Learning Dynamical Systems with Conic Optimization

Amir Ali Ahmadi
Princeton University
Operations Research and Financial Engineering
Affiliated member of ECE, COS, MAE, PACM, CSML

CWI Workshop on Polynomial Optimization and Applications to Control and Energy, December 2022

PRINCETON UNIVERSITY ORFE

1) Learning Dynamical Systems with Side Information

To appear in *SIAM Review*, 2022

with **Bachir El Khadir** (Two Sigma)

2) Learning Dynamical Systems with Safety Constraints

In preparation;
shorter version in *Learning for Dynamics and Control*, 2021

with **Abraar Chaudhry** (Princeton), **Vikas Sindhwani** (Google), **Stephen Tu** (Google)

PRINCETON UNIVERSITY ORFE

Learning dynamical systems with side information

- Goal is to learn a dynamical system

  $$\dot{x}(t) = f(x(t)) \quad (\text{where } f\colon \mathbb{R}^n \to \mathbb{R}^n)$$

  from a *limited* number of *noisy* measurements of its trajectories.

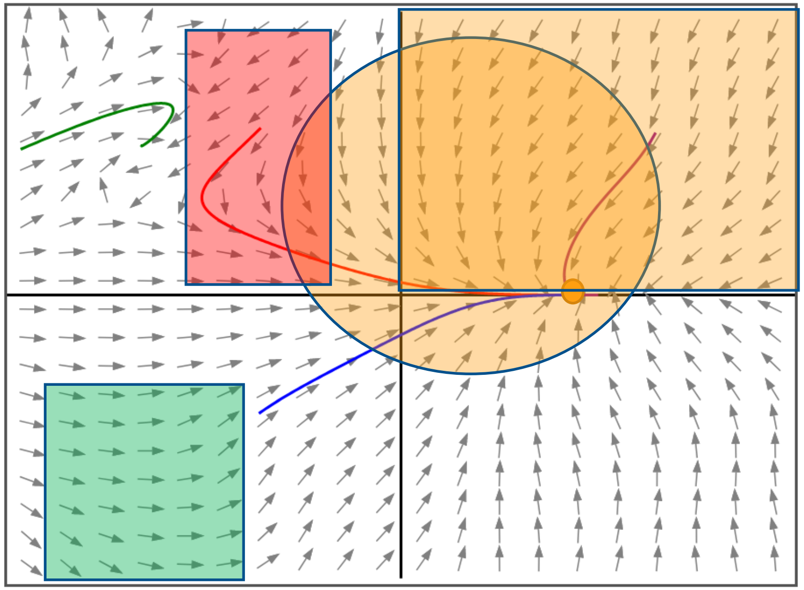

- How can side information (physical laws/contextual knowledge) be leveraged to help with learning?

Examples of "side information":

- Equilibrium points (and their stability)
- Invariance of certain sets
- Decrease of certain energy functions
- Sign conditions on derivatives of states
- Monotonicity conditions
- Having gradient/Hamiltonian structure
- (Non)reachability of a set B from a set A
- ...

Learning dynamical systems with side information

$$\dot{x}(t) = f(x(t))$$

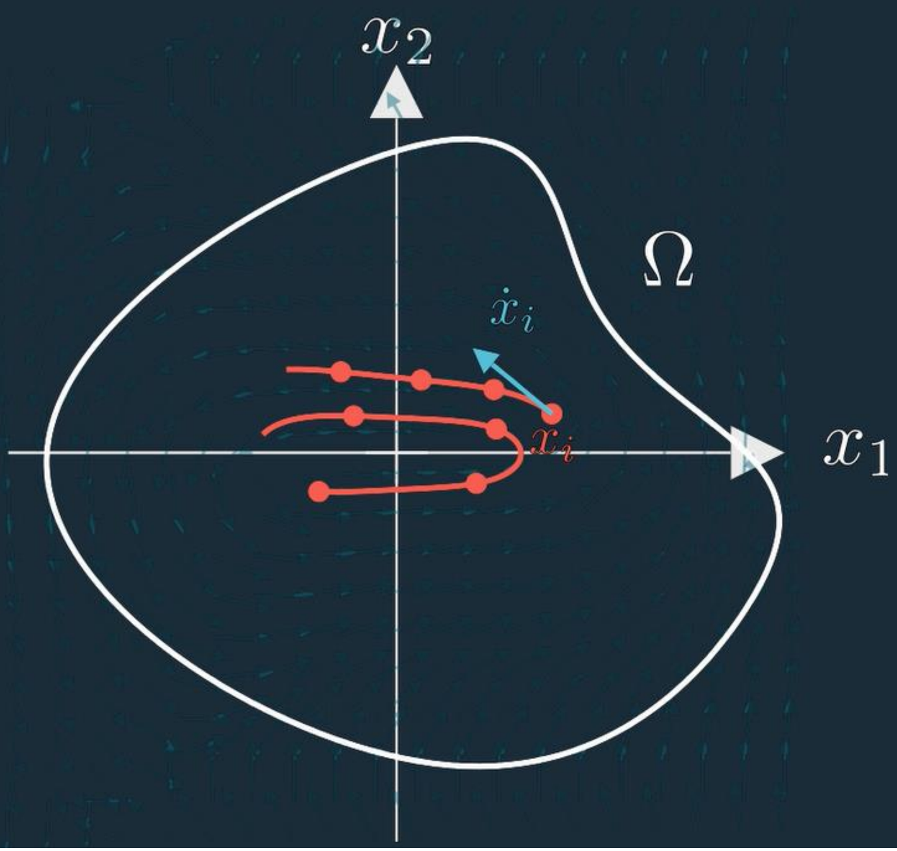

$$f : \mathbb{R}^n \rightarrow \mathbb{R}^n \quad \text{(unknown)}$$

Learn f on a compact set Ω from trajectories

$$\mathcal{D} = \{(x_i \in \mathbb{R}^n, \dot{x}_i \in \mathbb{R}^n)\}_{i=1,2,\ldots,N}$$

Our approach:

$$\min_{p : \mathbb{R}^n \rightarrow \mathbb{R}^n} \sum_i ||p(x_i) - \dot{x}_i||^2$$

Take p to be a polynomial

$$p(x_1, x_2) = \begin{pmatrix} a_1 x_1^4 + a_2 x_2^4 + a_3 x_1 x_2 + \ldots \\ b_1 x_1^4 + b_2 x_2^4 + b_3 x_1 x_2 + \ldots \end{pmatrix}$$

How to impose side information on a candidate polynomial vector field?

PRINCETON UNIVERSITY ORFE

Let's see a refresher on "SOS" & dynamical systems…

Optimization over nonnegative polynomials

Definition by example: How to pick c_1, c_2, c_3 so to make

$$p(x_1, x_2) = c_1 x_1^4 - 6x_1^3 x_2 - 4x_1^3 + c_2 x_1^2 x_2^2 + 10x_1^2 + 12x_1 x_2^2 + c_3 x_2^4$$

nonnegative over a given basic semialgebraic set?

Basic semialgebraic set: $\{x \in \mathbb{R}^n \mid g_i(x) \geq 0, i = 1, \ldots, m\}$

Ex: $x_1^3 - 2x_1 x_2^4 \geq 0$
$x_1^4 + 3x_1 x_2 - x_2^6 \geq 0$

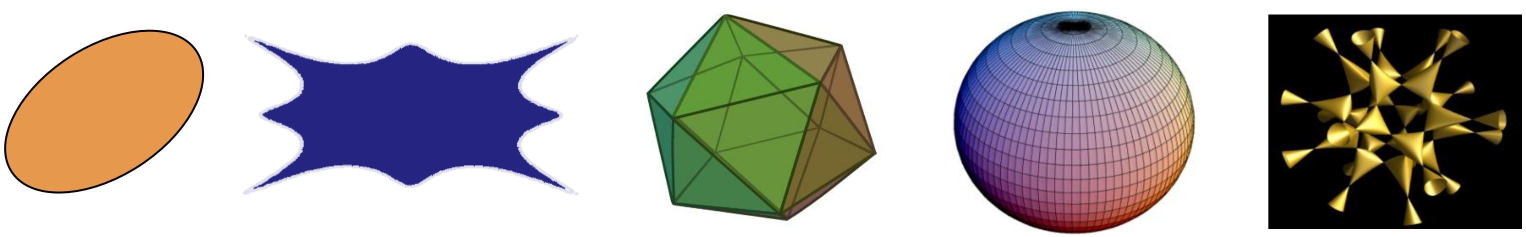

-This problem is fundamental to many areas of applied/computational mathematics.
-It is the problem that “SOS optimization” is designed to solve.

How to prove nonnegativity over a basic semialgebraic set?

Putinar's Psatz (1993):

$$\boldsymbol{p(x) \geq 0 \text{ on } K = \{x \in \mathbb{R}^n \mid g_i(x) \geq 0\}}$$

easy direction $\Uparrow$ $\Downarrow$ under positivity + the "Archimedean condition"

$\exists$SOS polynomials $\sigma_0(x), \ldots, \sigma_m(x)$ such that

$$\boldsymbol{p(x) = \sigma_0(x) + \textstyle\sum_i \sigma_i(x) g_i(x).}$$

Search for σ_i of bounded degree is an SDP!

This approach can be put to great use for analysis of dynamical systems!

[Parrilo,PhD thesis], [Henrion, Garulli, Positive polynomials in control],…

Certifying collision avoidance

$$\dot{x} = f(x)$$

(vector valued polynomial)

$\mathcal{S}$: needs safety verification

$\mathcal{U}$: unsafe (or forbidden) set

(both sets basic semialgebraic)

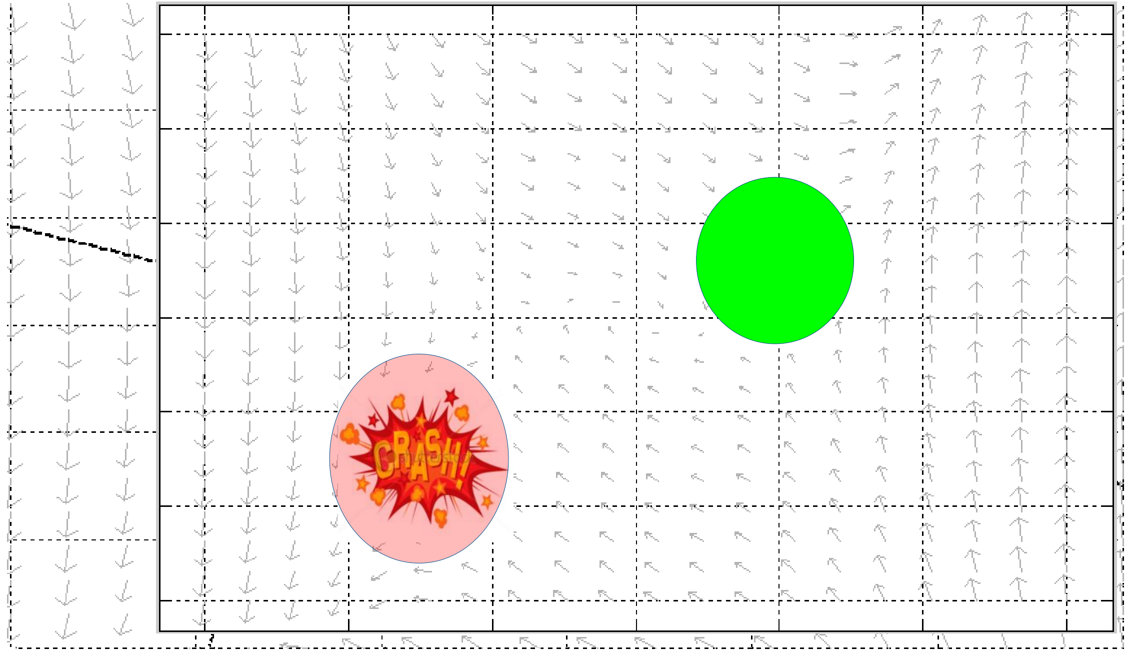

Safety guaranteed if we find a "Lyapunov function" such that:

$$B(\mathcal{S}) < 0 \qquad \dot{B} = \langle \nabla B(x), f(x) \rangle \leq 0$$
$$B(\mathcal{U}) > 0$$

Example: certifying stability

$$\dot{x} = f(x) \qquad f : \mathbb{R}^n \to \mathbb{R}^n$$

Ex. $\dot{x}_1 = -x_2 + \frac{3}{2}x_1^2 - \frac{1}{2}x_1^3 x_2$

$\dot{x}_2 = 3x_1 - x_1 x_2$

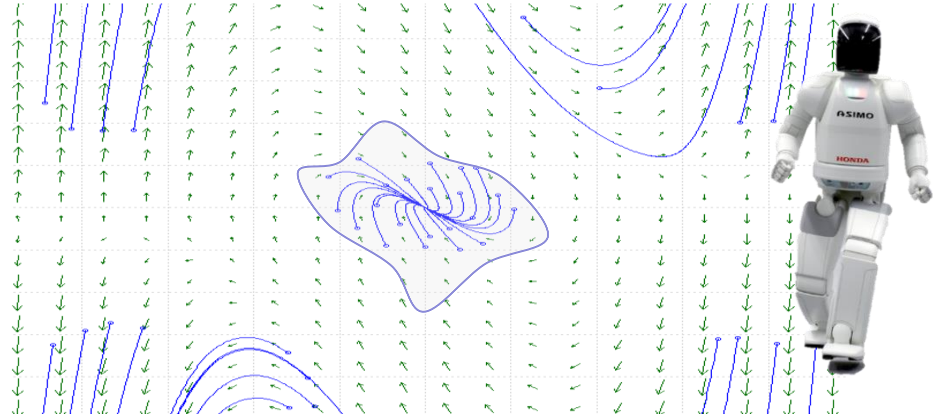

Locally asymptotic stability (LAS) of equilibrium points

Lyapunov's theorem (and its converse):

The origin is LAS if and only if there exists a C^1 function $V\colon \mathbb{R}^n \to \mathbb{R}$ that vanishes at the origin and a scalar $\beta > 0$ such that

$$V(x) > 0$$
$$V(x) \leq \beta \Rightarrow \dot{V}(x) = \nabla V(x)^T f(x) < 0$$

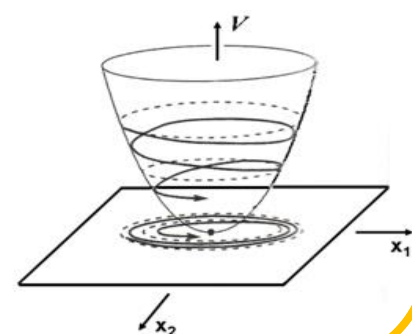

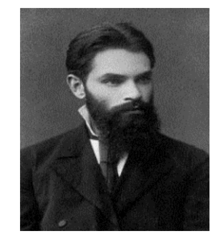

Let SOS optimization find a polynomial V and certify its inequalities.

Complexity of testing LAS for $\dot{x} = f(x)$

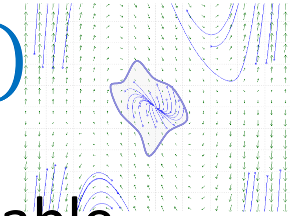

- If deg(f)=1, LAS ⇔ quadratic Lyapunov fn. Poly-time checkable.
- **Conjecture of Arnol'd (1976): LAS is undecidable when deg(f)>1.**

Existence of a **polynomial Lyapunov function**, together with a **computable upper bound** on its degree would imply decidability (e.g., by quantifier elimination).

Thm: Deciding local asymptotic stability of cubic vector fields is **strongly NP-hard**.

[AAA, *American Control Conference*]

Thm: The origin of the following vector field is LAS but the **there is no polynomial Lyapunov function** (of any degree):

$$f(x,y) = \begin{pmatrix} -2y(-x^4+2x^2y^2+y^4) \\ 2x(x^4+2x^2y^2-y^4) \end{pmatrix} - (x^2+y^2)\begin{pmatrix} 2x(x^4+2x^2y^2-y^4) \\ 2y(-x^4+2x^2y^2+y^4) \end{pmatrix}$$

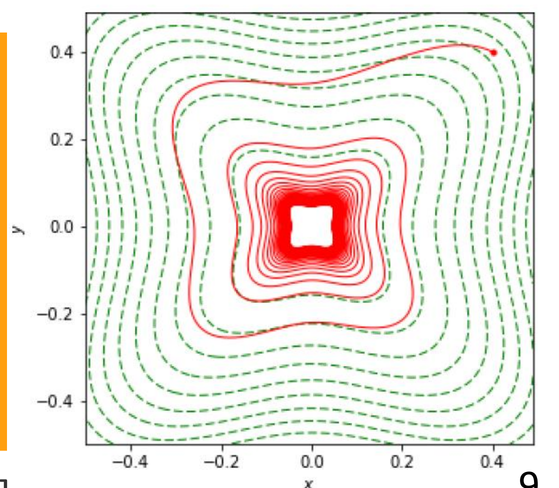

[AAA, El Khadir, *Systems & Control Letters*]

Stability does not imply polynomial Lyapunov function

$$f(x,y) = \begin{pmatrix} -2y(-x^4+2x^2y^2+y^4) \\ 2x(x^4+2x^2y^2-y^4) \end{pmatrix} - (x^2+y^2)\begin{pmatrix} 2x(x^4+2x^2y^2-y^4) \\ 2y(-x^4+2x^2y^2+y^4) \end{pmatrix}$$

Claim 1: System is GAS.

Claim 2: No polynomial Lyapunov function (of any degree) **even locally**!

Proof:

$$\boxed{W(x,y) = \frac{x^4+y^4}{x^2+y^2}}$$

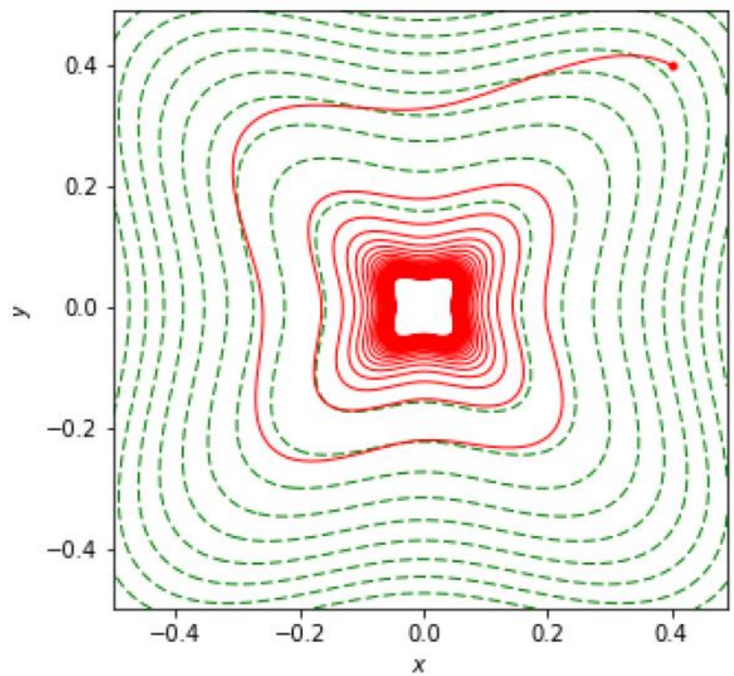

[AAA, El Khadir, *Systems & Control Letters* '18]

Proof outline

$$f(x,y) = \underbrace{\begin{pmatrix} -2y(-x^4+2x^2y^2+y^4) \\ 2x(x^4+2x^2y^2-y^4) \end{pmatrix}}_{f_0(x,y)} \underbrace{- (x^2+y^2) \begin{pmatrix} 2x(x^4+2x^2y^2-y^4) \\ 2y(-x^4+2x^2y^2+y^4) \end{pmatrix}}_{f_1(x,y)}$$

Claim 2: No polynomial Lyapunov function (of any degree) **even locally**!

Proof idea:

Suppose we had one: $p = \sum_{k=0}^{\infty} p_k$

➔ $\langle \nabla p_{k_0}(x,y), f_0(x,y) \rangle \leq 0$

➔ $\langle \nabla p_{k_0}(x,y), f_0(x,y) \rangle = 0.$

➔ A polynomial must be constant on the unit level set of $W(x,y) = (x^4+y^4)/(x^2+y^2)$

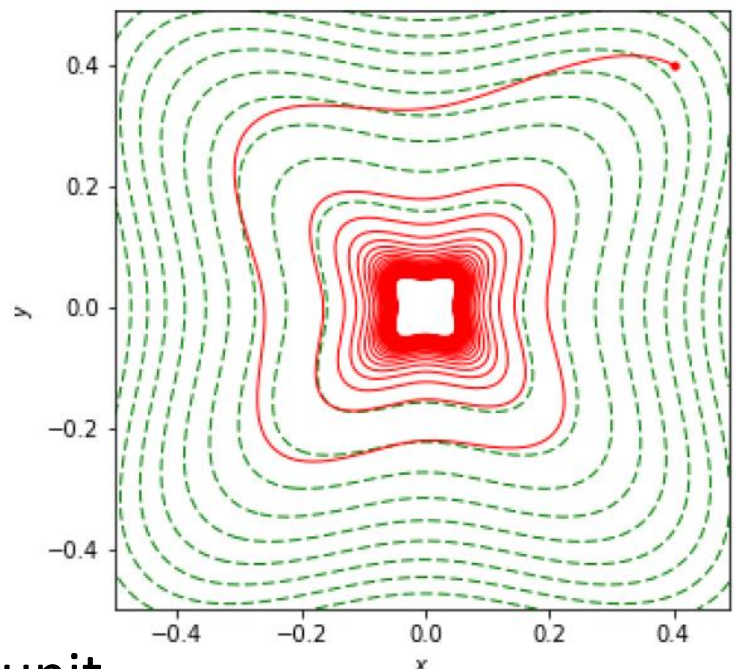

In practice however…

The nonlinear control community has had great success with low-degree polynomial Lyapunov functions and SOS

SOS for stabilizing a humanoid robot on one foot

$\dot{x} = f(x, u)$, 30 states, cubic dynamics

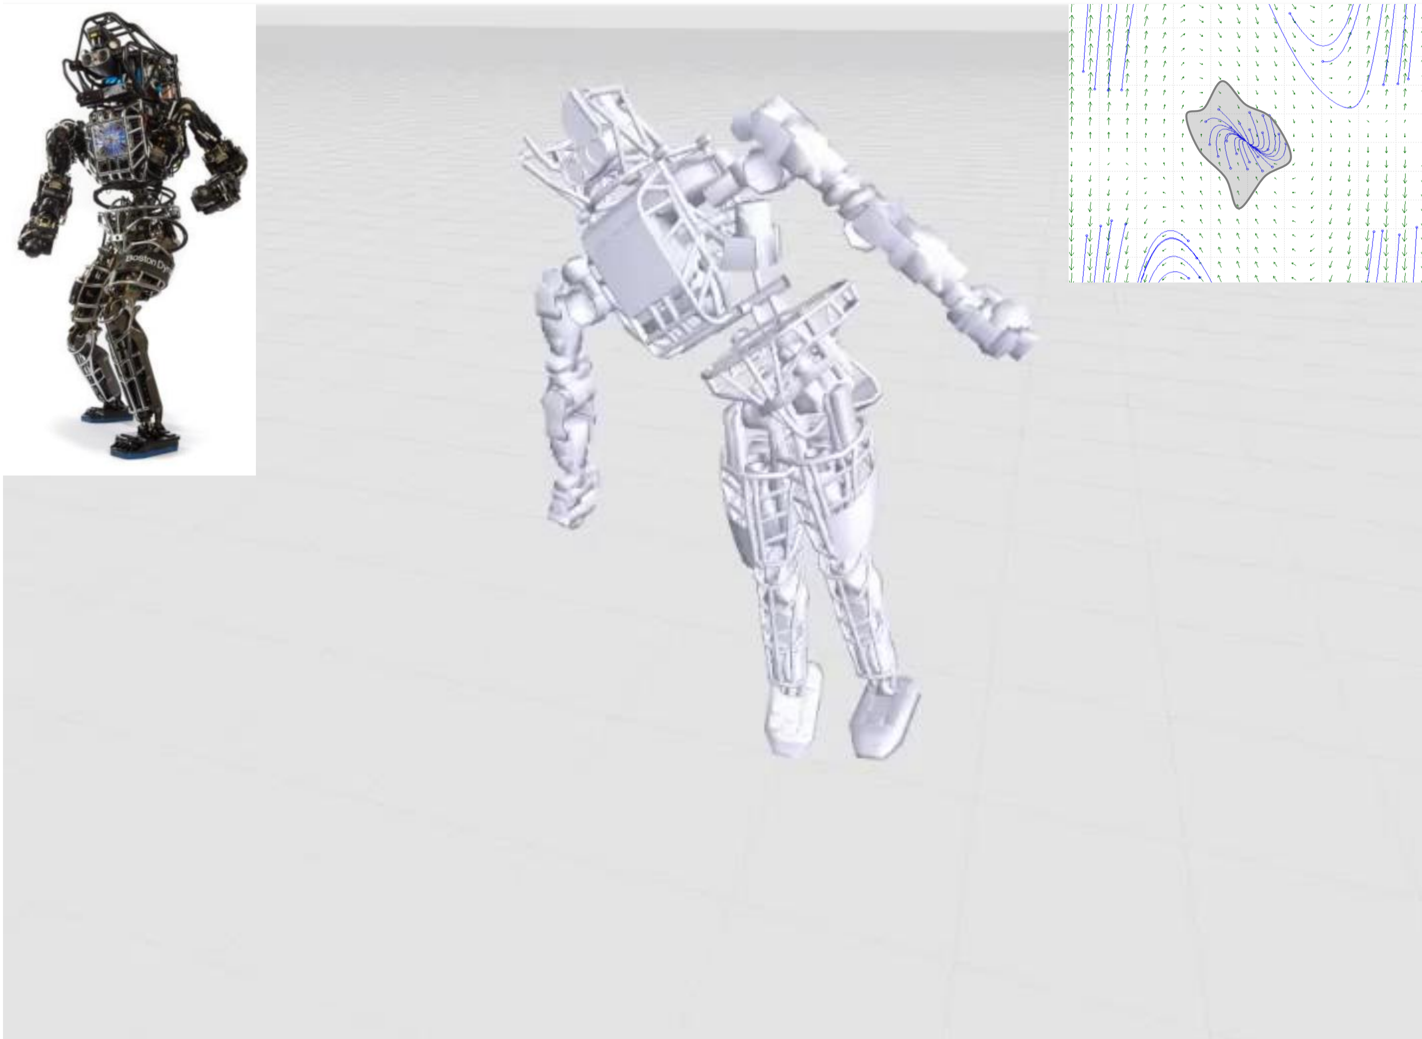

(w/ Majumdar and Tedrake)

Back to learning dynamical systems with side information

- Goal is to learn a dynamical system

  $\dot{x} = f(x)$ (where $f\colon \mathbb{R}^n \to \mathbb{R}^n$)

  from a *limited* number of *noisy* measurements of its trajectories.

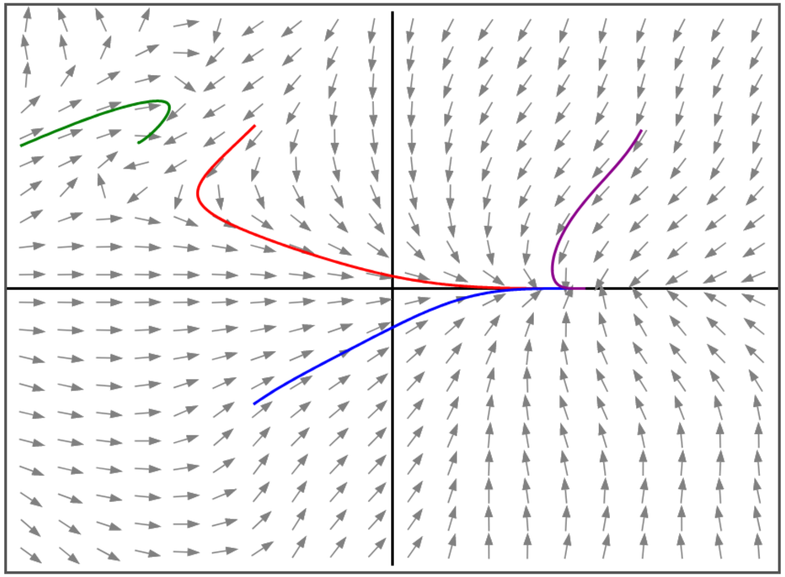

Examples of "side information":

- Equilibrium points (and their stability)
- Invariance of certain sets
- Decrease of certain energy functions
- Sign conditions on derivatives of states
- Having gradient structure
- Monotonicity conditions
- (Non)reachability of a set B from a set A
- …

> - Parametrize a polynomial vector field $p\colon \mathbb{R}^n \to \mathbb{R}^n$.
> - Use SOS optimization to impose side information as constraints on p.
> - Pick the p that best explains the data.

An epidemiology example

A model from the epidemiology literature for spread of Gonorrhea in a heterosexual population:

$$\dot{x} = f_1(x, y) = -a_1 x + b_1(1 - x)y$$

$$\dot{y} = f_2(x, y) = -a_2 y + b_2(1 - y)x$$

$x(t)$: fraction of infected males at time t

$y(t)$: fraction of infected females at time t

a_1: recovery rate of males

a_2: recovery rate of females

b_1: infection rate of males

b_2: infection rate of females

For our experiments:

$a_1 = a_2 = .1; b_1 = b_2 = .05.$

This is taken to be "the ground truth".

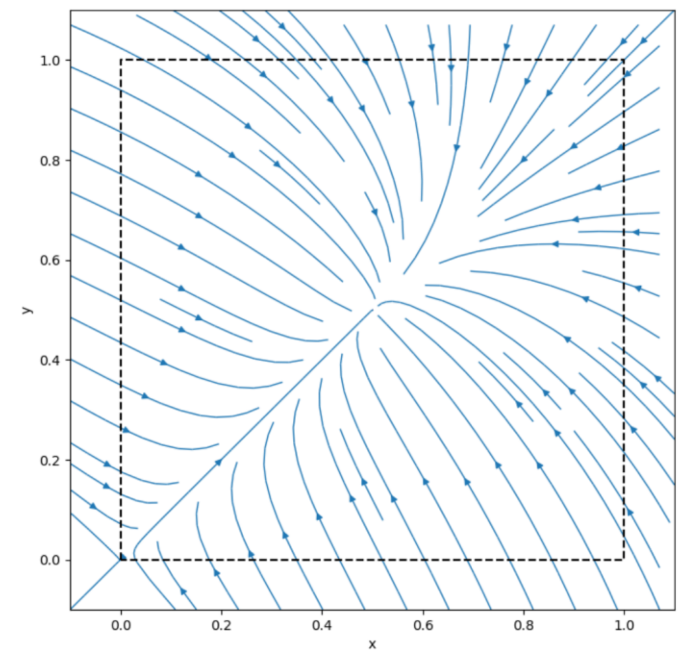

- The dynamics (both its parameters and its special structure) is unknown to us.
- We only get to observe noisy trajectories of this dynamical system.

The setup

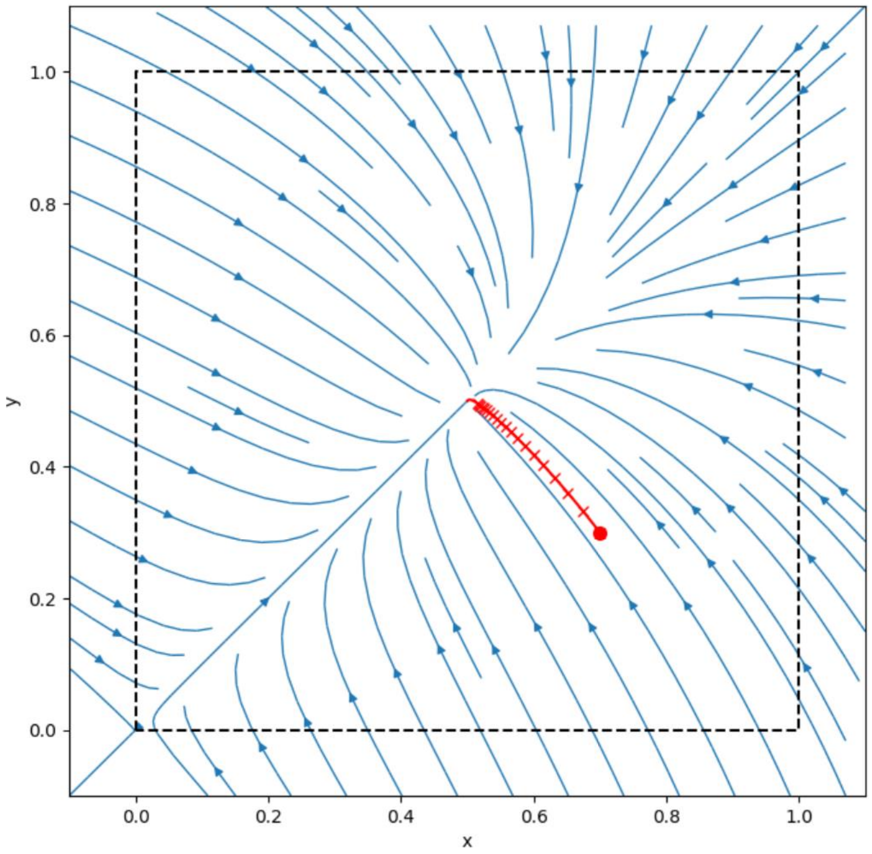

- The true dynamics f is unknown
- **What we observe:**
  Noisy measurements of the vector field on 20 points from a single trajectory starting from [0.7;0.3]
- **Goal:**
  - Learn a polynomial vector field p (of degree 2 or 3) that best agrees with the observed trajectory
  - Incorporate side information to generalize better to unobserved trajectories

Learning p of degree 2

The true dynamics f (unknown)

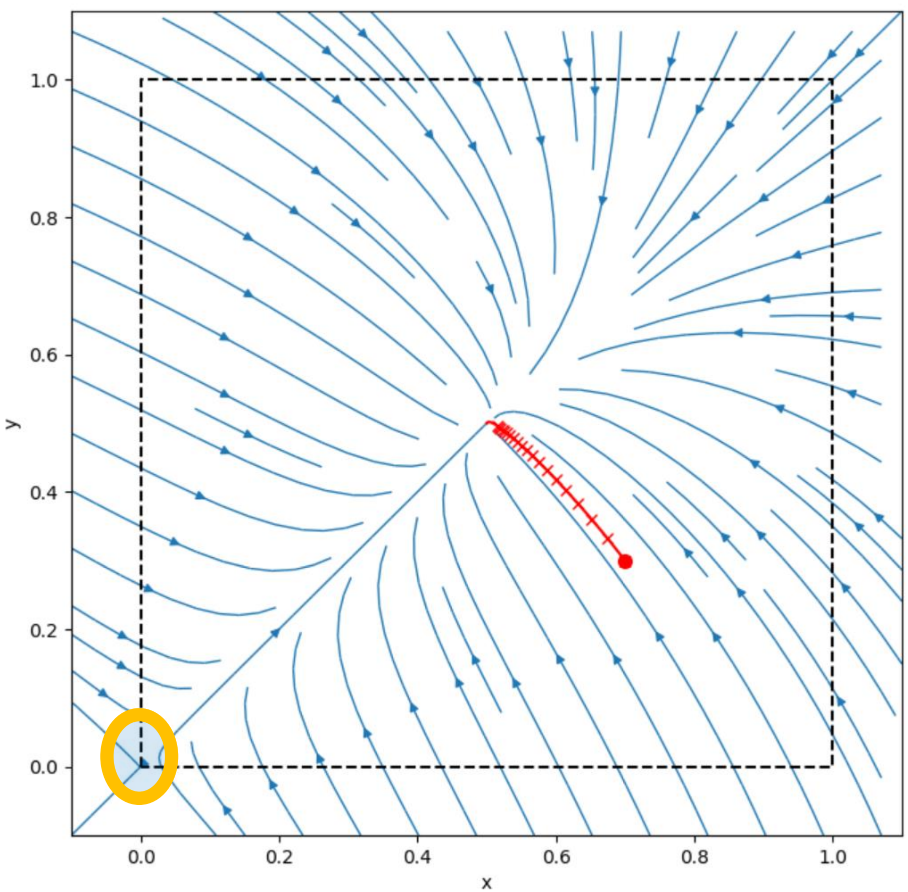

Least squares solution

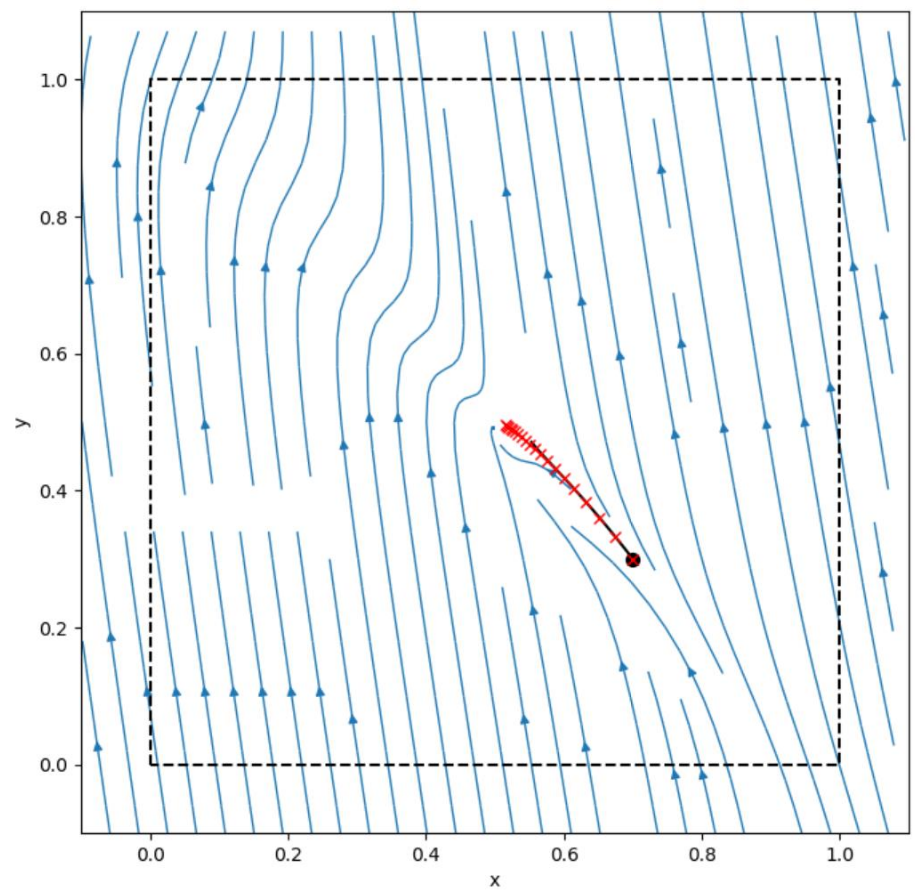

- Good performance on the observed trajectory. Horrible elsewhere.
- What side information can you think of in the context of this problem?

$f(0) = 0$!!

Learning p of degree 2

The true dynamics f (unknown)

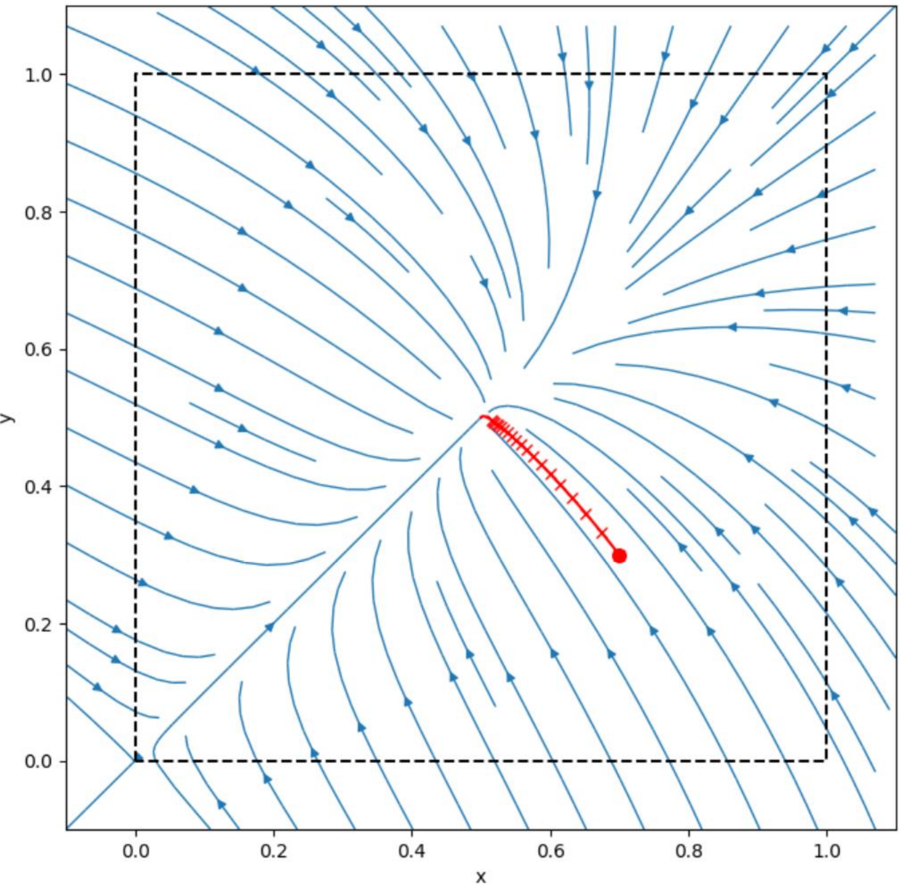

Least squares solution subject to $p(0) = 0$

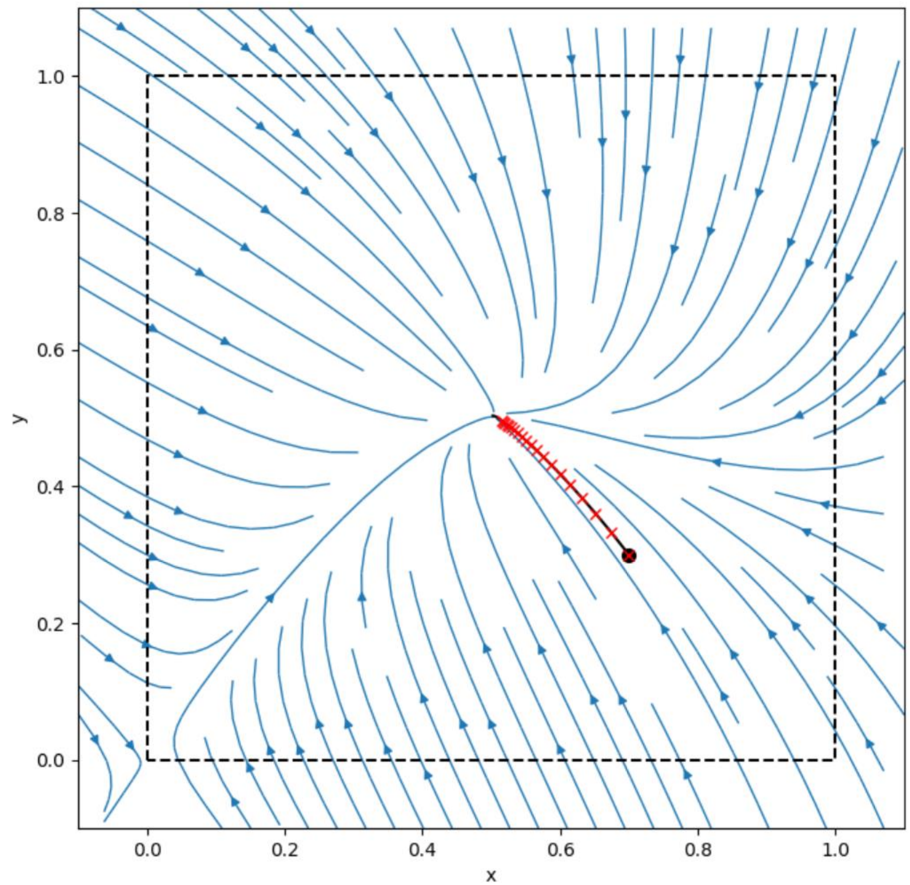

- Already gets the qualitative behavior on unobserved state space correctly!

Maybe we are helping ourselves by taking p to have degree 2 (same as f)?
Let's try higher degree p to assume a richer model than the truth.

Learning p of degree 3

The true dynamics f (unknown)

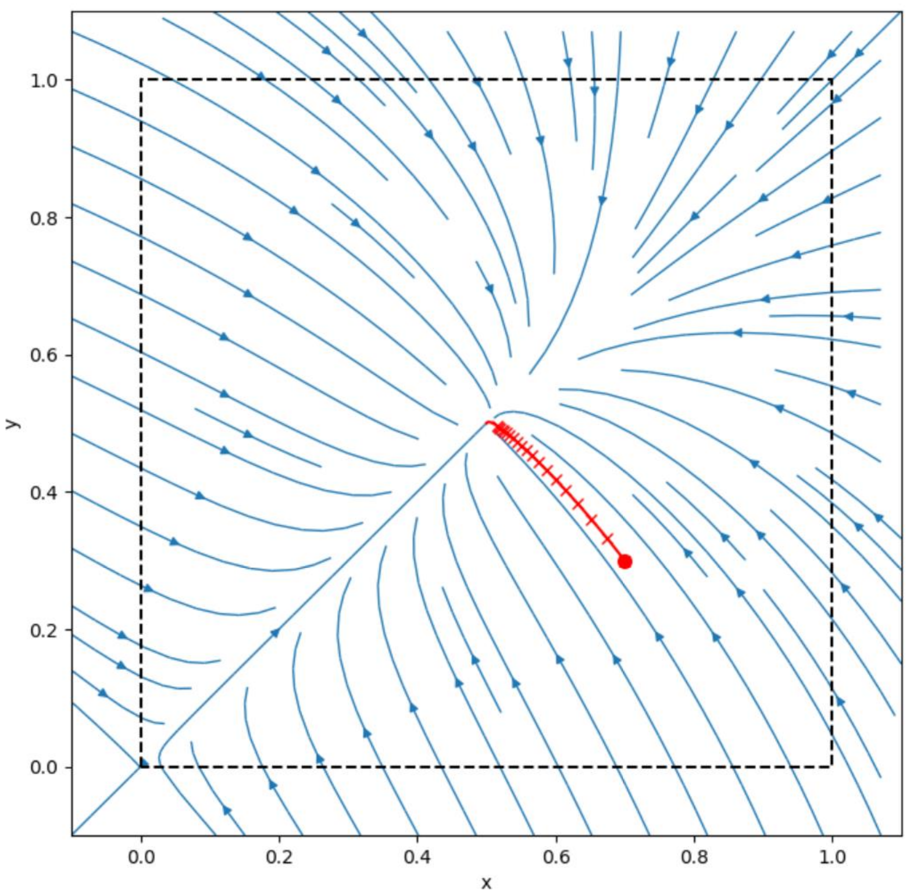

Least squares solution

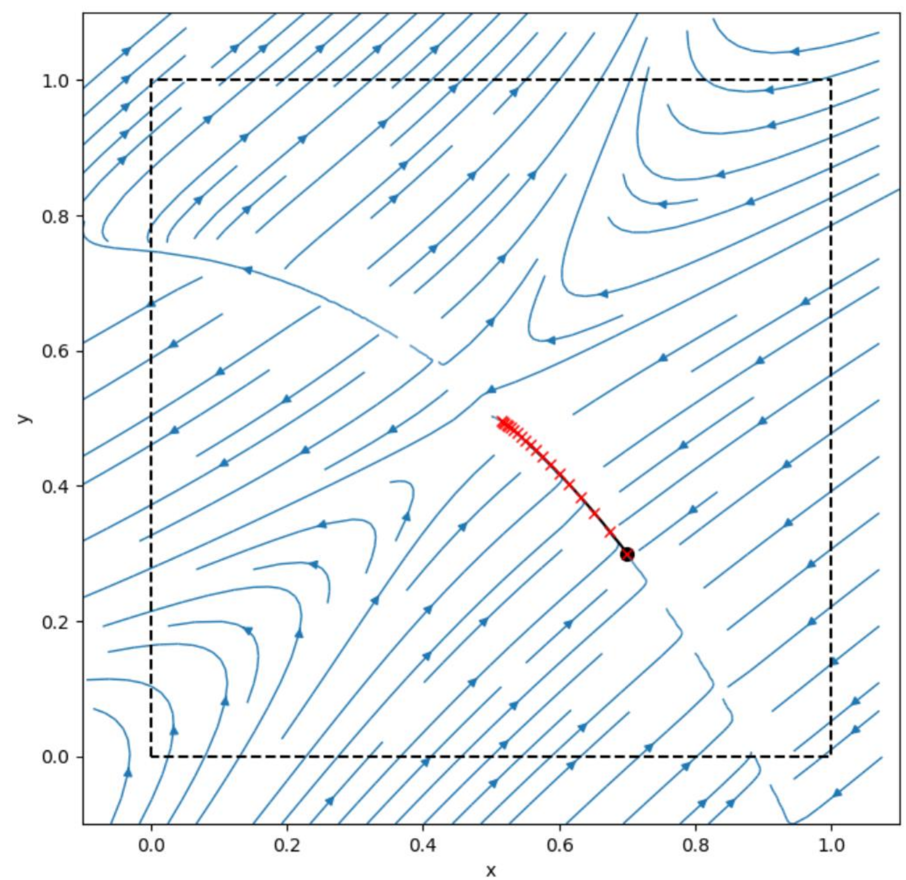

- Good performance on the observed trajectory. Again horrible elsewhere.

Learning p of degree 3

The true dynamics f (unknown)

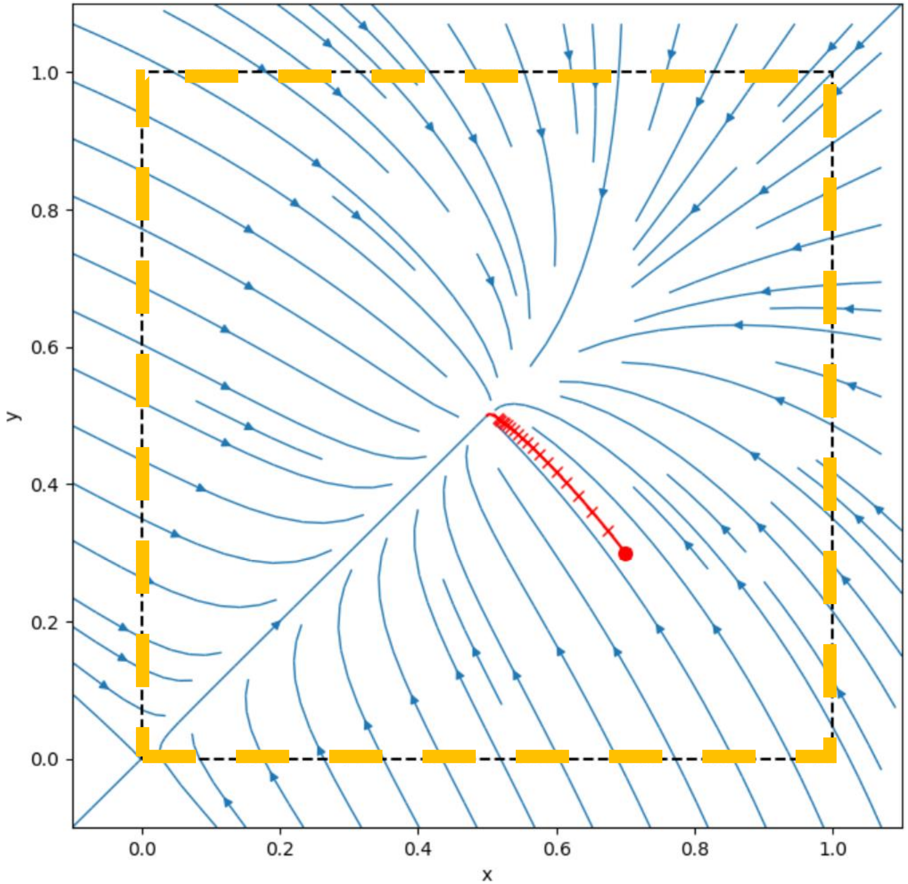

Least squares solution subject to $p(0) = 0$

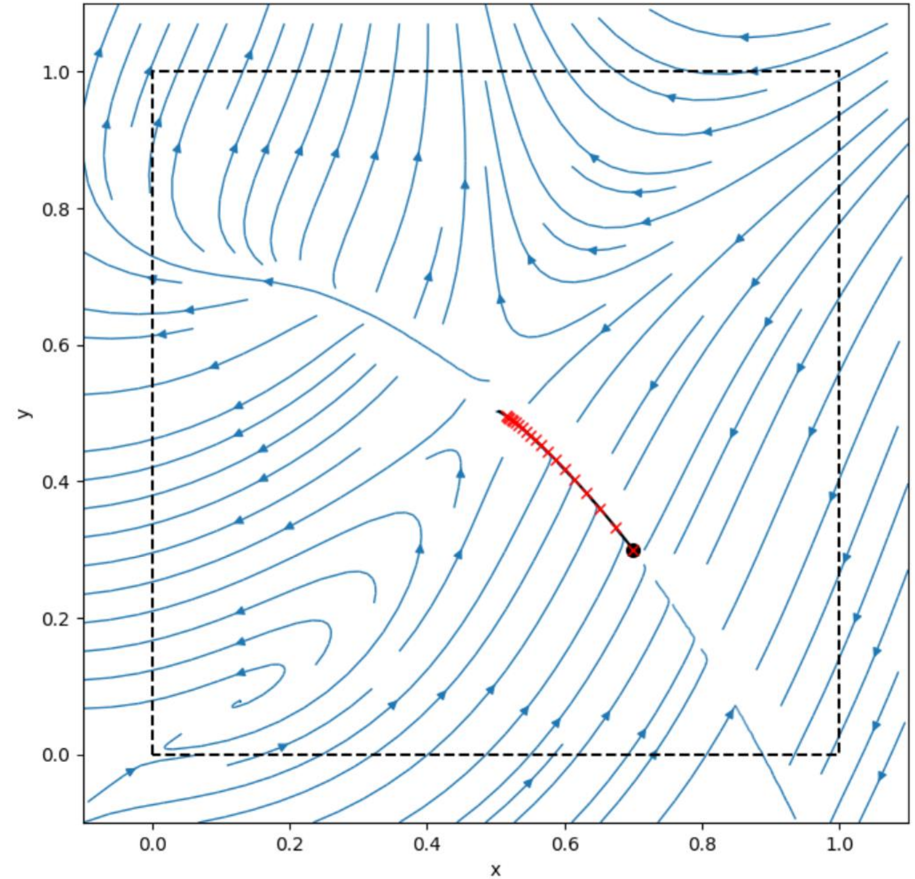

- Still pretty bad. What other side information can you think of?

Fraction of infected individuals cannot go negative or more than one!

The unit square must be an invariant set!!

Learning p of degree 3

The true dynamics f (unknown)

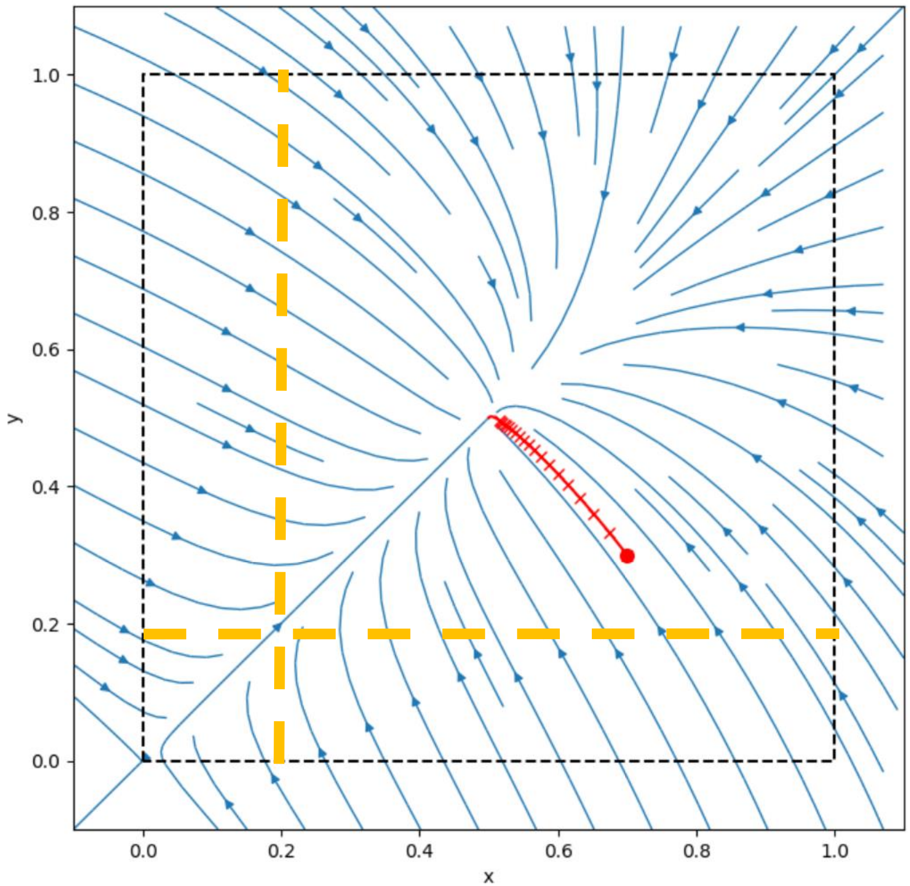

Least squares solution subject to $p(0) = 0$, unit square invariant

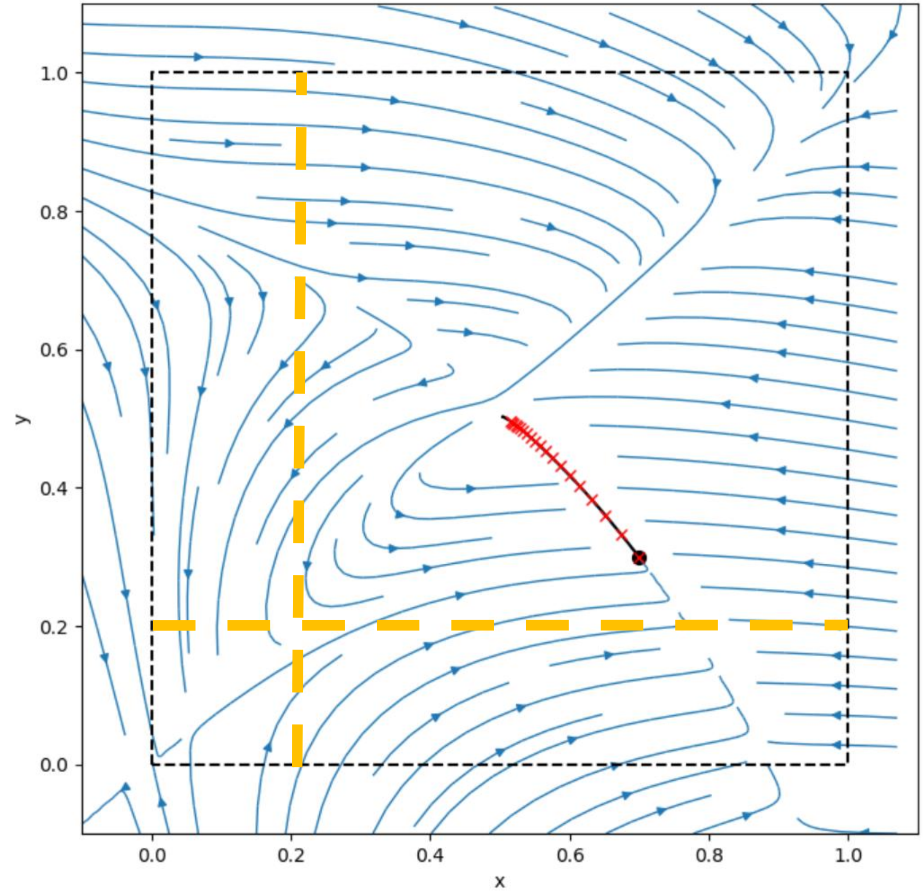

- Better, but not perfect. What other side information can you think of?

More infected females should imply higher infection rate for males! (and vice versa)

Side information: directional monotonicity

The true dynamics f (unknown)

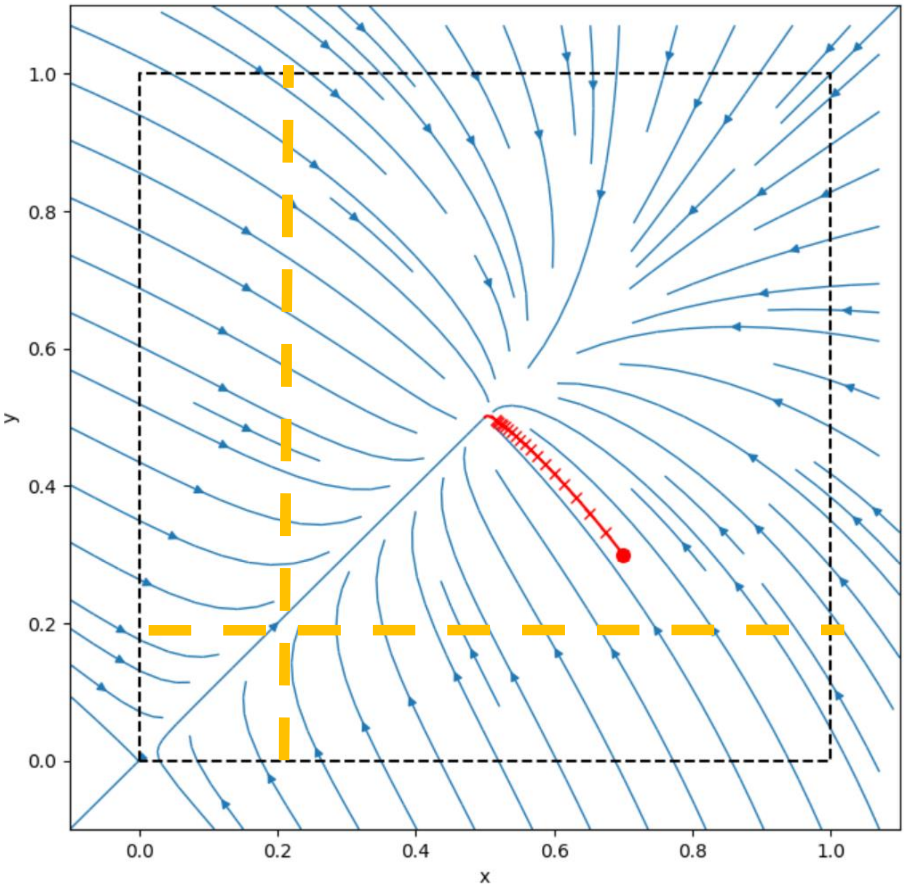

$$\frac{\partial f_1(x,y)}{\partial y} \geq 0, \forall (x,y) \in [0,1]^2$$

$$\frac{\partial f_2(x,y)}{\partial x} \geq 0, \forall (x,y) \in [0,1]^2$$

We want p to satisfy the same constraints!

Learning p of degree 3

The true dynamics f (unknown)

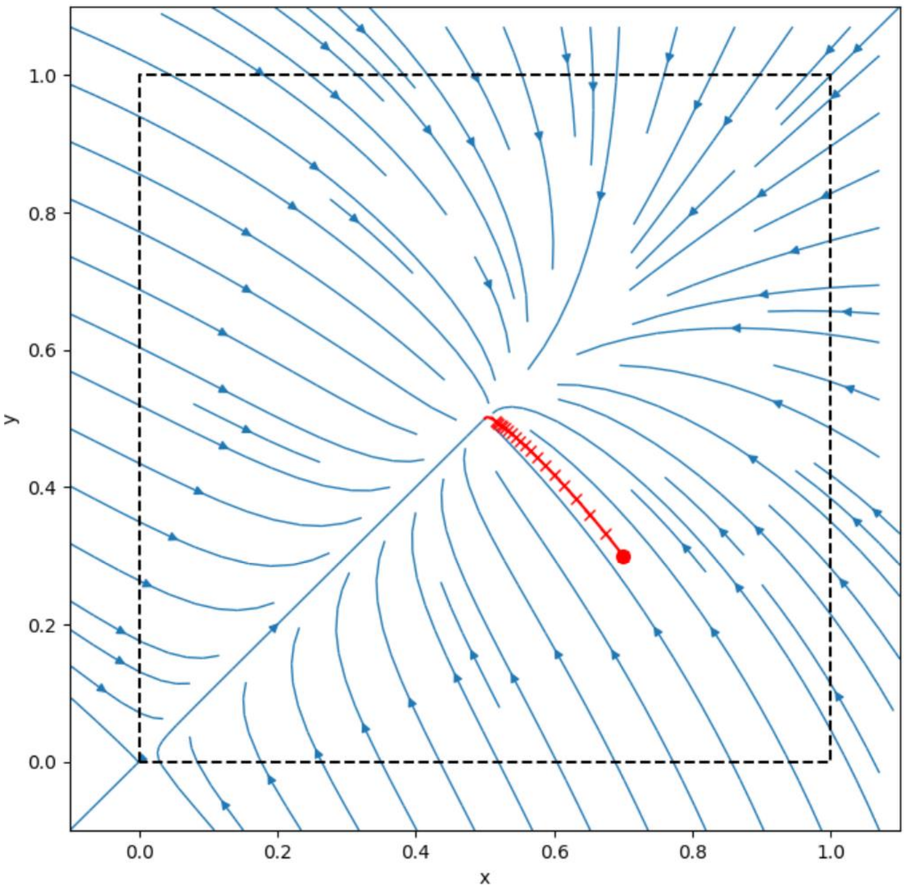

Least squares solution subject to $p(0) = 0$, unit square invariant, directional monotonicity

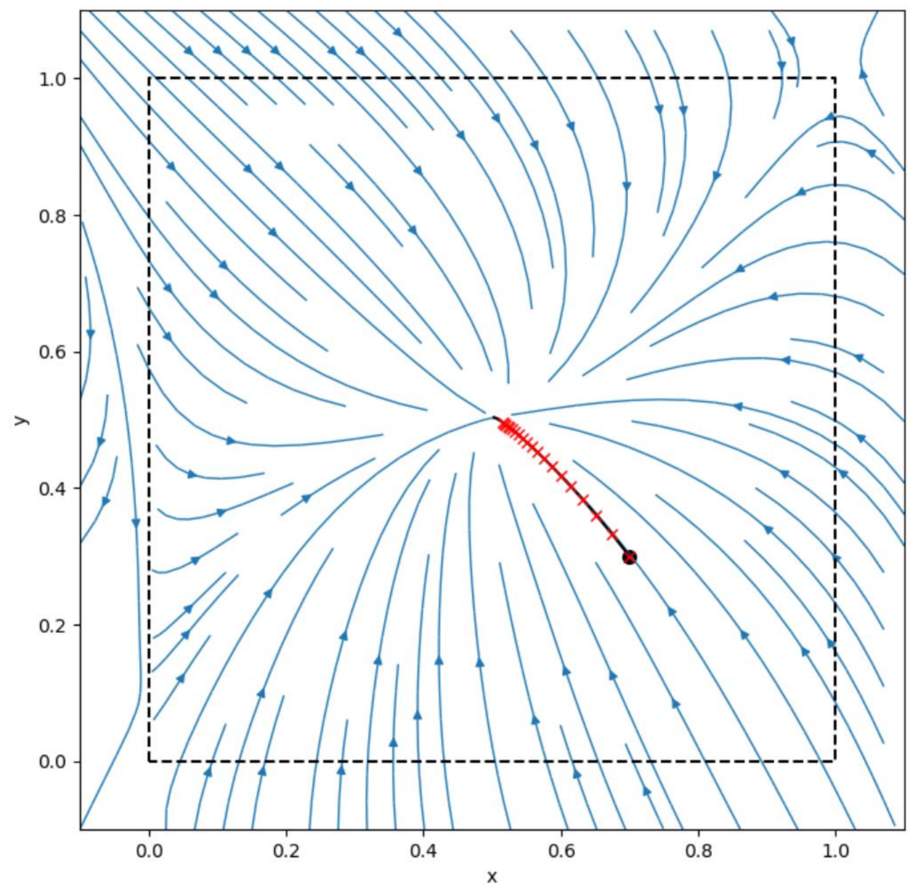

- Now we are getting the qualitative behavior correct everywhere!

Let's learn p of degree 2 again just for fun

The true dynamics f (unknown)

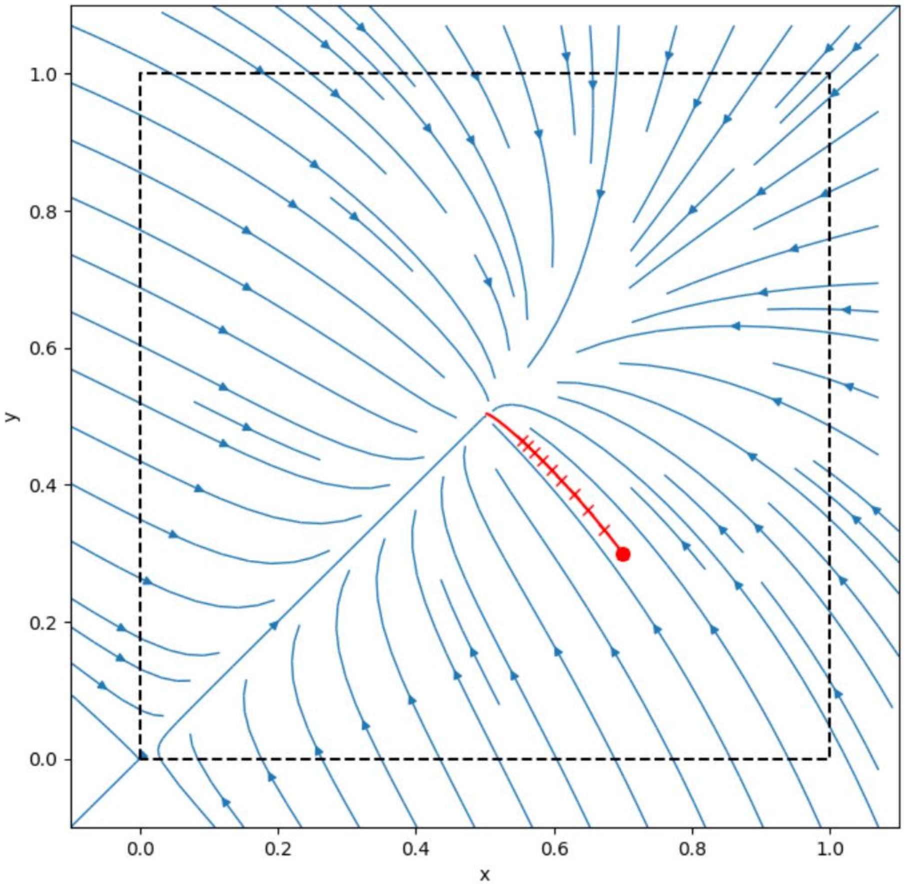

Least squares solution subject to $p(0) = 0$, unit square invariant, directional monotonicity

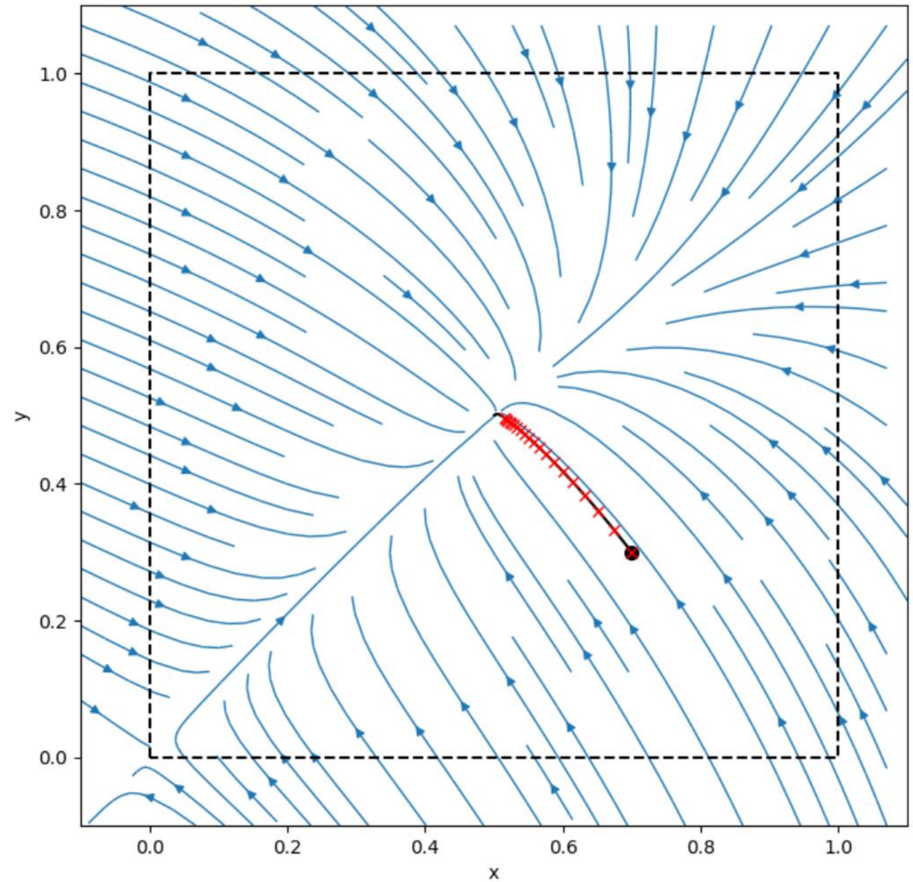

- p is pretty much dead on everywhere even though it was trained on a single trajectory!

The SDP that is being solved in the background

$$\min_{\substack{p=\begin{pmatrix}p_1\\p_2\end{pmatrix},\ \deg(p)\leq 3 \\ \sigma_0,\sigma_1,\ \deg(\sigma_i)\leq 2 \\ \hat{\sigma}_0,\hat{\sigma}_1,\hat{\sigma}_2 \\ \deg(\hat{\sigma}_0)\leq 2}} \sum_{i=1}^{20} \left(p(x^i,y^i)-\hat{f}(x^i,y^i)\right)^2$$

s.t. $p_1(0,0)=0, \quad p_2(0,0)=0$

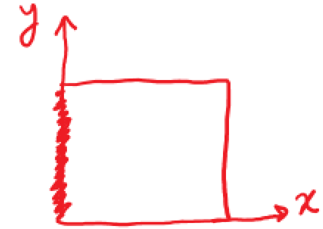

$$\left.\begin{array}{l} p_1(0,y)=y\,\sigma_0(y)+(1-y)\,\sigma_1(y) \\ \sigma_0,\ \sigma_1 \ \text{SOS} \end{array}\right\} \Rightarrow \left[x=0,\ 0\leq y\leq 1 \Rightarrow \dot{x}\geq 0\right]$$

(+ three similar constraints)

$$\left.\begin{array}{l} \frac{\partial p_1}{\partial y}(x,y)=\hat{\sigma}_0(x,y)+\hat{\sigma}_1(1-x)x+\hat{\sigma}_2(1-y)y \\ \hat{\sigma}_0 \ \text{SOS},\ \hat{\sigma}_1\geq 0,\ \hat{\sigma}_2\geq 0 \end{array}\right\} \Rightarrow \left[0\leq x\leq 1,\ 0\leq y\leq 1 \Rightarrow \frac{\partial p_1}{\partial y}(x,y)\geq 0\right]$$

(similarly for $\frac{\partial p_2}{\partial x}(x,y)$)

Output of SDP solver:

p1=0.2681*x^3 - 0.0361*x^2*y - 0.095*x*y^2 + 0.1409*y^3 - 0.4399*x^2 + 0.0956*x*y - 0.0805*y^2 + 0.1232*x + 0.0201*y

p2=0.1188*x^3 + 0.2606*x^2*y + 0.2070*x*y^2 + 0.0005*y^3 - 0.3037*x^2 - 0.4809*x*y -0.099*y^2+ 0.2794*x+0.01689*y

Following learning with optimal control

The true dynamics f (unknown)

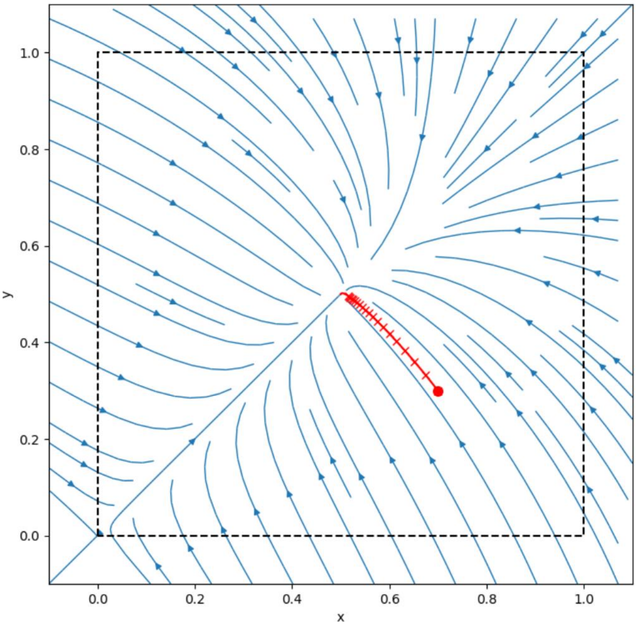

$$\dot{x}(t) = f(x(t)) - \begin{pmatrix} u_1 x_1 \\ u_2 x_2 \end{pmatrix}$$

$\boldsymbol{u_1}$: fraction of males to test

$\boldsymbol{u_2}$: fraction of females to test

As before: f is unknown; we only see 20 points from a single trajectory starting from $(0.7,0.3)^T$

Goal is to choose u_1, u_2 to minimize:

$$c(u_1, u_2) := x_1(T, \hat{x}_{\text{init}}) + x_2(T, \hat{x}_{\text{init}}) + \alpha(u_1 + u_2)$$

$$\hat{x}_{\text{init}} = (0.5, 0.4)^T \;\; T = 20, \, \alpha = 0.4$$

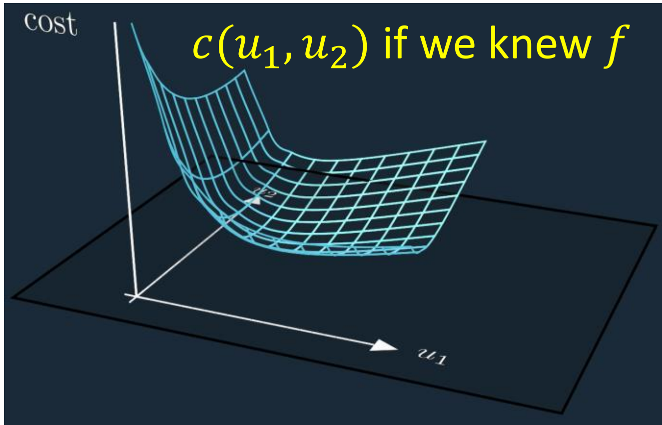

Applying optimal control to dynamics learned from data

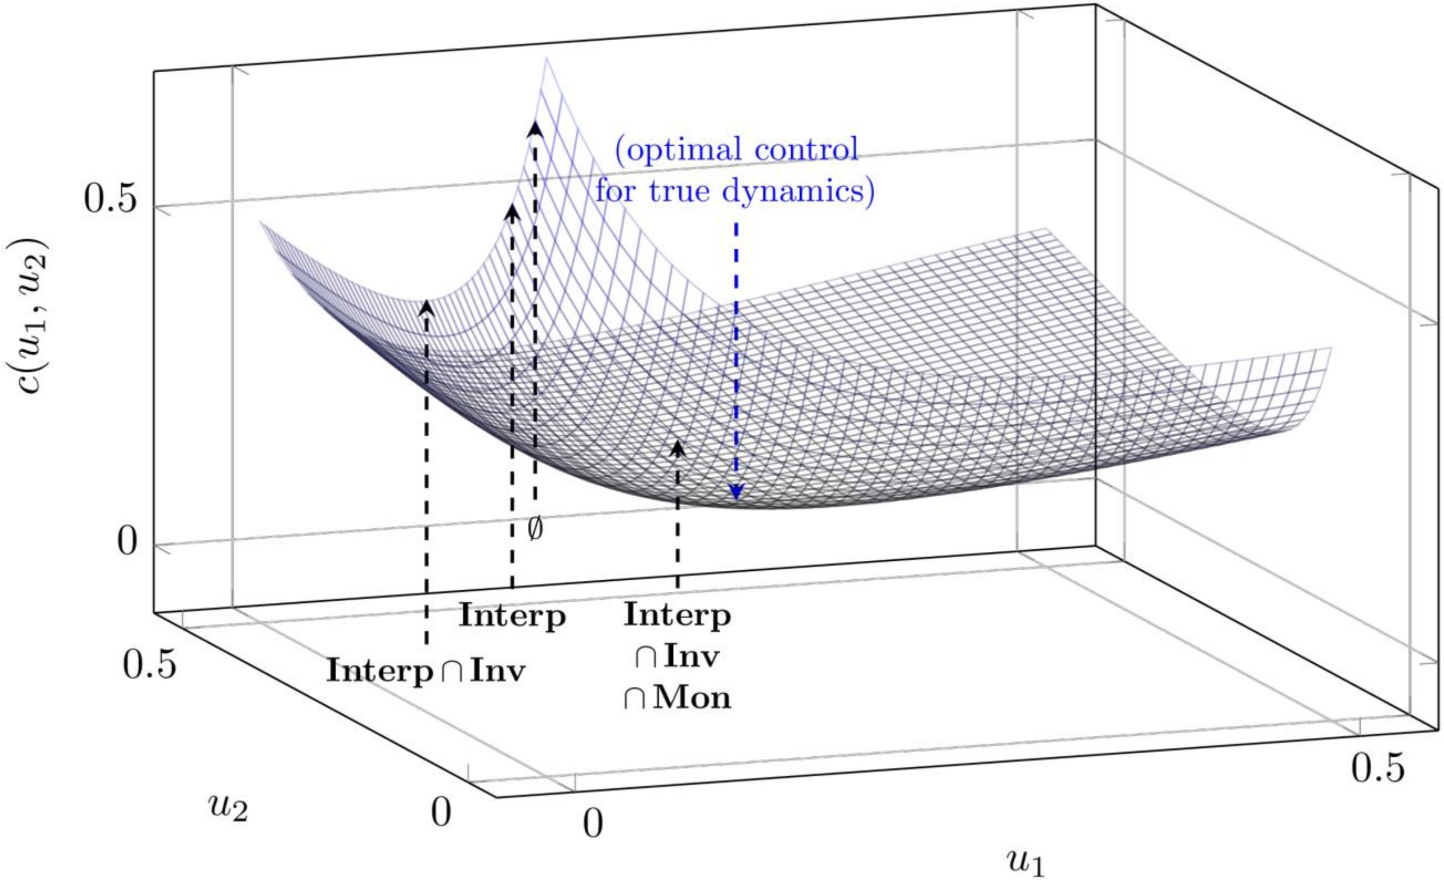

Side information	$x_1(T, \hat{x}_{\text{init}})$	$x_2(T, \hat{x}_{\text{init}})$
None	0.45	0.40
Interp	0.41	0.29
Interp $\cap$ **Inv**	0.31	0.12
Interp $\cap$ **Inv** $\cap$ **Mon**	0.01	0.01
True vector field	0.00	0.00

Side information useful in many other domains

See paper for an example in cell biology, or for learning a chaotic system.

A physical example: **the simple pendulum**

$$(\theta_1, \theta_2) = (\theta, \dot{\theta})$$

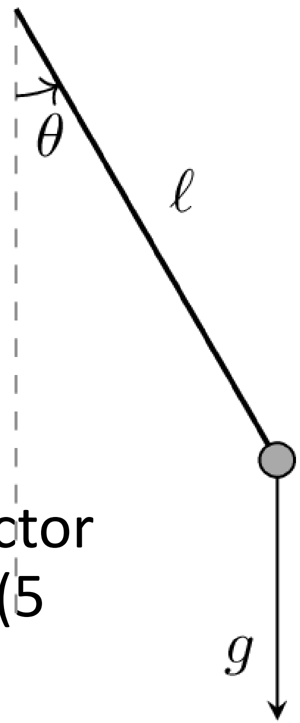

The true dynamics f (unknown)

$$\begin{pmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \end{pmatrix} = f(\theta_1, \theta_2) = \begin{pmatrix} \theta_2 \\ -\frac{g}{l}\sin(\theta_1) \end{pmatrix}$$

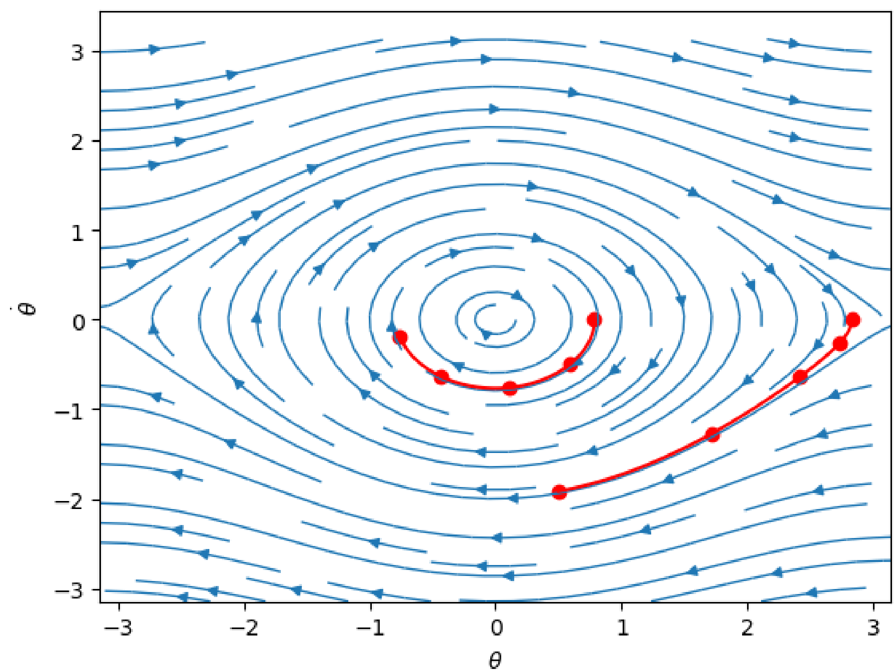

- **What we observe:** Noisy measurements of the vector field on two trajectories (5 noisy points from each)

- **Goal:** Learn a polynomial vector field p that is close to f over $[-\pi, \pi] \times [-\pi, \pi]$

PRINCETON UNIVERSITY ORFE

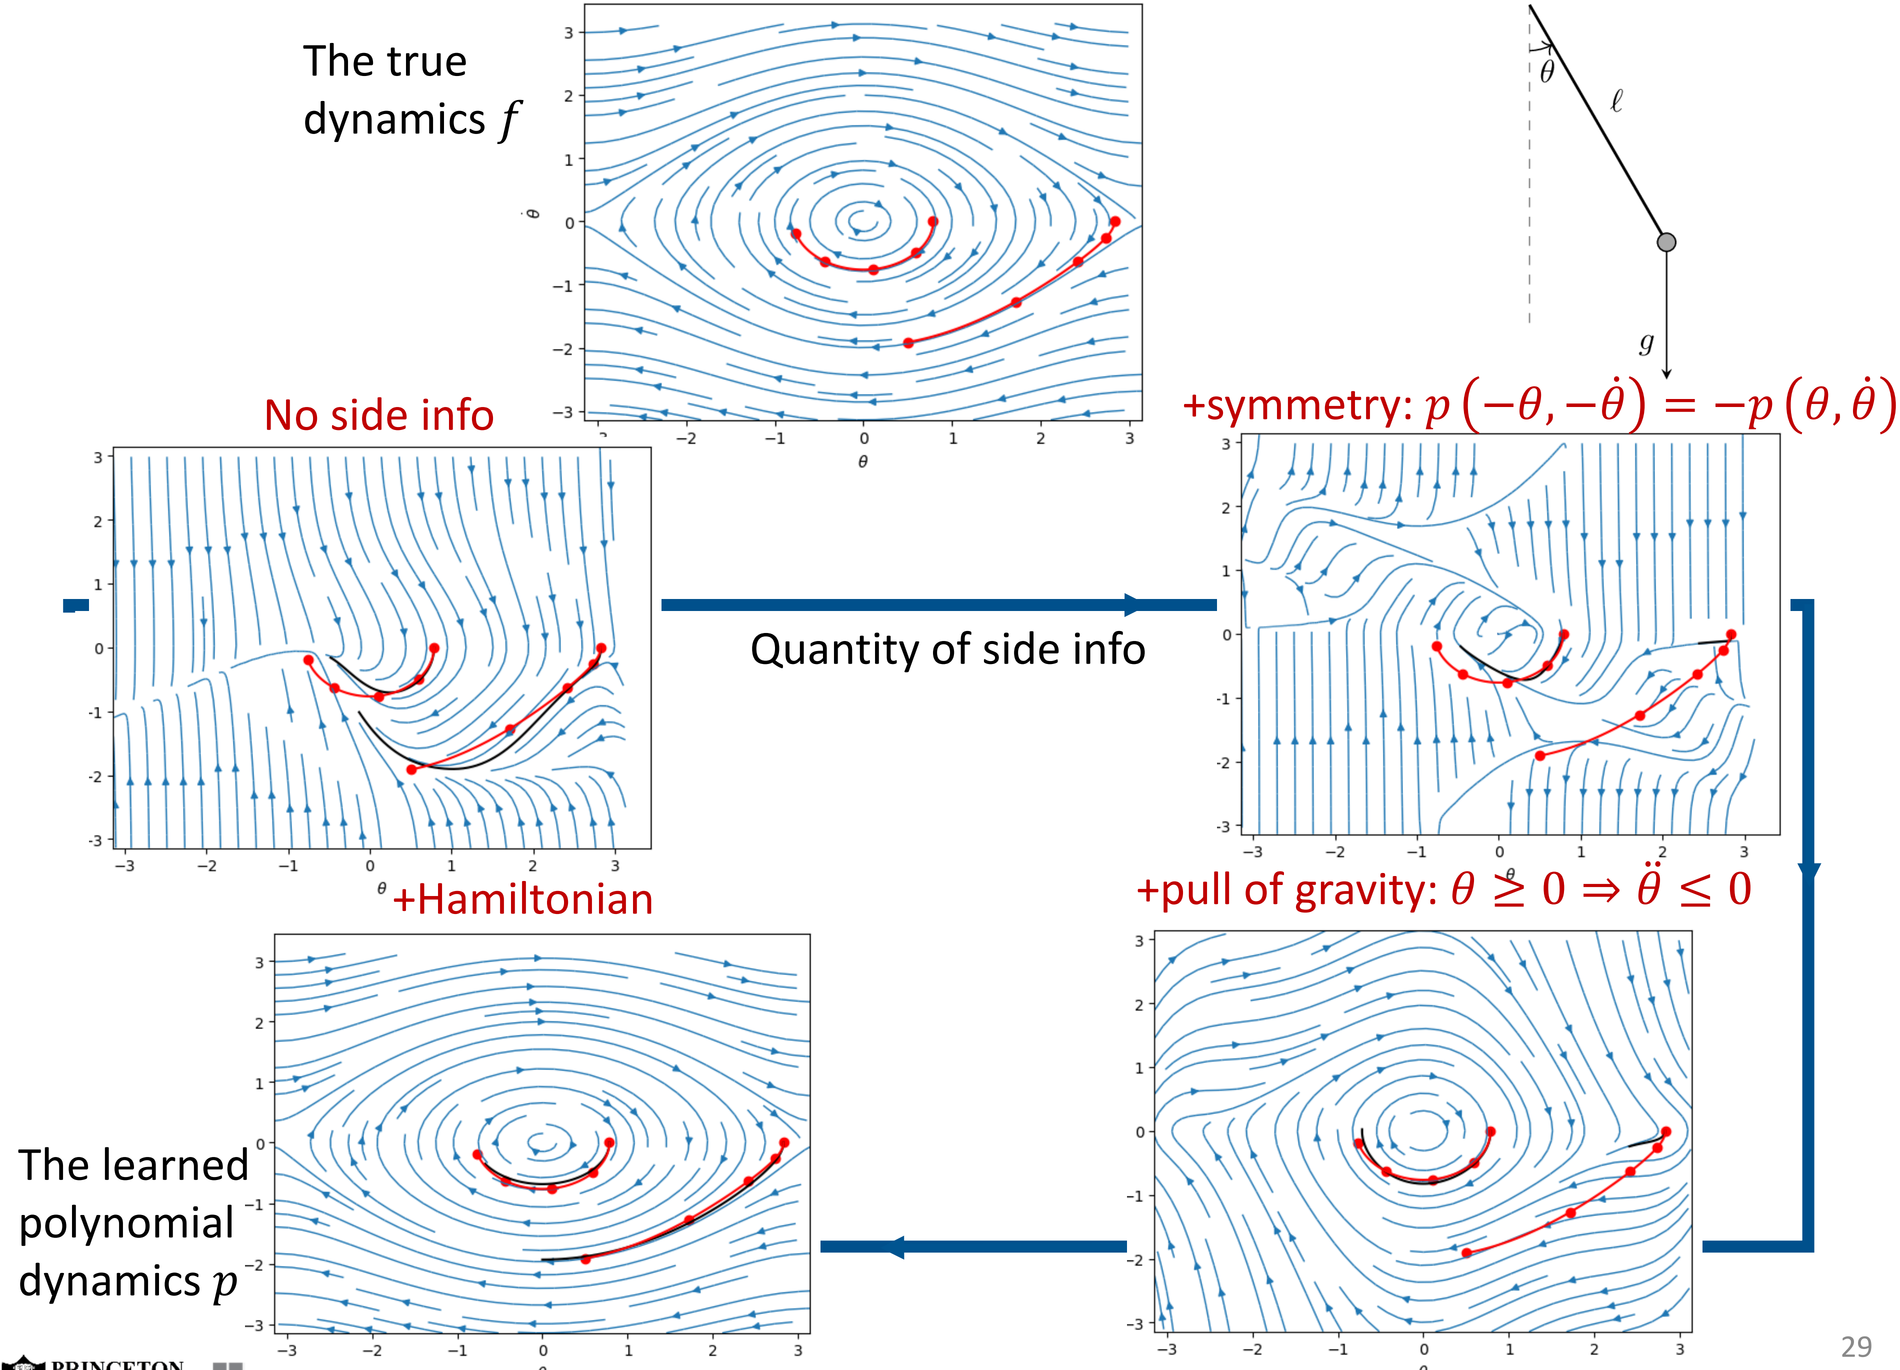
The true dynamics f
No side info
Quantity of side info
+symmetry: $p\left(-\theta,-\dot{\theta}\right)=-p\left(\theta,\dot{\theta}\right)$
+pull of gravity: $\theta \geq 0 \Rightarrow \ddot{\theta} \leq 0$
+Hamiltonian
The learned polynomial dynamics p
θ
ℓ
g

Density of polynomial vector fields satisfying side information

Thm **[AAA, El Khadir].** For time horizon $T > 0$, desired accuracy $\epsilon > 0$, compact set $\Omega \subseteq \mathbb{R}^n$, and any C^1 vector field $f: \mathbb{R}^n \to \mathbb{R}^n$ that satisfies one of the following side information constraints:

i. equilibria at a given finite set of points ($f(v_i) = 0$),
ii. group symmetry ($f(\sigma(g)x) = \rho(g)f(x)\ \forall x \in \Omega, \forall g \in \mathrm{G}$),
iii. invariance of a (full-dimensional) convex set,
iv. directional monotonicity ($\frac{\partial f(x)}{\partial x_j} \geq 0, \forall x \in P_{ij}$),
v. nonnegativity ($f_i(x) \geq 0, \forall x \in P_i$),
vi. gradient or Hamiltonian structure (e.g., $f(x) = -\nabla V(x)\ \forall x \in \Omega$),

there exists a polynomial vector field $p: \mathbb{R}^n \to \mathbb{R}^n$ such that

1) **trajectories of f and p starting from any initial conditions $x_0 \in \Omega$ remain within ϵ for all time $t \in [0, T]$** (while they stay in Ω),

2) **p satisfies the same side information as f.**

What about multiple side information constraints?

Are polynomial vector fields still dense in the space of C^1 vector fields?

In general, no!

$$f(x) := \begin{cases} 0 & x \geq 0 \\ -e^{-\frac{1}{x^2}} & x < 0 \end{cases}$$

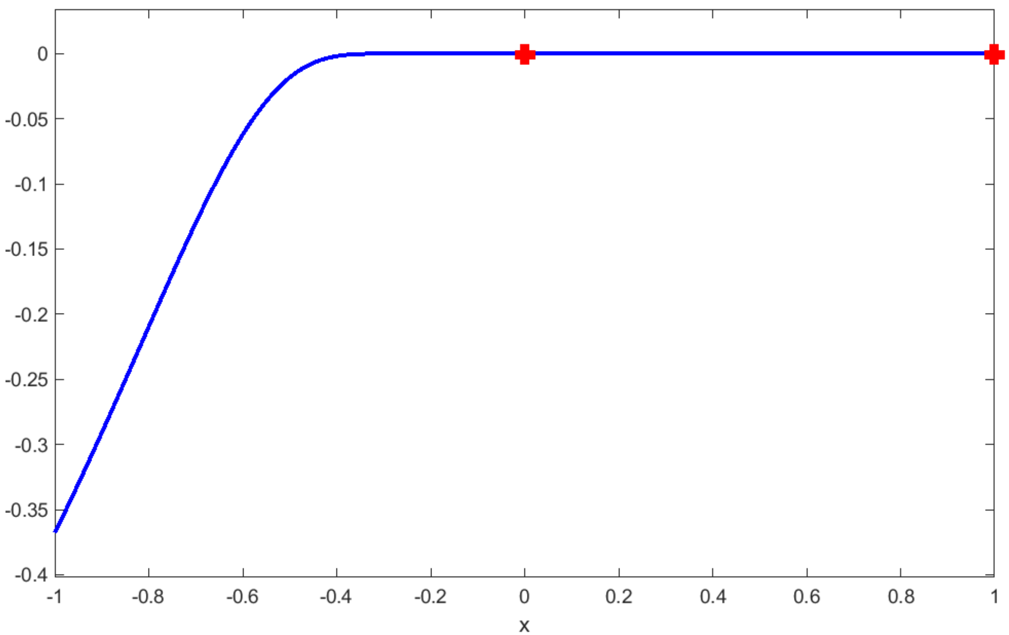

$\Omega = [-1,1]$

Side information:

$f(0) = f(1) = 0$

f nondecreasing over $[-1,1]$

Density with approximate satisfaction of side information

Thm [AAA, El Khadir]. For any continuously differentiable vector field $f: \mathbb{R}^n \to \mathbb{R}^n$, any $T > 0, \epsilon > 0$, $\boldsymbol{\delta > 0}$, and any compact set $\Omega \subseteq \mathbb{R}^n$,

there exists a polynomial vector field $p: \mathbb{R}^n \to \mathbb{R}^n$ such that

1) **trajectories of $\boldsymbol{f}$ and $\boldsymbol{p}$ starting from any initial conditions $\boldsymbol{x_0} \in \boldsymbol{\Omega}$ remain within $\boldsymbol{\epsilon}$ for all time $\boldsymbol{t} \in [\boldsymbol{0}, \boldsymbol{T}]$**(while they stay in Ω),

2) **$\boldsymbol{p}$ $\boldsymbol{\delta}$-satisfies* any combination of the constraints that $\boldsymbol{f}$ satisfies.**

Moreover, all such properties of p come with an **SOS certificate** (and such polynomial vector field can be found by **semidefinite programming**).

***δ-satisfies:**

	$f(v_i) = 0$	$\|\|p(v_i)\|\| \leq \delta$
Invariance: (n(x) is the outgoing normal at the boundary point x)	$\langle n(x), f(x) \rangle \leq 0$	$\langle n(x), p(x) \rangle \leq \delta$
	$\dfrac{\partial f(x)}{\partial x_j} \geq 0$	$\dfrac{\partial p(x)}{\partial x_j} \geq -\delta$
	$f(x) = -\nabla V(x)$	$\left\|\|\mathrm{p}(x) + \nabla W(x)\right\|\| \leq \delta$

etc.

Takeaways

- When data is limited, side information can help you learn better models.
- SOS is a powerful method for imposing side information on polynomial vector fields.

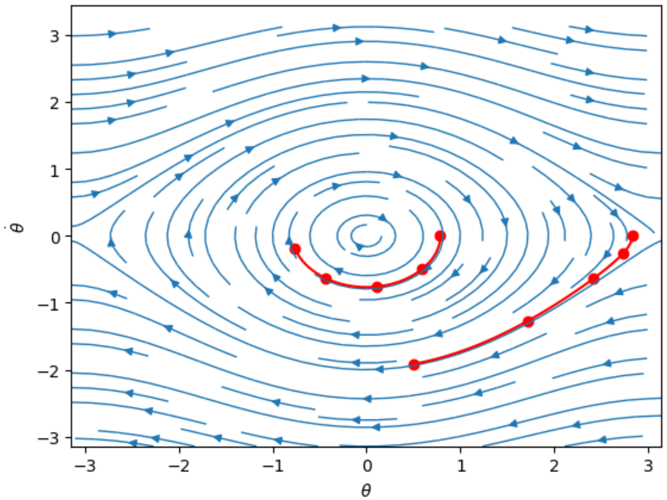

the true dynamics
(not polynomial)

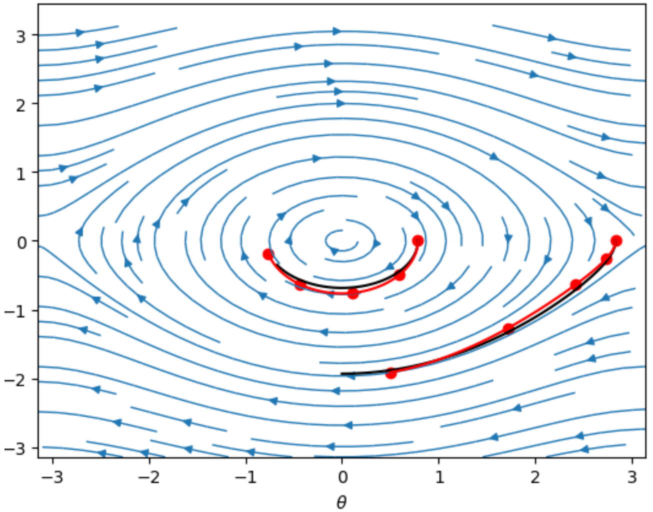

the polynomial dynamics
learned by exploiting side info
using SOS techniques

Safely Learning Dynamical Systems

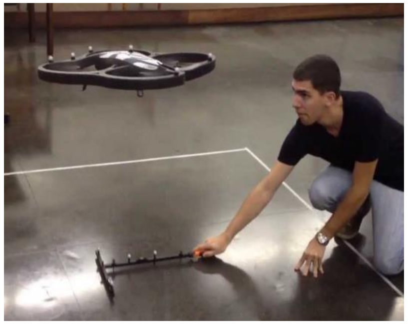

with **Abraar Chaudhry** (Princeton), **Vikas Sindhwani** (Google), **Stephen Tu** (Google)

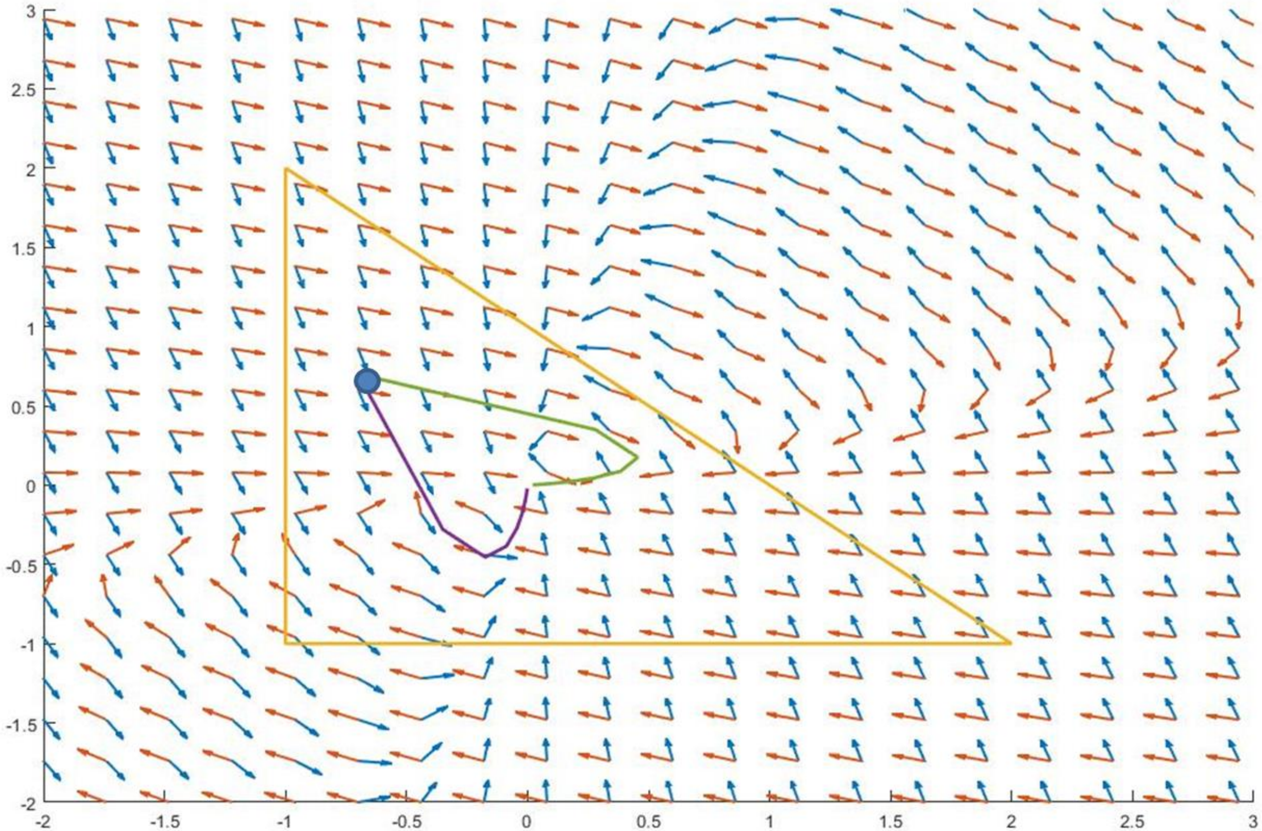

Problem setup

- Our goal is to safely learn a dynamical system

$$x_{t+1} = f_\star(x_t), \quad f_\star : \mathbb{R}^n \to \mathbb{R}^n$$

  by querying appropriate trajectories
- There is a safety region $S \subseteq \mathbb{R}^n$; the state should never go outside of S during our learning process
- The map $f_\star$ is known to be in some initial uncertainty set: $f_\star \in U_0$

- When we query a point $x \in S$, we observe the trajectory of the true dynamics starting from x for $\boldsymbol{T}$ **time steps**
- The "T-step safety region":

$$S^T(f) := \{x \in S \mid f^{(i)}(x) \in S, i = 1, \ldots, T\}$$

- As we (safely) gather more data, the uncertainty over $f_\star$ can shrink

PRINCETON UNIVERSITY ORFE

Safe Learning

- Let $\phi_{f,T}(x)$ represent the trajectory of length T starting from x under the dynamics f
- Suppose we have already queried $\phi_{f_\star,T}(x_j) \quad j = 1, \ldots, k$
- Then we can update out uncertainty set of $f_\star$ as follows:

$$U_k := \{f \in U_0 \mid \phi_{f,T}(x_j) = \phi_{f_\star,T}(x_j)\,, j = 1, \ldots, k\}$$

- With this information, we know that T-step safety is guaranteed for points in the following set: $S_k^T := \bigcap_{f \in U_k} S^T(f)$
- In our learning process we want to accomplish two goals:

1. $\textbf{(Safety)}$ for each $k = 1, \ldots, m$, $x_k \in S_{k-1}^T$,
2. $\textbf{(Learning)}$ $U_m \to \{f_\star\}$

- **For which S, U_0, and T is this possible?**
- If safe learning is possible, may also be concerned with learning with low "query cost"

Conic optimization for safe learning

Content of our paper:

Safe Learning of

1. Linear Systems, One-Step Safety (focus of today)
2. Nonlinear Systems, One-Step Safety→ SOCP (exact)
3. Linear Systems, Two-Step Safety → SDP (exact)
4. Linear Systems, Infinite-Step Safety → SDP (approximation)
5. Nonlinear Systems, Infinite-Step Safety → SDP (approximation)

Linear dynamics, polytopic uncertainty set, 1-step safety

- Safety region is a polytope:

$$S = \{x \in \mathbb{R}^n \mid h_i^T x \leq b_i \quad i = 1, \dots, r\}$$

- We want safety for one time step, i.e. $T = 1$
- Dynamics linear, i.e. $f(x) = A_\star x$ ($A_\star \in \mathbb{R}^{n \times n}$ unknown)
- Polytopic initial uncertainty set:

$$A_\star \in U_0 = \{A \in \mathbb{R}^{n \times n} \mid \mathrm{Tr}(V_j^T A) \leq v_j \quad j = 1, \dots, s\}$$

- Linear query cost: $c^T x$

Definition of Safe Learning for this specific case

Can we devise an algorithm to safely learn whenever possible and otherwise certify the impossibility of safe learning?

Recall:

$$U_k = \{A \in U_0 \mid Ax_l = y_l \quad l = 1, \ldots, k\}$$

$$S_k^1 = \{x \in S \mid Ax \in S, \forall A \in U_k\}$$

We say that one-step safe learning is possible if for some nonnegative integer m, we can sequentially choose vectors $x_k \in S$, for $k = 1, \ldots, m$, and observe measurements $y_k = A_\star x_k$ such that:

1. **(Safety)** for $k = 1, \ldots, m$, we have $Ax_k \in S \quad \forall A \in U_{k-1}$,
2. **(Learning)** the set of matrices U_m is a singleton.

Searching over one-step safe region

- We need a way to search over S_k^1
- Suppose we have thus far collected pairs (x_l, y_l) with $y_l = A_\star x_l$
- The next cheapest one-step safe point is a solution to the following program:

$$\begin{aligned} \min_x \quad & c^T x \\ \text{s.t.} \quad & h_i^T x \leq b_i \quad i = 1, \ldots, r \\ & \begin{bmatrix} \max_A & h_i^T A x \\ \text{s.t.} & \mathrm{Tr}(V_j^T A) \leq v_j \quad j = 1, \ldots, s \\ & A x_l = y_l \quad l = 1, \ldots, k \end{bmatrix} \leq b_i \quad i = 1, \ldots, r. \end{aligned}$$

- Can we solve this problem?

Duality

- Let us take the dual of the inner problems:

$$\begin{bmatrix} \max_A & h_i^T A x \\ \text{s.t.} & \mathrm{Tr}(V_j^T A) \leq v_j \; j = 1, \ldots, s \\ & A x_l = y_l \; l = 1, \ldots, k \end{bmatrix} = \begin{bmatrix} \min_{\mu^{(i)}, \eta^{(i)}} & \sum_{l=1}^k y_l^T \eta_l^{(i)} + \sum_{j=1}^s \mu_j^{(i)} v_j \\ \text{s.t.} & x h_i^T = \sum_{l=1}^k x_l \eta_l^{(i)T} + \sum_{j=1}^s \mu_j^{(i)} V_j^T \\ & \mu^{(i)} \geq 0 \end{bmatrix}$$

- We can proceed by taking the dual of all the inner problems
- Turns min/max into a min/min → overall problem becomes an LP

LP formulation for searching over the set S_k^1 :

$$\begin{aligned} \min_{x, \mu, \eta} \quad & c^T x \\ \text{s.t.} \quad & h_i^T x \leq b_i \quad i = 1, \ldots, r \\ & \sum_{l=1}^k y_l^T \eta_l^{(i)} + \sum_{j=1}^s \mu_j^{(i)} v_j \leq b_i \quad i = 1, \ldots, r \\ & x h_i^T = \sum_{l=1}^k x_l \eta_l^{(i)T} + \sum_{j=1}^s \mu_j^{(i)} V_j^T \quad i = 1, \ldots, r \\ & \mu^{(i)} \geq 0 \quad i = 1, \ldots, r. \end{aligned}$$

Naïve Algorithm

- Now that we can search over S_k^1 , can we just solve this LP iteratively?
- The algorithm can go like this:

1. Initialize Data = []
2. If U_k is a singleton, return $A_\star$
3. Solve LP with Data, get optimal x_k
4. Query x_k and observe $y_k := A_\star x_k$
5. Add (x_k, y_k) to data
6. Goto 2

- Problem: this may "stall" i.e. the optimal x from the LP may not add information
- This happens when the optimal x is in the linear span of our previous x's

Revised Algorithm

1. Initialize Data = []
2. If U_k is a singleton, return
3. Solve LP with Data, get optimal x_k
4. If x_k is in the linear span of previous x's, add a "valid" perturbation to x_k, if there is no valid perturbation, return safe learning impossible
5. Query x_k and observe
6. Add (x_k, y_k) to data
7. Goto 2

- Perturbation is "valid" if it maintains one-step safety and lowers $\dim(U_k)$
- Can test whether there is a valid perturbation by finding a basis for $\mathrm{span}(S_k^1)$ that is inside S_k^1; this can be done with some LPs

Theorem: Safe Learning is possible iff the above algorithm succeeds.

Corollary: If safe learning is possible, then it is possible with at most n queries.

One-step safety – a numerical experiment

- $S = \{x \in \mathbb{R}^4 \mid \|x\|_\infty \leq 1\}$
- $U_0 = \{A \in \mathbb{R}^{4\times 4} \mid |A_{ij}| \leq 4 \quad \forall i,j\}$
- It takes 4 steps to learn $A_\star$ exactly
- The (projection of) optimal query points are marked
- The last (largest) set is the "true 1-step safety region", i.e.

$$S^1(A_\star) = \{x \in S \mid A_\star x \in S\}$$

$$A_\star = \begin{bmatrix} 2 & 1 & 4 & 2 \\ 2 & -3 & -1 & -2 \\ -2 & -3 & 1 & 0 \\ 2 & 0 & -2 & 2 \end{bmatrix}$$

$$c = (-1,-1,0,0)^T$$

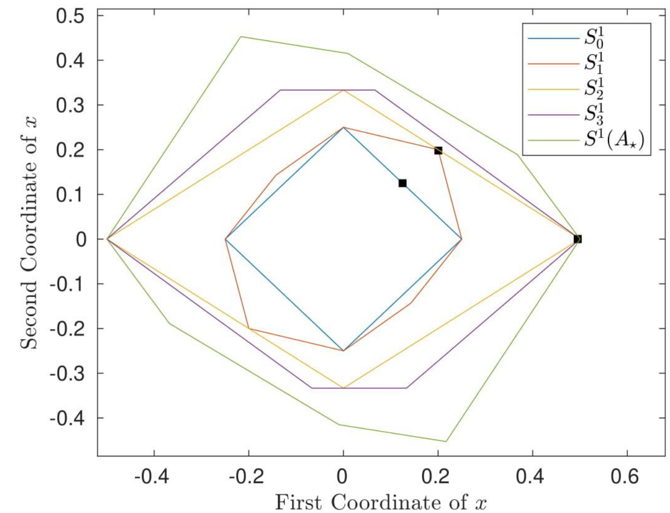

- Cost of safe learning $\sum_{k=1}^{4} c^T x_k$: -1.6385
- Cost of safe learning without exploiting information on the fly: -1.0000
  - This comes from optimizing over S_0^1
- Lower bound: -2.2264
  - This comes from optimizing over $S^1(A_\star)$

PRINCETON UNIVERSITY ORFE

Uncertainty shrinks over time, safety sets grow

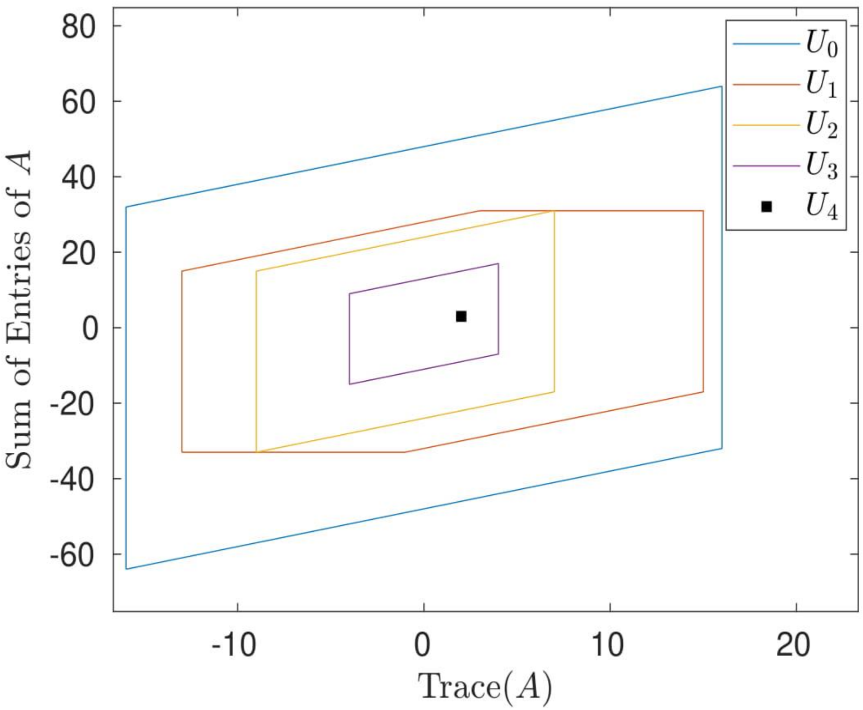

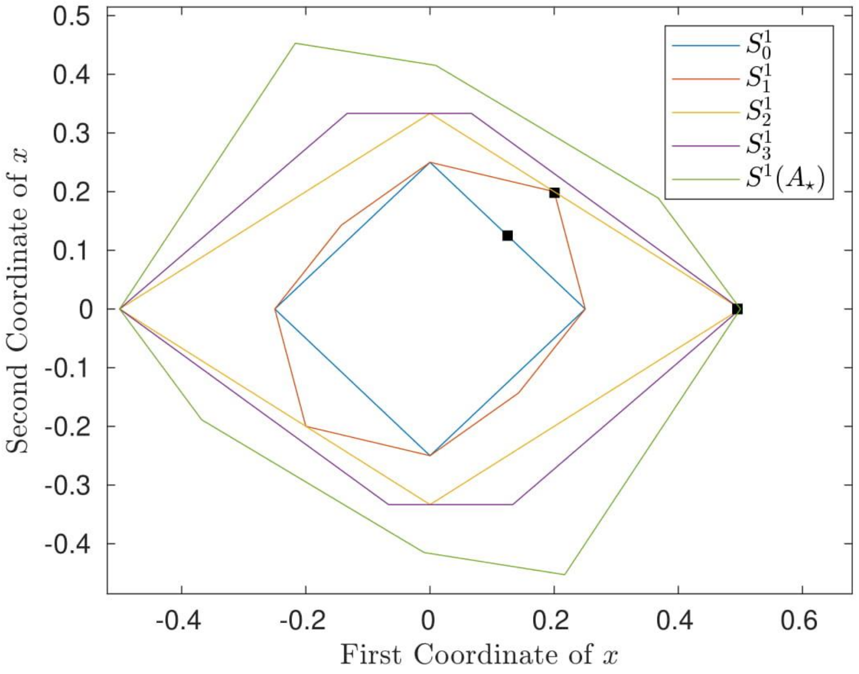

In summary…

- There are interesting optimization problems at the interface of learning and dynamical systems!
- Conic programming (especially semidefinite optimization) proves to be a powerful tool in this area.
- We saw two examples of this today for handling safety constraints or exploiting side information.

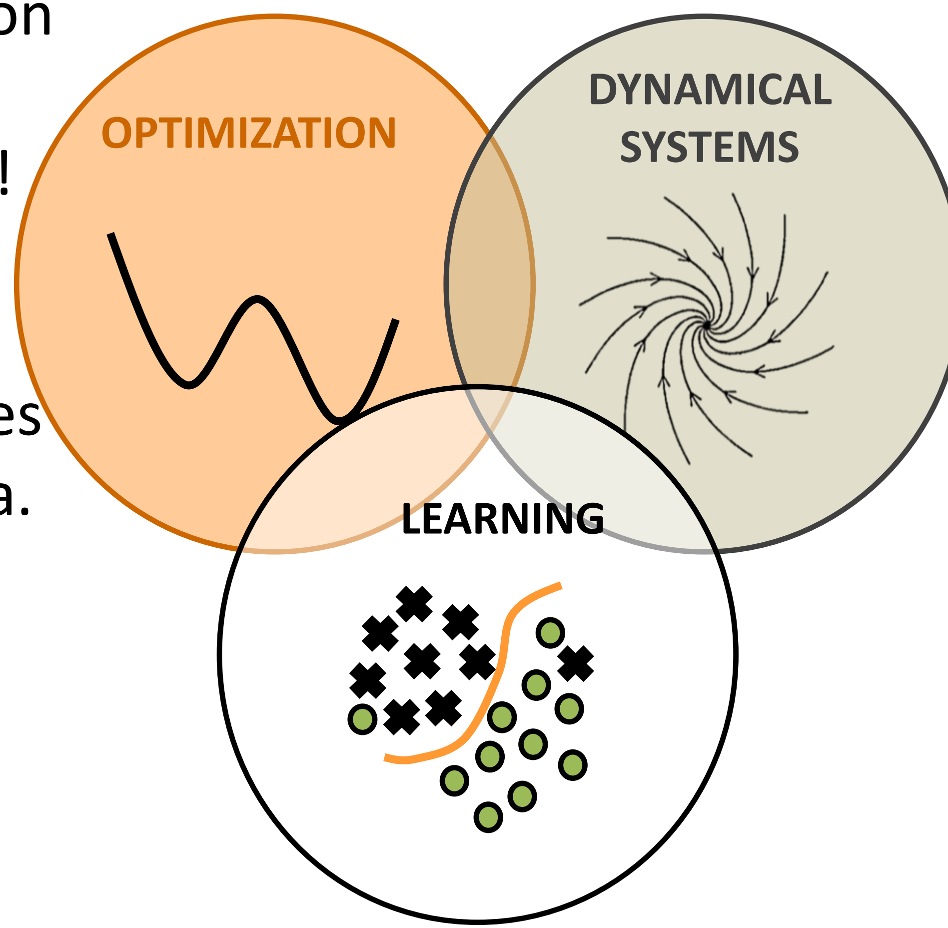

Thank you for listening.

Want to know more? http://aaa.princeton.edu