Glimmers of Autonomy: Structure-Aware Reachability Analysis and Control Synthesis for Unknown Nonlinear Systems

Melkior Ornik

ILLINOIS

AFOSR DC Review (FA9550-23-1-0131)
August 28, 2024

Melkior Ornik – mornik@illinois.edu

Maximal Understanding in the Face of Minimal Knowledge

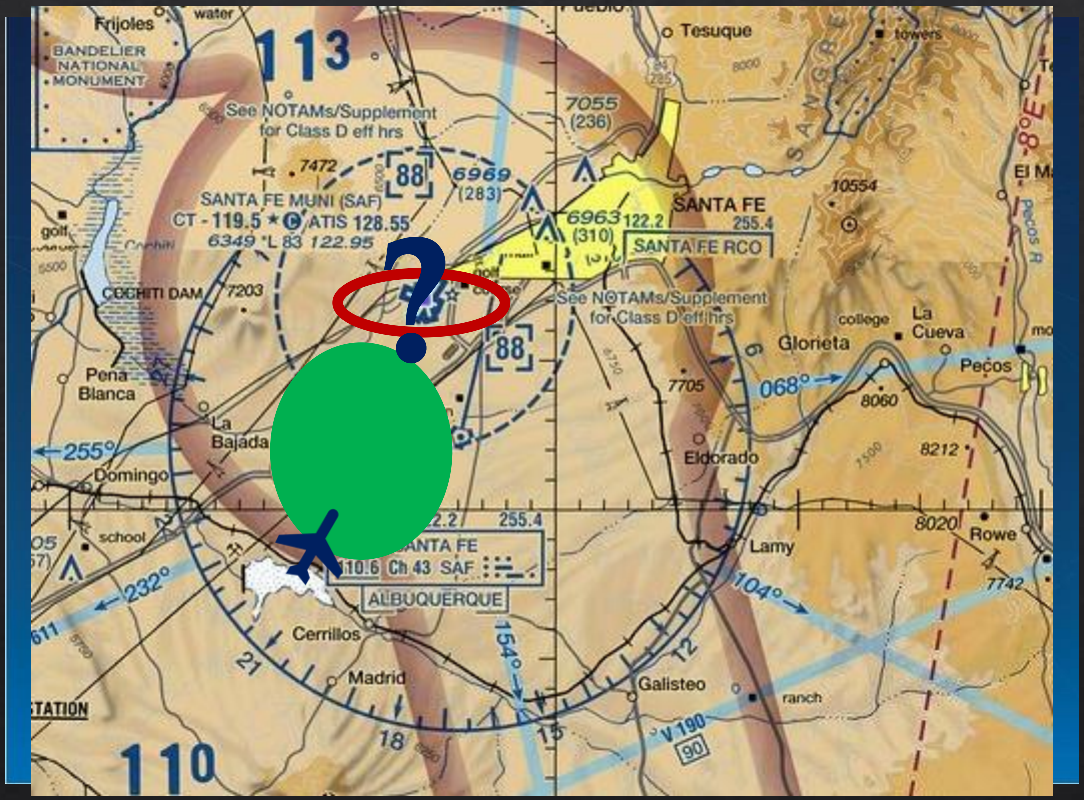

Adaptive / Robust Control

Reach the objective under (some) lack of knowledge of dynamics

Objective might not be reachable

First decide on a reachable target, *then* plan how to get there

Certifiable Capabilities

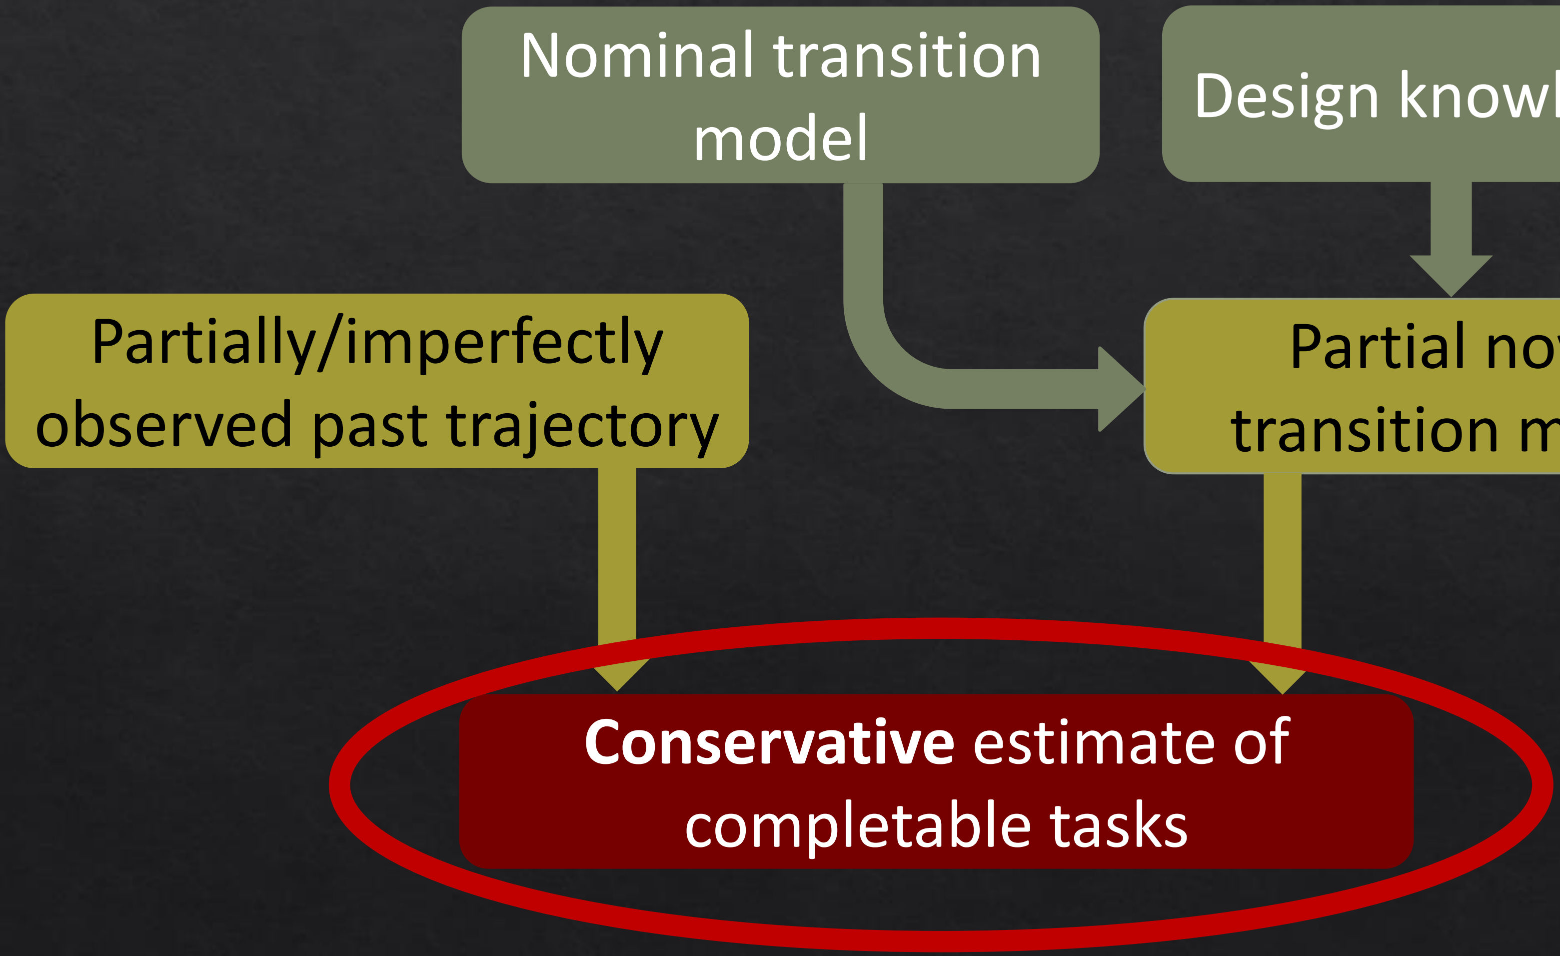

Before/during **provably learning online** how to complete a task (and ideally getting less conservative)

Disasters

$$\dot{x} = f(x, u), \quad u \in U$$

Change in dynamics (e.g., physical damage)

Partial loss of control (e.g., adversarial takeover)

Actuator degradation

$$\dot{x} = \hat{f}(x, u, v), \quad (u, v) \in \hat{U}$$

Loss of Control Authority

$$\dot{x} = f(x, u, v), \quad (u, v) \in U$$

Uncontrolled system input
(not a disturbance)

Two-player game:

P1 (“controller”) wishes to reach a state
P2 (“environment”/”adversary”) wishes to obstruct P1

Can P1 win? (Can P1 win for any state? Can P1 win if there is a time limit?)

“resilience”

Quantitative Resilience

Intuitively, a system is resilient if a player can counteract the adversary and still have some meaningful control authority

If the target state is in the guaranteed reachable set of the initial state (*guaranteed* with respect to all possible adversarial control inputs):

it also matters how hard it is to reach the state compared to the system with nominal dynamics.

$$r_q = \frac{\inf_{\bar{u}} T_R^{\bar{u}}(x_0, x_{goal})}{\sup_{v} \inf_{u} T_R^{(u,v)}(x_0, x_{goal})}$$

The adversary chooses an input such that whatever the controller chooses, performance will be bad

Reach-Time Resilience

First idea: measure of performance – reach time

$$r_q = \frac{\inf_{\bar{u}} T_R^{\bar{u}}(x_0, x_{goal})}{\sup_v \inf_u T_R^{(u,v)}(x_0, x_{goal})}$$

Resilience quotient = 0: no resilience; **resilience quotient = 1:** perfect resilience

Determining minimal time to reach a target even for nominal linear dynamics with input constraints is not simple *(see last year)*

Energy Resilience

Determining minimal energy to reach a target (soft input constraints) for nominal dynamics $\underbrace{\text{is simple:}}_{\text{for linear systems}}$ controllability Gramian

Step 0: Determine worst-case minimal energy for reachability at a particular time for *disturbed dynamics*

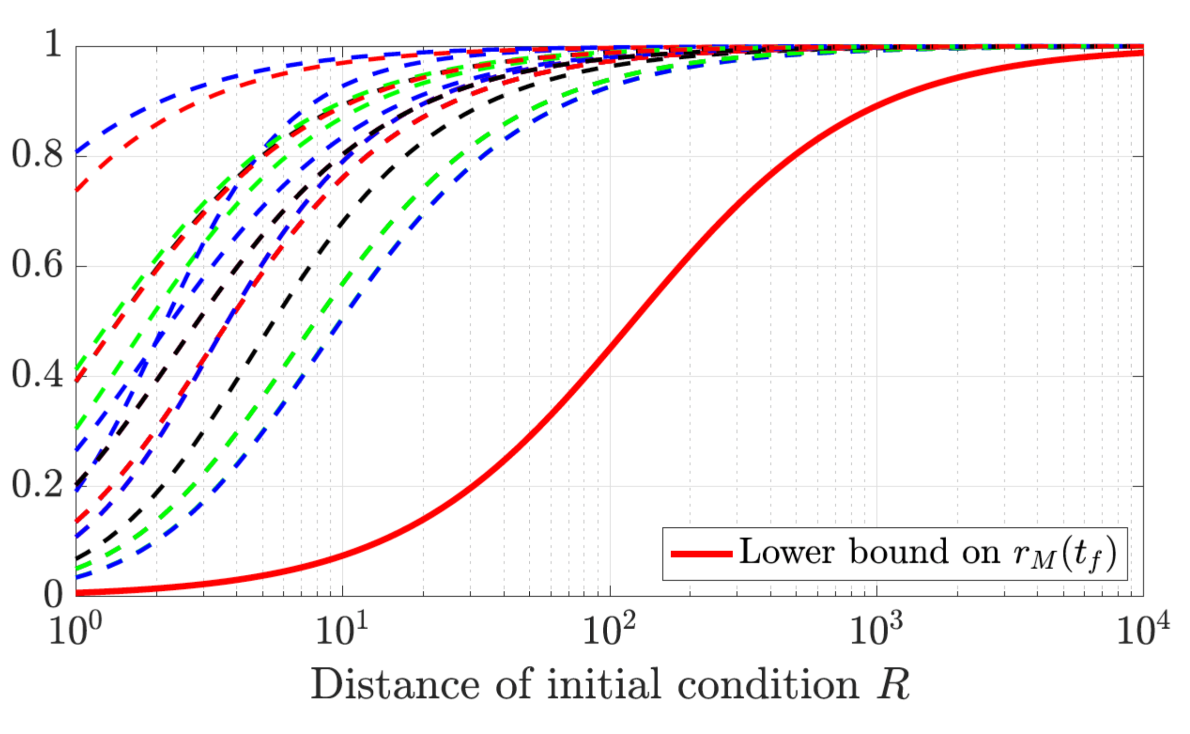

Slight problem: quotient of energies goes to zero the closer we start to equilibrium (player constantly fights the adversary) – difference of energies

R. Padmanabhan, C. Bakker, S. A. Dinkar, M. Ornik. How much reserve fuel: Quantifying the maximal energy cost of system disturbances. ArXiv preprint arXiv:2408.10913 [eess.SY], 2024.

Energy Resilience for Driftless Systems

"Disturbances" in our scenario have structure, *and* faulty/hostile actuators still expend energy

For d

- O
- W

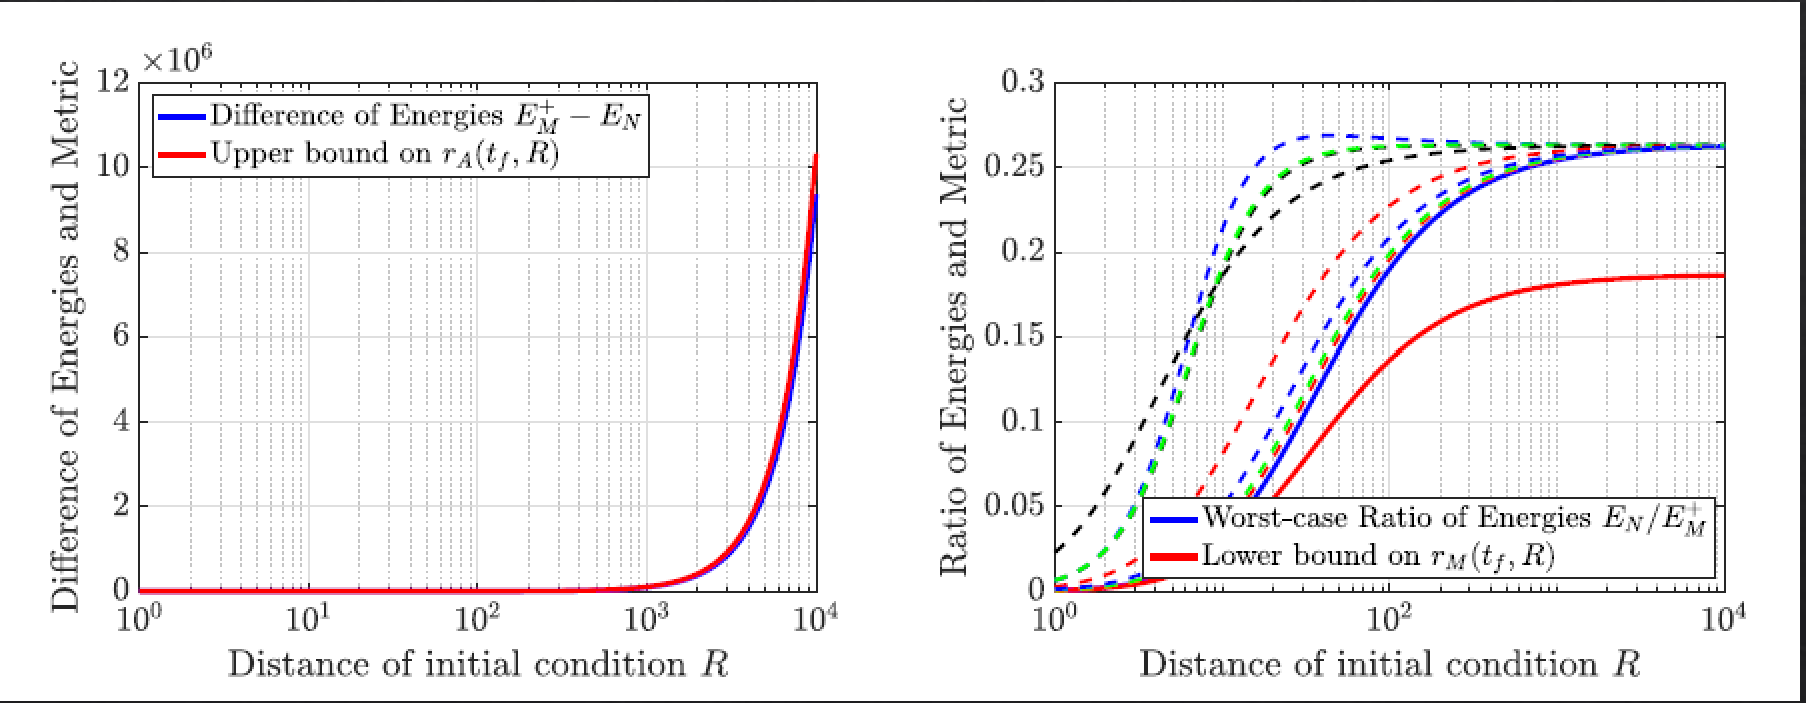

The stronger the adversary, the larger the energy
The stronger the controller, the smaller the energy

Towards Nonlinearity

Driftless → Almost driftless → Almost linear → Nonlinear

Driftless → General linear → Almost linear

Almost driftless: growth or magnitude bound on a nonlinear drift term

- Idea: difficult to express optimal control input, possible to bound it *(first results this week)*

Nonlinear (on a $\underbrace{\text{compact set)}}_{\text{Switched systems?}}$ = linear + bound on a nonlinear term

Melkior Ornik – mornik@illinois.edu

Disasters

$$\dot{x} = f(x, u), \quad u \in U$$

Change in dynamics
(e.g., physical damage)

Partial loss of control
(e.g., adversarial takeover)

$$\dot{x} = \hat{f}(x, u, v), \quad (u, v) \in \hat{U}$$

Guaranteed Reachability on Manifolds

$$\dot{x} = f^{?}(x, u), \quad u \in U$$

$\dot{x} \in T_x\mathcal{M}$

$x \in \mathcal{M}$

Physical knowledge,
a priori safety constraint?

After a change in dynamics, the system might not be able to reach its target using **any** control law

What is it **certifiably capable** of doing (even if we don't yet know how)?

$$R^{\mathcal{G}}(x_0) = \bigcap_{\tilde{f} \in D_{con}} R^{\tilde{f}}(x_0)$$

T. Shafa, M. Ornik. Guaranteed reachability on Riemannian manifolds for unknown nonlinear systems. ArXiv preprint arXiv:2404.09850 [eess.SY], 2024.

Melkior Ornik – mornik@illinois.edu

Idea in Euclidean Space

- In theory, guaranteed reachability set is well-defined
- In practice, how do compute it?

Idea: By finding all trajectories obtained by integrating **guaranteed velocities**

Difficult, but one level easier to underapproximate

$V^{\mathcal{G}}(t) = \bigcap_{\tilde{f} \in D_{con}} V^{\tilde{f}}(t)$, we will obtain at least some (if not all) guaranteed reachable states

Online “local System ID”

Velocities guaranteed at time = ~~velocities guaranteed at start time~~ “modulo” ~~maximal system wildness~~

“Physics and design” = Lipschitz

Melkior Ornik – mornik@illinois.edu

Challenge with Manifolds

- Velocities guaranteed at different times are no longer in the same space ($T_x M$)
- Lipschitz constant on M is nontrivial to define

Let V be a continuous vector field on M and τ be the parallel transport. Then L is the *classical Lipschitz constant* on V if

$$L = \sup_{\gamma} \frac{|\tau_\gamma V(\gamma(0)) - V(\gamma(1))|_{h_x}}{\mathrm{Length}(\gamma)}$$

where $\gamma : [0,1] \to M$ varies over all $\mathcal{C}^1$-paths and τ_γ is shorthand for the parallel transport along the curve γ from $\gamma(0)$ to $\gamma(1)$.

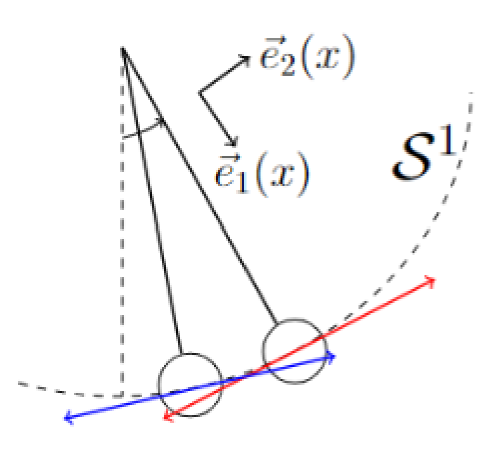

Why not just embed everything into Euclidean space?

- Doesn't feel right
- Worse Lipschitz
- **No full actuation**

T. Shafa, M. Ornik. Guaranteed reachability on Riemannian manifolds for unknown nonlinear systems. ArXiv preprint arXiv:2404.09850 [eess.SY], 2024.

Velocities on Manifolds

Idea still the same, but underapproximation of guaranteed velocities will depend on the:

Riemannian metric tensor/
Covariant derivatives/
Choice of connection

For flat manifolds, we recover previous results for Euclidean spaces

Theorem 1. *Let $f(x_0), G(x_0), L_f, L_G, H_x, \Gamma^k_{ij}$, and g^Γ_l for $l \in [m]$ be defined as above. Let $\gamma : [0,1] \to M$ define a geodesic curve from x_0 to x. Let $\tilde{\tau}$ define the parallel transport using the flat connection. Set $a(x) = (\|H_x^{-1}\|\|H_x\|)^{\frac{1}{2}}\|H_x\|^{\frac{1}{2}}\left\|\begin{bmatrix} g^\Gamma_1 & \dots & g^\Gamma_m \end{bmatrix}\right\|$, $b(x) = (\|H_x^{-1}\|\|H_x\|)^{\frac{1}{2}}\left\|\sum_{i,j,k}\dot{\gamma}^i\Gamma^k_{ij}f^j(x_0)\vec{e}_k\right\|$, $c(x) = (\|H_x^{-1}\|\|H_x\|)^{\frac{1}{2}}\left(L_g + \|H_x\|^{-\frac{1}{2}}L_f\right)d(x_0,x)$ and*

$$\overline{d}(x_0,x) = \frac{\|\tilde{\tau}^x_{x_0}G^\dagger(x_0)\|^{-1} - a(x) - b(x)}{c(x)}.$$

If $d(x_0,x) \le \overline{d}(x_0,x)$,

$$\overline{\mathcal{V}}^{\mathcal{G}}_x = \mathbb{B}^n\left(\tilde{\tau}^x_{x_0}f(x_0); \alpha(x_0,x)\right) \cap \mathrm{Im}(\tilde{\tau}^x_{x_0}G(x_0))$$

where $\overline{\mathcal{V}}^{\mathcal{G}}_x \in T_xM$, and

$$\begin{gathered}\alpha(x_0,x) = \|\tilde{\tau}^x_{x_0}G^\dagger(x_0)\|^{-1} - \\ (\|H_x^{-1}\|\|H_x\|)^{\frac{1}{2}}\Bigg(\|H_x\|^{\frac{1}{2}}\left\|\begin{bmatrix} g^\Gamma_1 & \dots & g^\Gamma_m \end{bmatrix}\right\| + \\ \left\|\sum_{i,j,k}\dot{\gamma}^i\Gamma^k_{ij}f^j(x_0)\vec{e}_k\right\| + \left(L_g + \|H_x\|^{-\frac{1}{2}}L_f\right)d(x_0,x)\Bigg),\end{gathered}$$

then $\overline{\mathcal{V}}^{\mathcal{G}}_x \subseteq \mathcal{V}^{\mathcal{G}}_x$.

T. Shafa, M. Ornik. Guaranteed reachability on Riemannian manifolds for unknown nonlinear systems. ArXiv preprint arXiv:2404.09850 [eess.SY], 2024.

Numerical Example

Operating on $SO(3)$: difficult to meaningfully think about in Euclidean space, possibly useful for applications, and difficult to even draw the guaranteed reachable set

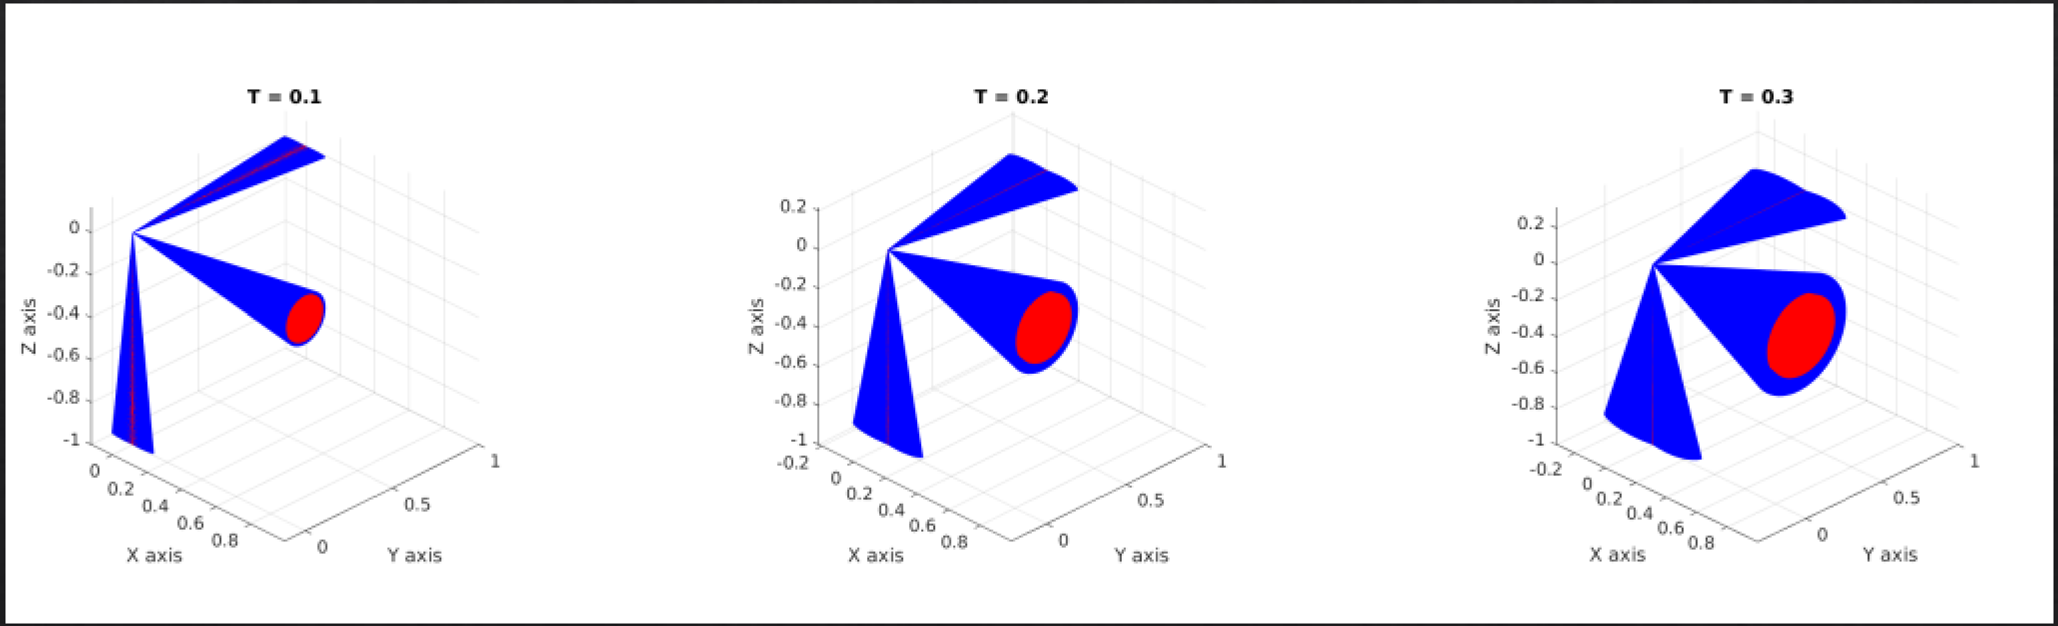

T. Shafa, M. Ornik. Guaranteed reachability on Riemannian manifolds for unknown nonlinear systems. ArXiv preprint arXiv:2404.09850 [eess.SY], 2024.

Learn-Control Pipeline: "Glimmers of Autonomy"

Once we have established what the system is capable of doing, we still do not know *how* to do it

- **End-to-end planning, learning and control**:
  - **Task assignment**: what task can be provably completed
  - **Real-time learning**: what do we need to know in order to be able to complete it: persistent excitation allows us to learn local dynamics
  - **Assured control:** complete the task

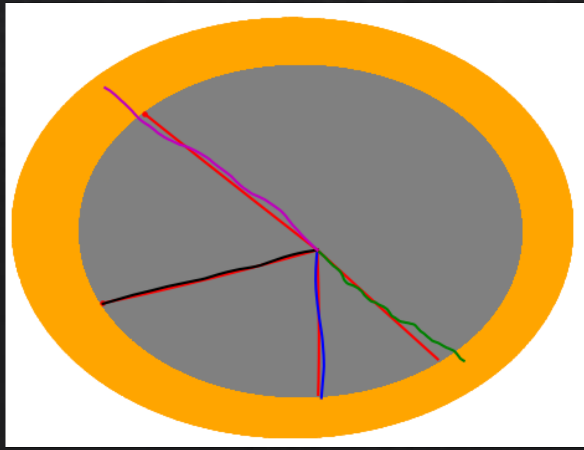

Melkior Ornik – mornik@illinois.edu

Previous Work: Online "Local System ID"

When the system is at state x, applying a constant control u^+ and observing the system response provides an approximation for $f(x, u^+)$

Control-affine systems: quickly sequentially applying affinely independent controls provides an approximation for the system dynamics at x

Let ϕ be the system trajectory. Let $u^0, u^1, \ldots, u^m$ be affinely independent with $\|u^j\| \leq \delta$ for all j. Define $x_0 = x = \phi(T)$, and $x_j = \phi(T + j\varepsilon)$, where input $u(t) = u^j$ is applied for $t \in [T + j\varepsilon, T + (j+1)\varepsilon)$. Then

$$\left\| f(x, \bar{u}) - \sum_{j=0}^{m} \lambda_j \frac{x_{j+1} - x_j}{\varepsilon} \right\| \leq K\varepsilon \frac{4m^{3/2} + \delta}{\delta},$$

for all $\bar{u}$, where $\bar{u} = \lambda_0 u^0 + \ldots + \lambda_m u^m$.

Guarantees in Online Planning

Previously: If we know the local dynamics, we can choose a control input that **appears to** work well right now. Goodness function encodes trajectory quality: the higher its value, the better the direction of the system.

No guarantees! How to even choose the goodness function?

Step 1: choose a target guaranteed to be reachable

Step 2: guaranteed reachability computations also produce a **guaranteed trajectory**

Step 3: **approximately** follow the guaranteed trajectory

Control Design

Step 3a: perform online learning (a previous result bounds the distance from the trajectory)

Step 3b: choose a point on the trajectory and use a control input that currently approximately (with known error bounds) moves the system towards that point

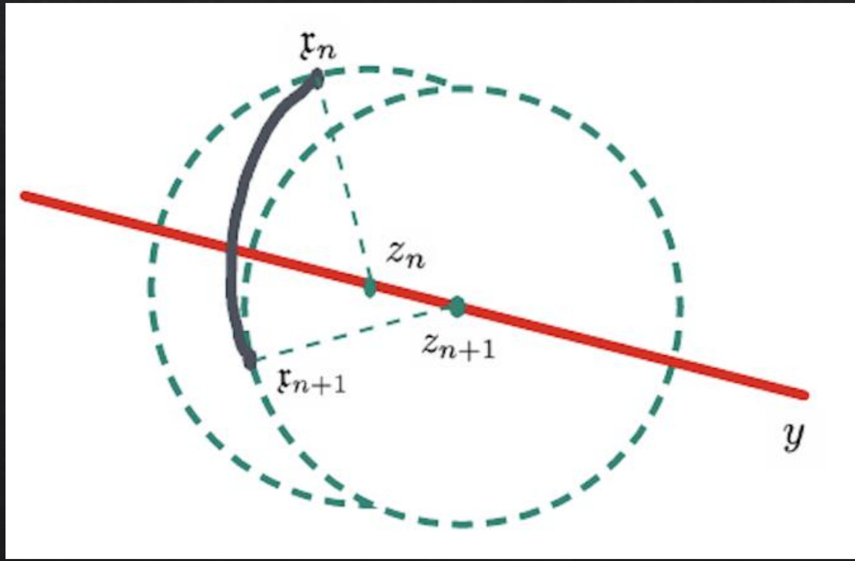

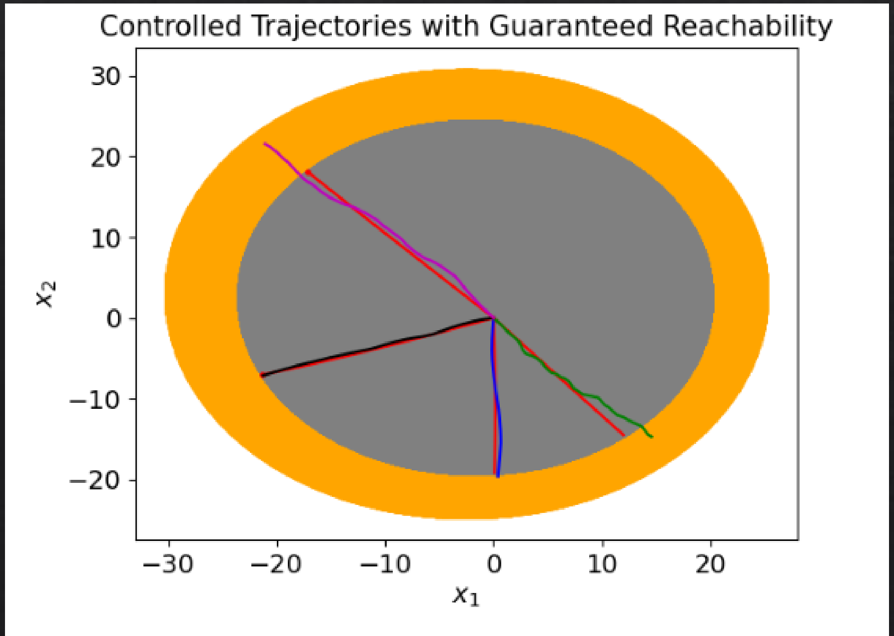

Y. Meng, T. Shafa, J. Wei, M. Ornik. Online learning and control synthesis for reachable paths of unknown nonlinear systems. ArXiv preprint arXiv:2403.03413 [math.OC], 2024.

Teaser: Learning from Multiple Partial Trajectories

Multiple agents operating at the same time can collect data in parallel.

Can we obtain information about global (or “less local”) dynamics?

Arbitrarily many agents sampling arbitrarily close: trivial

Bounded number of agents:

- **Where to place them?**
- **Interpolation with error bound quantification**

With good placement, uncertainty does not grow indefinitely – long-term guarantees?

Teaser: Complicated Objectives

Only objective currently: **reachability/stabilization**

In progress (in some sense): **safety, trajectory tracking**

As means towards reachability, through self-designed **waypoints**

More complicated missions: **hybrid** (in some sense; even if the dynamics are not hybrid, the reachability specification/coordinate system might change)

First approach: set waypoints such that at each waypoint, the next waypoint is provably reachable (local guarantees, global heuristics)

Medium-Term Goal (Last Year): Capabilities

“Almost driftless” with disturbance

- **Combination of scenarios**: e.g., partially unknown dynamics with partial loss of control (Some work already done: actuator degradation with disturbances)

- **End-to-end planning, learning and control**:
  - **Task assignment**: what task can be provably completed
  - **Real-time learning**: what do we need to know in order to be able to complete it
  - **Assured control:** complete the task

Multiple trajectories, manifolds

- (More) **physics-based and design-based results**: exploitation of significant unchanged prior knowledge for better estimation

Melkior Ornik – mornik@illinois.edu

Long-Term Goals: Validation

- **Using sensors and perception** to recognize fault type and gather information
  - Fault detection; sensor fusion; state estimation
- **Complex missions** in high-fidelity simulation
  - Concurrent sensing and actuation faults
  - Noise+hostile action
- **Onboard implementation**
  - Time-delayed control
  - Real-time computation

ONR

Acknowledgments

AFOSR YIP 2023 (FA9550-23-1-0131, Frederick Leve), NASA (AVIATE ULI)

Real research (among others, for this year only):

- **Partial loss of control:** Ram Padmanabhan, Siddharth Dinkar, Craig Bakker (PNNL)
- **Unknown dynamics:** Taha Shafa, Yiming Meng, Jesse Wei, collaborations on the side with X. Li (UWisc), J. Liu (Waterloo)

mornik@illinois.edu