


Learning to Accelerate Optimizers with Guarantees

MIT REALM Talk 2024
Rajiv Sambharya




Tracking a reference trajectory with a quadcopter



Success!

(If given enough time)



Failure: not enough time to solve

Model predictive control
optimize over a smaller horizon (T steps),
implement first control,
repeat

Model predictive controller



minimize $\sum_{t=1}^T \|x_t - x_t^{\text{ref}}\|_2^2$

subject to $x_{t+1} = Ax_t + Bu_t$

$$x_t \in \mathcal{X}, \quad u_t \in \mathcal{U}$$

$$x_0 = x_{\text{init}}$$

Current state,
reference trajectory



Control
inputs

Challenge: we need faster methods for optimization

Claim: real-world optimization is parametric


Robotics and control



Energy



Signal processing



Can machine learning speed up parametric optimization?

Goal: Do mapping quickly and accurately


Parameter

$\theta \longrightarrow$

minimize $f_{\theta}(z)$
subject to $g_{\theta}(z) \leq 0$

Optimal solution

$\longrightarrow z^*(\theta)$



Learning to Optimize

The learning to optimize paradigm

Goal: solve the parametric optimization problem fast


$$\begin{aligned} &\text{minimize} && f_{\theta}(z) \\ &\text{subject to} && g_{\theta}(z) \leq 0 \end{aligned}$$

Offline

Data collection




Parameters



Solve

Optimal solutions




Training

Training parameter θ

θ

Learnable Optimizer with weights w



Learn

Candidate solution

$\hat{z}_w(\theta)$


Loss

Deploy

Online evaluation

Unseen parameter θ

Learned Optimizer



High-quality solution

Challenges in learning to optimize methods

- I: Lack convergence guarantees
- II: Lack generalization guarantees
- III: Hard to integrate with state-of-the-art solvers

We need **reliable** L2O methods

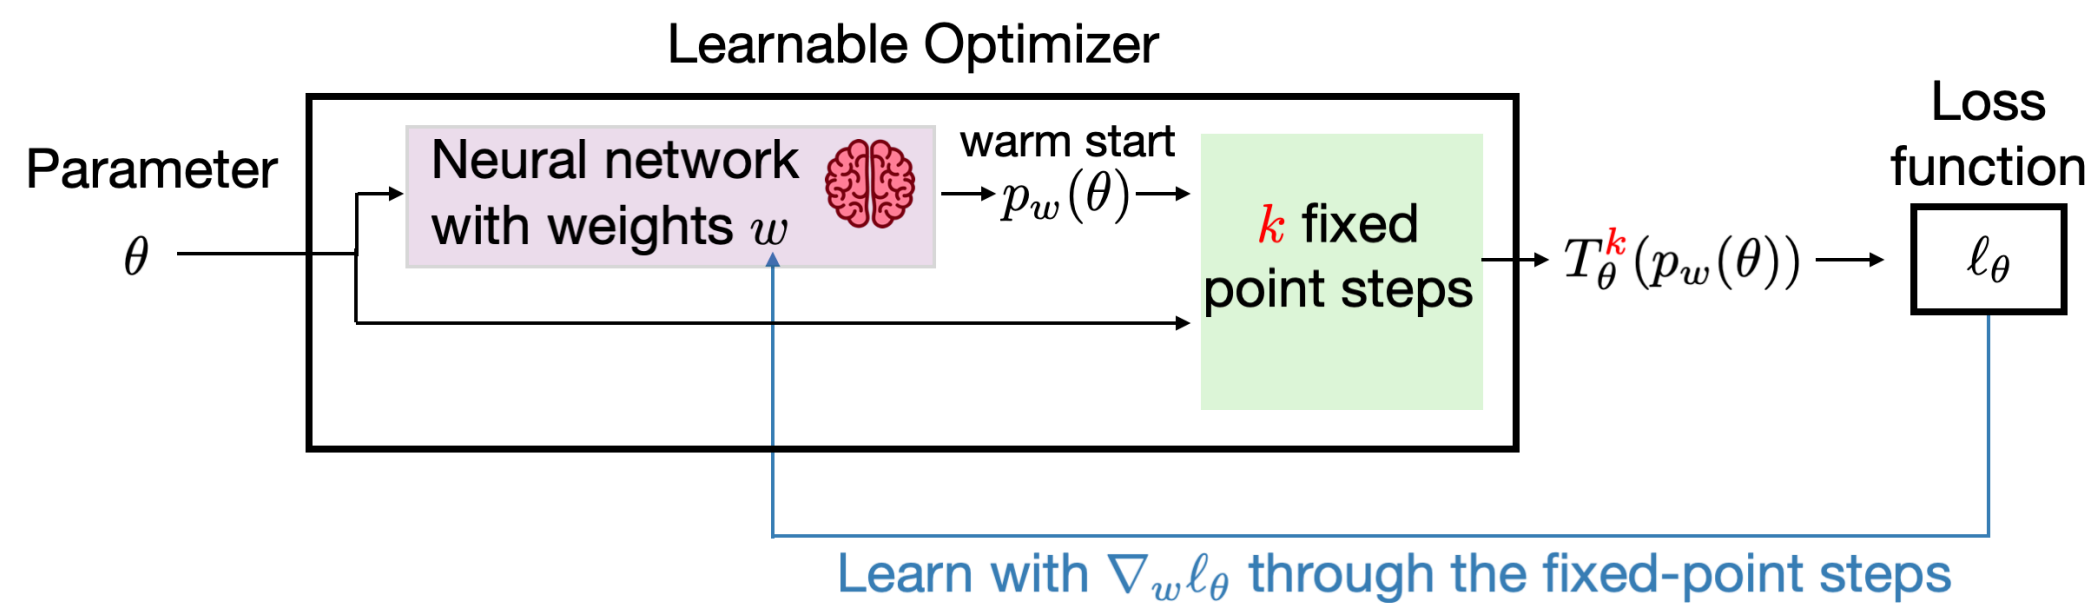


Learning to Optimize: A Primer and A Benchmark [Chen. et al 2021]


“So, to conclude this article, let us quote Sir Winston Churchill: ‘Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning.’”

Talk Outline

- **Part 1: Learning to Warm-Start Fixed-Point Optimization Algorithms**



- **Part 2: Practical Performance Guarantees for Classical and Learned Optimizers**



Collaborators



Georgina
Hall



Brandon
Amos




Bartolomeo
Stellato




Talk Outline

- **Part 1: Learning to Warm-Start Fixed-Point Optimization Algorithms**



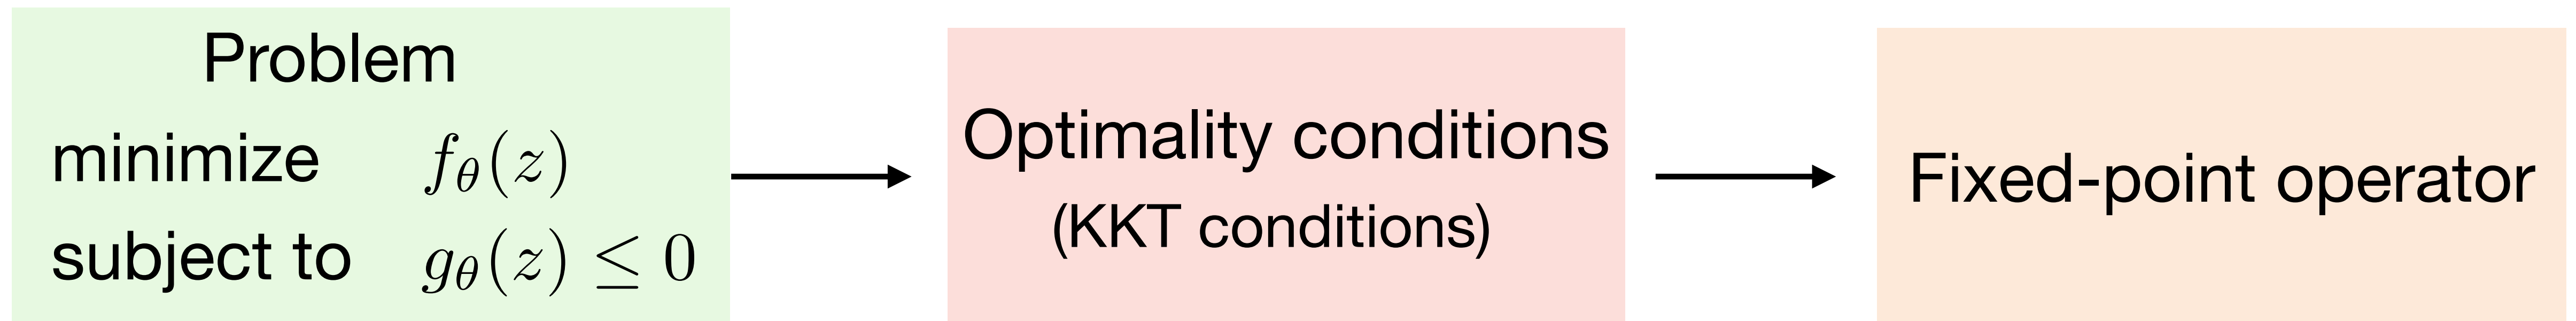
- **Part 2: Practical Performance Guarantees for Classical and Learned Optimizers**




Fixed-point optimization problems are ubiquitous

Parametric fixed-point problem: find z such that $z = T_{\theta}(z)$

Convex optimization




Unconstrained, smooth convex optimization



Many optimization algorithms are fixed-point iterations

Fixed-point iterations: $z^{i+1} = T_\theta(z^i)$



Initialize with z^0 (a warm-start)

Terminate when $\|T_\theta(z^j) - z^j\|_2$ is small

Fixed-point residual

Example: Proximal gradient descent

$$\text{minimize } \underbrace{g_\theta(z)}_{\substack{\text{Convex} \\ \text{Smooth}}} + \underbrace{h_\theta(z)}_{\substack{\text{Convex} \\ \text{Non-smooth}}}$$

Iterates $z^{i+1} = \text{prox}_{\alpha h_\theta}(z^i - \alpha \nabla g_\theta(z^i))$

$$\text{prox}_s(v) = \arg \min_x \left(s(x) + \frac{1}{2} \|x - v\|_2^2 \right)$$



Problem: limited iteration budget



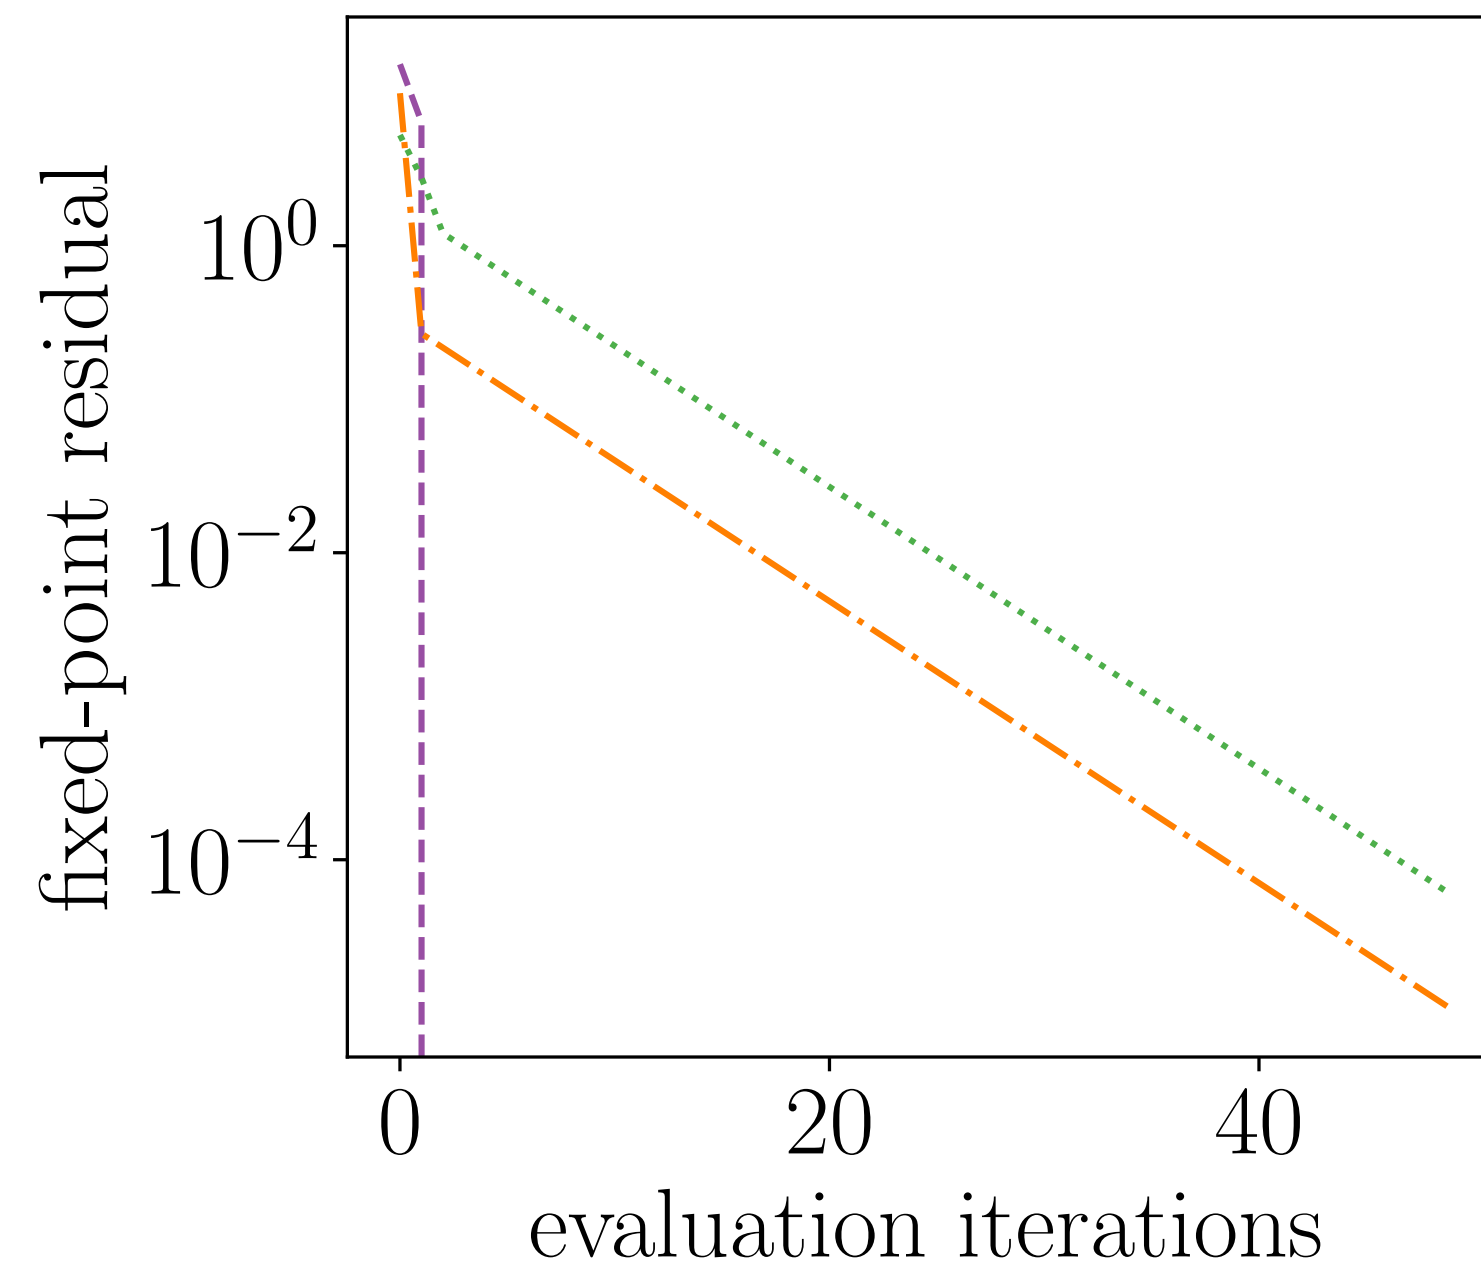

Solution: learn the warm-start to improve the solution within budget

Some warm starts are better than others

$$\begin{aligned} &\text{minimize} && 10z_1^2 + z_2^2 \\ &\text{subject to} && z \geq 0 \end{aligned}$$

★ Optimal solution at the origin

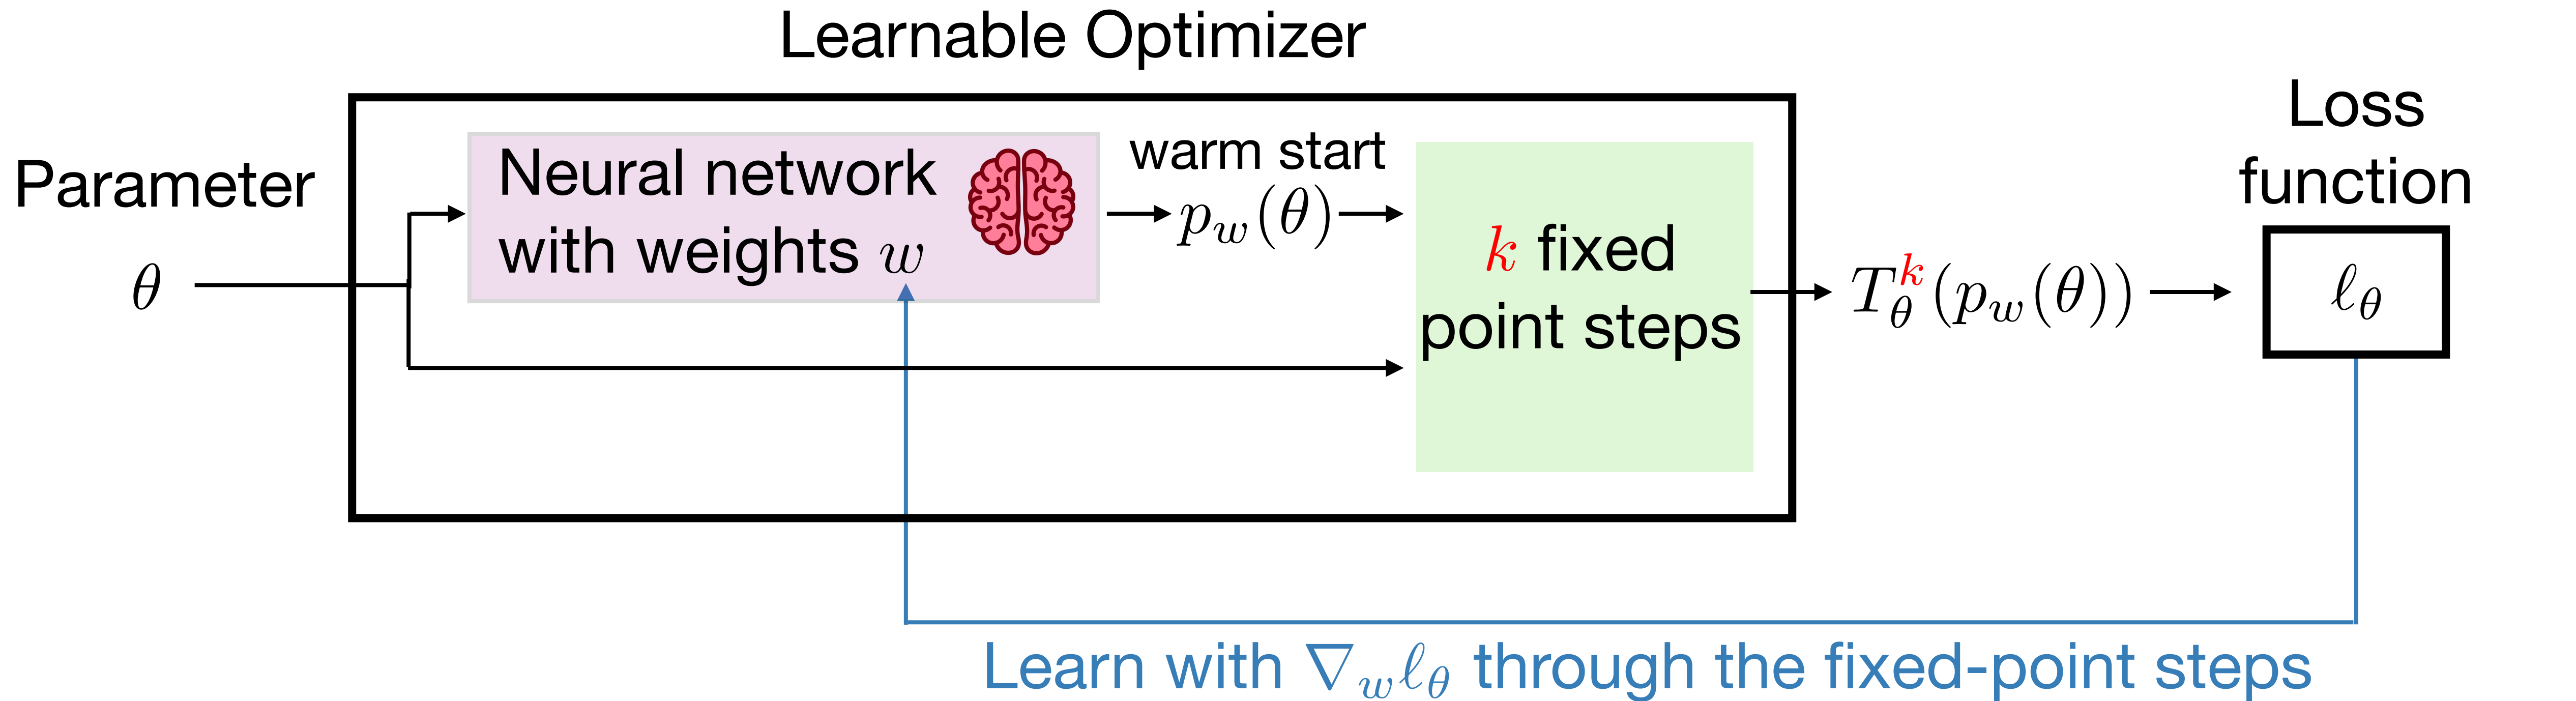
Run proximal gradient descent to solve



All three warm starts appear to be equally suboptimal but converge at very different rates

The quality of the warm start depends on the algorithm

End-to-end learning architecture




Loss function: $l_\theta(z) = \|z - z^*(\theta)\|_2$ Ground truth solution

Learned warm start tailored for downstream algorithm

Benefits of our learning framework


End-to-end learning: warm-start predictions tailored to downstream algorithm




Guaranteed convergence



Learned solver with convergence



Generalization guarantees



- I. Guarantees from **k** training steps to **t** evaluation steps
- II. Guarantees to unseen data

Easy integration with popular solvers



$$\begin{aligned} &\text{minimize} && (1/2)x^T P x + c^T x \\ &\text{subject to} && Ax + s = b \\ &&& s \in \mathcal{K} \end{aligned}$$

Conic programs


```
sol = scs_solver.solve(warm_start=True,
                      x=x0, y=y0, s=s0)
```


Allows us to quantify solve time in seconds

Numerical Experiments


Comparing our learned warm starts  against

Baseline initializations


1. Cold-start: initialize at zero 
2. Nearest neighbor: initialize with solution of nearest training problem



Robust Kalman filtering



Robust Kalman filtering



Second-order cone program

$\theta = \{y_t\}_{t=0}^{T-1}$
Noisy trajectory

minimize $\sum_{t=0}^{T-1} \|w_t\|_2^2 + \mu\psi_\rho(v_t)$
subject to $x_{t+1} = Ax_t + Bw_t \quad \forall t$
 $y_t = Cx_t + v_t \quad \forall t$


$\{x_t^*, w_t^*, v_t^*\}_{t=0}^{T-1}$
Recovered trajectory

Dynamics matrices: A, B


Observation matrix: C

Huber loss: ψ_ρ






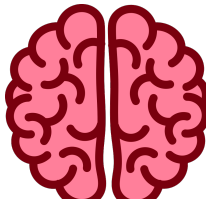
Robust Kalman filtering visuals



-  Noisy trajectory
-  Optimal solution



Solution after 5 fixed-point steps with different initializations

-  Nearest neighbor 
-  Previous solution 
-  Learned: $k = 5$ 

With learning, we can estimate the state well

Model predictive control (MPC) of a quadcopter

Current state,
previous control
reference trajectory



Controller



Control
inputs

$$\theta = (x_{\text{init}}, u_{\text{prev}}, \{x_t^{\text{ref}}\}_{t=1}^T)$$



Quadratic program

minimize $\sum_{t=1}^T (x_t - x_t^{\text{ref}})^T Q (x_t - x_t^{\text{ref}}) + \sum_{t=0}^{T-1} u_t^T R u_t$

subject to $x_{t+1} = A(\theta)x_t + B(\theta)u_t$

$$u_{\min} \leq u_t \leq u_{\max}$$
$$x_{\min} \leq x_t \leq x_{\max}$$
$$|u_{t+1} - u_t| \leq \Delta u$$
$$x_0 = x_{\text{init}}$$
$$u_{-1} = u_{\text{prev}}$$




$$\{x_t^*, u_t^*\}_{t=0}^T$$

Linearized dynamics

MPC of a quadcopter in a closed loop


Budget of 15 fixed-point steps






Nearest neighbor



Previous solution



Learned: $k = 5$



With learning, we can track the trajectory well

Image deblurring





Image deblurring



$\theta = b$
Blurred image




Quadratic program
minimize $\|Ax - b\|_2^2 + \lambda \|x\|_1$
subject to $0 \leq x \leq 1$



x^*
Deblurred image

A : blur operator

Image deblurring





percentile optimal blurred cold-start nearest neighbor learned

50 fixed-point steps



Distance to nearest neighbor increases








10th






50th



90th




99th




With learning, we can deblur all of the images quickly

Talk Outline

- **Part 1: Learning to Warm-Start Fixed-Point Optimization Algorithms**

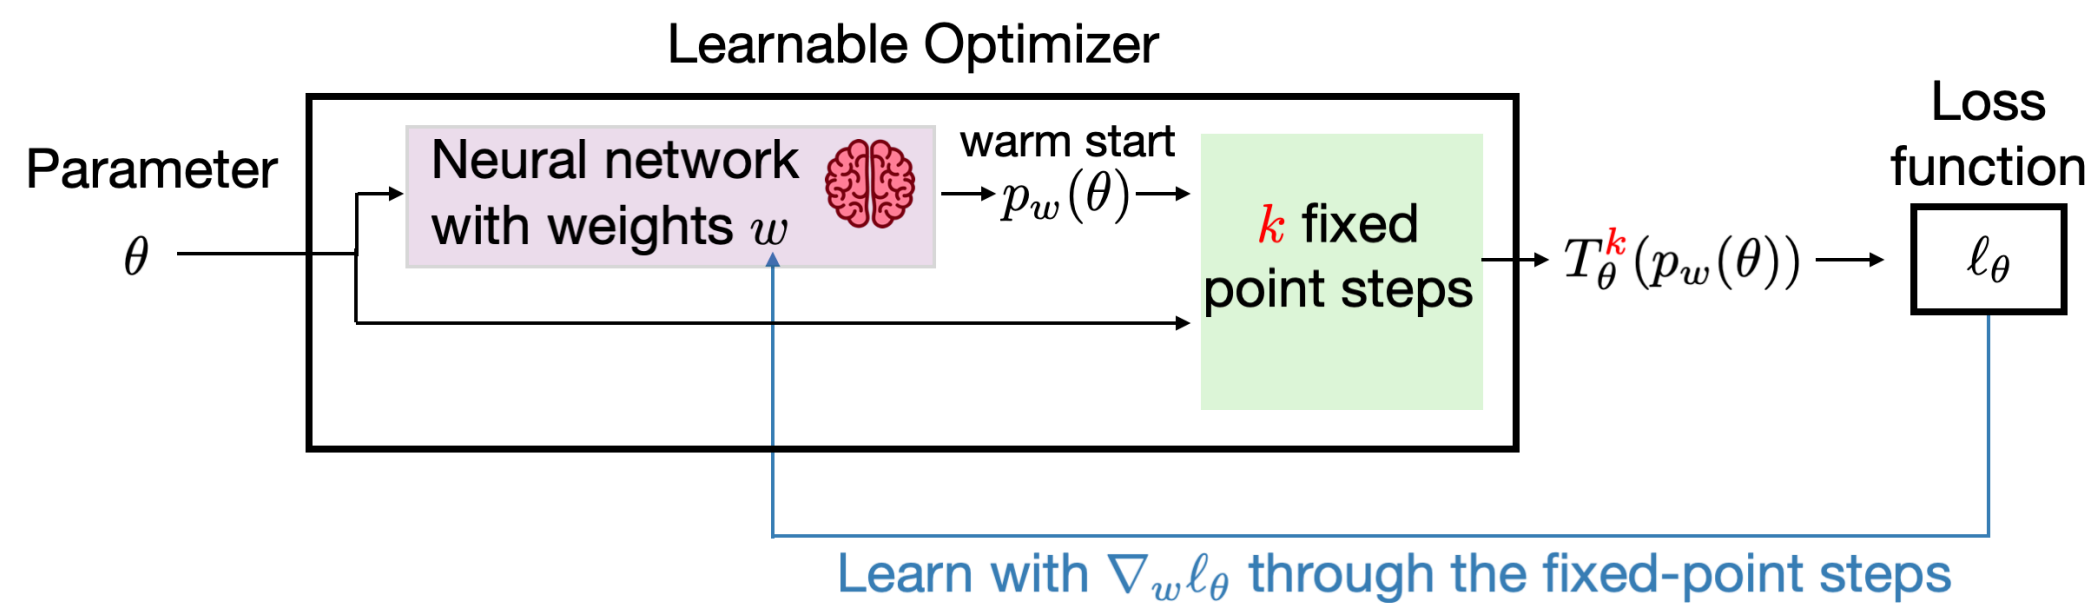


- **Part 2: Practical Performance Guarantees for Classical and Learned Optimizers**




Talk Outline

- **Part 1: Learning to Warm-Start Fixed-Point Optimization Algorithms**




- **Part 2: Practical Performance Guarantees for **Classical** and Learned Optimizers**

Classical = no learning



Worst-case bounds can be very loose



Example: robust Kalman filtering

Second-order cone program

$$\begin{aligned} &\text{minimize} && \sum_{t=0}^{T-1} \|w_t\|_2^2 + \mu \psi_\rho(v_t) \\ &\text{subject to} && x_{t+1} = Ax_t + Bw_t \quad \forall t \\ & && y_t = Cx_t + v_t \quad \forall t \end{aligned}$$

▼ SCS empirical average performance over 1000 parametric problems

◀ Worst-case bound

In practice: **linear** convergence over the parametric family

Worst-case analysis: **sublinear** convergence

Worst-case bounds do not consider the **parametric** structure

Approach: solve N problems and then bound

We will bound 0-1 error metrics

We will provide guarantees for
any measured quantity

algorithm steps tolerance

$$e(\theta) = \mathbf{1}(\ell^k(\theta) > \epsilon)$$

Standard metrics

e.g., fixed-point residual

algorithm steps cold start tolerance


$$e(\theta) = \mathbf{1}(\ell_{\theta}^{\text{fp}}(T_{\theta}^k(\mathbf{0}))) > \epsilon$$

Task-specific metrics:

e.g., quality of extracted states
in robust Kalman filtering

recovered state optimal state

$$e(\theta) = \mathbf{1} \left(\max_{t=1, \dots, T} \|x_t - x_t^*\|_2 > \epsilon \right)$$



Background: Kullback-Liebler Divergence

KL divergence: measures distance between distributions

$$\text{KL}(q \parallel p) = \sum_{i=1}^m q_i \log \left(\frac{q_i}{p_i} \right)$$

Our bounds on the risk will take the form

$$\text{KL}(\text{empirical risk} \parallel \text{risk}) \leq \text{regularizer}$$

Invert these bounds by solving

$$\text{risk} \leq \text{KL}^{-1}(\text{empirical risk} \mid \text{regularizer})$$

1D convex optimization problem

$$\begin{aligned} \text{KL}^{-1}(q \mid c) = & \text{maximize } p \\ & \text{subject to } q \log \frac{q}{p} + (1 - q) \log \frac{1-q}{1-p} \leq c \\ & 0 \leq p \leq 1 \end{aligned}$$

Statistical learning theory can provide probabilistic guarantees

algorithm steps

tolerance

$$e(\theta) = \mathbf{1}(\ell^k(\theta) > \epsilon)$$

Sample convergence bound: with probability $1 - \delta$ [Langford et. al 2001]


$$\mathbf{E}_{\theta \sim \mathcal{X}} e(\theta) \leq \text{KL}^{-1} \left(\frac{1}{N} \sum_{i=1}^N e(\theta_i) \mid \frac{\log(2/\delta)}{N} \right)$$

Number of problems

$$\mathbf{P}(\ell^k(\theta) > \epsilon) = \text{risk} \leq \text{KL}^{-1} (\text{empirical risk} \mid \text{regularizer})$$


”With probability $1 - \delta$, 90% of the time the fixed-point residual is below $\epsilon = 0.01$ after $k = 20$ steps”

Robust Kalman filtering guarantees







With 1000 samples, we provide strong probabilistic guarantees on the 99th quantile

Visualizing Robust Kalman filtering guarantees



Task-specific error metric


$$e(\theta) = \mathbf{1} \left(\max_{t=1, \dots, T} \|x_t - x_t^*\|_2 > \epsilon \right)$$

-  Noisy trajectory
-  Optimal solution
-  Solution after 15 steps
-  Region with guarantee


“With high probability, 90% of the time, all of the recovered states after 15 steps of problems drawn from the distribution will be within the correct ball with radius 0.1”

Talk Outline

- **Part 1: Learning to Warm-Start Fixed-Point Optimization Algorithms**



- **Part 2: Practical Performance Guarantees for Classical and **Learned** Optimizers**



Tutorial on Amortized Optimization [Amos 2023]

“Despite having the capacity of surpassing the convergence rates of other algorithms, oftentimes in practice amortized optimization methods can deeply struggle to generalize and converge to reasonable solutions.”

PAC-Bayes guarantees for learned optimizers

algorithm steps

tolerance

$$e_w(\theta) = \mathbf{1}(\ell_w^k(\theta) > \epsilon)$$

learnable weights

McAllester bound: given posterior and prior distributions [McAllester et. al 2003]

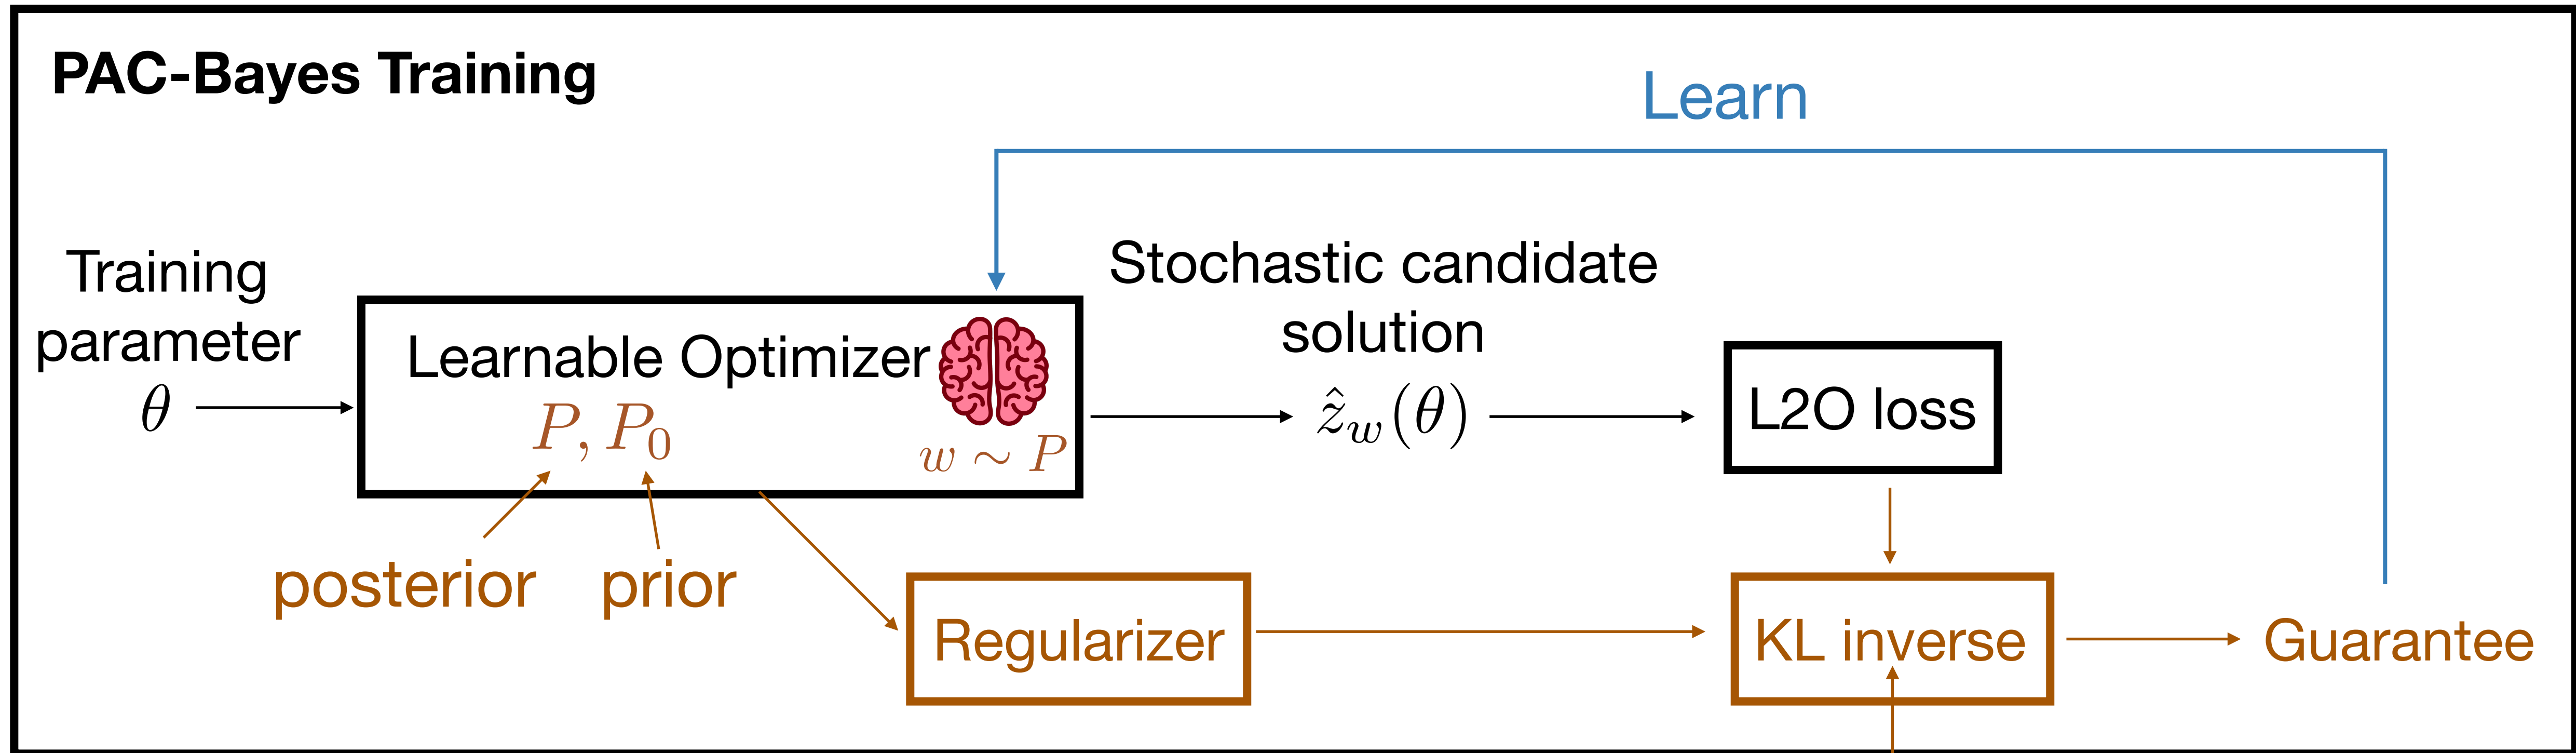
P and P_0 , with probability $1 - \delta$

$$\mathbf{E}_{\theta \sim \mathcal{X}} \mathbf{E}_{w \sim P} e_w(\theta) \leq \text{KL}^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{E}_{w \sim P} e_w(\theta_i) \middle| \frac{1}{N} (\text{KL}(P \parallel P_0) + \log(N/\delta)) \right)$$

risk \leq KL⁻¹ (empirical risk | regularizer)

Optimize the bounds directly

PAC-Bayes training architecture to optimize the guarantees

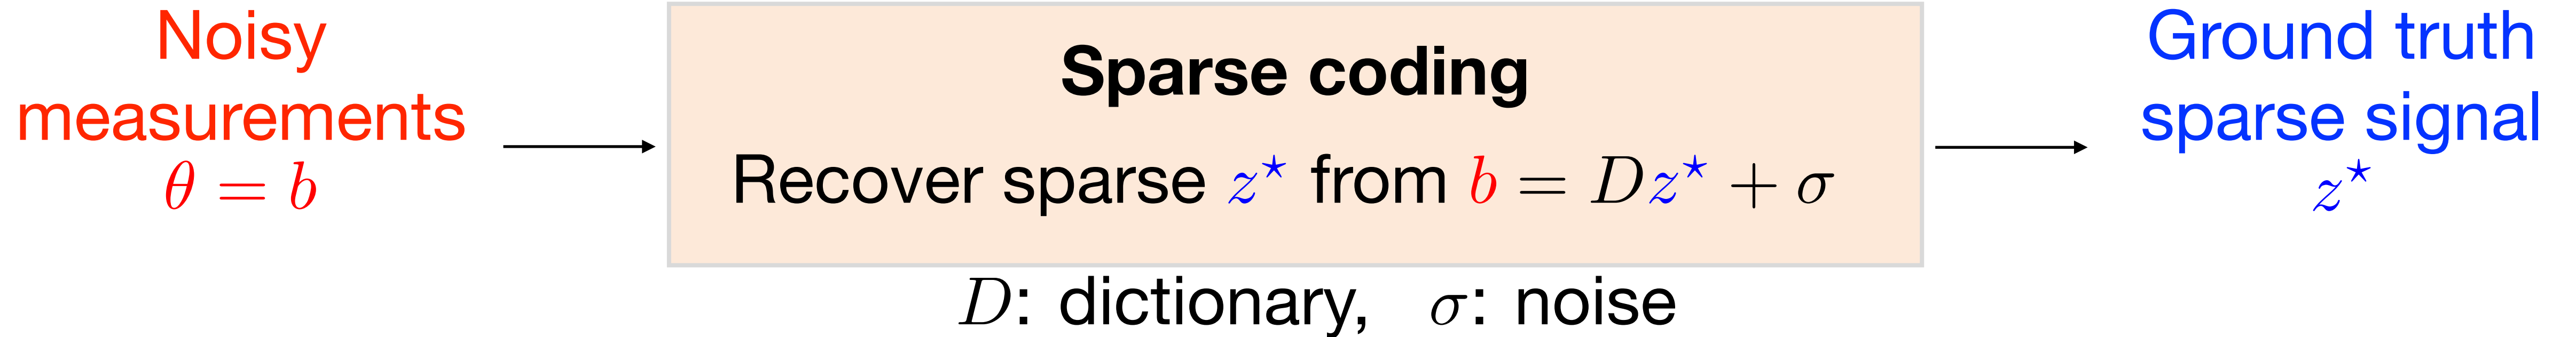


Use differentiable optimization

We show that the derivative always exists

We implement the learnable optimizer and train with this architecture

Learned algorithms for sparse coding



Standard technique

minimize $\|Dz - b\|_2^2 + \lambda \|z\|_1$

ISTA (iterative shrinkage thresholding algorithm)
(Classical optimizer)

$$z^{j+1} = \text{soft threshold}_{\frac{\lambda}{L}} \left(z^j - \frac{1}{L} (Dz^j - b) \right)$$


Learned ISTA
(Learned optimizer)

$$z^{j+1} = \text{soft threshold}_{\psi^j} \left(W_1^j z^j + W_2^j b \right)$$


+ variants [Gregor and LeCun 2010, Liu et. al 2019]

$$\text{soft threshold}_{\psi}(z) = \mathbf{sign}(z) \max(0, |z| - \psi)$$

Learned ISTA results for sparse coding



Baseline: Classical Optimizer



ISTA

Bound

LISTA




ALISTA




TiLISTA



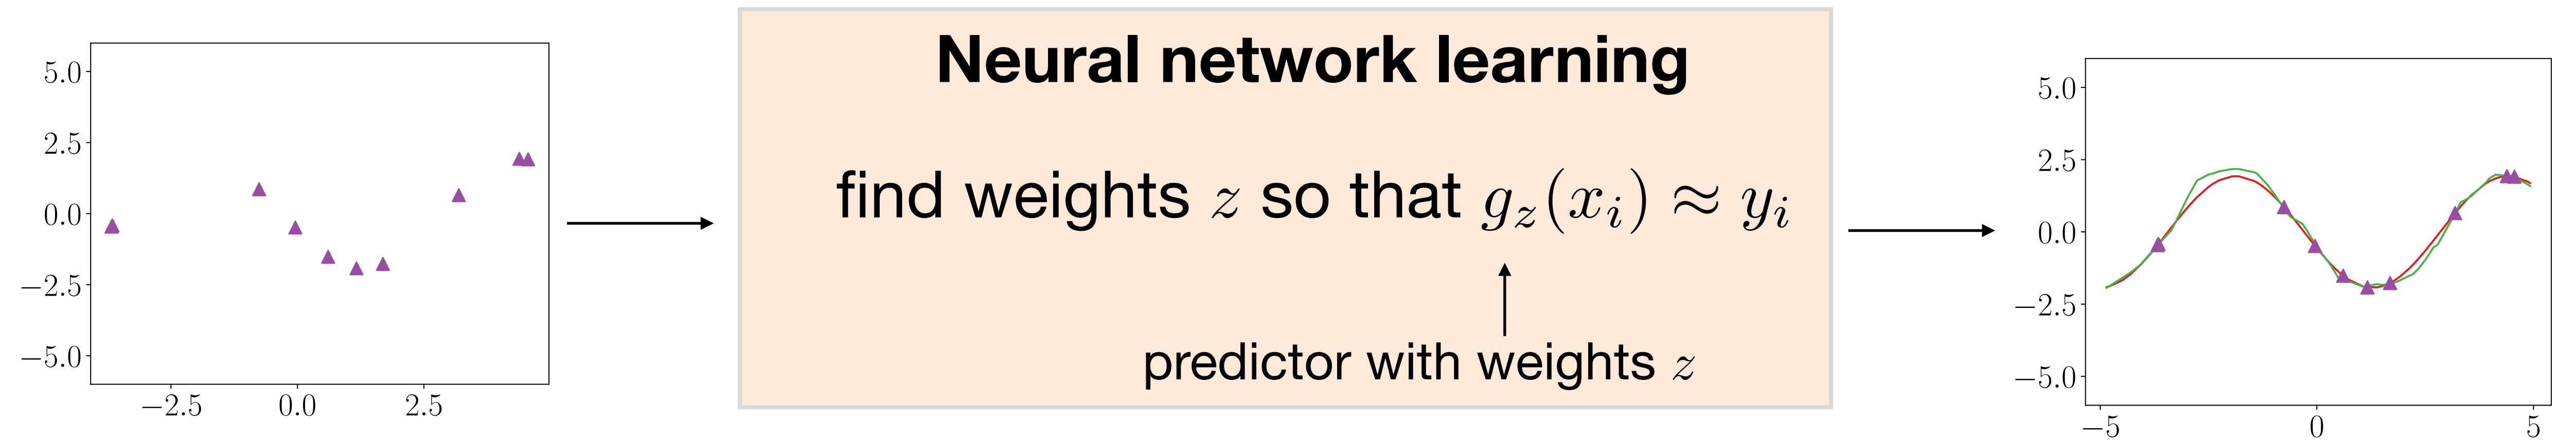
GLISTA



Empirical



K-shot Meta-Learning for Sine Curves



Training dataset
with K points

$\mathcal{D}^{\text{train}}$

Gradient step

$$\hat{z} = z - \alpha \nabla_z \mathcal{L}(z, \mathcal{D}^{\text{train}})$$


Weights that generalize
to new points quickly

\hat{z}

Model-Agnostic Meta-Learning (MAML) [Finn et. al 2017]

MAML learns a shared initialization z so that \hat{z} performs well on test data

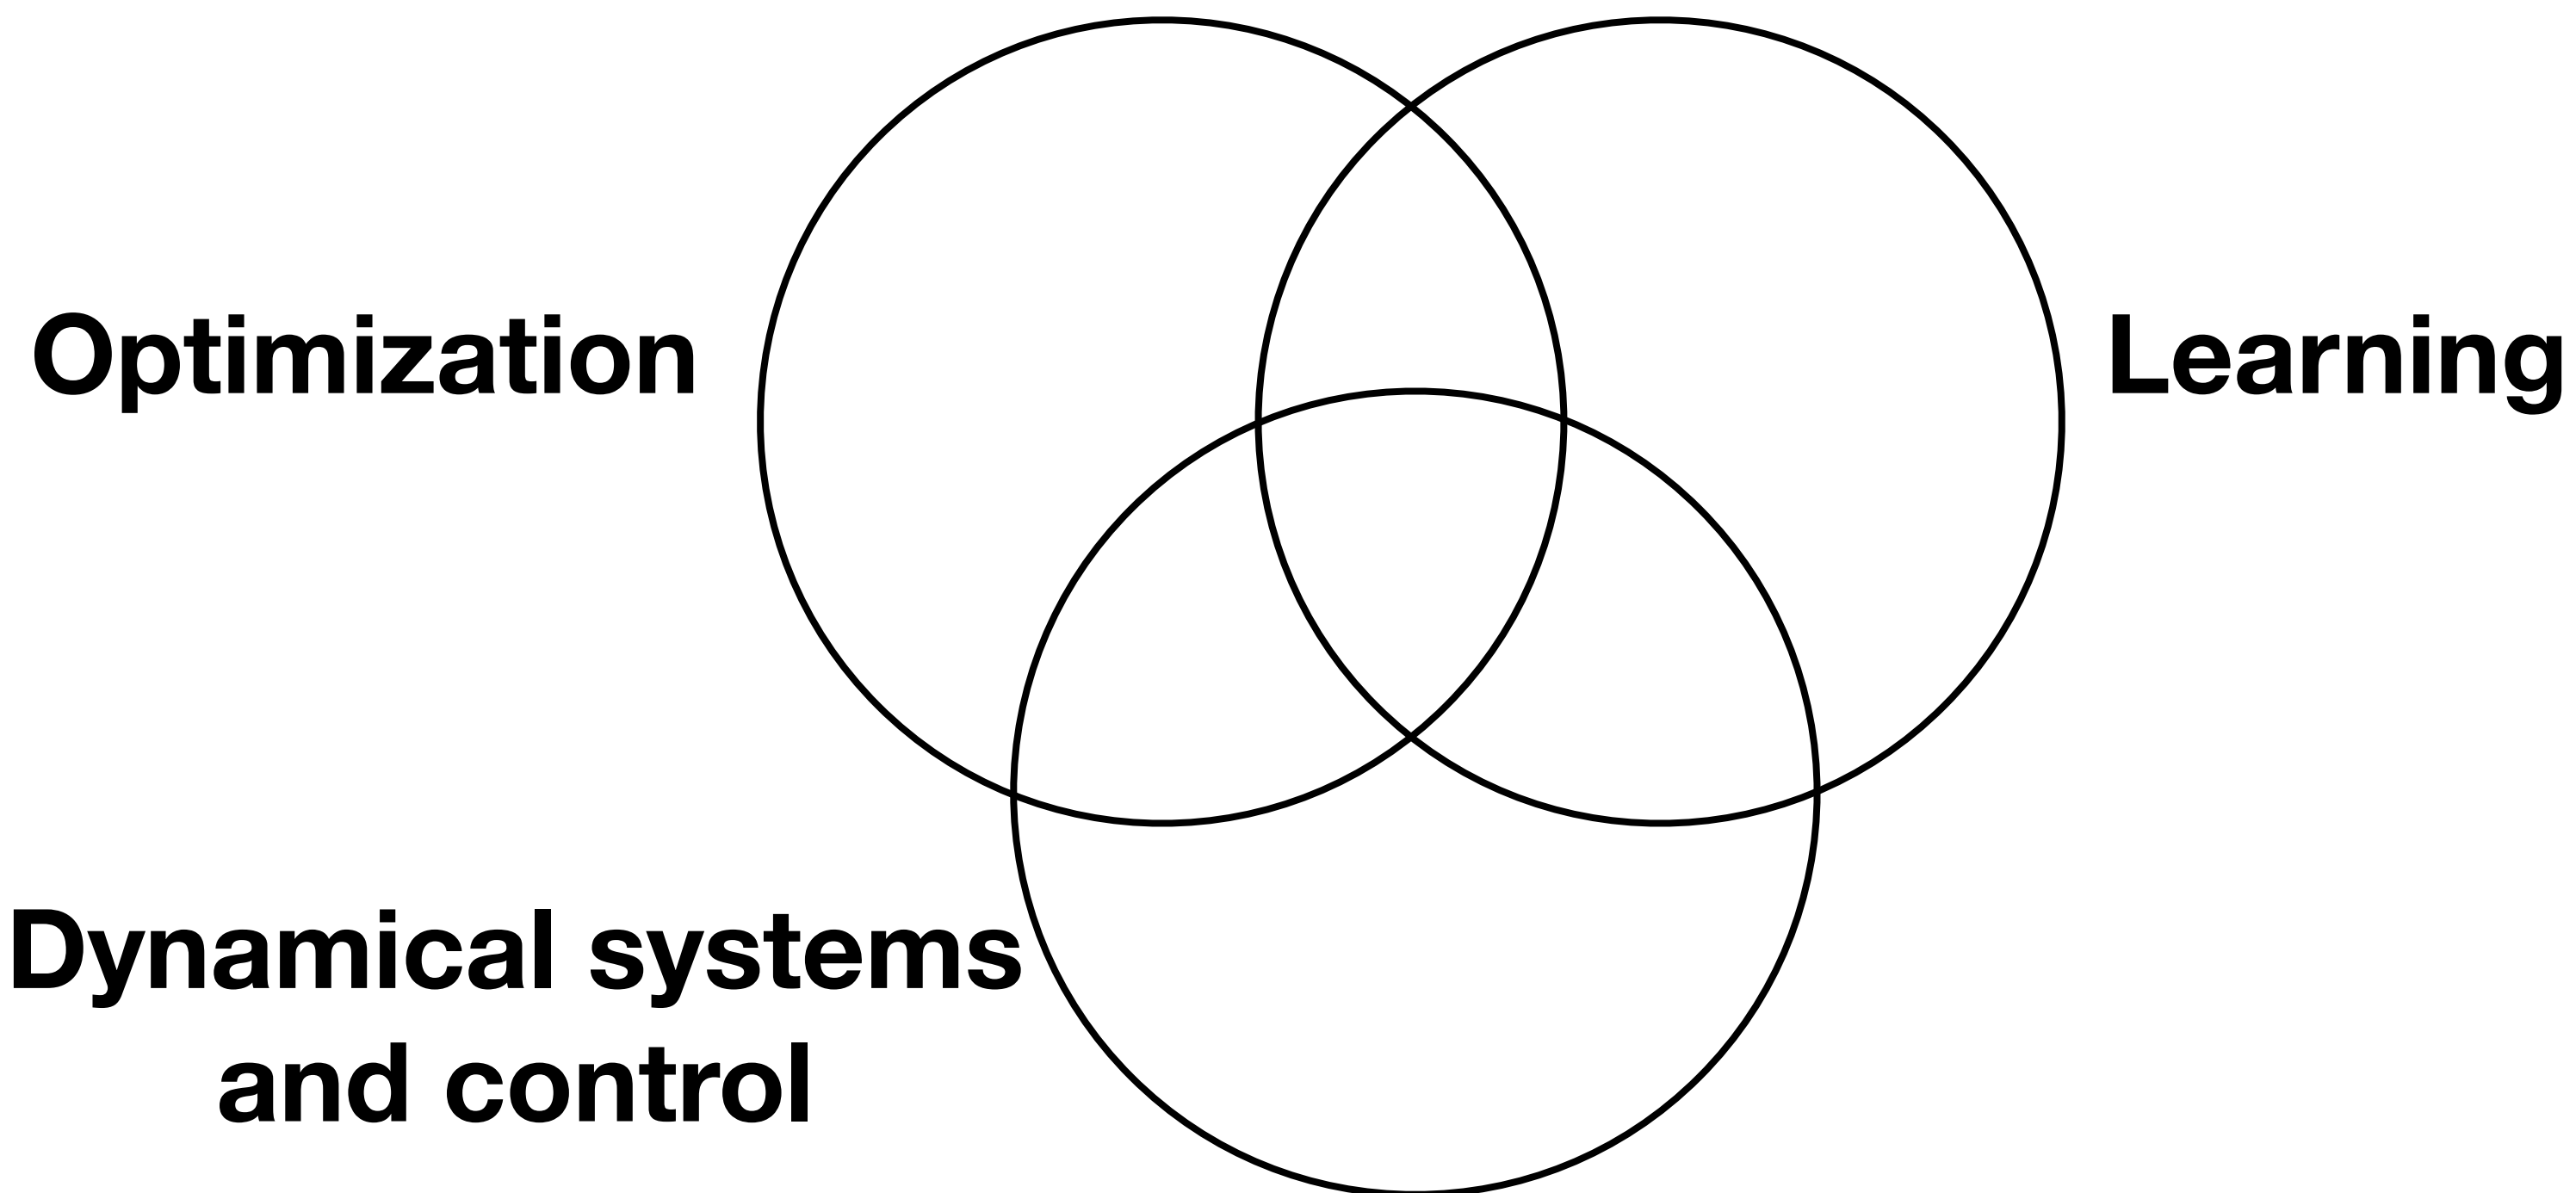
Visualizing Guarantees: K-shot Meta-Learning for Sine Curves



With high probability, 90% of the time stochastic MAML after 10 steps will stay within the band

The pretrained baseline only stays within the band 30% of the time

Future directions



Connections with REALM

Learning dynamical systems,
certificates for stability and safety

Focus on guarantees

Safe Control with Learned Certificates: A Survey of Neural Lyapunov, Barrier, and Contraction Methods for Robotics and Control [Dawson et. al 2023]

“Closely related is the issue of generalization error, which relates a learned certificate’s performance on a finite training set with its performance on the full state space... Some works have [obtained] probabilistic upper-bounds on the generalization error, but these bounds tend to be conservative.”

Conclusions

We do not need to sacrifice **guarantees for learning-based systems**

Learning to Warm-Start
Fixed-Point Optimization Algorithms

Journal of Machine Learning Research
(accepted conditioned
on minor revision)

<https://arxiv.org/pdf/2309.07835.pdf>



End-to-End Learning to Warm-Start for
Real-Time Quadratic Optimization

Learning for Dynamics and
Control Conference

<https://arxiv.org/pdf/2212.08260.pdf>



Practical Performance Guarantees
for Classical and Learned Optimizers

To be on Arxiv soon!



rajivs@princeton.edu



rajivsambharya.github.io