Sparse Koopman Autoencoders Identify Local Dynamical Regimes in Multibasin Systems

Aidan Li[1,2,3] Uday Kiran Reddy Tadipatri[4] Mahan Fathi[5]
Sarath Chandar[1,2,6,7] Ross Goroshin[2,3]

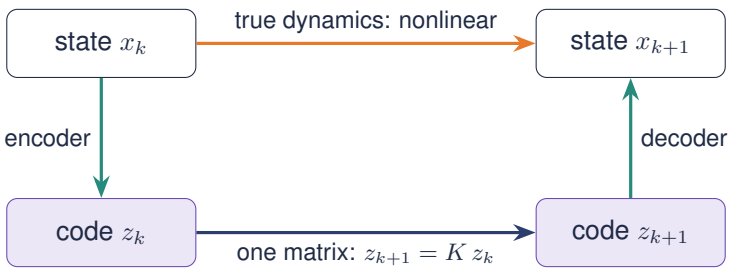
[1]Chandar Research Lab [2]Mila – Quebec AI Institute [3]Université de Montréal [4]University of Pennsylvania
[5]NVIDIA [6]Polytechnique Montréal [7]Canada CIFAR AI Chair

Paper

Code

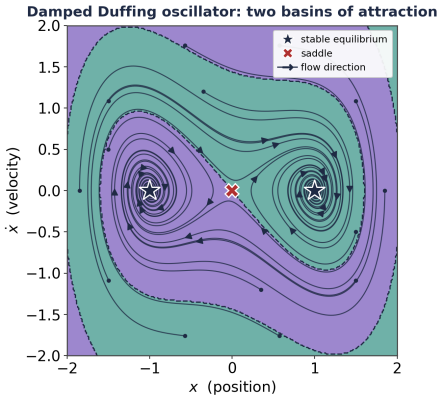
Koopman autoencoders

Learn a code in which nonlinear dynamics become **linear**.



The same matrix K is used for every state and at every step.

Many systems have several basins of attraction



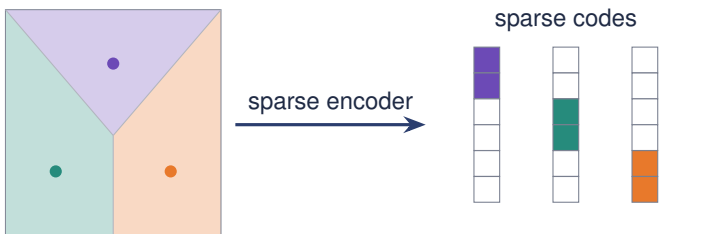
- Different starting states settle at different attractors.
- Inside one basin, the dynamics can be linearized.
- In general, **no single linear model** of finite dimension, with a continuous encoder and decoder, covers all basins.

Brunton et al. 2016 · Pan & Duraisamy 2024 · Liu, Ozay & Sontag 2025

A sparse code can represent several regimes

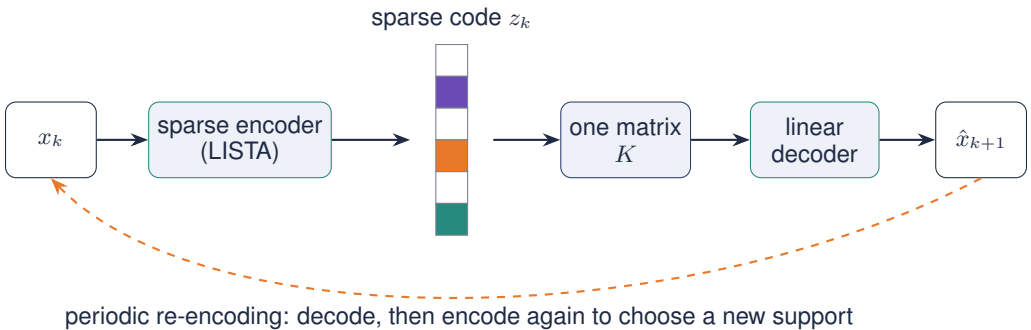
the **support** (which coordinates are nonzero) $\Rightarrow$ *which* regime

the **nonzero values** $\Rightarrow$ linear coordinates within that regime



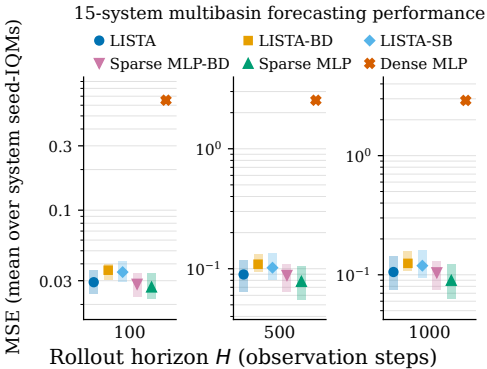
Each region of state space uses its own coordinates,
so the one matrix K acts as a different linear map in each.

Sparse Koopman Autoencoder (SKAE)



Trained on trajectories only, without any basin labels.

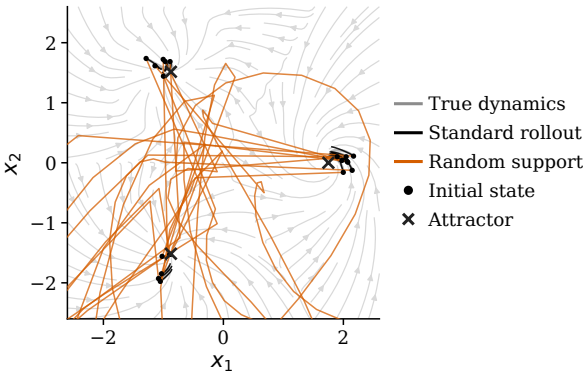
Result 1: sparse codes forecast better



- 15 multibasin systems, 15 seeds each
- **Orange cross:** dense code
  Other markers: sparse codes
- 10 chaotic flows: sparse codes are also better, by 1.2–1.6× at 4,000 steps

23–32× lower error than the dense code on multibasin systems, at 1,000 steps

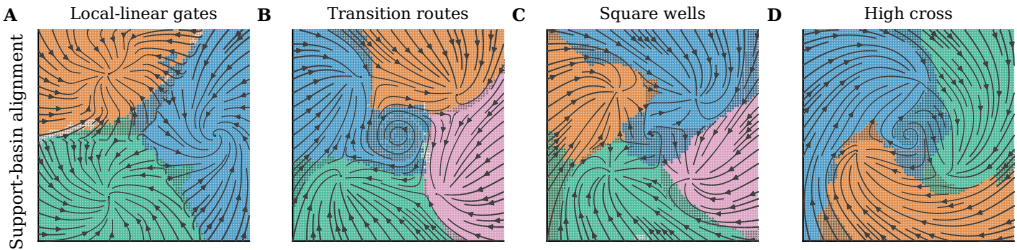
Result 2: the forecasts depend on the support



- Take a trained SKAE and encode a starting state.
- Keep the nonzero values, but move them to **random coordinates**.
- Forecast from the edited code.

21-step forecast error: 0.0112 → 2060

Result 3: the supports match the basins



Dots: states, colored by their support family's basin. Dark dots: mismatches. Lines: true dynamics.

Support families match the basins one-to-one in 175 of 225 runs (dense code: 0)

Summary

- One Koopman autoencoder with a **sparse code** holds several local linear regimes.
- It forecasts better than a dense code, and the forecasts depend on the support.
- The supports identify the basins without basin labels.

Paper

Code