How to provide performance guarantees for data-driven models?

Prof. Alan Kaptanoglu and many others

Courant Institute of Mathematical Sciences, NYU

16th Plasma Kinetics Working Meeting

System ID/Operator Learning: Turning time series data into ODEs/PDEs

- Goal: Take some data obtained from dynamical system and either:
  - *Predict the dynamical evolution* (neural ODEs, FNOs, RNNs, etc.) $\dot{\boldsymbol{X}} = N_{\boldsymbol{\theta}}(\boldsymbol{X})$
  - *Infer the equations of motion* (symbolic regression, sparse regression, etc.) represented by system of ordinary differential equations (ODEs) or partial differential equations (PDEs), (SINDy, OpInf, etc.)

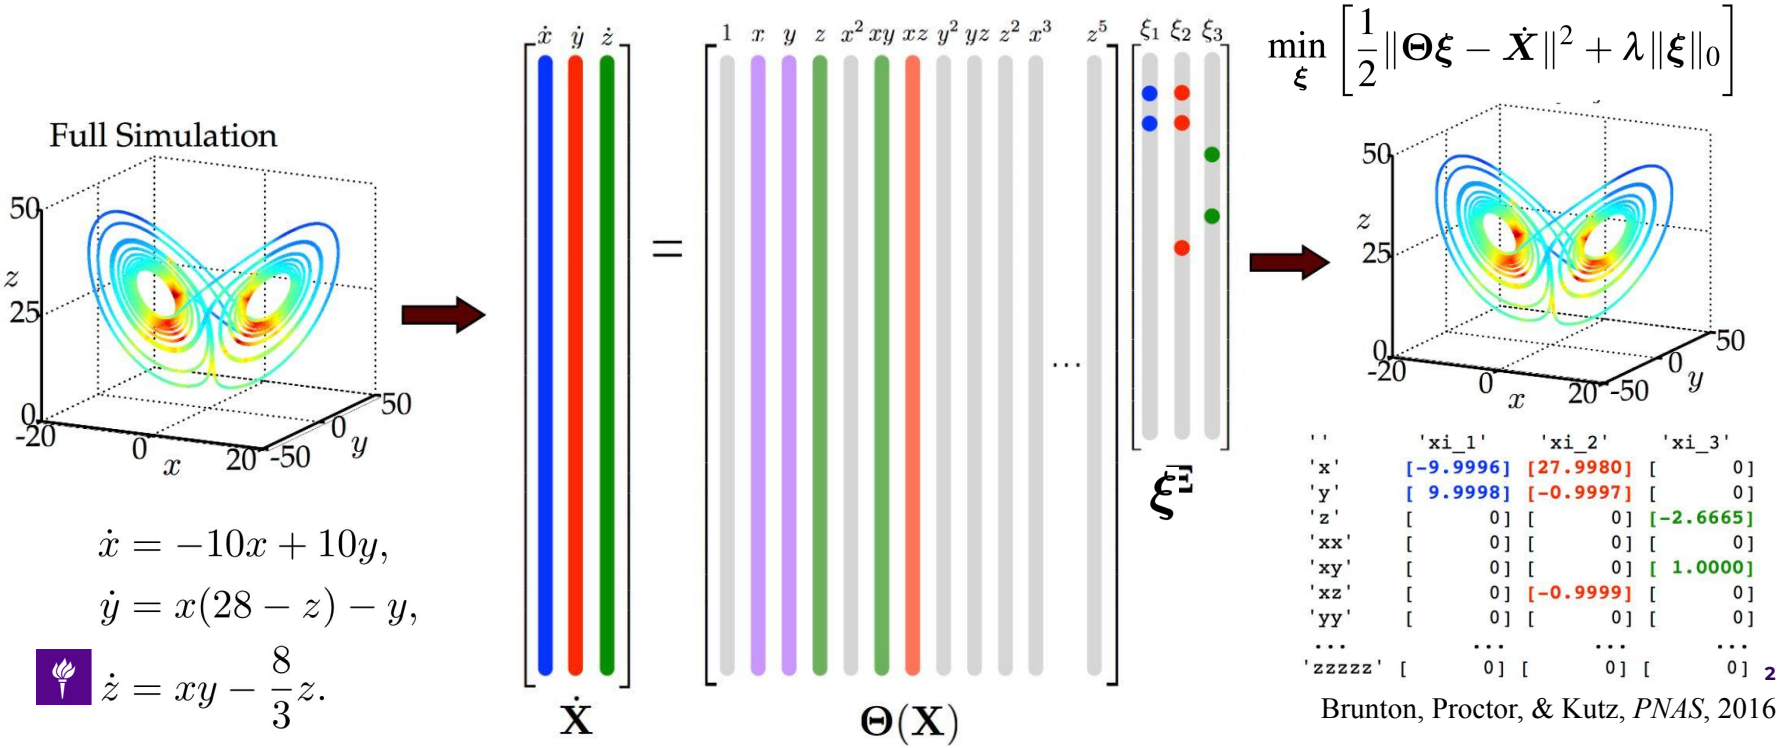

Brunton, Proctor, & Kutz, *PNAS*, 2016

Be careful with naive system identification

- Some potential issues to notice:
  - **Errors in model identification are inevitable** from finite sampling rates, noise in the data, imperfect optimization, and so on.
  - **Physically stable ODEs are roughly a set of measure zero** in space of all ODEs.
  - Integration of your continuous ODE may fail when it is discretized (e.g. with Forward/Backward Euler)
  - If identifying PDEs, not clear a-priori how to solve the PDE you identified.
  - Finding accurate model requires **good** data. What does good mean?
  - What if you only have partial state measurements?

What dynamical properties should ML-learned operators have?

Continuous-time stability of $\dot{x} = \mathscr{L}x$

Discrete-time stability of $\dot{x}_i = \mathscr{L}_{ij}x_j$

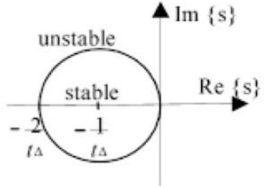

explicit Euler

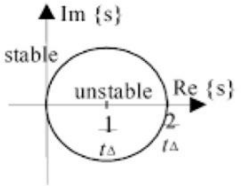

implicit Euler

Local asymptotic stability of a fixed point ($\dot{x} = 0$)

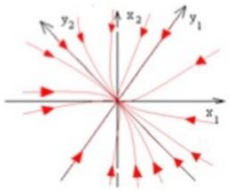

Global asymptotic stability of a fixed point ($\dot{x} = 0$)

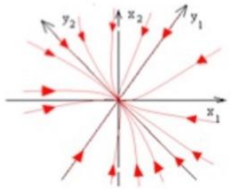

Periodic orbits, energy-preserving structure, strange attractors, etc.

Global boundedness

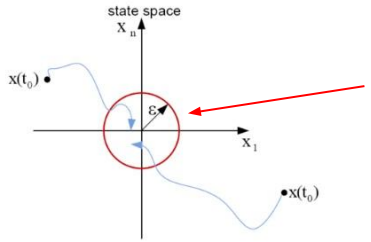

"Attracting" or "Trapping" region

Why is dynamical stability very hard to get right?

- True Lorenz system exhibits a monotonically attracting trapping region

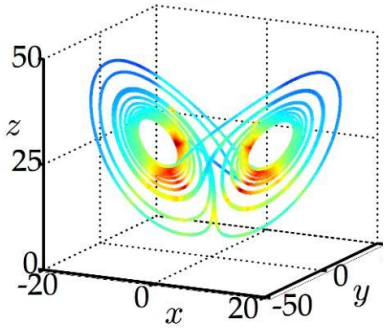

$$\dot{x} = -10x + 10y,$$
$$\dot{y} = x(28 - z) - y,$$
$$\dot{z} = xy - \frac{8}{3}z.$$

- *Arbitrarily* small perturbation $\varepsilon > 0$ can produce unbounded trajectories:

$$\dot{x} = -10x + 10y,$$
$$\dot{y} = x(28 - z) - y + \varepsilon y^3$$
$$\dot{z} = xy - \frac{8}{3}z.$$

Define

$V(x) = x^2 + y^2 + z^2 = ||\mathbf{x}||^2$

$dV/dt \sim \varepsilon y^4$ for $||\mathbf{x}||^2 >> 1$

$d^2V/d^2t \sim \varepsilon y^6$ for $||\mathbf{x}||^2 >> 1$

…

- Actually this is a nice example that *sparse* regression can give benefits beyond reducing overfitting!
- Potential rebuttal – sometimes we don't care about the very far field. But in practice $\varepsilon >>$ machine precision

NYU

Two paths to promoting/guaranteeing stability

- Path 1: Forget about stability theory, just optimize the *multi-time-step prediction error**:

$$\min_{\xi} \left[\sum_{j=n_t+1}^{N-n_t} \sum_{i=-n_t}^{n_t} w_i \| \boldsymbol{x}_{t_{j+i}} - \left(\boldsymbol{x}_{t_j} + \boldsymbol{\xi}^T \int_{t_j}^{t_{j+i}} \boldsymbol{\Theta}(\boldsymbol{X}_{t'}) dt' \right) \|_2^2 + R(\boldsymbol{\xi}) \right]$$

  - Self-consistently trains the model by integrating forward and backward n_t time steps and comparing with the data. Requires a differentiable ODE solver like Jax's Runge-Kutta, and by definition then we are now analyzing the stability of the *discretized* system of ODEs.
  - Gives accurate predictions and retains basic notions of stability, given lots of good data.
  - Potential issue #1: if the dynamical system exhibits chaos, $n_t < \boldsymbol{\tau}_L$ is recommended
  - Potential issue #2: process *promotes* but does not *guarantee* model stability.
  - Potential issue #3: Global boundedness is very strong notion – trajectories in the *entire* state space should be bounded. For model to learn this feature may require huge amount of data that samples all the dynamical variety in the full state space.

*Chen, Ricky TQ, et al. *Advances in neural information processing systems* 31, 2018

Two paths to promoting/guaranteeing stability

- Path 2: Build analytic stability theorems directly into the model-learning process.
  - *Guarantees* rather than just *promotes* desired stability properties.
  - **Primary issue:** nonlinear stability theorems are extremely hard to come by. Most proofs of nonlinear stability come from *Lyapunov's direct method*[1]. Sufficient condition to show stability is to construct a *Lyapunov function (LF)* $V(x) > 0$ s.t. $dV/dt < 0$ on some domain Ω. It is very hard to find general Lyapunov functions for general nonlinear dynamical systems.
  - **Some good news:** there is a lot to say for polynomial dynamical systems, e.g. we can certify local/global stability of a FP of arbitrary order polynomial systems by searching for a polynomial LF. This is called the *polynomial sum-of-squares approach*[2].
  - Some bad news again: All of the dynamical systems theory assumes that *the model is known*, and only searching for a LF to certify stability. Data-driven methods **need to jointly learn the model AND a LF**[3] to certify it at the same time!
  - Growing literature[4] on colearning a control policy and certificate (LF), given a model.

[1]Khalil, Hassan K., and Jessy W. Grizzle. *Nonlinear systems*. Prentice hall, 2002.
[2]Chernyshenko, Sergei I., et al. *Philosophical Transactions of the Royal Society A*, 2014.
[3]Kolter, J. Zico, and Gaurav Manek. *Advances in neural information processing systems*, 2019.
[4]Dawson, Charles, Sicun Gao, and Chuchu Fan. *IEEE Transactions on Robotics*, 2023.

Path 2 is feasible for certain nonlinear systems

- What if a good Lyapunov function is known for a class of nonlinear models?
- Common class of systems is quadratic in nonlinearity.

$$\dot{a}_i(t) = E_i + \sum_{j=1}^{r} L_{ij} a_j + \sum_{j,k=1}^{r} Q_{ijk} a_j a_k \qquad Q_{ijk} + Q_{jik} + Q_{kji} = 0$$

- Under two additional assumptions (1) energy-preserving nonlinearities (EPN), and (2) effective nonlinearity (EN), there is a necessary & sufficient theorem*.
- Actually these assumptions are well-satisfied by lots of reduced-order and data-driven fluid and plasma models!
- Notice Q_{ijk} already $O(n^3)$ values if dense

"effective nonlinearity" (EN)

$$\|\boldsymbol{a}\| \gg 1 \longrightarrow \dot{a}_i(t) \approx Q_{ijk} a_j a_k$$

*Schlegel, Michael, and Bernd R. Noack. *Journal of Fluid Mechanics*, 2015

Schlegel and Noack's trapping theorem

$$\dot{a}_i(t) = E_i + \sum_{j=1}^{r} L_{ij} a_j + \sum_{j,k=1}^{r} Q_{ijk} a_j a_k$$

$$\boldsymbol{y}(t) = \boldsymbol{a}(t) - \boldsymbol{m}$$

$$K = \frac{1}{2}\boldsymbol{y}^T\boldsymbol{y} \qquad \frac{d}{dt}K = \boldsymbol{y}^T\boldsymbol{A}^S\boldsymbol{y} + \boldsymbol{d}_m^T\boldsymbol{y}$$

$$\boldsymbol{A}^S = \boldsymbol{L}^S - \boldsymbol{m}^T\boldsymbol{Q}, \;\; \boldsymbol{L}^S = \frac{1}{2}(\boldsymbol{L} + \boldsymbol{L}^T),$$

$$\boldsymbol{d}_m = \boldsymbol{E} + \boldsymbol{L}\boldsymbol{m} + \boldsymbol{Q}\boldsymbol{m}\boldsymbol{m}.$$

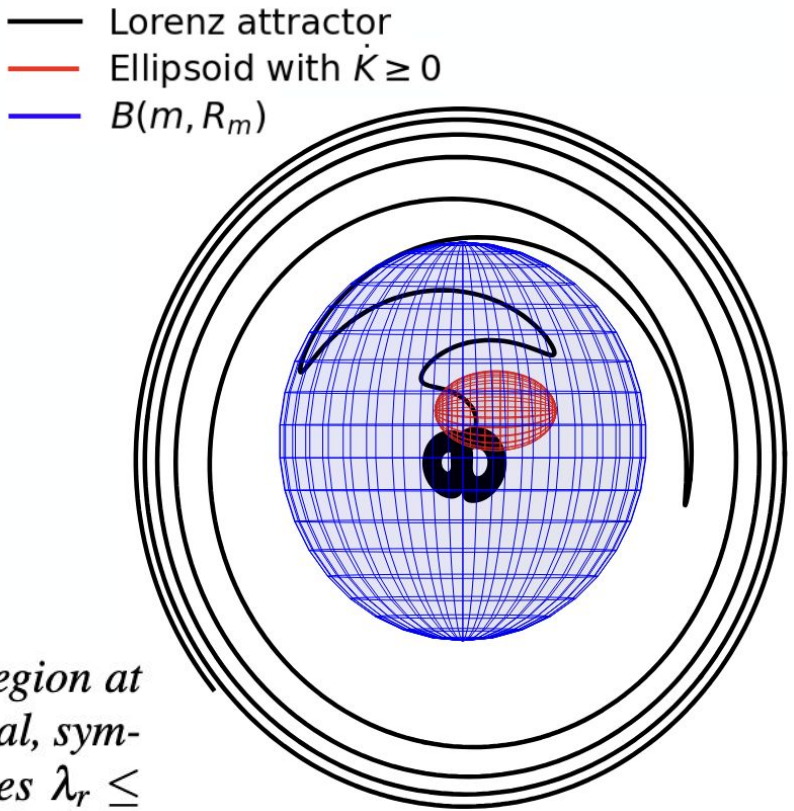

Theorem 1. *There exists a monotonically trapping region at least as small as the ball $B(\boldsymbol{m}, R_m)$ if and only if the real, symmetric matrix $\boldsymbol{A}^S$ is negative definite with eigenvalues $\lambda_r \leq \cdots \leq \lambda_1 < 0$; the radius is then given by $R_m = \|\boldsymbol{d}_m\|/|\lambda_1|$.**

Nonlinear stability theorems can be incorporated into optimization!

Provably globally-stable model trained on *one* noisy Lorenz trajectory, predicts new trajectories starting from 20 random initial conditions.

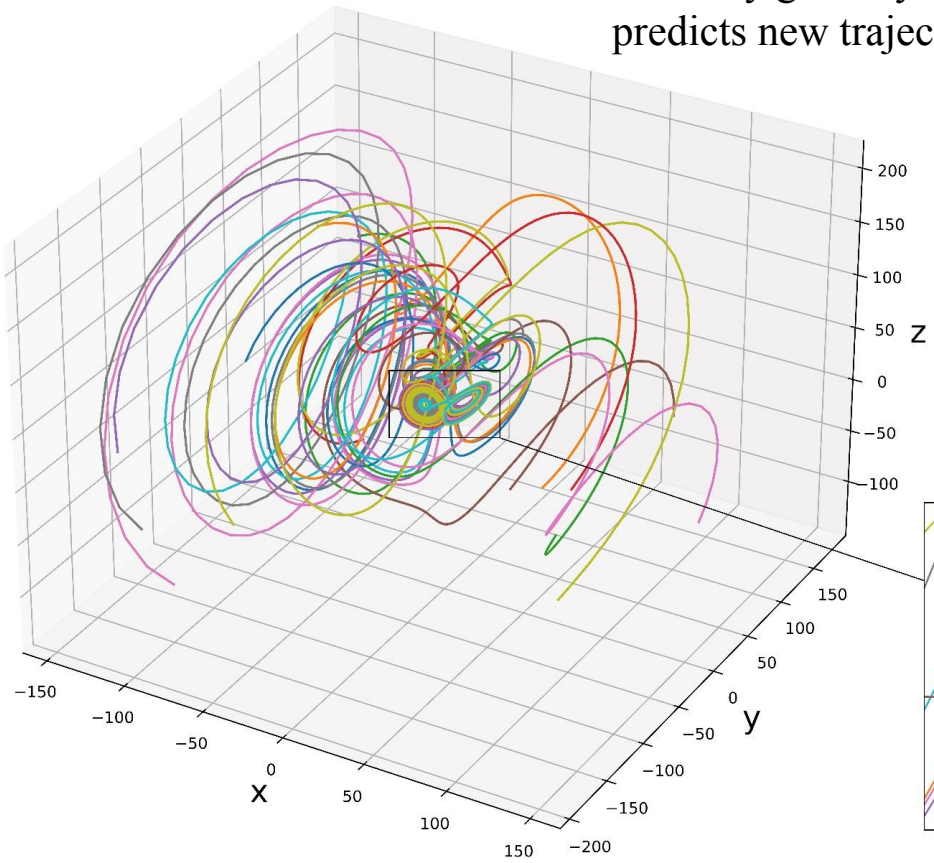

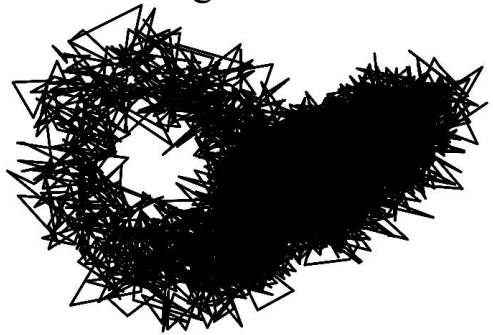

(b)

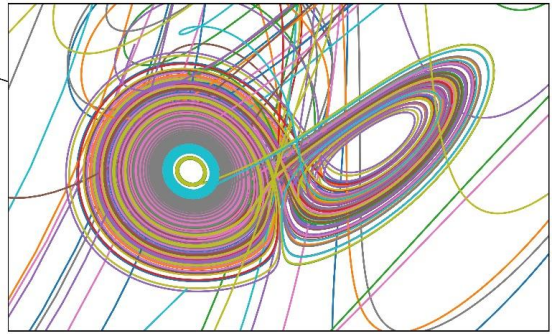

Generalizations to weakly broken EPN

Mai Peng

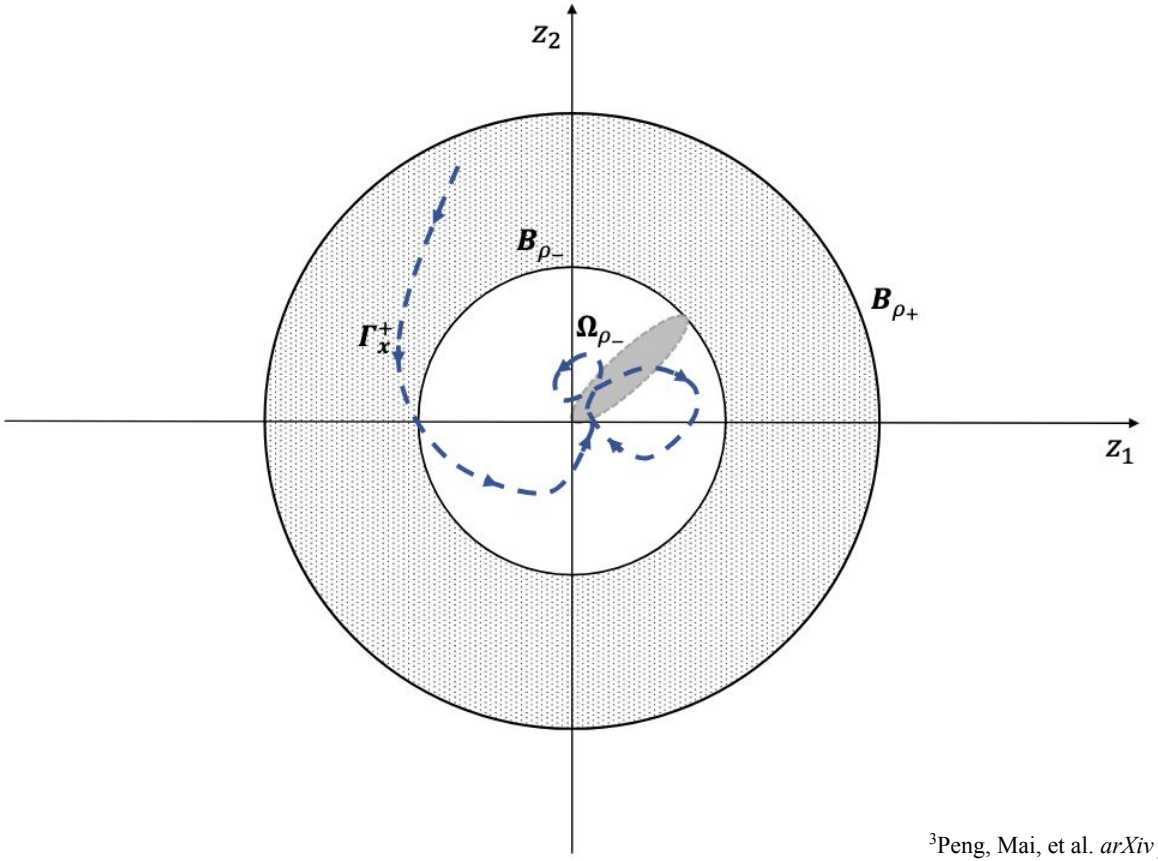

NYU

[3]Peng, Mai, et al. *arXiv preprint arXiv:2403.00324v3*, 2023.

Trapping theorem can be enforced during optimization

- The LF and trapping theorem is *known*, so can just impose it during optimization that (1) nonlinearities are EP and (2) A^S is negative definite[1-7].
- Example below for the lid-cavity flow, weakly breaks EPN

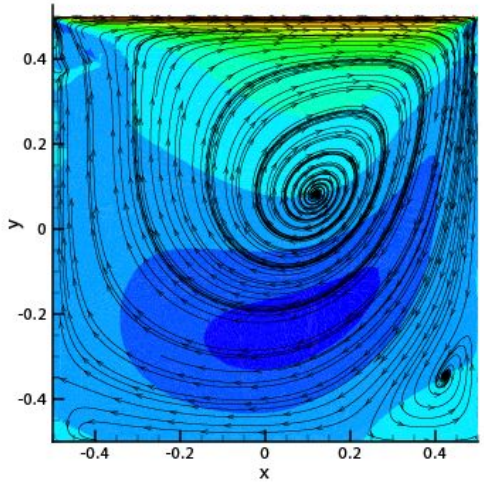

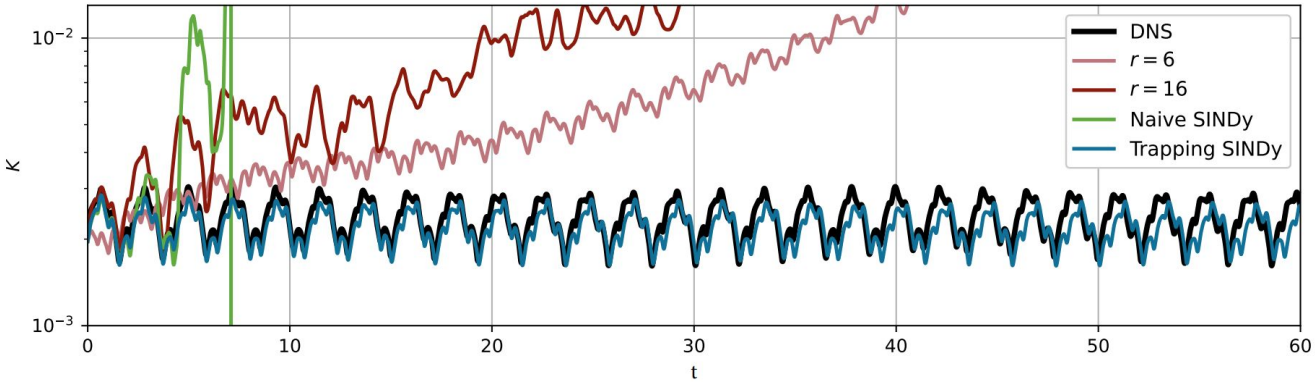

NYU

[1]Kaptanoglu, Alan A., et al. *Physical Review Fluids*, 2021.
[2]Goyal, Pawan, et al., *arXiv preprint arXiv:2308.13819*, 2023.
[3]Peng, Mai, et al. *arXiv preprint arXiv:2403.00324v3*, 2023.
[4]Ouala et al., *Physica D*, 2023.
[5]Sawant, Kramer, & Peherstorfer, *CMAME*, 2023.
[6]Duff, Igor Pontes, Pawan Goyal, and Peter Benner. *IFAC-PapersOnLine*, 2024.
[7]Koike, Tomoki, and Elizabeth Qian. *AIAA SCITECH 2024 Forum*, 2024.

Big picture future questions

- Can we understand when multi-step prediction error is good enough?
- Can we find robust ways to co-learn dynamical systems and Lyapunov functions that certify them?
- Should we be asking for low-polynomial-order high-dimensional models, or high-order low-dimensional models?
- Can we extend data-driven modeling for polynomial systems and other systems where a Lyapunov function is already constructed?
  - What if NEP is broken?
    - Is the Schlegel and Noack trapping theorem the “only game in town” for EN systems?
    - Can we more robustly solve the optimization problem?
  - Higher-order polynomial nonlinearities in EN systems?
  - Effective nonlinearity vs sparsity?

Conclusion and discussion

- **Takeaways:**
  - Guaranteeing dynamical stability when identifying ODEs is very challenging but can be critical to performance.
  - Two (three?) paths identified to promote/guarantee notions of dynamical stability but all paths are challenging.
  - For effectively nonlinear polynomial systems of arbitrary order, there is probably an analytic necessary and sufficient theorem for global boundedness and similar notions.

Please reach out if you want to collaborate on this class of problems!

Extra Slides

An extremely simple GAS quadratic polynomial system with no polynomial Lyapunov function

$$\begin{aligned}\dot{x} &= -x + xy \\ \dot{y} &= -y.\end{aligned}$$

$$V(x,y) = \ln(1+x^2) + y^2$$

Pop quiz: tell me why the trapping theorem cannot be applied here

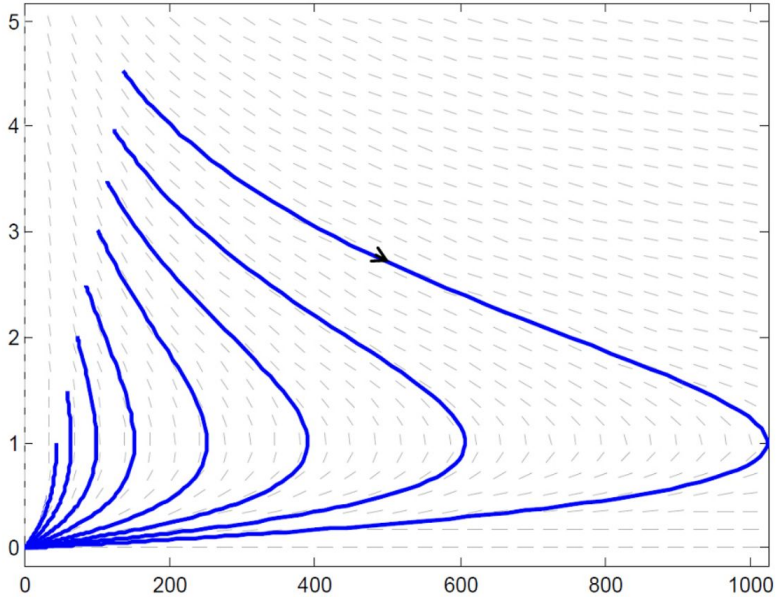

NYU

Ahmadi, Amir Ali, et al. 2*011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE, 2011.

Generalizations to weakly broken EPN

Mai Peng

Theorem 2. *Regarding the system* (20) *and a Lyapunov matrix* $\boldsymbol{P}$*, there exists a trapping region if the following statements are true:*

a. $\tilde{\boldsymbol{A}}^S = \frac{1}{2}\left(\boldsymbol{P}^{\frac{1}{2}}\boldsymbol{A}\boldsymbol{P}^{-\frac{1}{2}} + \boldsymbol{P}^{-\frac{1}{2}}\boldsymbol{A}^T\boldsymbol{P}^{\frac{1}{2}}\right)$ *is Hurwitz, i.e.* $\lambda_r(\tilde{\boldsymbol{A}}^S) \leq \ldots \leq \lambda_1(\tilde{\boldsymbol{A}}^S) < 0$.

b. $\lambda_1^2 - \frac{4\epsilon_Q}{3}\|\tilde{\boldsymbol{d}}\|_2 \geq 0$*, where* $\epsilon_Q = \|\tilde{\boldsymbol{H}}_0\|_F$*,* $\tilde{\boldsymbol{H}}_0 = \boldsymbol{P}^{-\frac{1}{2}}\mathcal{T}(\tilde{Q}_{ijk}+\tilde{Q}_{ijk}+\tilde{Q}_{kji})(\boldsymbol{P}^{-\frac{1}{2}}\otimes\boldsymbol{P}^{-\frac{1}{2}})$*,* $\tilde{\boldsymbol{Q}} = \boldsymbol{P}\boldsymbol{Q}$ *and* $\tilde{\boldsymbol{d}} = \boldsymbol{P}^{\frac{1}{2}}\boldsymbol{d}$.

Then the smallest radius of the trapping region $\mathbf{B}_{\boldsymbol{y}}(\rho_-)$ *is given by*

$$\rho_- = \frac{3|\lambda_1|}{2\epsilon_Q}\left(1 - \sqrt{1 - \frac{4\epsilon_Q}{3\lambda_1^2}\|\tilde{\boldsymbol{d}}\|_2}\right), \tag{28}$$

while all trajectories starting inside $\mathbf{B}_{\boldsymbol{y}}(\rho_+)$ *are bounded where*

$$\rho_+ = \frac{3|\lambda_1|}{2\epsilon_Q}\left(1 + \sqrt{1 - \frac{4\epsilon_Q}{3\lambda_1^2}\|\tilde{\boldsymbol{d}}\|_2}\right). \tag{29}$$

$$\boldsymbol{\xi}, \boldsymbol{\zeta}, \boldsymbol{m} = \underset{\boldsymbol{\xi}', \boldsymbol{m}}{\operatorname{argmin}} \ \max_{\boldsymbol{\zeta}} \left[\frac{1}{2}\|\dot{\boldsymbol{X}}[:] - \hat{\boldsymbol{\Theta}}(\boldsymbol{X})\boldsymbol{\xi}'\|_2^2 + \gamma\|\boldsymbol{\xi}'\|_1 + \boldsymbol{\zeta}^T(\boldsymbol{C}\boldsymbol{\xi}' - \boldsymbol{q}) + \frac{\boldsymbol{\lambda}_1(\tilde{\boldsymbol{A}}^S)}{\eta}\right],$$

[3]Peng, Mai, et al. *arXiv preprint arXiv:2403.00324v3*, 2023.