morari, G.J. Pappas   arXiv:1903.01287

2. Efficient and Accurate Estimation of ***Lipschitz Constants*** for Deep Neural Networks

   M. Fazlyab, A. Robey, H. Hassani , M. Morari, G.J. Pappas

   NeurIPS – Spotlight (2019)   arXiv:1906.04893

Lipschitz Constant for Robustness Analysis of Neural Networks

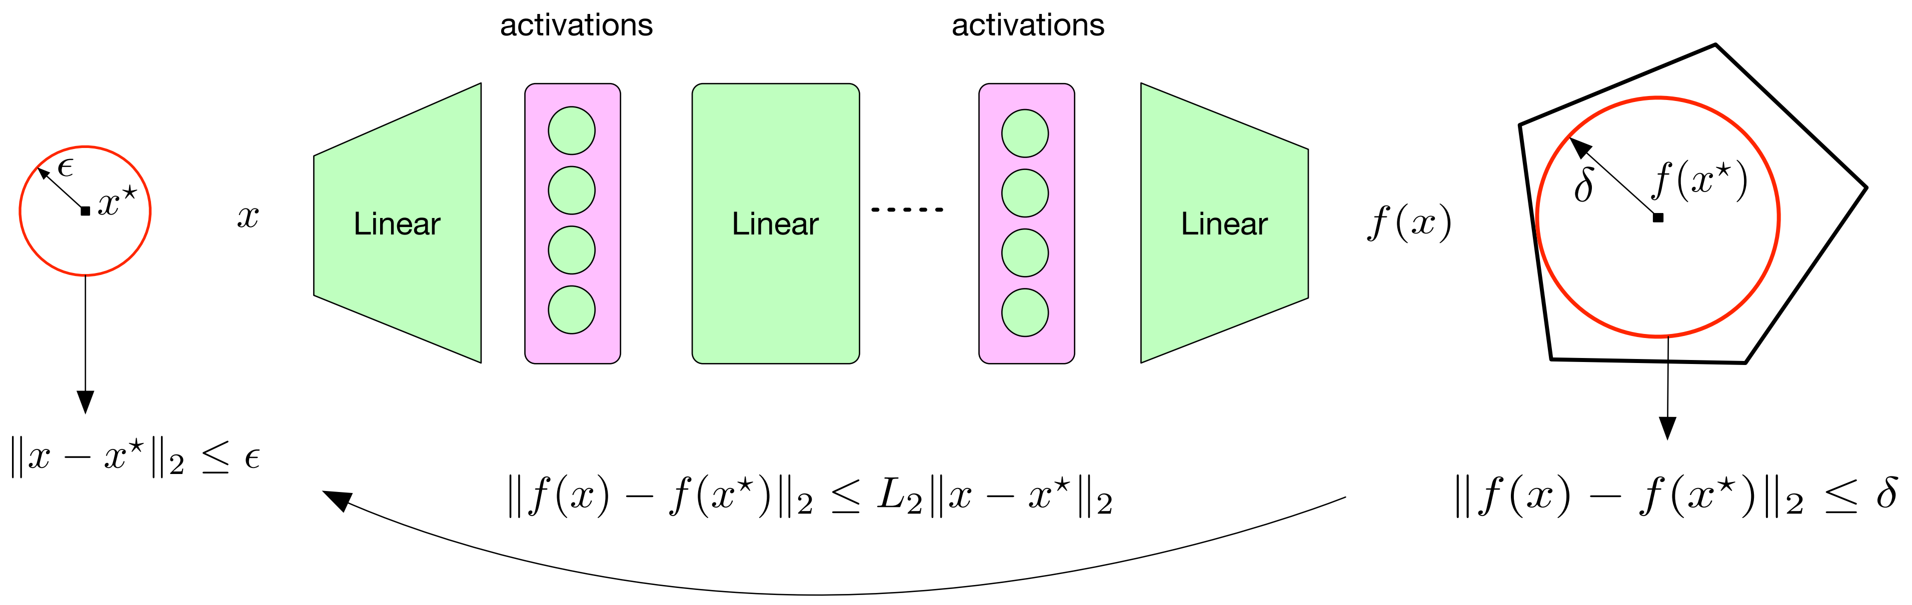

Numerical Experiments

- **Platform:** MATLAB, CVX toolbox, and MOSEK on a 9-core CPU with 16GB of RAM
- **LipSDP:** Lipschitz constant estimation using semidefinite programming
  - **LipSDP-Network**
  - **LipSDP-Neuron**
  - **LipSDP-Layer**
- **CPLip:** Combettes, Pesquet. "Lipschitz Certificates for Neural Network Structures Driven by Averaged Activation Operators." *arXiv:1903.01014*(2019).
- **SeqLip:** Virmaux, Scaman. "Lipschitz regularity of deep neural networks: analysis and efficient estimation." Advances in Neural Information Processing Systems. 2018.

Tightness of Lipschitz bound

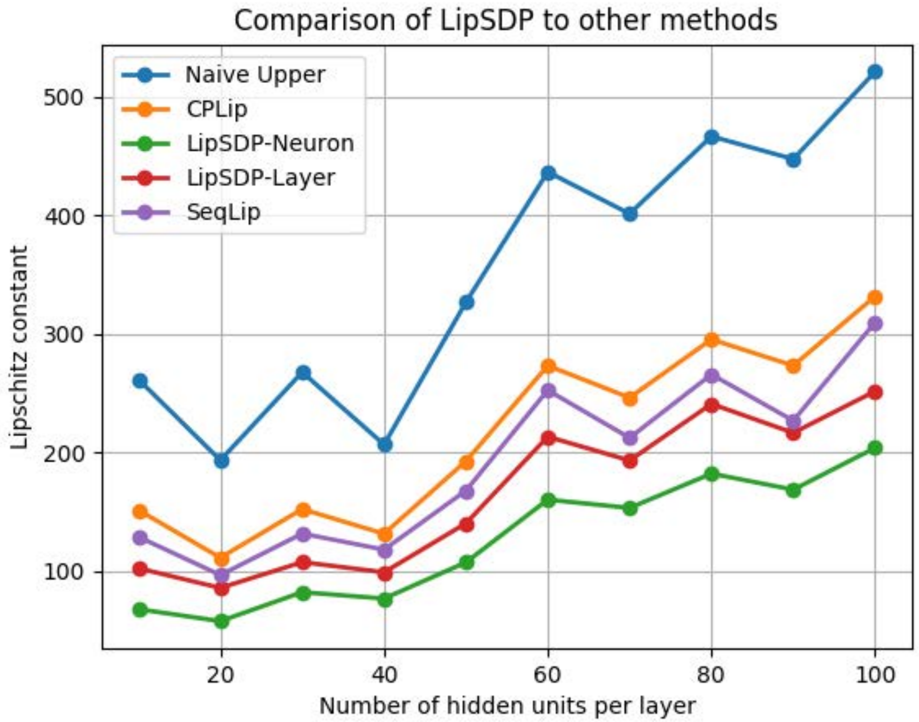

(a) Comparison of the Lipschitz constant found by various formulations for various five hidden-layer neural networks trained on the MNIST dataset with the Adam optimizer. Each network had a test accuracy above 97%.

Numerical Experiments

- **Training Methods:** two robust and one standard training procedures
  - **Adam:** Kingma, Ba. "Adam: A method for stochastic optimization." arXiv:1412.6980 (2014).
  - **LP-Train:** Wong, Kolter. "Provable defenses against adversarial examples via the convex outer adversarial polytope." arXiv:1711.00851 (2017).
  - **PGD-Train:** Madry et al. "Towards deep learning models resistant to adversarial attacks." arXiv:1706.06083 (2017).

Estimation of Input Perturbation ϵ leading to Misclassification

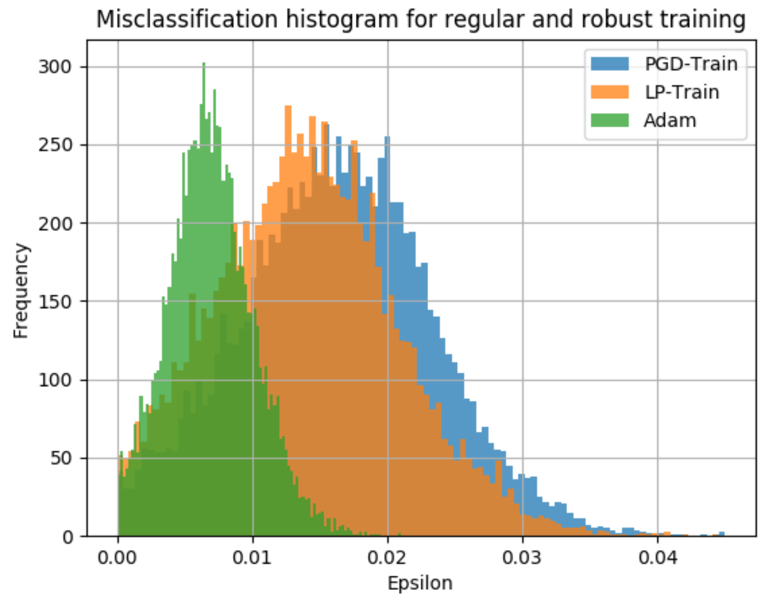

(b) Histograms showing the local robustness (in ℓ_∞ norm) around each correctly-classified test instance from the MNIST dataset. The neural networks had three hidden layers with 100, 50, 20 neurons, respectively. All classifiers had a test accuracy of 97%.

Conclusions: Quadratic Constraints to bound activation functions in NN

- Tight Safety Analysis via SDP
- Analysis of more general properties with NN in closed loop
- Computation costly

- Tightest reported bounds on Lipschitz constant
- Demonstrated Lipschitz constant as effective robustness measure
- Powerful Real time monitoring for possible misclassification via Lipschitz constant
- Computation cheap

Some Research Directions

- MPC Approximation via Neural Networks
- Robustness Analysis of Learning Enabled Components
- **Gaussian-Process based Model Predictive Control**

Gaussian-Process based Model Predictive Control

Prof. Melanie Zeilinger

Institute for Dynamic Systems and Control

ETH Zurich

Challenge: Model Uncertainty
Performance & Safety Require Good Model

$\dot{x} = f(x, u, t, d)$ Modeling challenged by complexity, variability, external disturbances

Goal: Data-driven model improvement

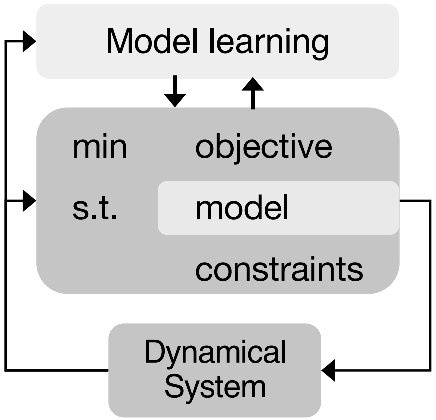

Example: Autonomous racing

- Difficult parameter tuning, in particular of tire models
- Properties of cars/track change over time

ORCA platform @IfA, ETH Zurich. Courtesy of Alex Liniger

Challenge: Model Uncertainty
Performance & Safety Require Good Model

$\dot{x} = f(x, u, t, d)$ ➡ Modeling challenged by complexity, variability, external disturbances

Goal: Data-driven model improvement

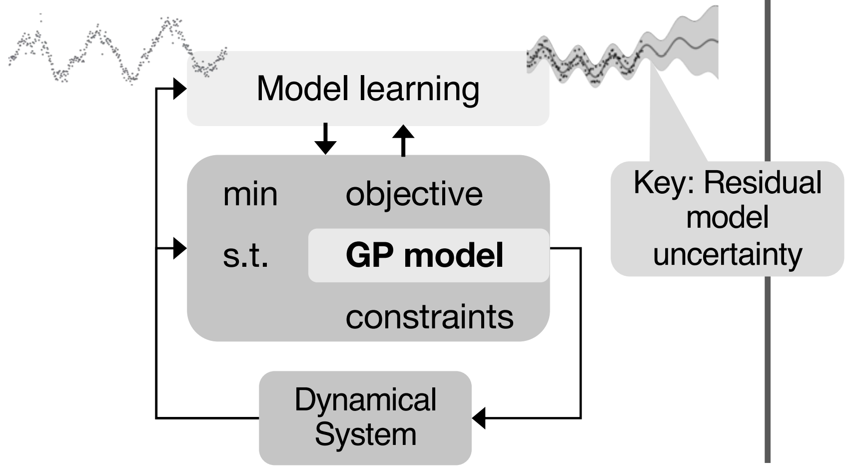

Model learning in MPC

- "Nominal" models
  (e.g. neural networks)
- Robust models
  (e.g. set-membership techniques)
- Stochastic models
  (e.g. Gaussian Processes)

Related work: Kocijan, Findeisen, Ostafew, Schöllig, Koller, Berkenkamp, Deisenroth, Borrelli, ...

Race Car Modeling with Gaussian Processes (GPs)

$$x_{k+1} = f(x_k, u_k) + B_d \, d(x_k, u_k)$$

Bicycle model with nonlinear tire forces

Uncertainty in velocity states

Model mismatch: Tire forces are complex and vary

States: Position, orientation, longitudinal and lateral velocity, yaw rate

Inputs: Motor duty cycle, steering angle

Constraints: Track boundaries, input constraints

GP model:

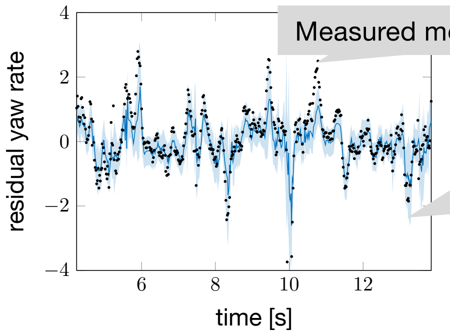

Measured model error

GP prediction with 2σ confidence interval

Bicycle model:

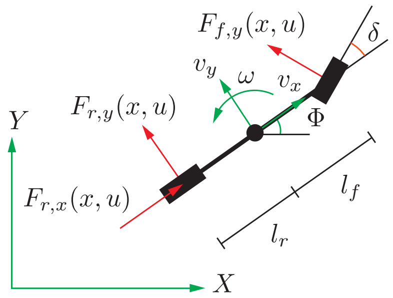

[Hewing & Zeilinger, ECC 2018]

GP-based MPC for Autonomous Miniature Race Cars

Experimental Results with ORCA Platform (@ IfA, ETHZ)

Nominal MPC (20 laps)

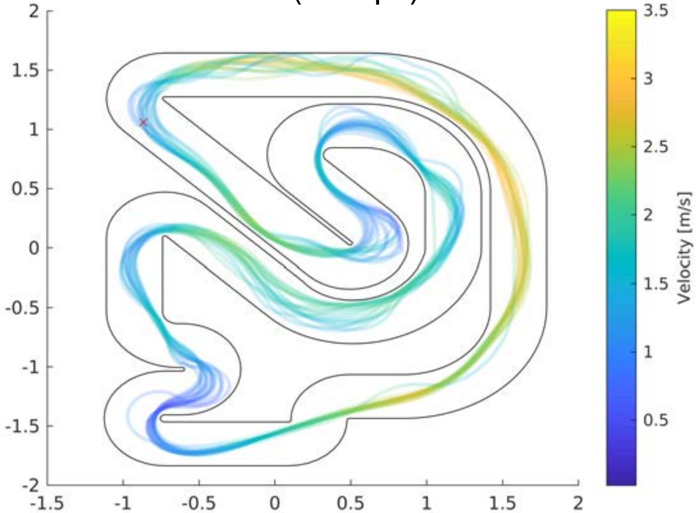

Laptime (mean): **10.1397s**
Solvetime (mean): **17.8 ms**, 96.87% <20ms

GP-based MPC (20 laps)

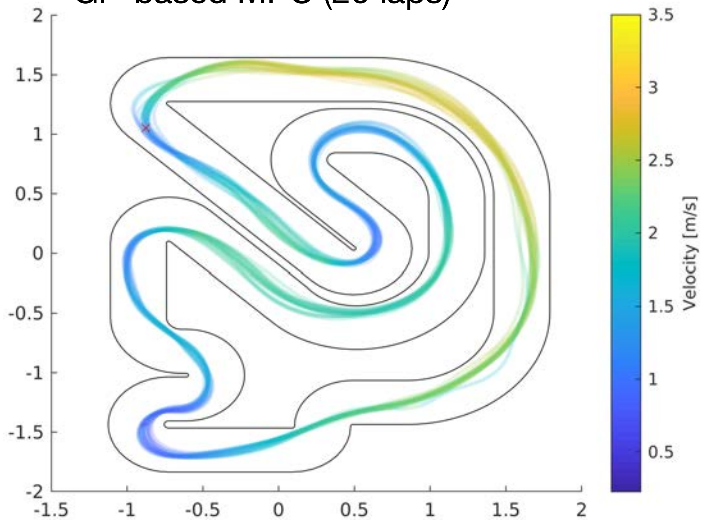

Laptime (mean): **9.6922s**
Solvetime (mean): **17.4 ms**, 96.57% <20ms

Controllers implemented with FORCES Pro: A. Domahidi and J. Jerez, embotech AG

Learning-based Motion Planning for AMZ Driverless Race Car

With Online Learning

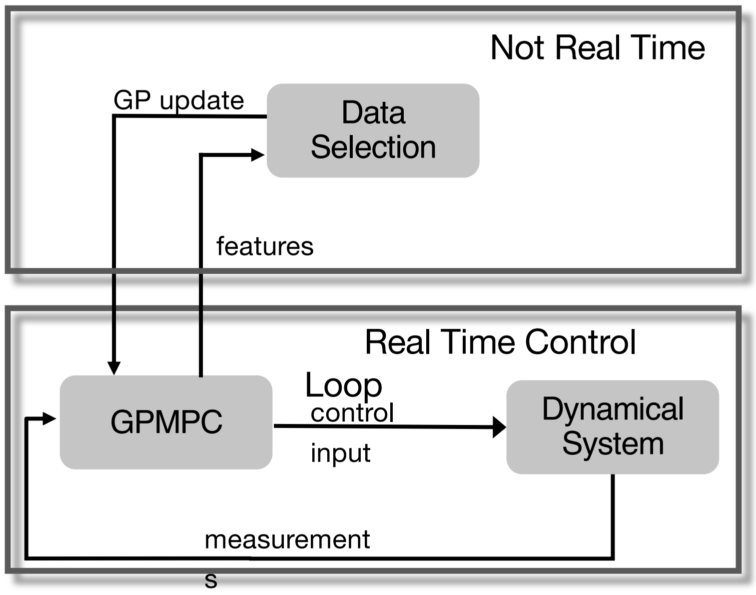

Collaboration with Academic Motorsports Club Zurich (AMZ)

Juraj Kabzan

Lukas Hewing

AMZ Electric Race Car

Learning-based Model Predictive Control for Autonomous Racing

Juraj Kabzan
Lukas Hewing
Alexander Liniger
Prof. Melanie Zeilinger

IDSC

driverless.amzracing.ch

https://ieeexplore.ieee.org/document/8754713

https://www.youtube.com/watch?v=bjlT-6KVQ7U&t

Conclusions

- In 50 years MPC has moved from PhD proposal to become the most widely used advanced high performance control method.
- The reasons for the success:
  - Intuitive concept
  - General applicability
  - Full use of designer (model) information
  - Complex specification guaranteed by design
- The remaining challenges:
  - Computation for stochastic, uncertain and switched systems
  - Learning and adaptation