Learning Control Laws

Stable Estimator of Dynamical Systems (SEDS)

Global Asymptotic Stability of Autonomous Dynamical System (DS)

Lyapunov's Theorem for Global Asymptotic Stability

Theorem: A DS is ***globally asymptotically stable*** at $\boldsymbol{x}^* \in \mathbb{R}^N$ if there exists a Lyapunov candidate function $V(x):\mathbb{R}^N \to \mathbb{R}^N$ that is radially unbounded; i.e. $V(\boldsymbol{x}) \to \infty$ as $||\boldsymbol{x}|| \Rightarrow \infty$, $\mathcal{C}^1$ and satisfies the following conditions:

$$(\mathbf{I})\ V(\boldsymbol{x}^*) = 0,\ (\mathbf{II})\ V(\boldsymbol{x}) > 0\ \forall\ \boldsymbol{x} \in \mathbb{R}^N \setminus \boldsymbol{x} = \boldsymbol{x}^*$$

$$(\mathbf{III})\ \dot{V}(\boldsymbol{x}^*) = 0,\ (\mathbf{IV})\ \dot{V}(\boldsymbol{x}) < 0\ \forall\ \boldsymbol{x} \in \mathbb{R}^N \setminus \boldsymbol{x} = \boldsymbol{x}^*$$

$$\dot{V}(\boldsymbol{x}) = \frac{\partial V}{\partial \boldsymbol{x}} \mathbf{f}(\boldsymbol{x}) < 0$$

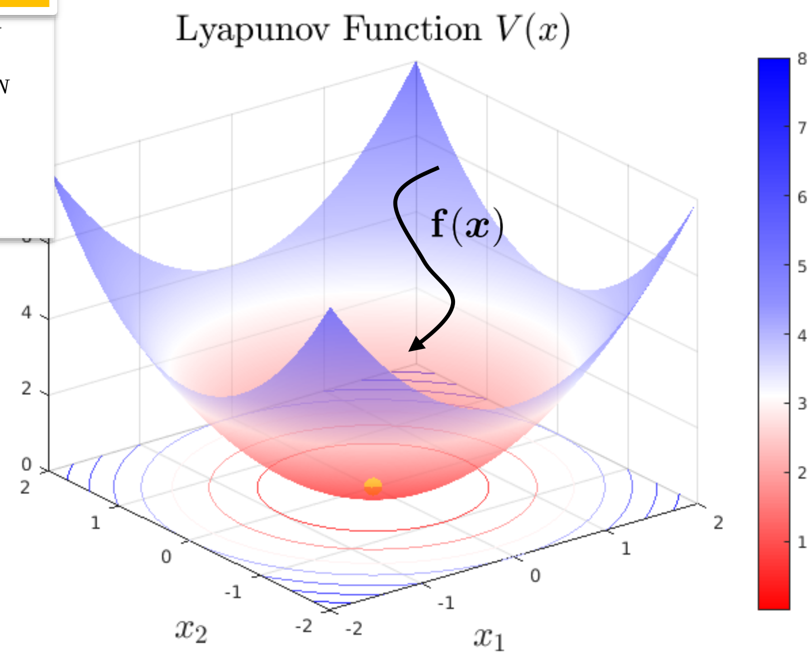

V should be non-increasing along all trajectories of $\mathbf{f}(x)$

Lyapunov Function ~ Energy-like Function

Global Asymptotic Stability of Autonomous Dynamical System (DS)

Lyapunov's Theorem for Global Asymptotic Stability

Theorem: A DS is ***globally asymptotically stable*** at $x^* \in \mathbb{R}^N$ if there exists a Lyapunov candidate function $V(x): \mathbb{R}^N \rightarrow \mathbb{R}^N$ that is radially unbounded; i.e. $V(x) \rightarrow \infty$ as $||x|| \Rightarrow \infty$, $\mathcal{C}^1$ and satisfies the following conditions:

$$(\mathbf{I})\ V(\boldsymbol{x}^*) = 0,\ (\mathbf{II})\ V(\boldsymbol{x}) > 0\ \forall\ \boldsymbol{x} \in \mathbb{R}^N \setminus \boldsymbol{x} = \boldsymbol{x}^*$$

$$(\mathbf{III})\ \dot{V}(\boldsymbol{x}^*) = 0,\ (\mathbf{IV})\ \dot{V}(\boldsymbol{x}) < 0\ \forall\ \boldsymbol{x} \in \mathbb{R}^N \setminus \boldsymbol{x} = \boldsymbol{x}^*$$

$$\dot{V}(\boldsymbol{x}) = \frac{\partial V}{\partial \boldsymbol{x}} \mathrm{f}(\boldsymbol{x}) < 0$$

V should be non-increasing along all trajectories of $\mathrm{f}(\boldsymbol{x})$

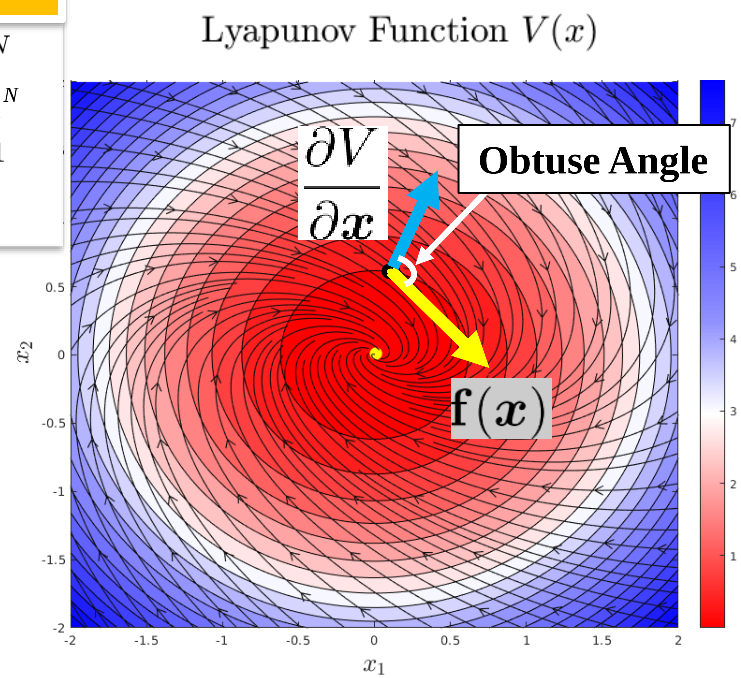

Level Sets of Lyapunov Function

Stability of a Linear Autonomous Dynamical System (DS)

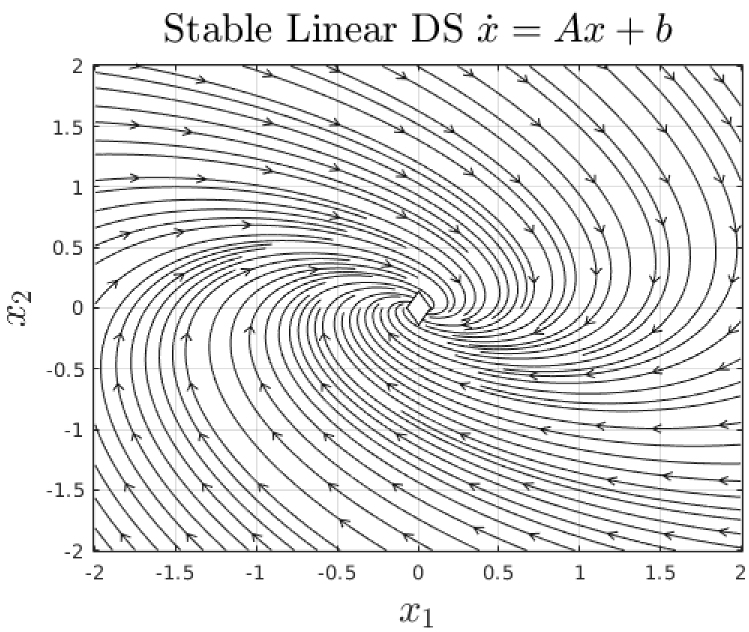

Quadratic Lyapunov Function (QLF)

$$V(\boldsymbol{x}) = (\boldsymbol{x} - \boldsymbol{x}^*)^T (\boldsymbol{x} - \boldsymbol{x}^*)$$

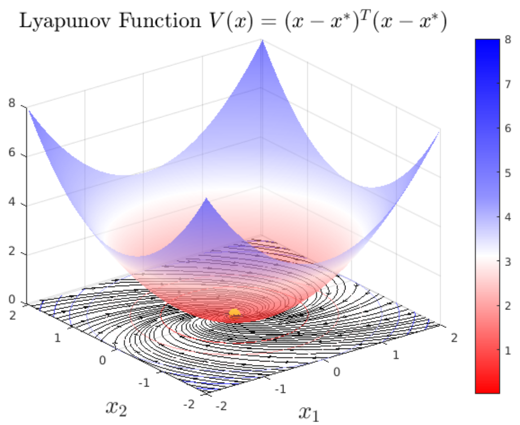

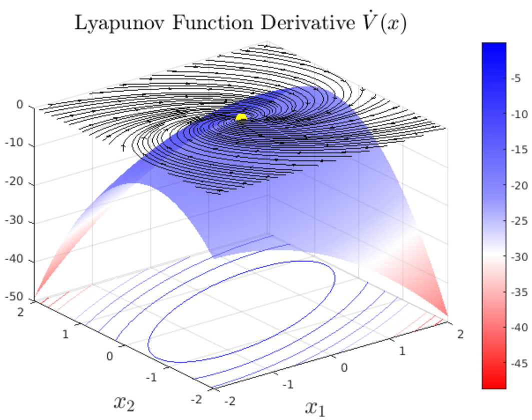

How to ensure $\dot{V}(\boldsymbol{x})$ is always negative?

$$\dot{V}(\boldsymbol{x}) = \frac{\partial V}{\partial \boldsymbol{x}} \mathbf{f}(\boldsymbol{x}) < 0$$

$$\mathbf{A}^T + \mathbf{A} \prec 0$$

Enforce the eigenvalues to be negative!

Stability of non-linear DS

What if *f(x)* is non-linear?

- Not easy to assess whether the system is stable.

- Traditionally, the following has been done:

    - local linearization;
    - numerical estimation of stability;
    - analytical solution in special cases.

Stable Estimator of Dynamical Systems (SEDS)

Khansari-Zadeh, S.M. and Billard, A., 2011. Learning stable nonlinear dynamical systems with gaussian mixture models. *IEEE Transactions on Robotics*, *27*(5), pp.943-957.

Mohi Khansari

Stable Estimator of Dynamical Systems (SEDS)

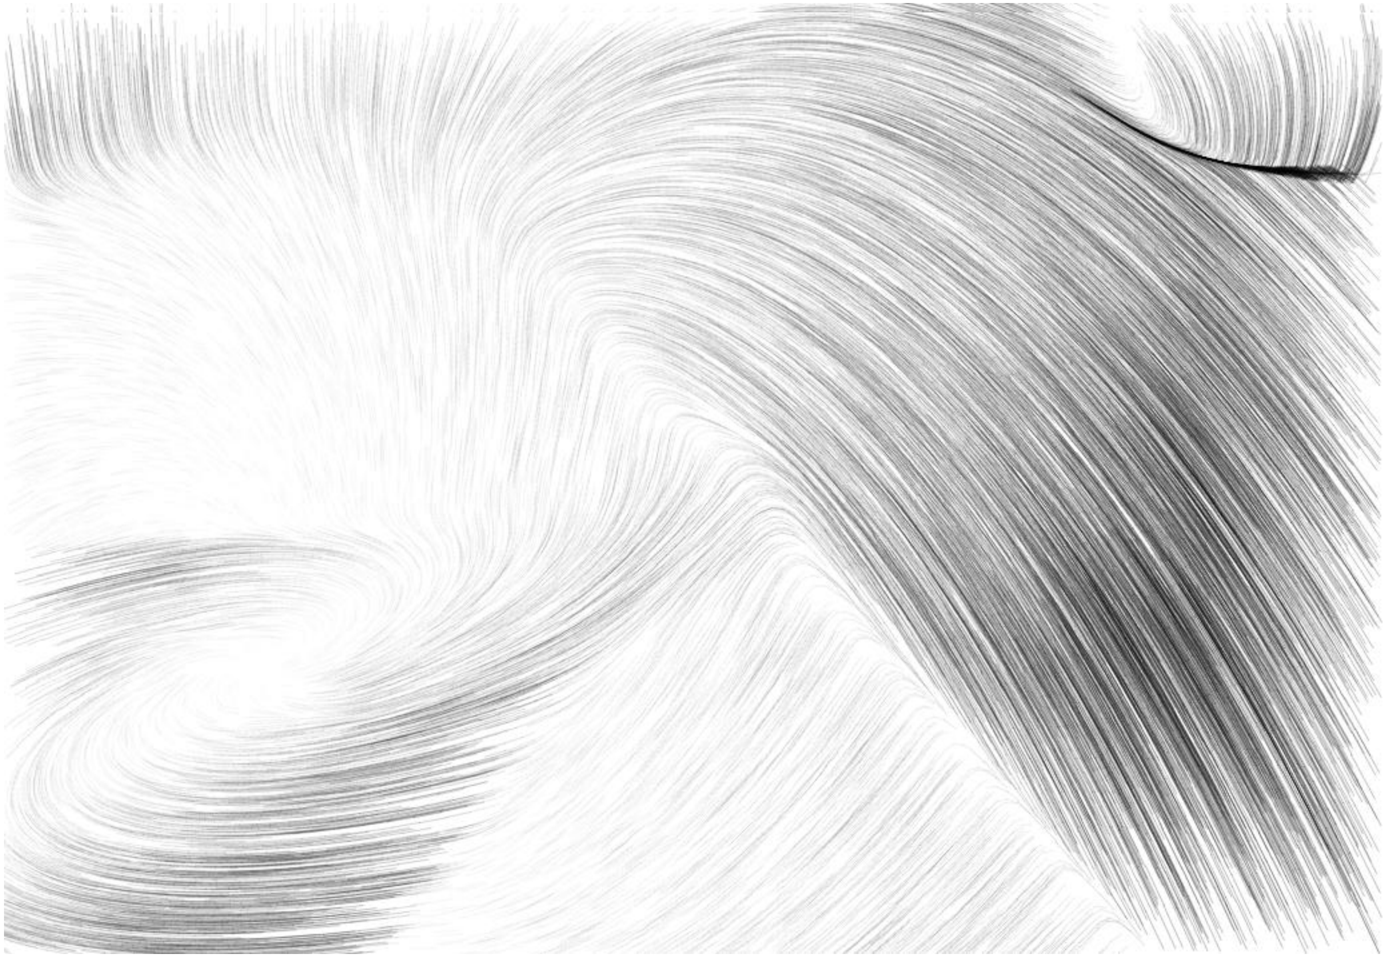

How to model this non-linear dynamical system?

SEDS starting point

DATA: set of M reference trajectories

$$\left\{X,\dot{X}\right\}=\left\{\left\{x_t^i,\dot{x}_t^i\right\}_{t=1}^{T_m}\right\}_{m=1}^{M}$$

T_m : Length of each trajectory

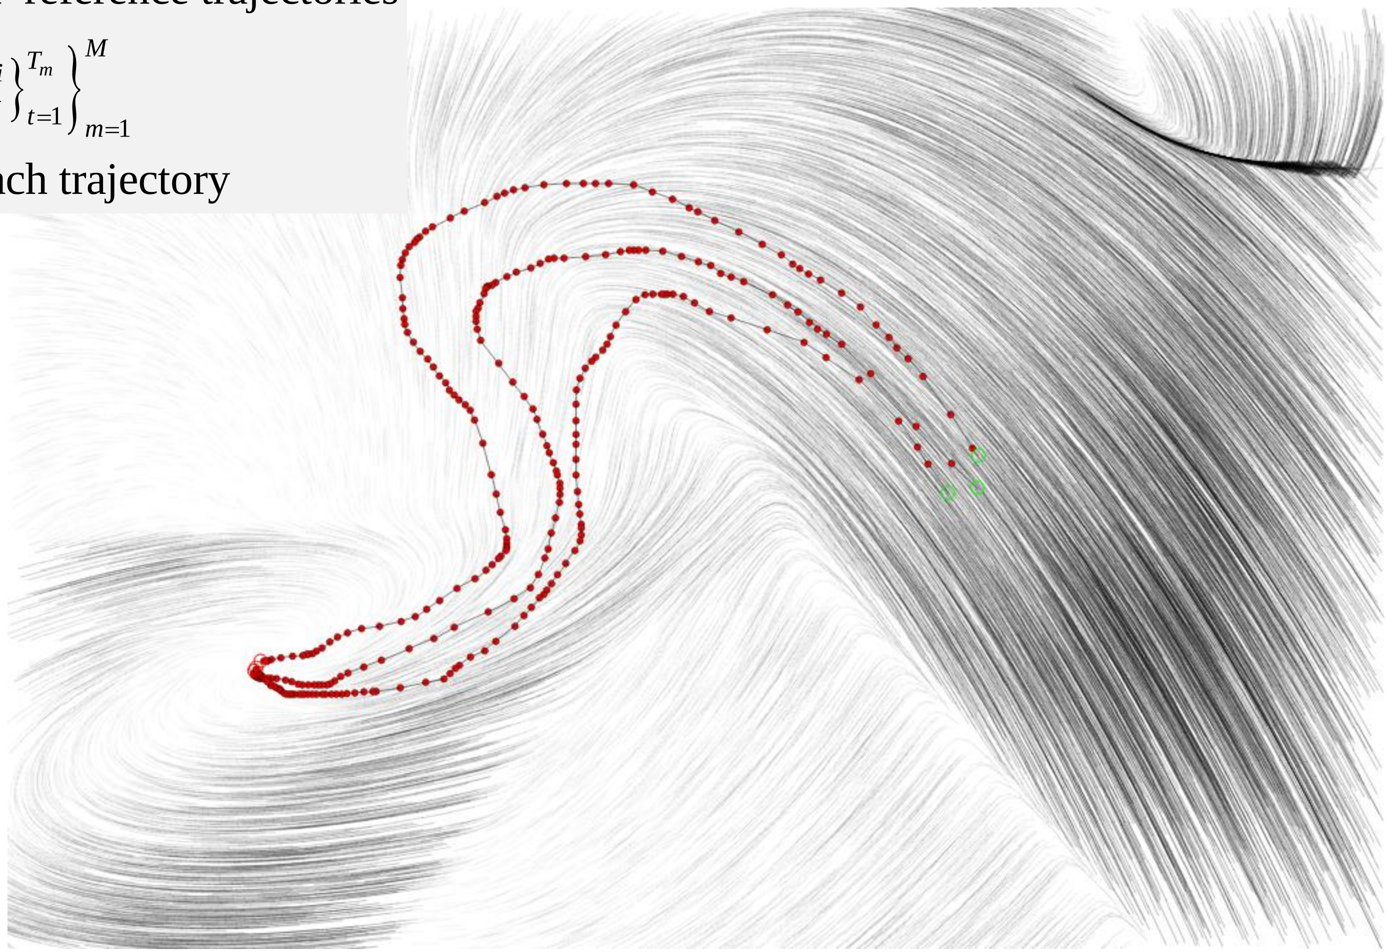

Start with sampled trajectories from a nonlinear DS

SEDS model

Model the data with a probabilistic model: $p\left(\dot{x}, x; \Theta\right)$

Θ : Model's parameters

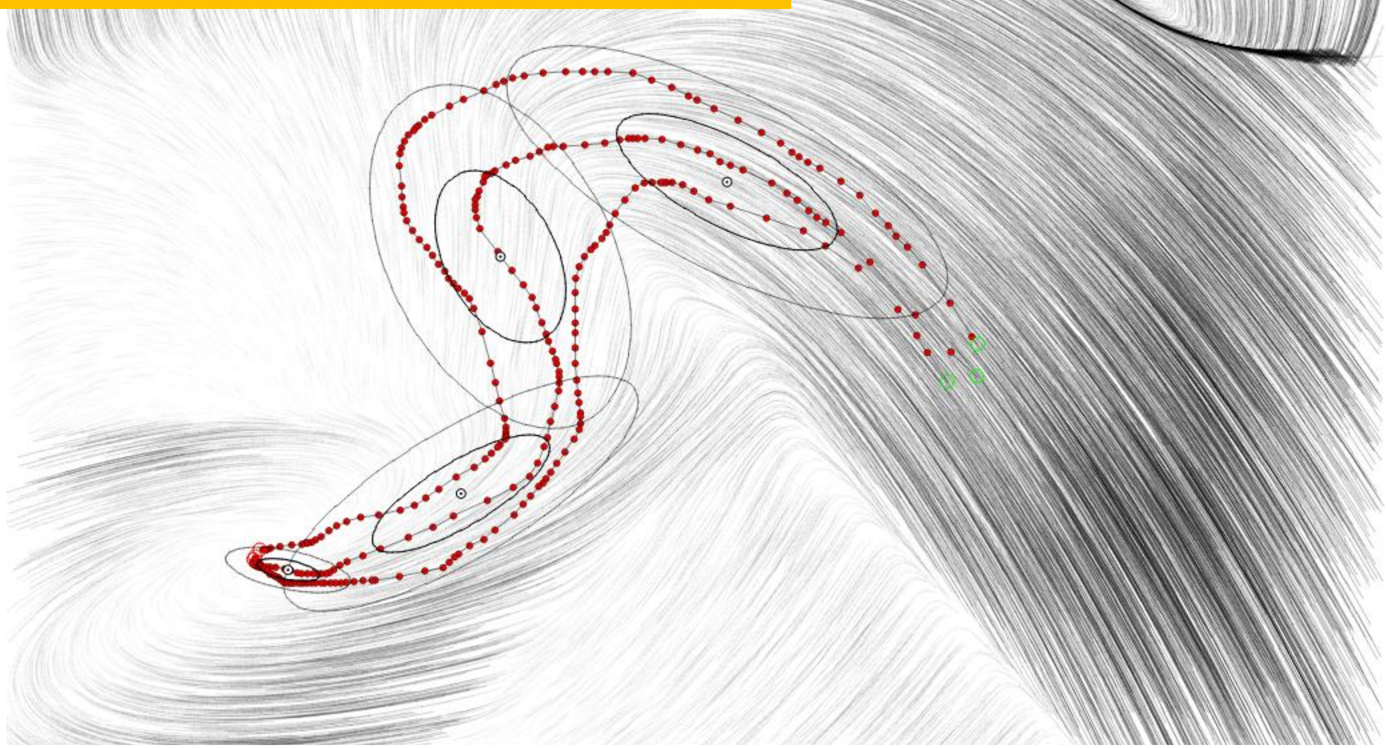

$$p\left(\dot{x}, x; \Theta\right) = \sum_{k=1}^{K} \pi_k \cdot p\left(\dot{x}, x; \mu^k, \Sigma^k\right), \qquad \text{with } p\left(\dot{x}, x; \mu^k, \Sigma^k\right) = N\left(\mu^k, \Sigma^k\right), \quad 0 < \pi_k \leq 1$$

$\Theta = \left\{\pi_k, \mu^k, \Sigma^k\right\}_{k=1}^{K}$: priors, means and covariance matrices of the K Gauss functions

SEDS model

Generate an estimate of the DS: $\dot{x} = f\left(x;\Theta\right) := E\left\{p\left(\dot{x} \mid x;\Theta\right)\right\}$

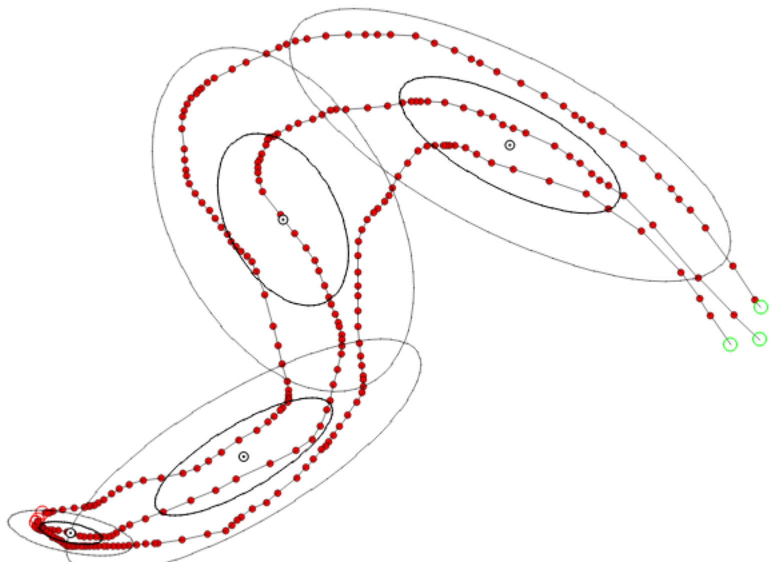

Nonlinearity comes from

$$\gamma_k(x) = \frac{\pi_k \cdot p\left(x;\mu_x^k,\Sigma_x^k\right)}{\sum_{l=1}^{K} \pi_l \cdot p(x;\mu_x^l,\Sigma_x^l)}$$

Gaussian Mixture Regression:

$$\dot{x} = \sum_{k=1}^{K} \gamma_k\left(x\right)\left(\underbrace{\Sigma_{\dot{x}x}^k\left(\Sigma_{xx}^k\right)^{-1}}_{A^k} x + \underbrace{\left(\mu_{\dot{x}}^k - \Sigma_{\dot{x}x}^k\left(\Sigma_{xx}^k\right)^{-1}\mu_x^k\right)}_{b^k}\right) = \sum_{k=1}^{K} \gamma_k\left(x\right)\left(A^k x + b^k\right)$$

K linear DS

SEDS as a mixture of linear DS

K linear DS

$$\dot{x} = \sum_{k=1}^{K} \gamma_k(x)\left(A^k x + b^k\right)$$

Nonlinearity comes from

$$\gamma_k(x) = \frac{\pi_k \cdot p\left(x; \mu_x^k, \Sigma_x^k\right)}{\sum_{k=1}^{K} \pi_k \cdot p(x; \mu_x^k, \Sigma_x^k)}$$

Mixing function

$$\gamma_k(\boldsymbol{x}) : \mathbb{R}^N \to \mathbb{R}$$
$$0 < \gamma_k(\boldsymbol{x}) < 1$$
$$\sum_{k=1}^{K} \gamma_k(\boldsymbol{x}) = 1$$

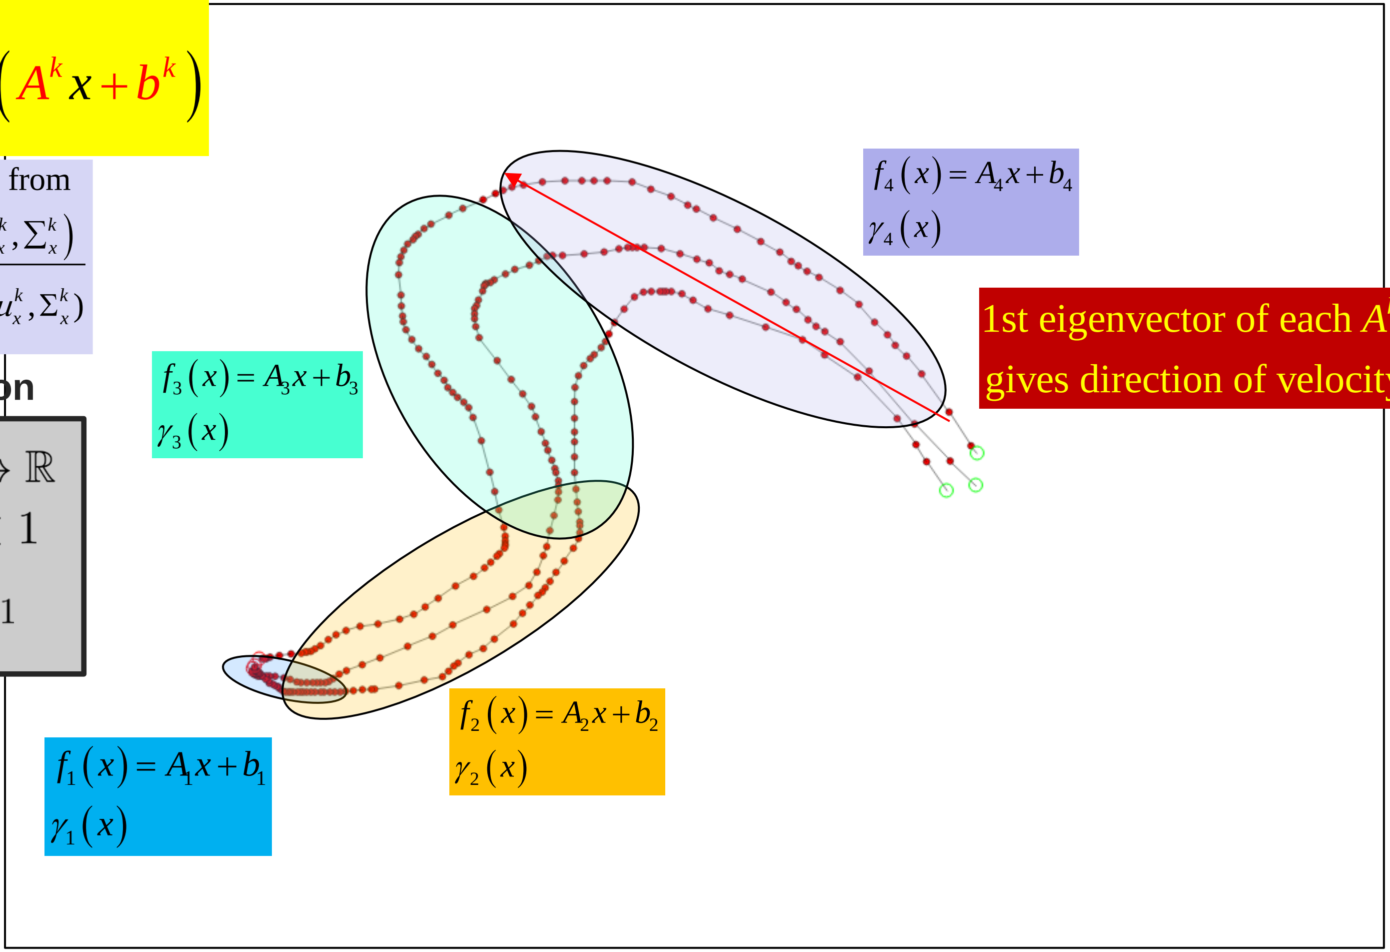

Conditions for SEDS stability

Model is parameterized only by the A^k matrices and b^k vectors.

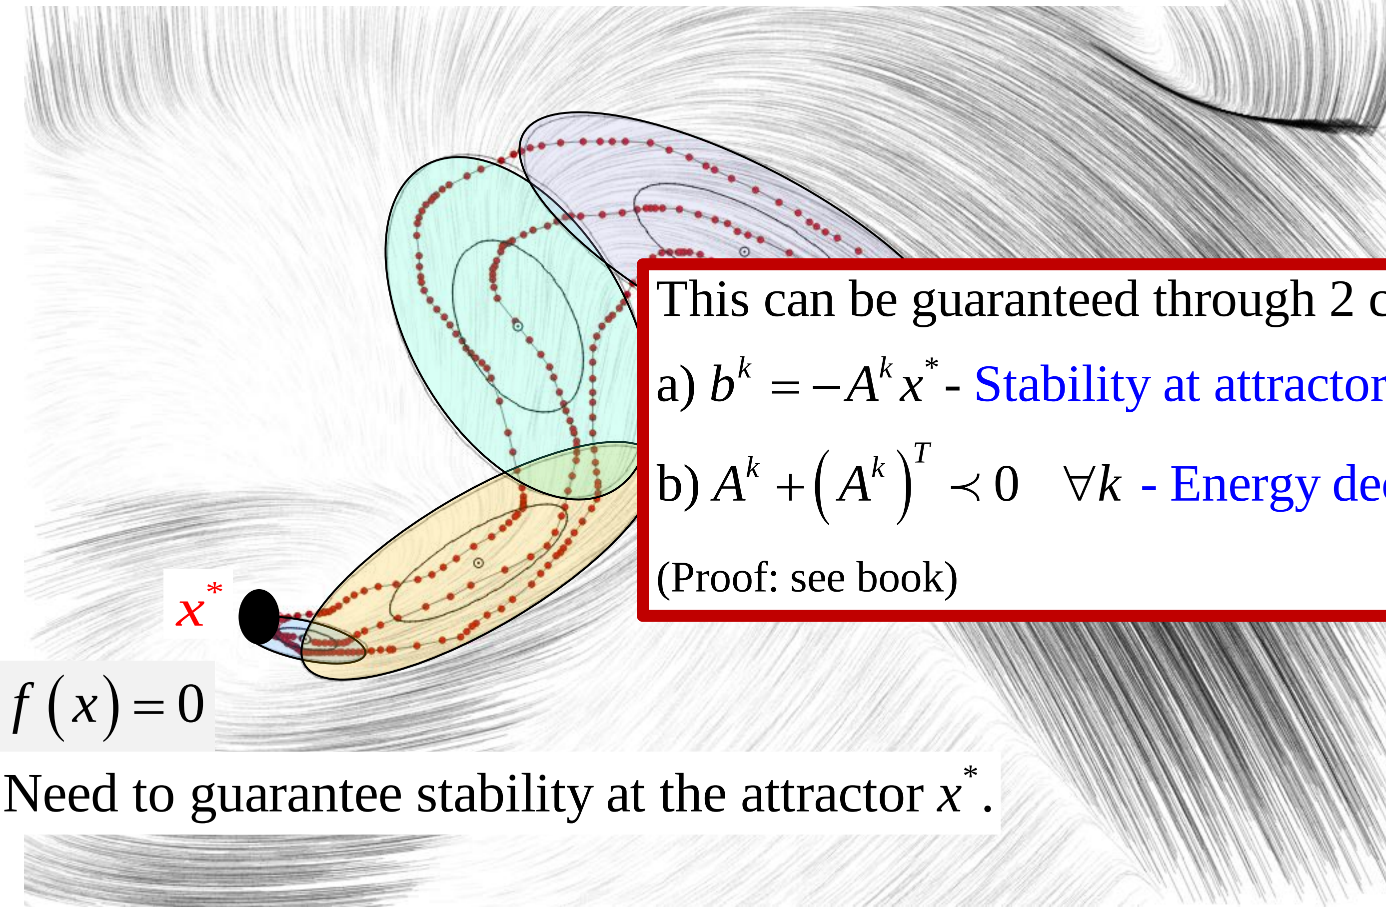

This can be guaranteed through 2 conditions:

a) $b^k = -A^k x^*$ - Stability at attractor

b) $A^k + \left(A^k\right)^T \prec 0 \quad \forall k$ - Energy decreases

(Proof: see book)

Need to guarantee stability at the attractor x^*.

Parametrization of SEDS

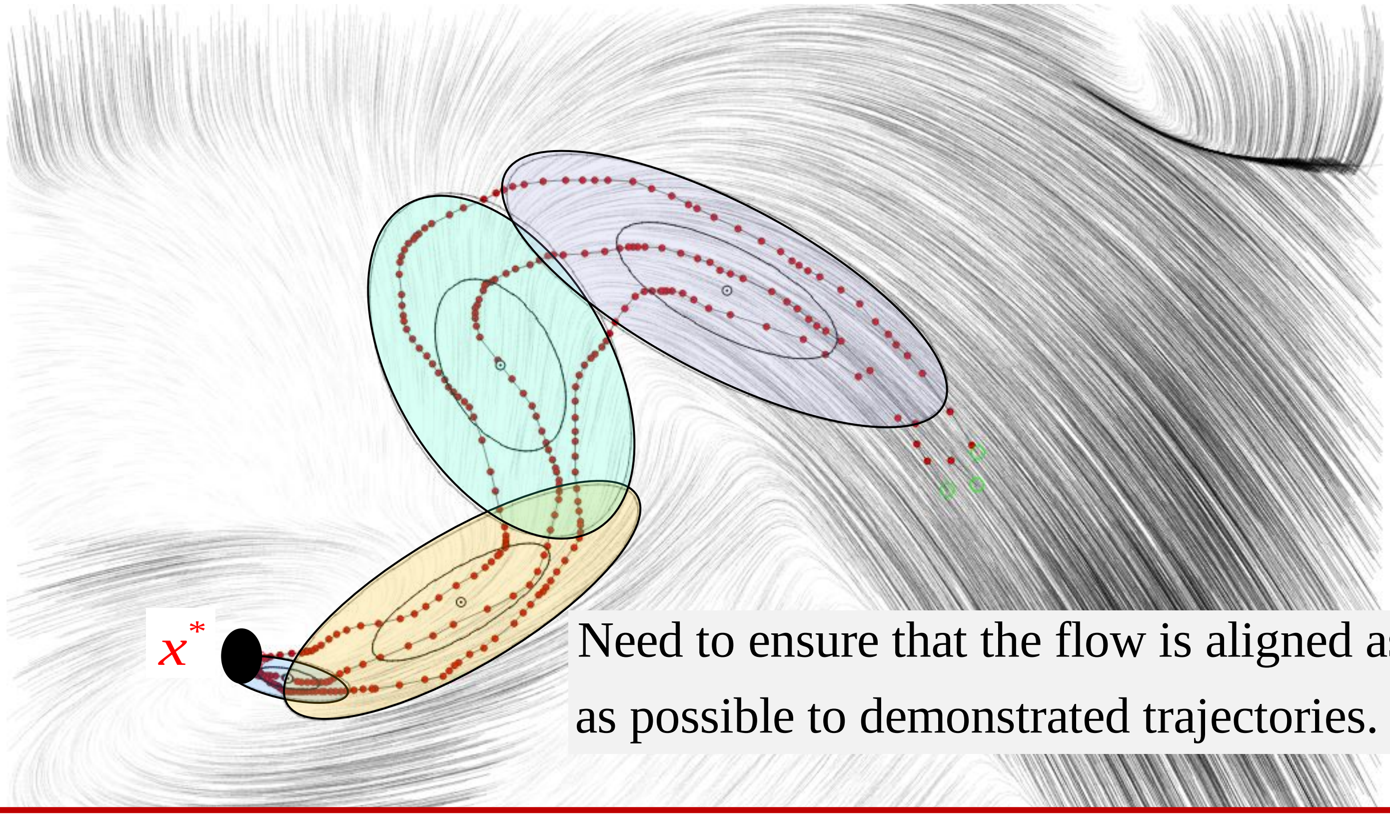

Need to ensure that the flow is aligned as closely as possible to demonstrated trajectories.

Two possible objective functions:

a) Maximum likelihood $\rightarrow$ fits at best the entire density

b) Mean-square error $\rightarrow$ fits at best the state space trajectories and velocities

Optimization of SEDS

Maximum likelihood

$$\min_{\Theta_{GMR}} J\left(\Theta_{GMR}\right) = -\frac{1}{M}\sum_{m=1}^{M}\sum_{t=0}^{T_m}\log p\left(x^{t,m}, \dot{x}^{t,m} \middle| \Theta_{GMR}\right)$$

Mean-square error

$$\min_{\Theta_{GMR}} J\left(\Theta_{GMR}\right) = -\frac{1}{2M}\sum_{m=1}^{M}\sum_{t=0}^{T_m}\left\| f\left(x^{t,m}\right) - \dot{x}^{t,m}\right\|$$

Non-convex optimization

Set of constraints

$$\begin{cases} \textbf{(a)}\ b^k = -A^k x^* \\ \textbf{(b)}\ A^k + (A^k)^T \prec 0 \\ \textbf{(c)}\ \Sigma^k \succ 0 \\ \textbf{(d)}\ 0 < \pi_k \leq 1 \\ \textbf{(e)}\ \sum_{k=1}^{K} \pi_k = 1, \end{cases} \quad \forall\, k = 1, \ldots, K \qquad \Sigma^k = \begin{bmatrix} \Sigma^k_{xx} & \Sigma^k_{x\dot{x}} \\ \Sigma^k_{\dot{x}x} & \Sigma^k_{\dot{x}\dot{x}} \end{bmatrix}$$

Maximum likelihood

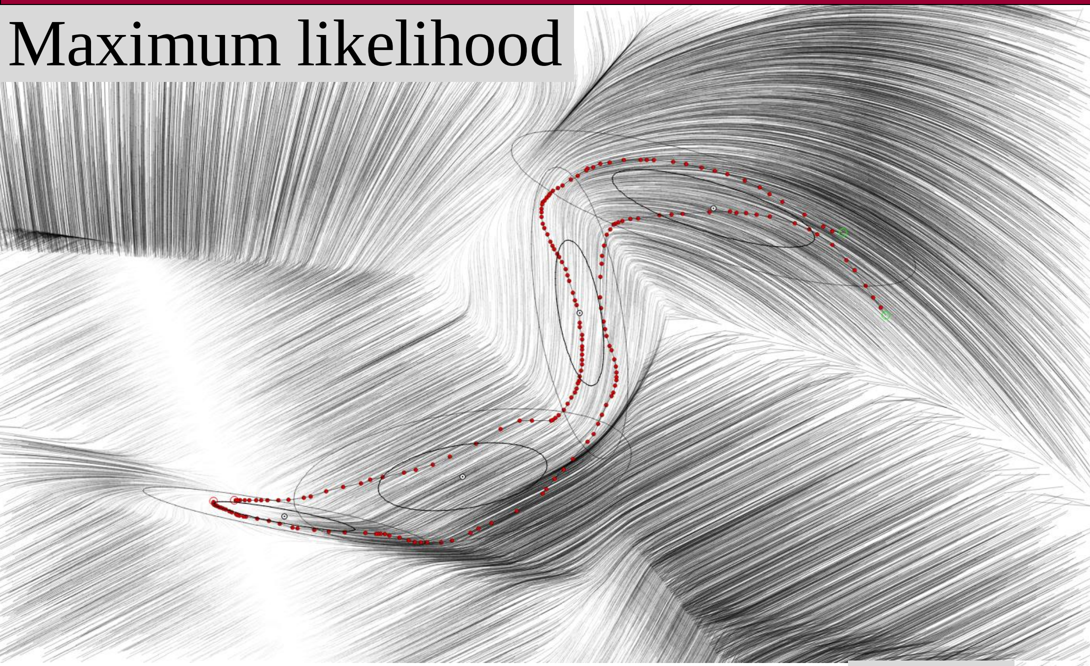

MSE

When trained with mean-square error (MSE) as objective function, the Gauss functions no longer need to fit the distribution of the data

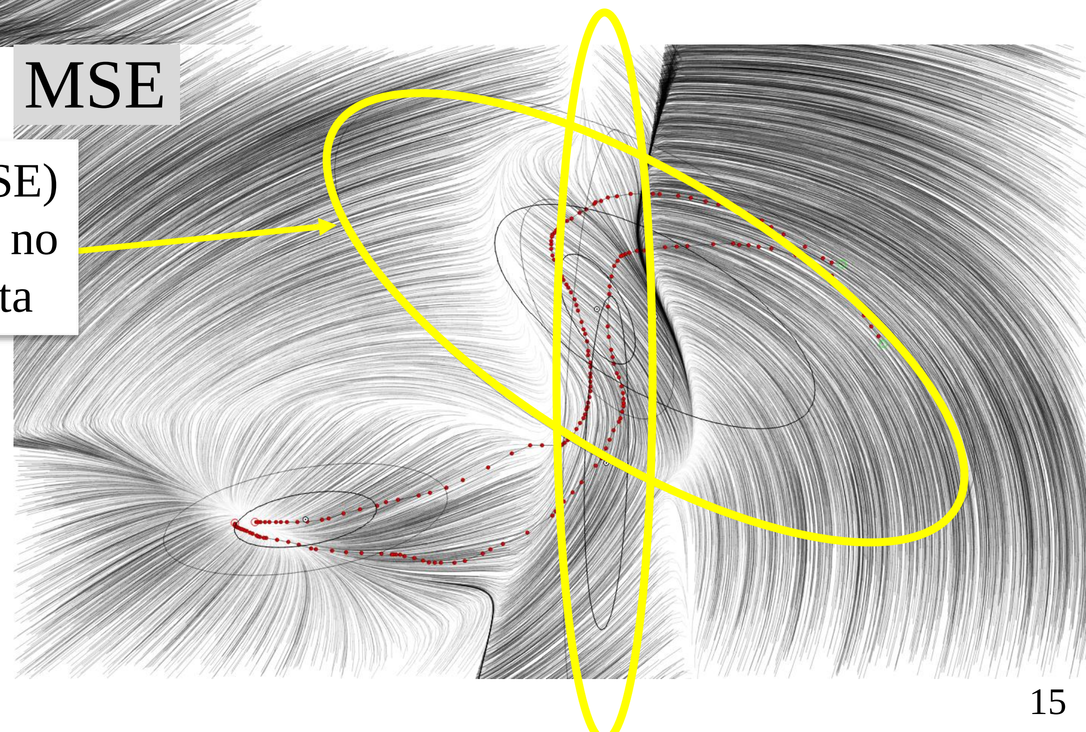

Hyperparameter and pre-selections for SEDS

Prior to training SEDS, the user must make several choices that will influence the quality of the learned model.

The choices are:

- Type of objective function
  - → This will affect the placement of the Gauss functions.
- Number of Gauss functions
  - → This can be automated by using the Bayesian Information Criterion (BIC) – BIC finds a balance between improved quality of the fit and increase in number of parameters.

LASA Handwriting Dataset - Benchmark

https://www.epfl.ch/labs/lasa/datasets/

EPFL

LASA Handwriting Dataset - Benchmark

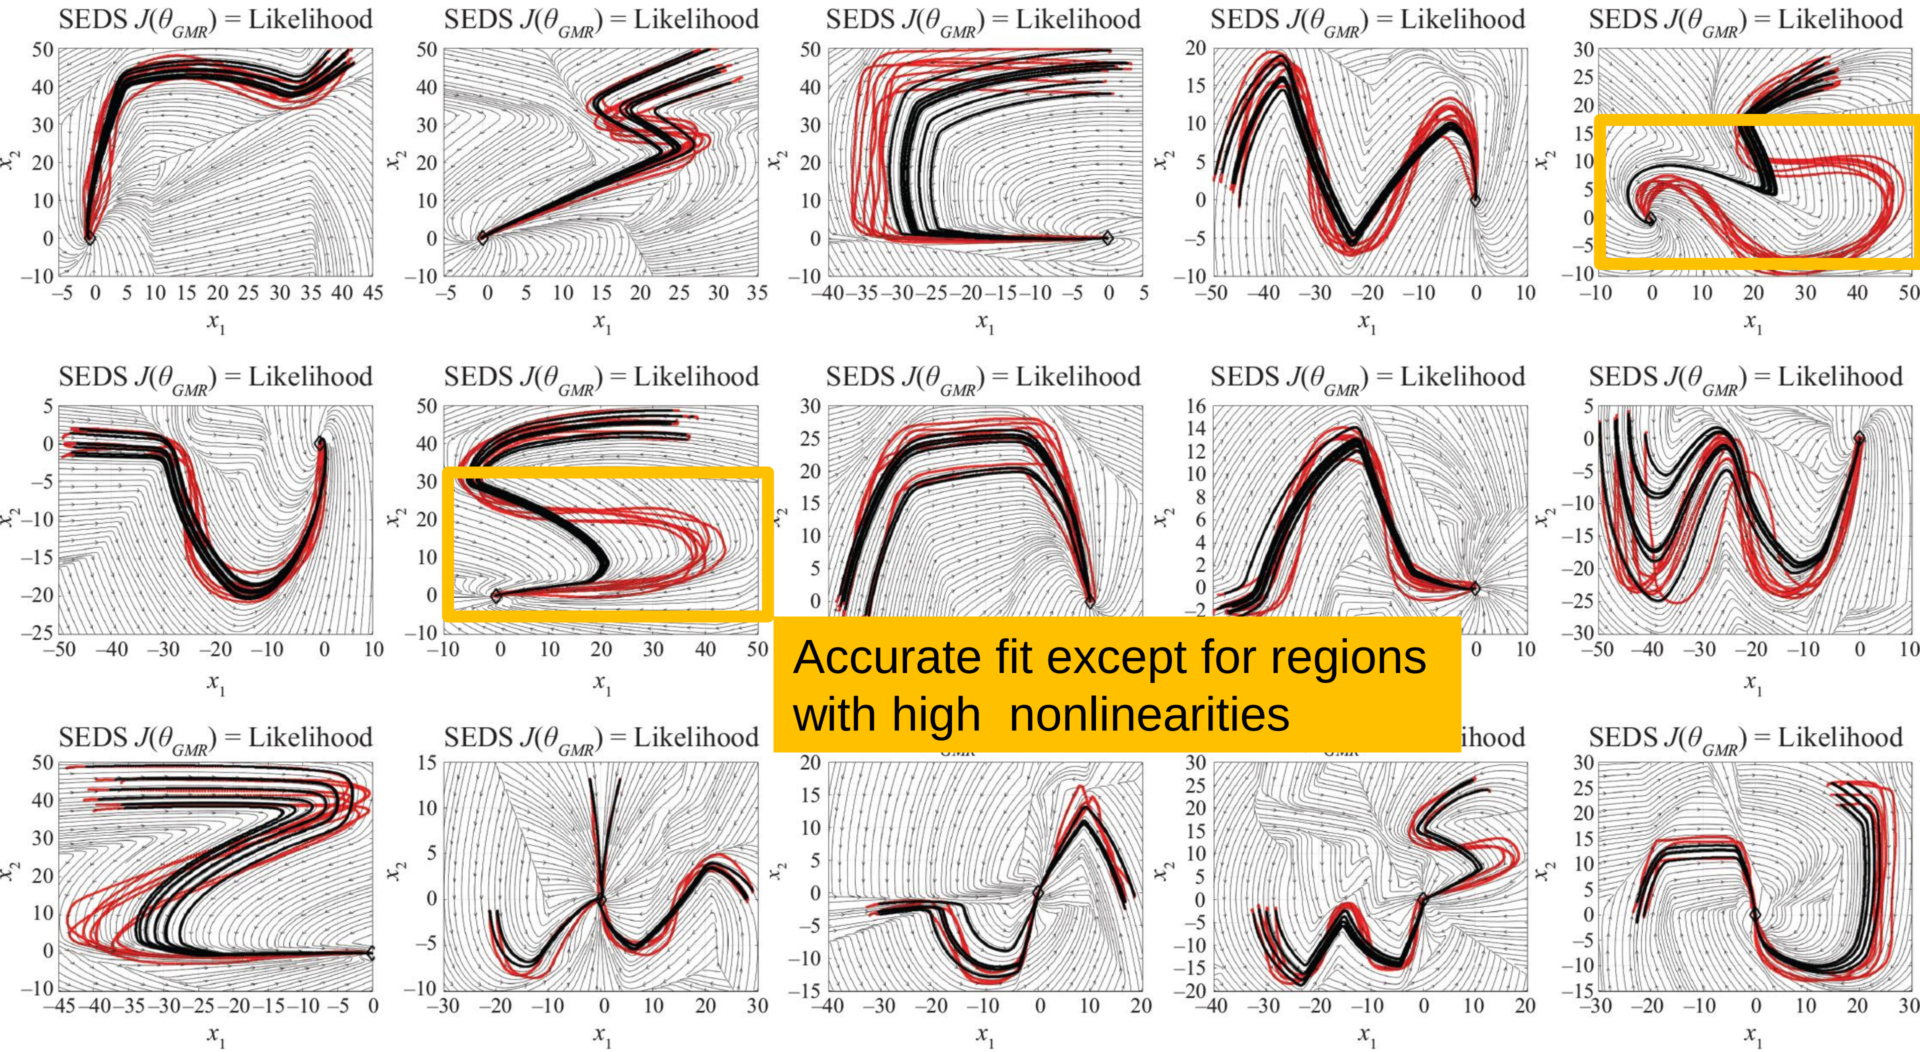

Accurate fit except for regions with high nonlinearities

Demonstrated trajectories

Reproduction from same initial position

Fit with maximum likelihood

LASA Handwriting Dataset - Benchmark

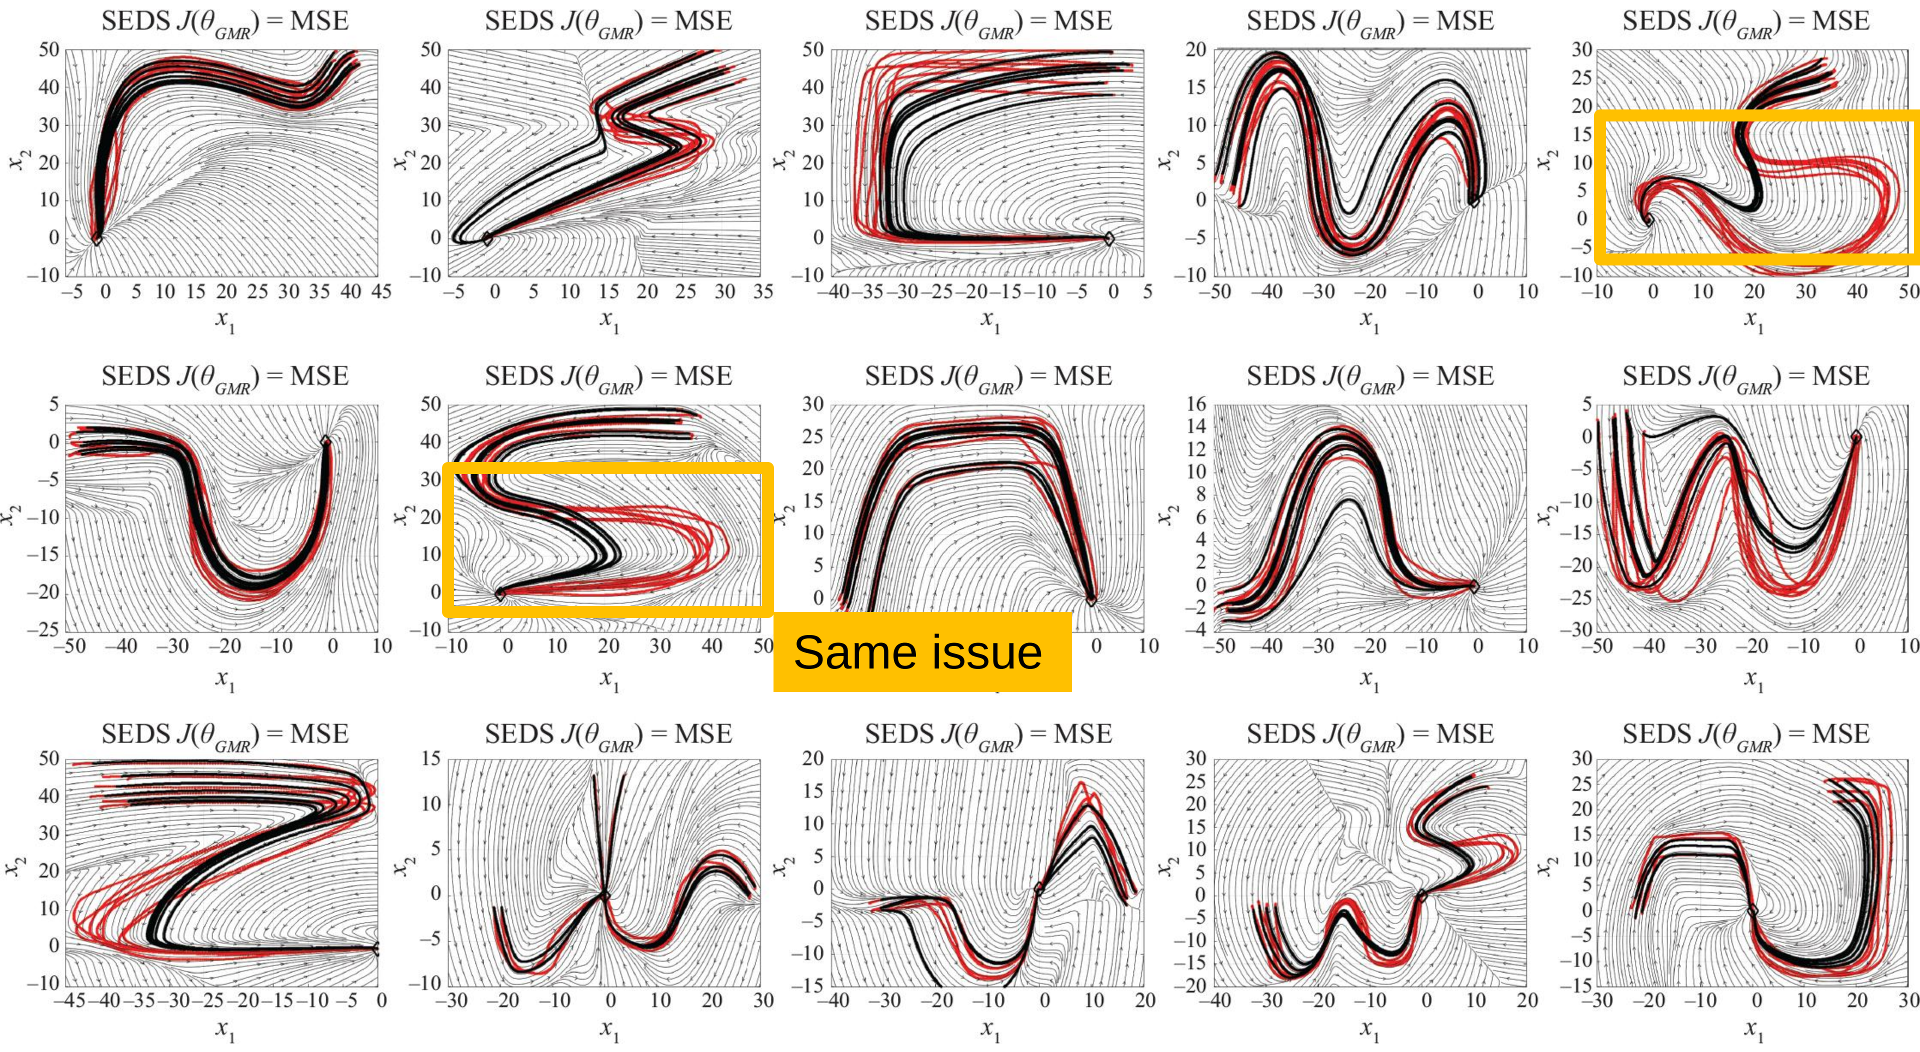

Fit with MSE

LASA Handwriting Dataset - Benchmark

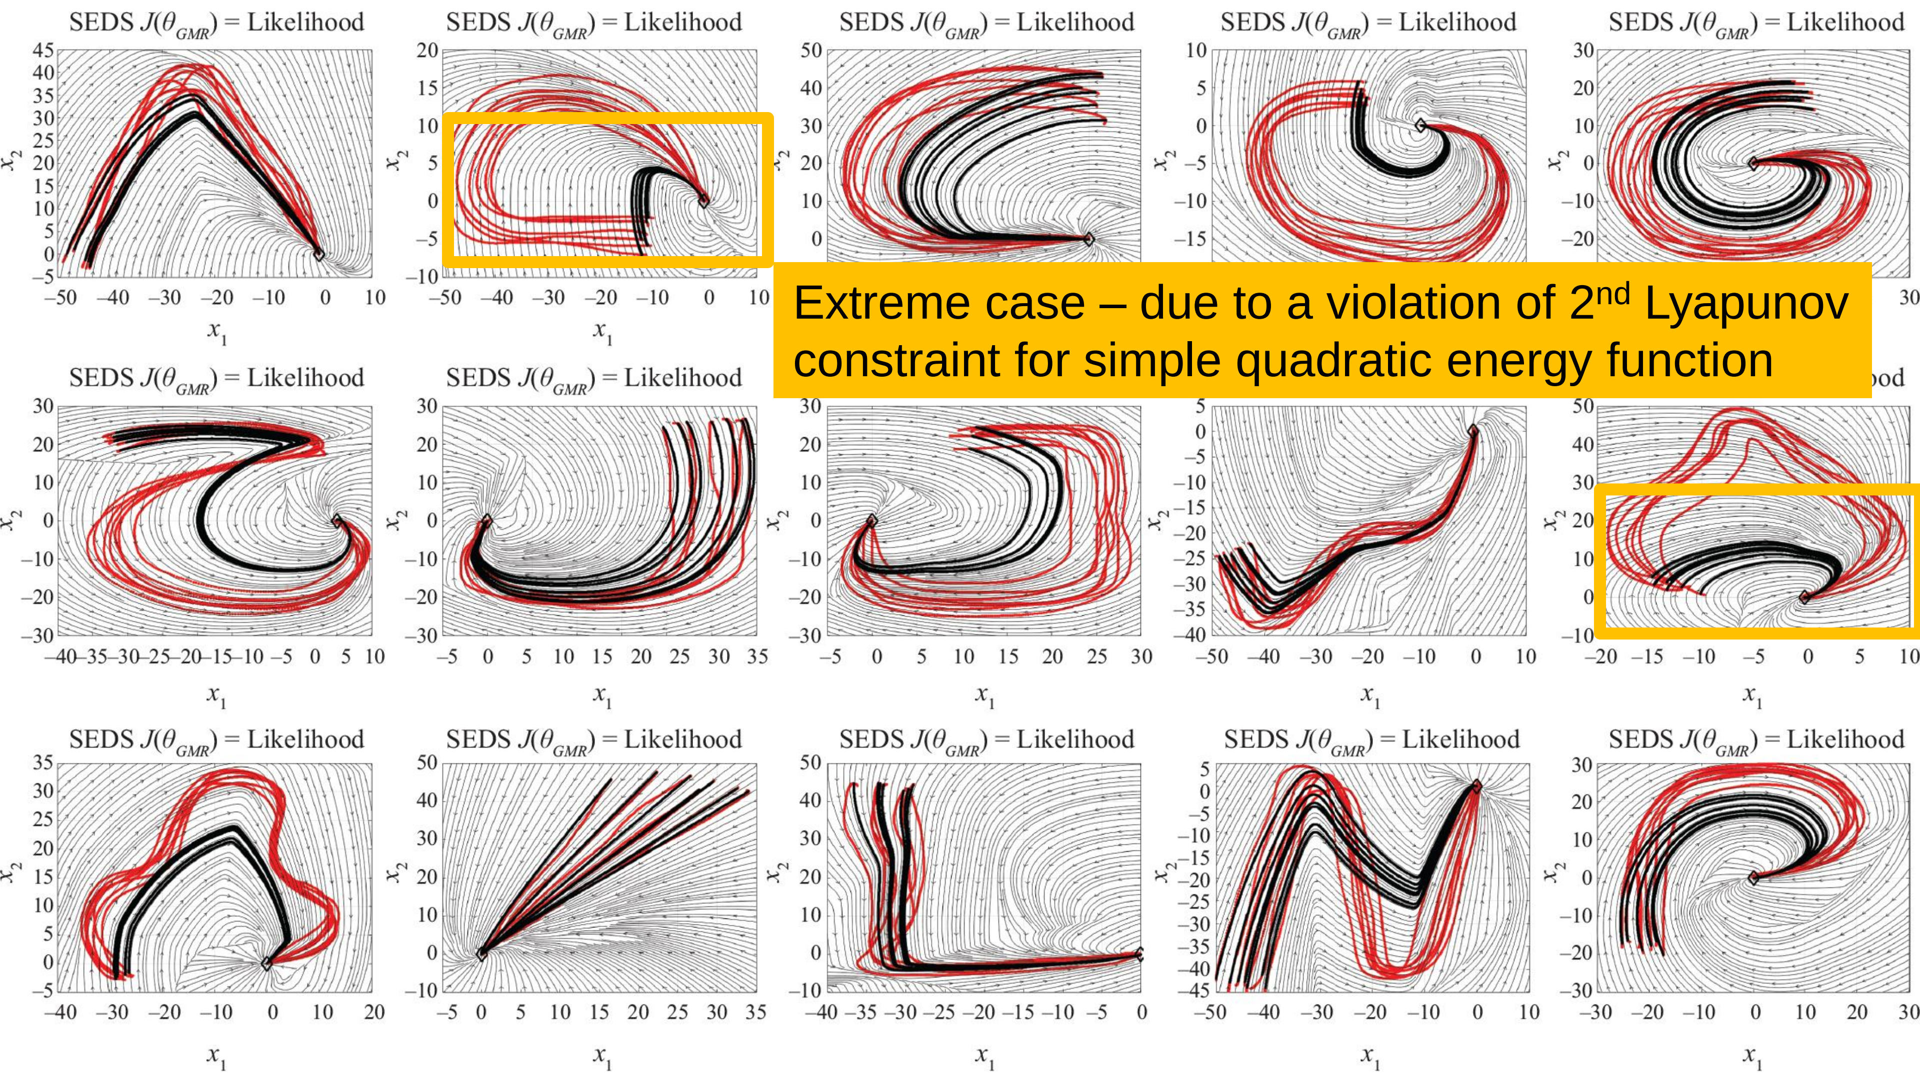

Extreme case – due to a violation of 2nd Lyapunov constraint for simple quadratic energy function

Fit with maximum likelihood

LASA Handwriting Dataset - Benchmark

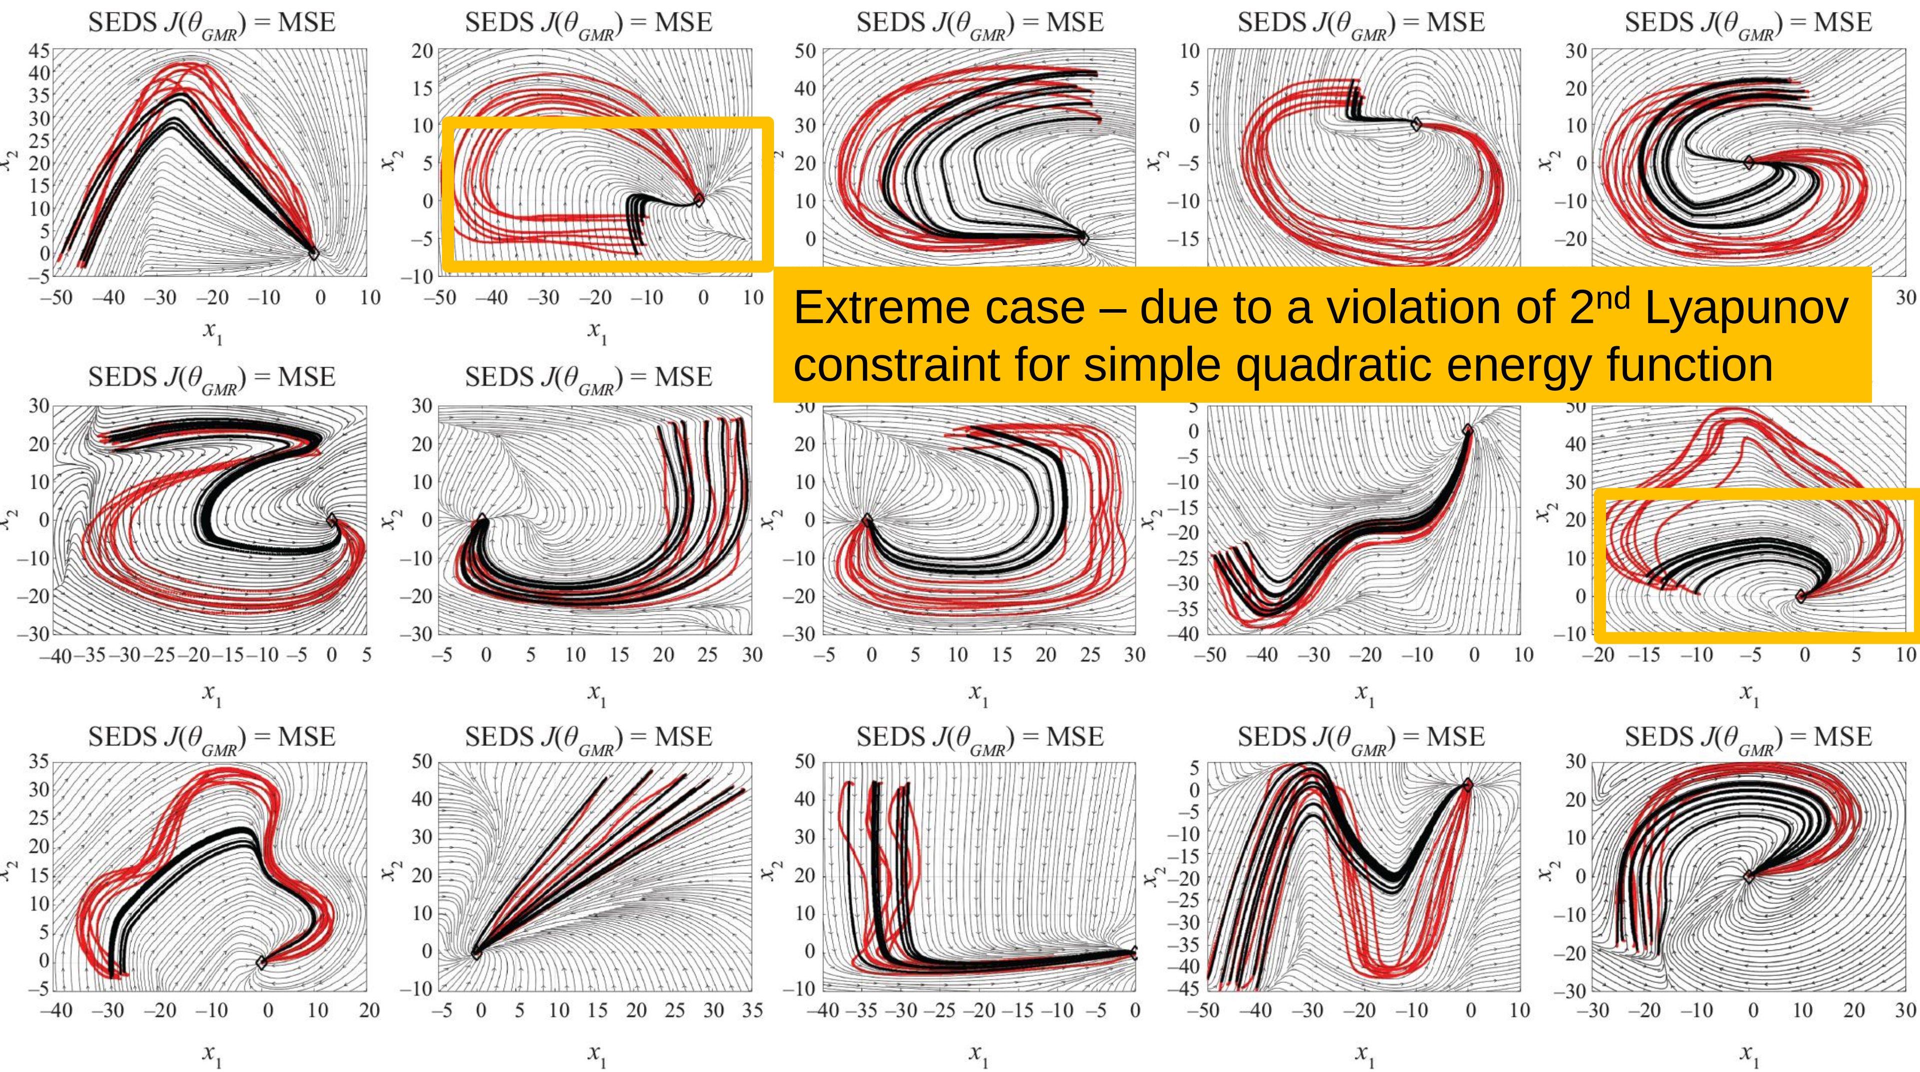

Extreme case – due to a violation of 2nd Lyapunov constraint for simple quadratic energy function

Fit with MSE

Kinesthetic teaching

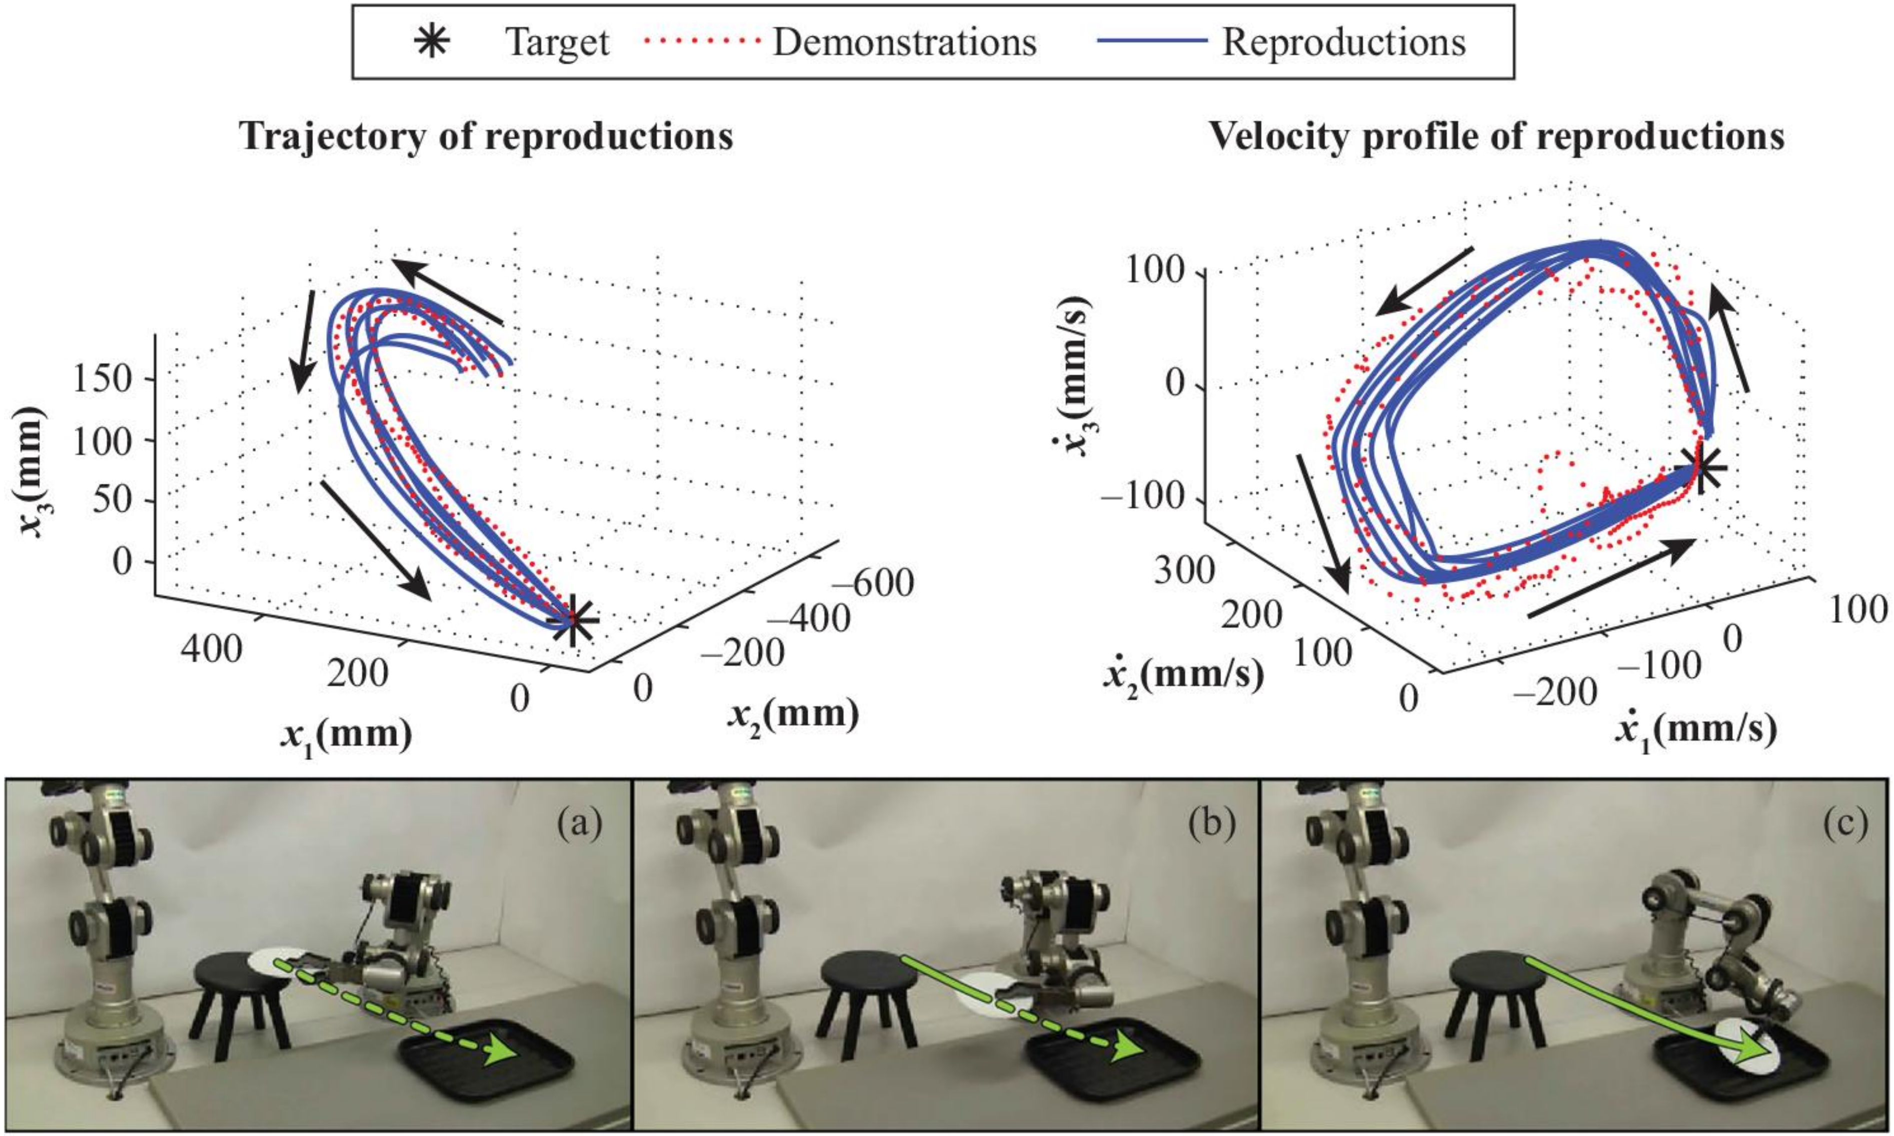
Target
Demonstrations
Reproductions
Trajectory of reproductions
Velocity profile of reproductions
150
100
50
0
x_3(mm)
400
200
0
x_1(mm)
0
−200
−400
−600
x_2(mm)
100
0
−100
$\dot{x}_3$(mm/s)
300
200
100
0
$\dot{x}_2$(mm/s)
−200
−100
0
100
$\dot{x}_1$(mm/s)
(a)
(b)
(c)

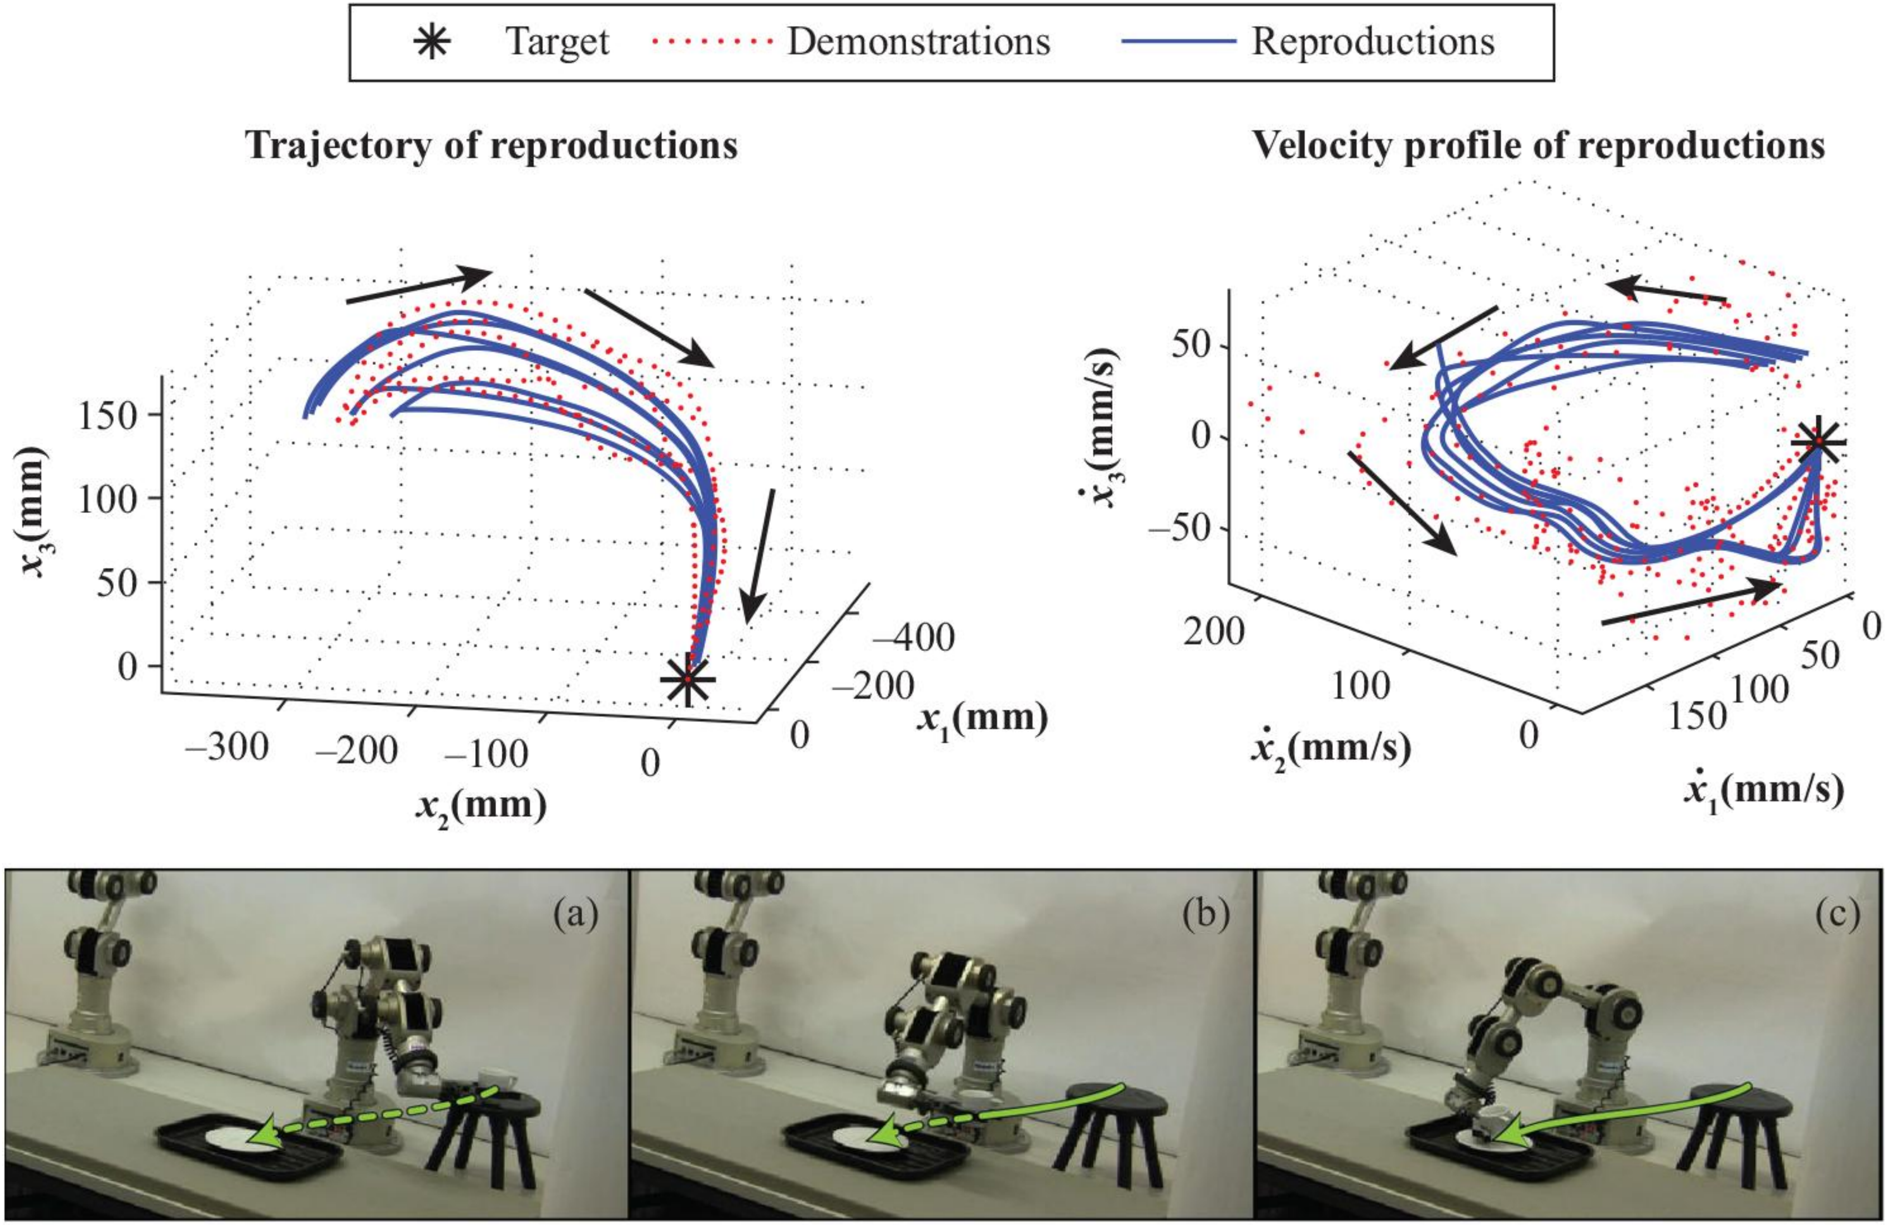
Target
Demonstrations
Reproductions
Trajectory of reproductions
Velocity profile of reproductions
x_3(mm)
x_2(mm)
x_1(mm)
$\dot{x}_3$(mm/s)
$\dot{x}_2$(mm/s)
$\dot{x}_1$(mm/s)
(a)
(b)
(c)

Reproduction


http://lasa.epfl.ch/
Applying Disturbance During Reproduction

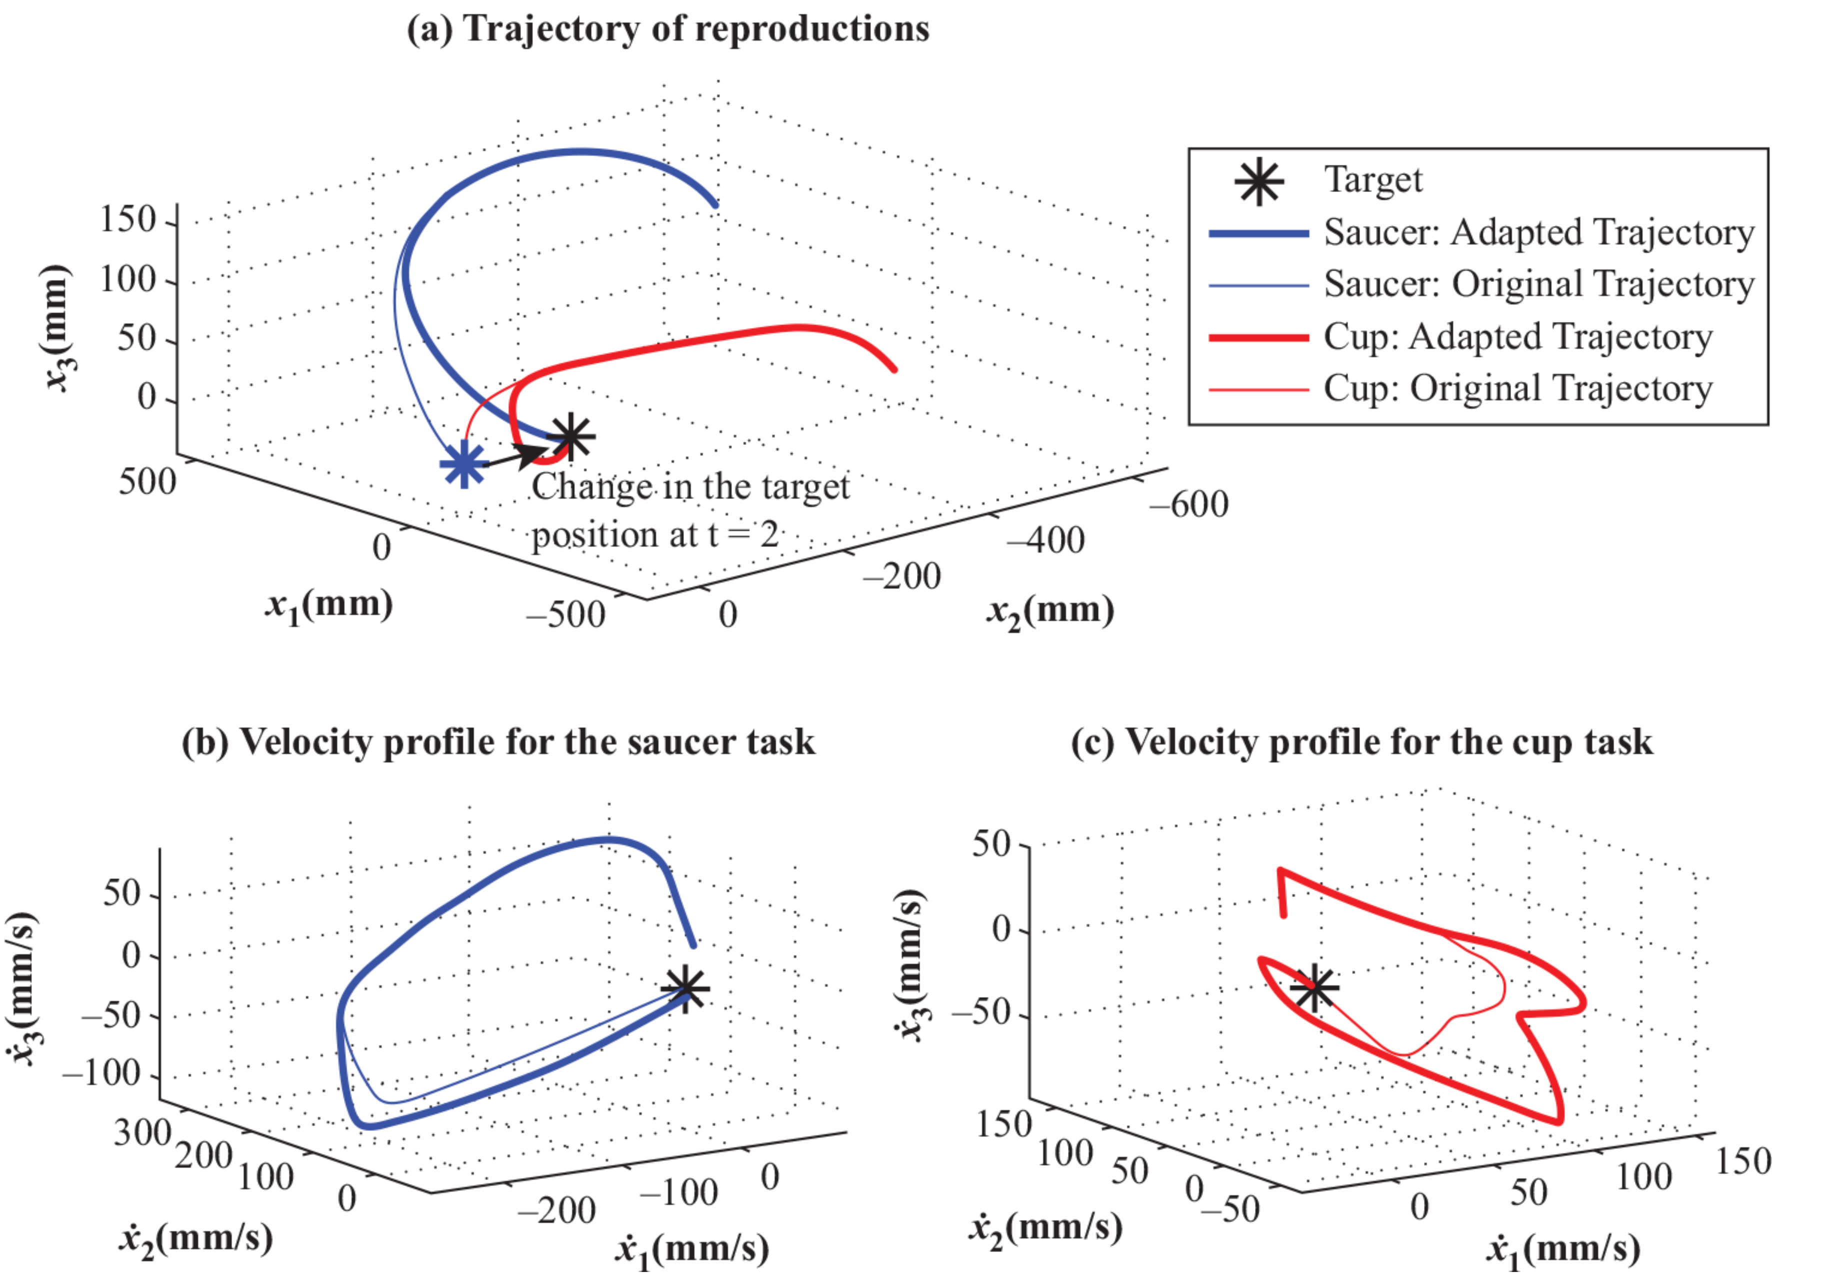
(a) Trajectory of reproductions
150
100
50
0
x_3(mm)
500
0
–500
x_1(mm)
0
–200
–400
–600
x_2(mm)
Change in the target
position at t = 2
Target
Saucer: Adapted Trajectory
Saucer: Original Trajectory
Cup: Adapted Trajectory
Cup: Original Trajectory
(b) Velocity profile for the saucer task
50
0
–50
–100
$\dot{x}_3$(mm/s)
300
200
100
0
$\dot{x}_2$(mm/s)
–200
–100
0
$\dot{x}_1$(mm/s)
(c) Velocity profile for the cup task
50
0
–50
$\dot{x}_3$(mm/s)
150
100
50
0
–50
$\dot{x}_2$(mm/s)
0
50
100
150
$\dot{x}_1$(mm/s)


EPFL
Robustness to Perturbations

SEDS Summary

- Automatically estimate *globally asymptotically stable* dynamical systems from sampled trajectories
- Extension of Gaussian Mixture Model
  - Uses same objective function (maximum likelihood)
  - Add new set of constraints to enforce stability
- Stability is guaranteed through Lyapunov stability constraints
  - Assumes a quadratic Lyapunov function
- High accuracy for a large number of nonlinear dynamics
- Limitations:
  - Non-convex optimization
  - Poor accuracy for highly nonlinear dynamics (high curvature)

Extensions to SEDS

Approach	Stability ensured via
SEDS (Constrained-GMR) [1]	QLF (Lyapunov)
Tau-SEDS (SEDS-extension) [2]	Complex (Lyapunov) Function + Diffeomorphic Transformation
CDSP (SEDS-extension) [3]	Partial Contraction Theory
LPV-DS (GMM-based) [4,5]	P-QLF (Lyapunov)
ANPs (NN-based) [6]	P-QLF (Lyapunov)

[1] S. Khansari-Zadeh and A. Billard. Learning stable nonlinear dynamical systems with Gaussian mixture models. IEEE Transactions on, 27(5):943–957, Oct 2011.
[2] K. Neumann and A. Billard. Learning robot motions with stable dynamical systems under diffeomorphic transformations. Robotics and Autonomous Systems. 2015
[3] H. Ravichandar, I. Salehi and A. Dani. Learning partially contracting dynamical systems from demonstrations. In Proc. of the 1st Conference on Robot Learning (CoRL). Nov. 2017.
[4] Figueroa N., and Billard, A. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. In Proc. of the 2nd Conference on Robot Learning. Oct 2018.
[5] Figueroa, Nadia, and Aude Billard. "Locally active globally stable dynamical systems: Theory, learning, and experiments." *The International Journal of Robotics Research* 41.3 (2022): 312-347.
[6] Totsila, Dionis, et al. "Sensorimotor Learning with Stability Guarantees via Autonomous Neural Dynamic Policies." *IEEE Robotics and Automation Letters* (2025).