Learning Control Laws

Linear Parameter Varying Dynamical Systems (LPVDS)

SEDS on Highly Non-Linear Trajectories

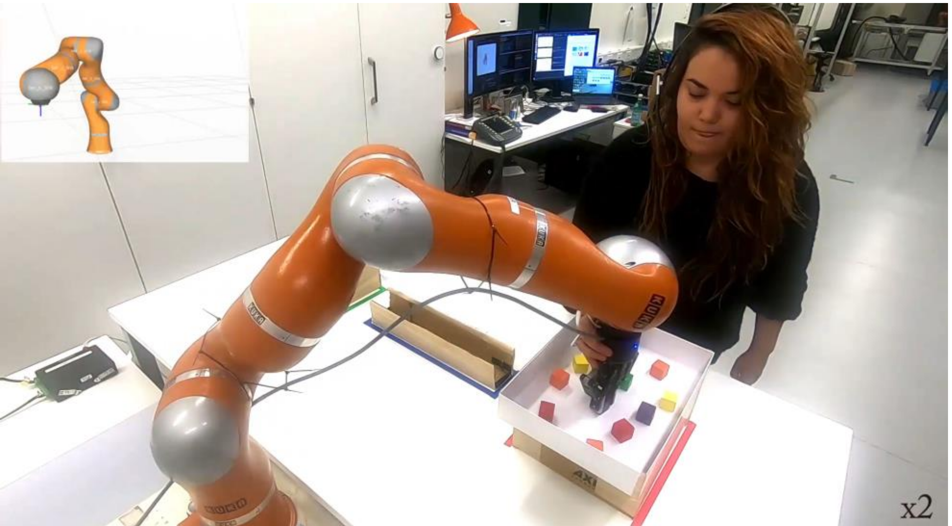

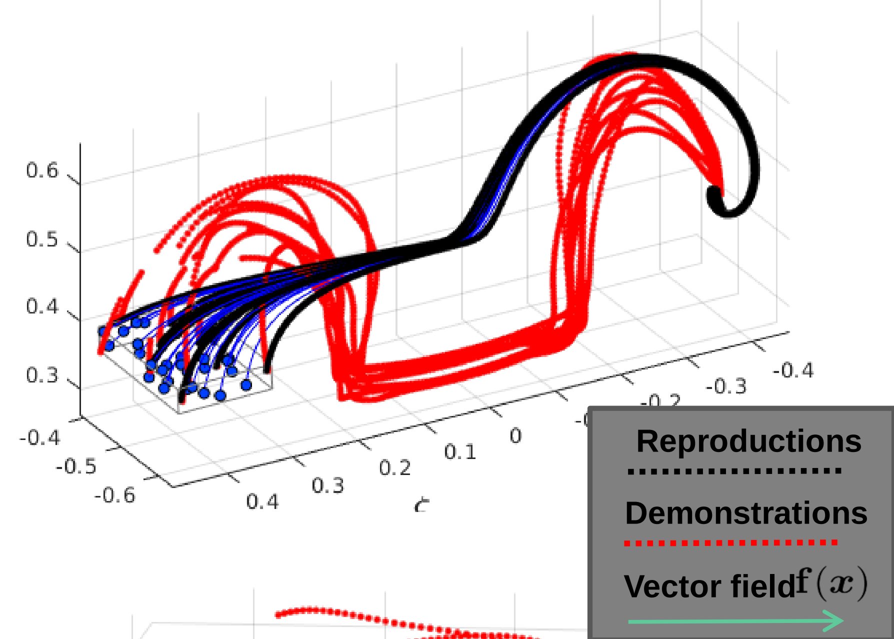

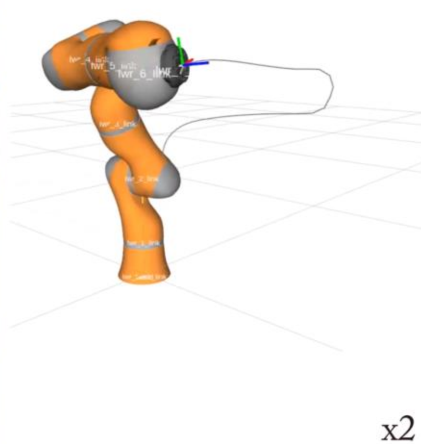

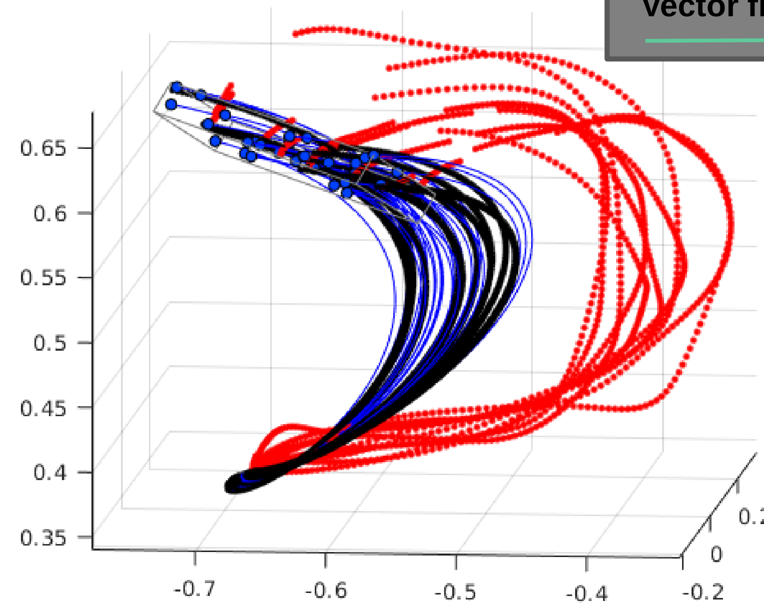

SEDS on Highly Non-Linear Trajectories

$$\dot{\boldsymbol{x}} = \mathbf{f}(\boldsymbol{x}) = \sum_{k=1}^{K} \gamma_k(\boldsymbol{x})(\mathbf{A}_k \boldsymbol{x} + \boldsymbol{b}_k)$$

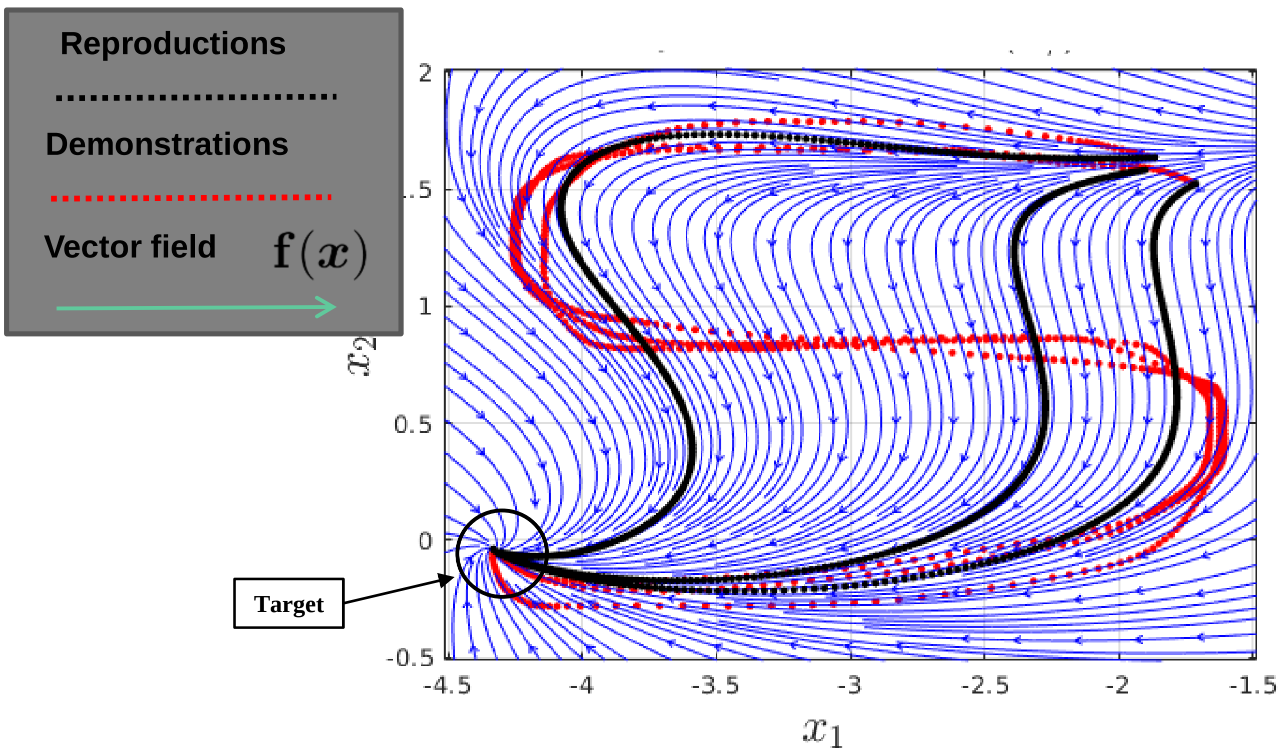

- ✓ **Convergence ensured**
- ➢ **Inaccurate Reproduction of highly non-linear motions**

Why?

SEDS on Highly Non-Linear Trajectories

$$\dot{\boldsymbol{x}} = \mathbf{f}(\boldsymbol{x}) = \sum_{k=1}^{K} \gamma_k(\boldsymbol{x})(\mathbf{A}_k \boldsymbol{x} + \boldsymbol{b}_k)$$

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

SEDS Lyapunov Function

Highly Non-linear trajectories violate stability condition

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

If V is too conservative.

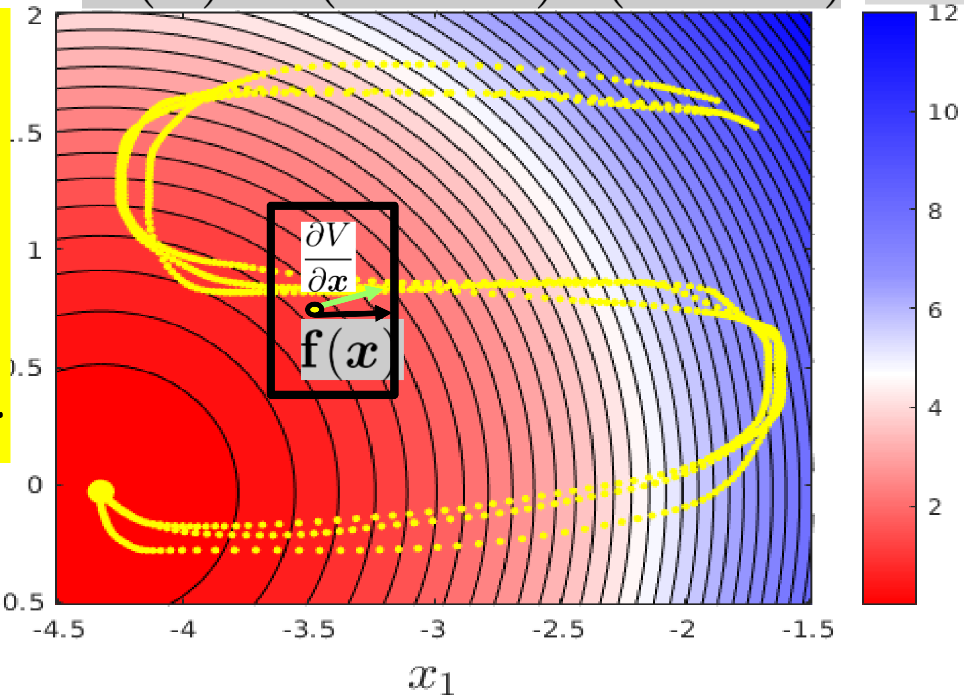

SEDS on Highly Non-Linear Trajectories

$$\dot{\boldsymbol{x}} = \mathbf{f}(\boldsymbol{x}) = \sum_{k=1}^{K} \gamma_k(\boldsymbol{x})(\mathbf{A}_k \boldsymbol{x} + \boldsymbol{b}_k)$$

State dependent parameter vector ($\gamma_k(\boldsymbol{x})$)

Linear Time-Invariant (LTI) DS ($\mathbf{A}_k \boldsymbol{x} + \boldsymbol{b}_k$)

Stability of LTI can be shown if ∃ a generic Lyapunov Function:

$$V(x) = (x - x^*)^T P (x - x^*), \quad P = P^T, P \succ 0$$

Theorem:

The nonlinear DS above is Globally Asymptotically Stable at x^* if $\exists \; P = P^T, P \succ 0$, with $V(x) = (x - x^*)^T P (x - x^*)$, such that:

$$\begin{cases} (A^k)^T P + PA^k = Q^k, \quad Q^k = (Q^k)^T \\ b^k = -A^k x^* \end{cases} \qquad \forall k = 1, \ldots. K$$

See Theorem 3.3 (Book)

Learning Non-linear DS with GMM's and P-QLF

Goal: Learn the parameters of a non-linear DS with P-QLF

Quadratic Lyapunov Function (QLF)

$$V(x) = (x - x^*)^T (x - x^*)$$

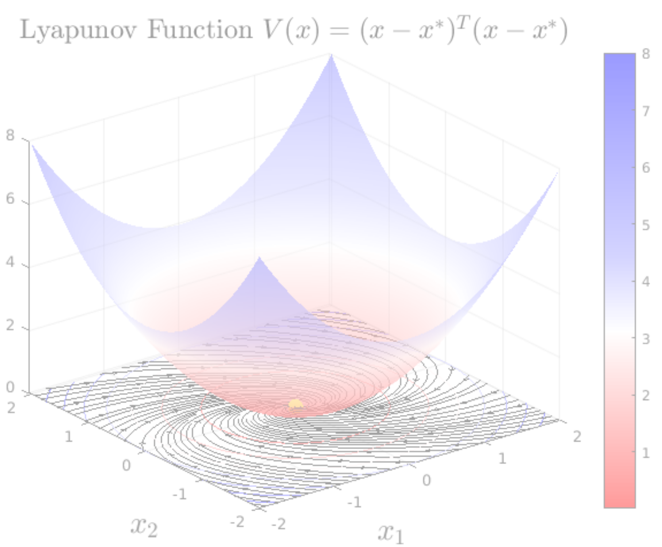

Parameterized Quadratic Lyapunov Function (P-QLF)

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

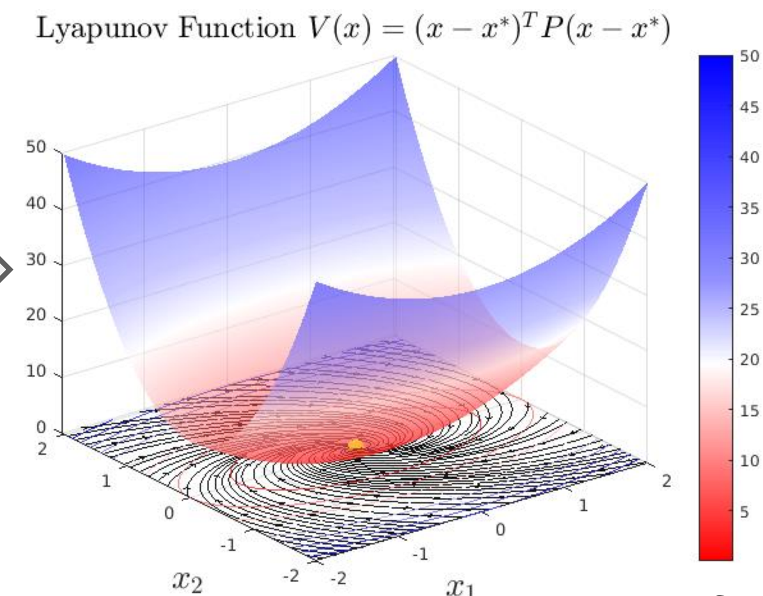

6

P's effect is of a reshaping of the Lyapunov function

P's effect is of a reshaping of the Lyapunov function

$$P = \begin{bmatrix} 1 & 0 \\ 0 & 20 \end{bmatrix}$$

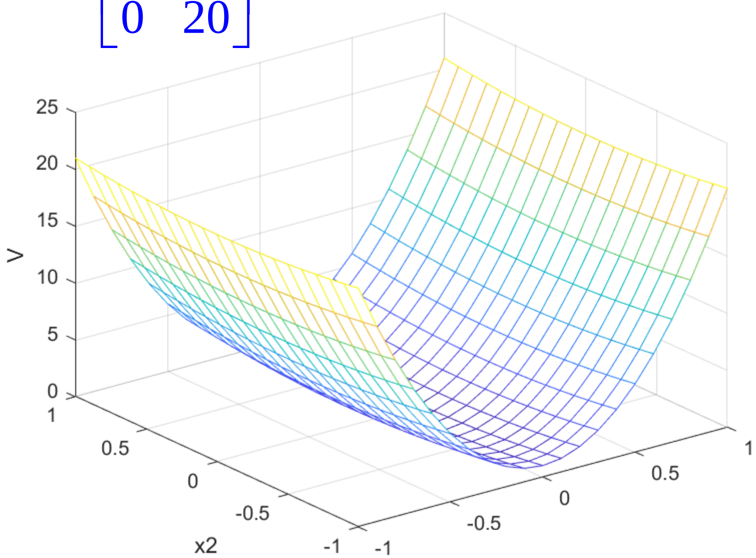

$$P = \begin{bmatrix} 20 & 0 \\ 0 & 1 \end{bmatrix}$$

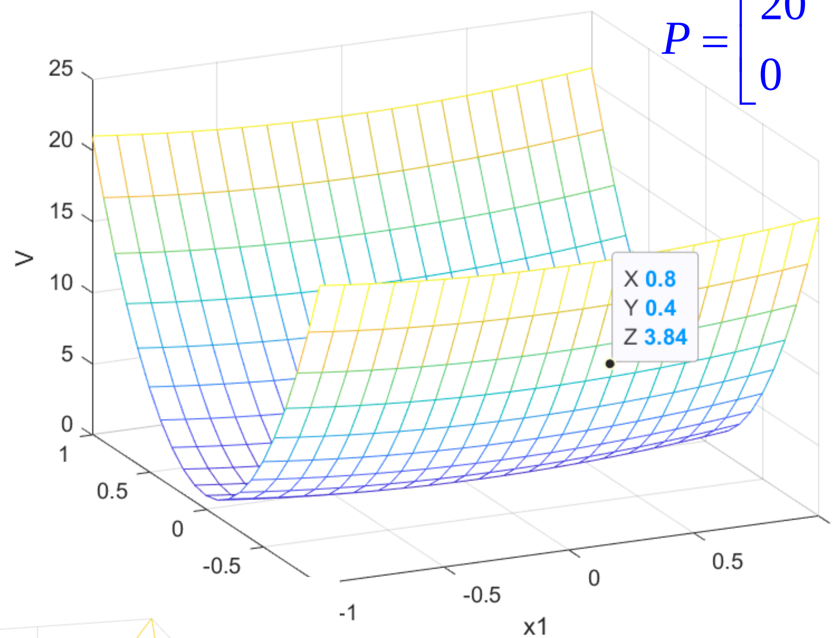

$$P = \begin{bmatrix} 2 & 1.5 \\ 1.5 & 2 \end{bmatrix}$$

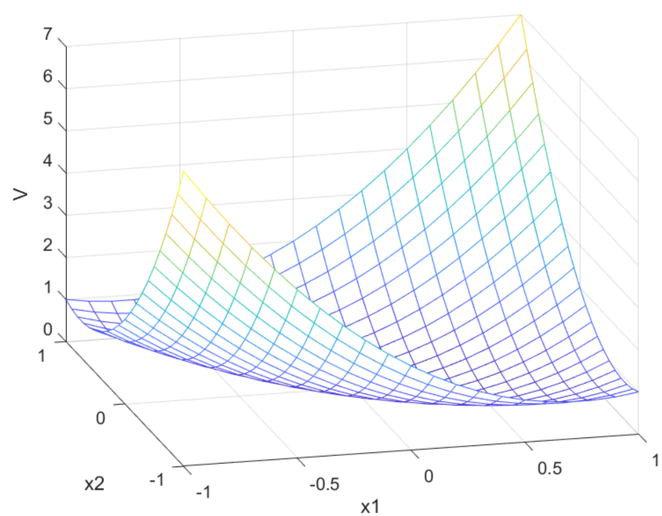

P-QLF Stability Condition

<u>Parameterized Quadratic Lyapunov Function (P-QLF)</u>

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

$$\mathbf{P} = \mathbf{P}^T \succ 0$$

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

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

$$\dot{V}(\boldsymbol{x}) = \frac{\partial V}{\partial \boldsymbol{x}} \mathbf{f}(\boldsymbol{x}) < 0 \longrightarrow \mathbf{A}^T \mathbf{P} + \mathbf{P} \mathbf{A} \prec 0$$

Enforce that the eigenvalues be negative!

Optimization of P-QLF – 1st formulation

Objective function: Maximum likelihood or Mean-square error

Constraints:

$$\begin{cases} b^k = -A^k x^* \\ \left(A^k\right)^T P + PA^k \prec 0 \end{cases} \quad \forall k = 1, \ldots K$$

Joint estimation of P and A makes the problem non-convex
Depends on good initial guess for P.

Idea: decouple the problem in two-steps:

1) Estimate the A^k matrices with standard GMM
2) Estimate P so as to enforce the stability constraints

Learning Non-linear DS with GMM's and P-QLF

(Proposed Approach) We decouple the density estimation from the DS parameters

$$\mathbf{f}(\boldsymbol{x}) = \sum_{k=1}^{K} \gamma_k(\boldsymbol{x})(\mathbf{A}_k\boldsymbol{x} + \boldsymbol{b}_k)$$

Step 1: **Learn the GMM density solely on position variables**

$$p(\boldsymbol{x}|\theta_\gamma) = \sum_{k=1}^{K} \pi_k \mathcal{N}(\boldsymbol{x}|\boldsymbol{\mu}^k, \boldsymbol{\Sigma}^k)$$
$$\theta_\gamma = \{\pi_k, \boldsymbol{\mu}^k, \boldsymbol{\Sigma}^k\}_{k=1}^{K}$$

$$\gamma_k(\boldsymbol{x}) = \frac{\pi_k p(\boldsymbol{x}|k)}{\sum_j \pi_j p(\boldsymbol{x}|j)}$$
$$\mathbf{A}_k = \boldsymbol{\Sigma}^k_{\boldsymbol{x}\dot{\boldsymbol{x}}}(\boldsymbol{\Sigma}^k_{\boldsymbol{x}})^{-1}$$
$$\boldsymbol{b}_k = \boldsymbol{\mu}^k_{\dot{\boldsymbol{x}}} - \mathbf{A}_k\boldsymbol{\mu}^k_{\boldsymbol{x}}$$

$p(x) \sim N(x;\mu,\Sigma)$

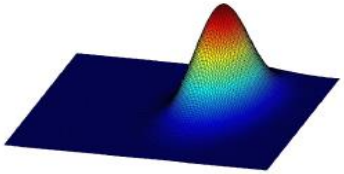

2D projection of a normal distribution

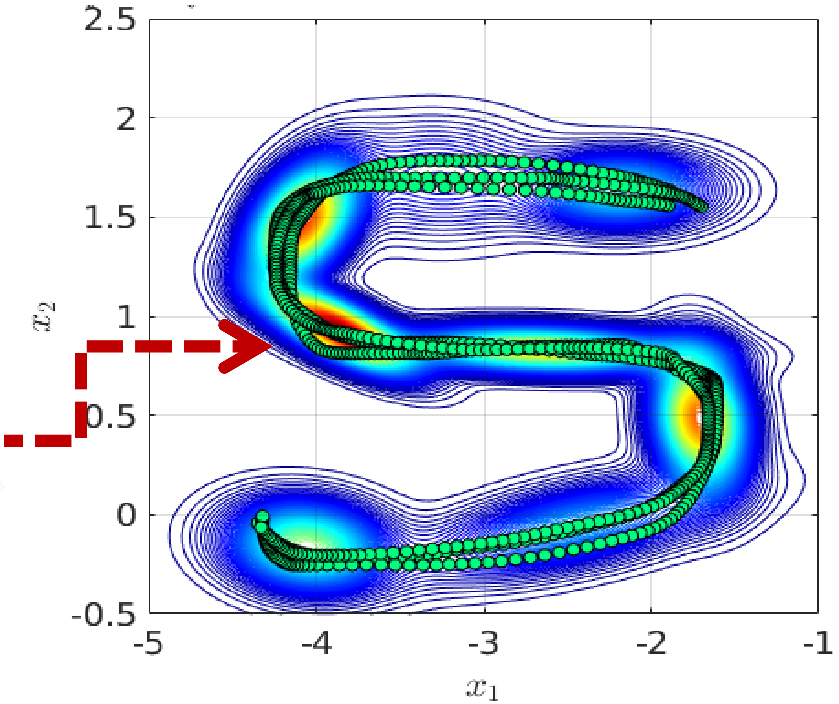

Learning Non-linear DS with GMM's and P-QLF

(Proposed Approach) We decouple the density estimation from the DS parameters

$$\mathbf{f}(\boldsymbol{x}) = \sum_{k=1}^{K} \gamma_k(\boldsymbol{x})(\mathbf{A}_k\boldsymbol{x} + \boldsymbol{b}_k) \qquad \theta_\gamma = \{\pi_k, \boldsymbol{\mu}^k, \boldsymbol{\Sigma}^k\}_{k=1}^{K}$$

Step 2: Estimate DS parameters via non-convex Semi-Definite Programming

$$\min_{\theta_f} J(\theta_f) = \textbf{MSE}$$

Stability Constraints

$$(\mathbf{A}_k)^T\mathbf{P} + \mathbf{P}\mathbf{A}_k \prec 0$$

$$\boldsymbol{b}_k = -\mathbf{A}_k\boldsymbol{x}^*$$

$$\mathbf{P} = \mathbf{P}^T \succ 0$$

$$\forall\ k = 1, \ldots, K$$

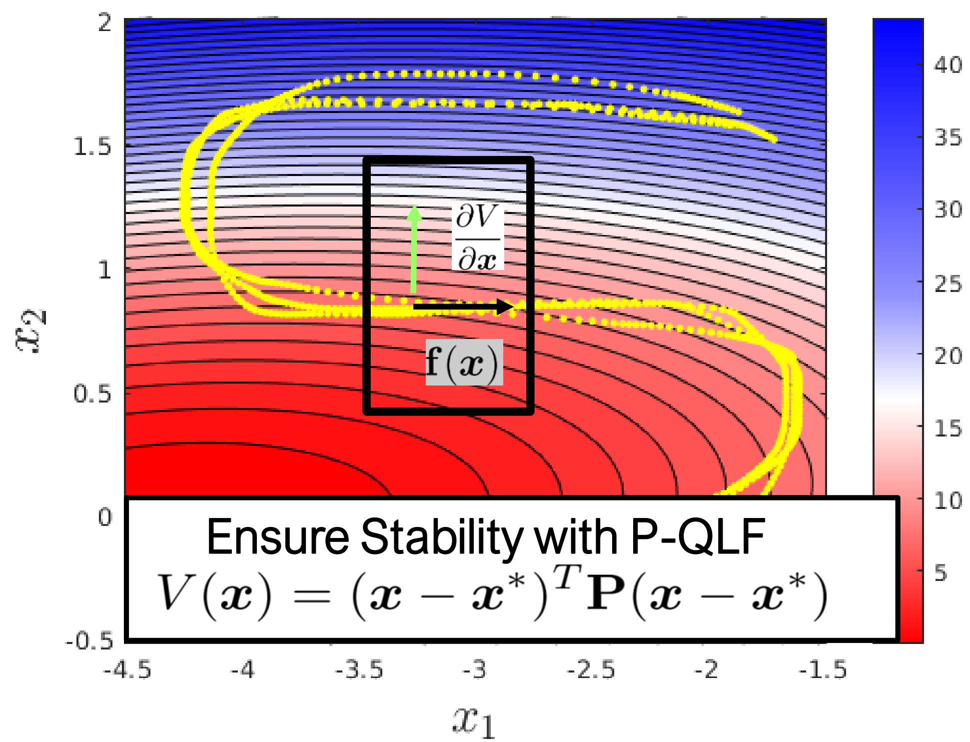

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Learning Non-linear DS with GMM's and P-QLF

(Caveat) Since the density estimation is decoupled, DS reproduction accuracy relies on how whether the mixture of Gaussians fits well the dynamics of the data.

→ Need to devise a new procedure to train GMM that is informed by the fact that data are samples of a DS.

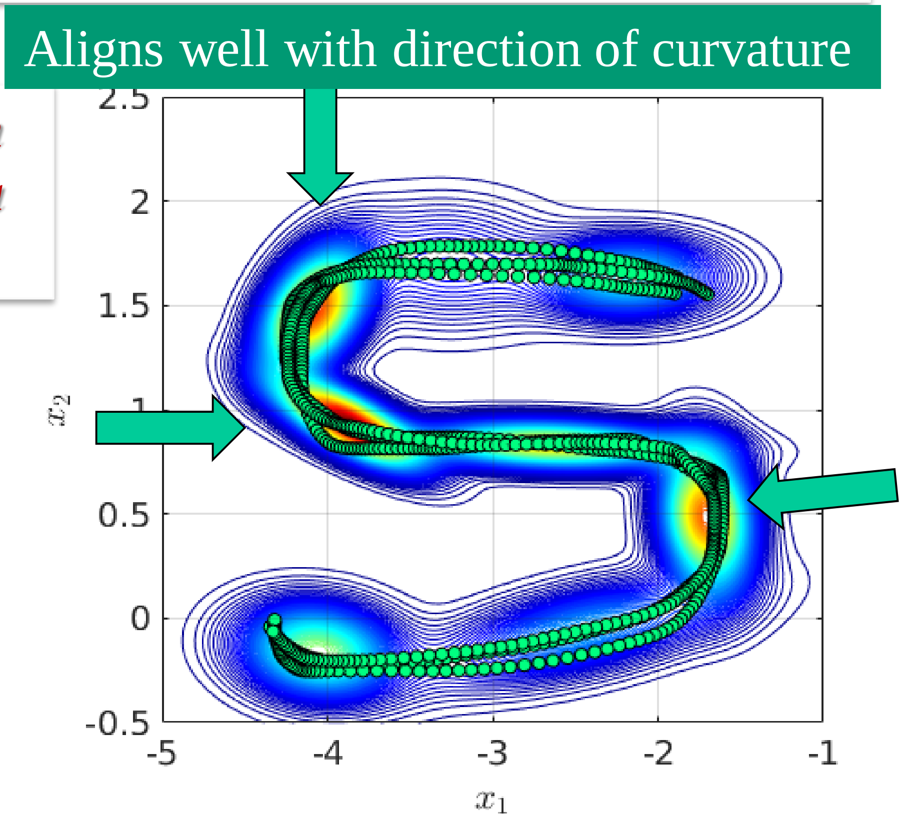

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Example: Training Dataset

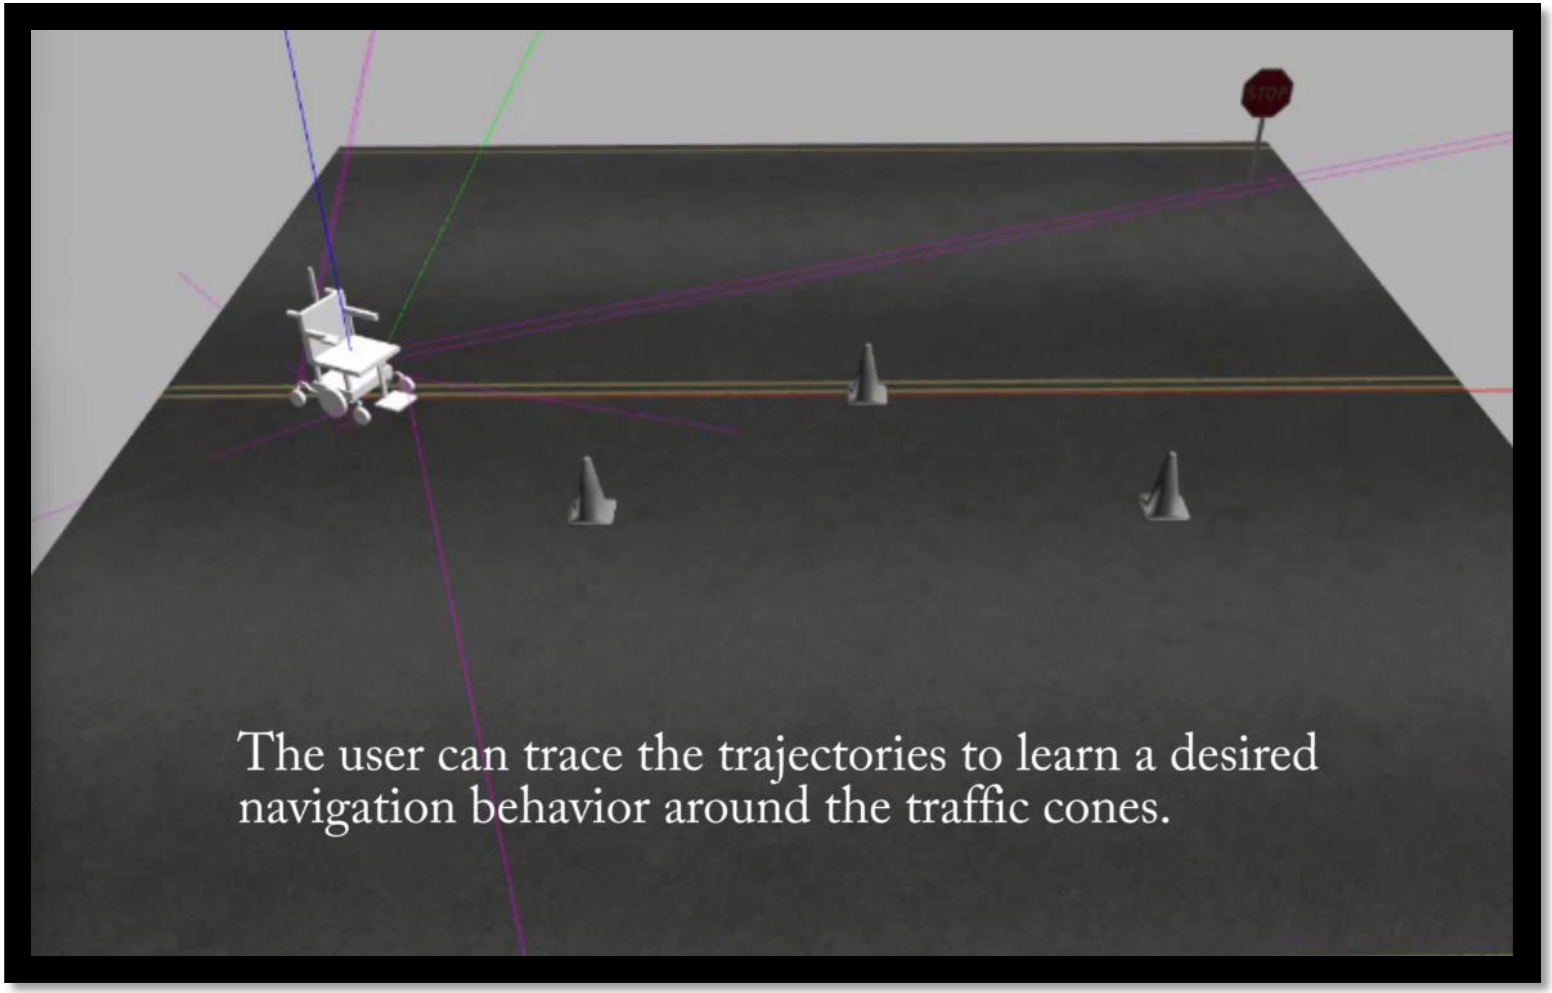

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Fit with traditional GMM training

Use classic EM estimation to fit the Gauss functions

Use Bayesian Information Criterion (BIC) to determine optimal number of Gauss functions.

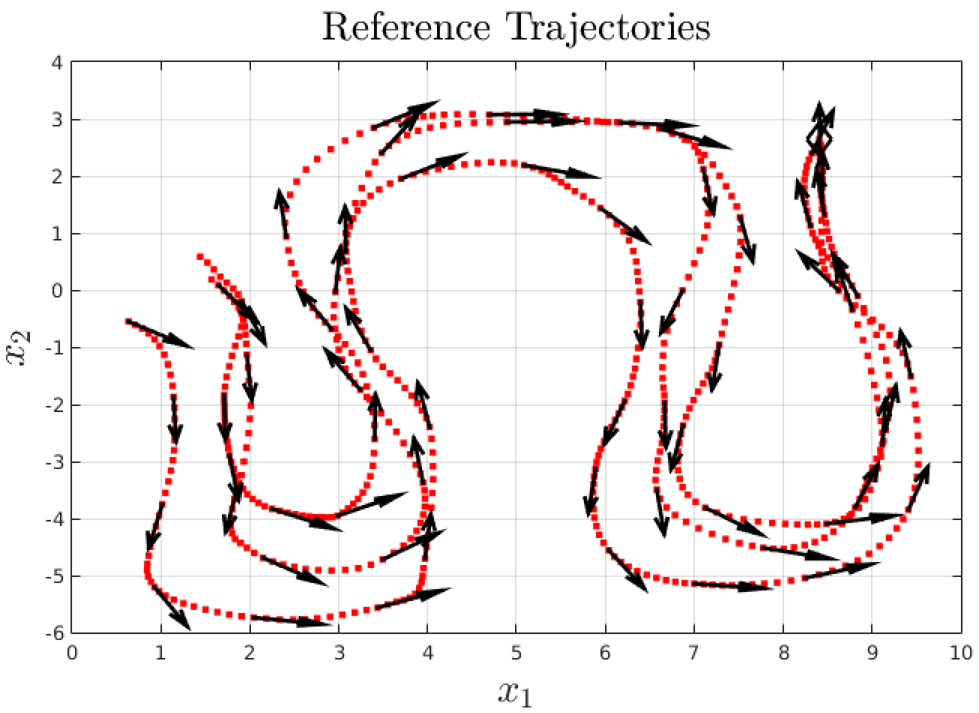

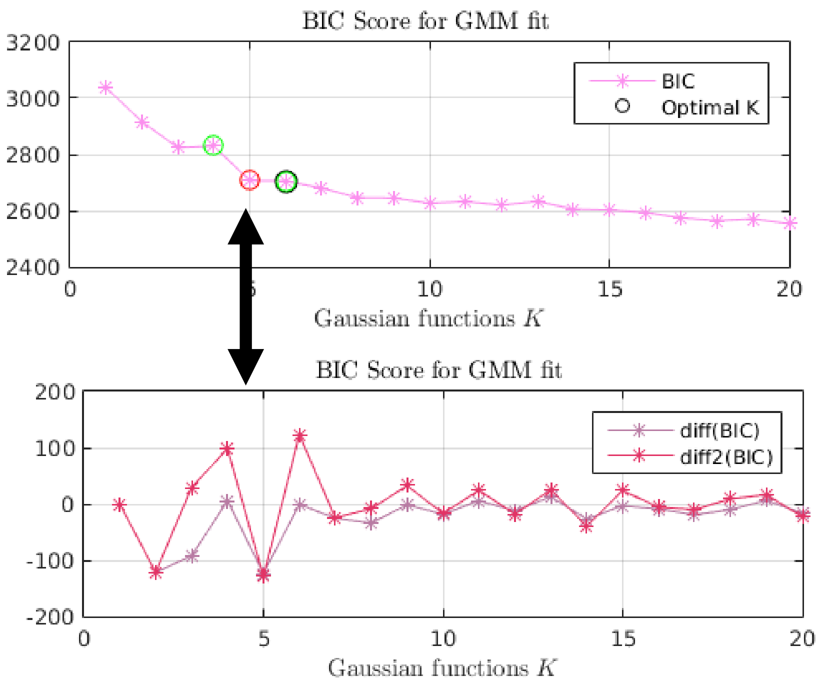

Repeat with different initial conditions and compare the fits.

N. Figueroa and A. Billard. . *A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.*

Result from traditional GMM fit

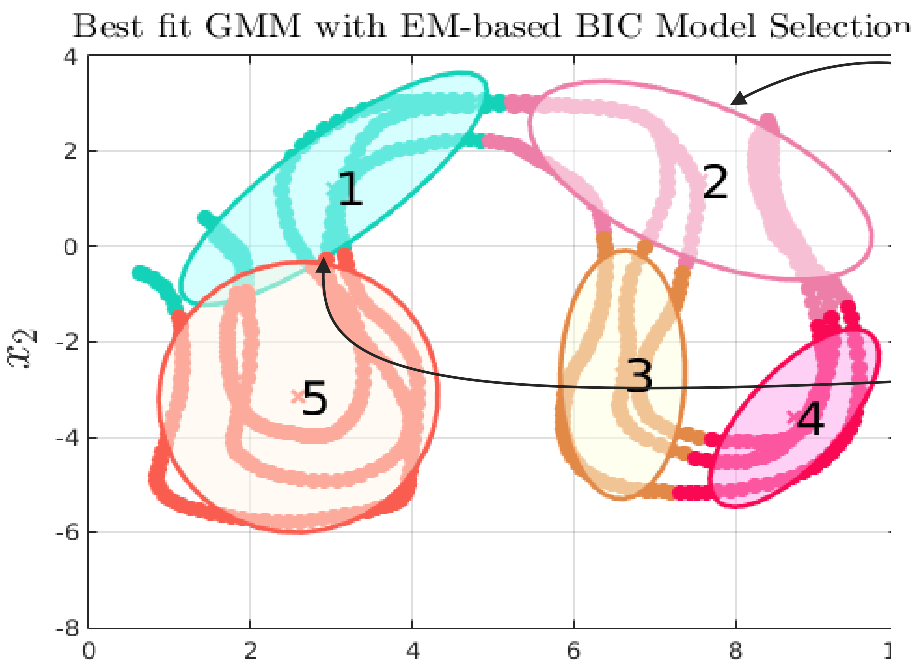

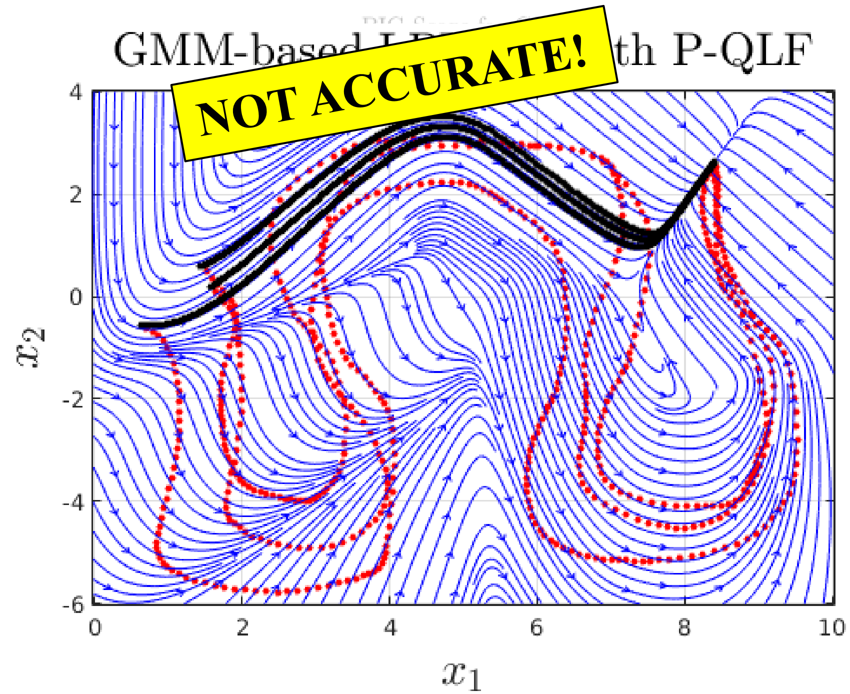

NOT PHYSICALLY CONSISTENT!

DO NOT FOLLOW ORDERING COMING FROM VELOCITY FLOW

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Physically-Consistent GMM

IDEA: ALIGN GAUSS FUNCTION WITH VELOCITY FLOW

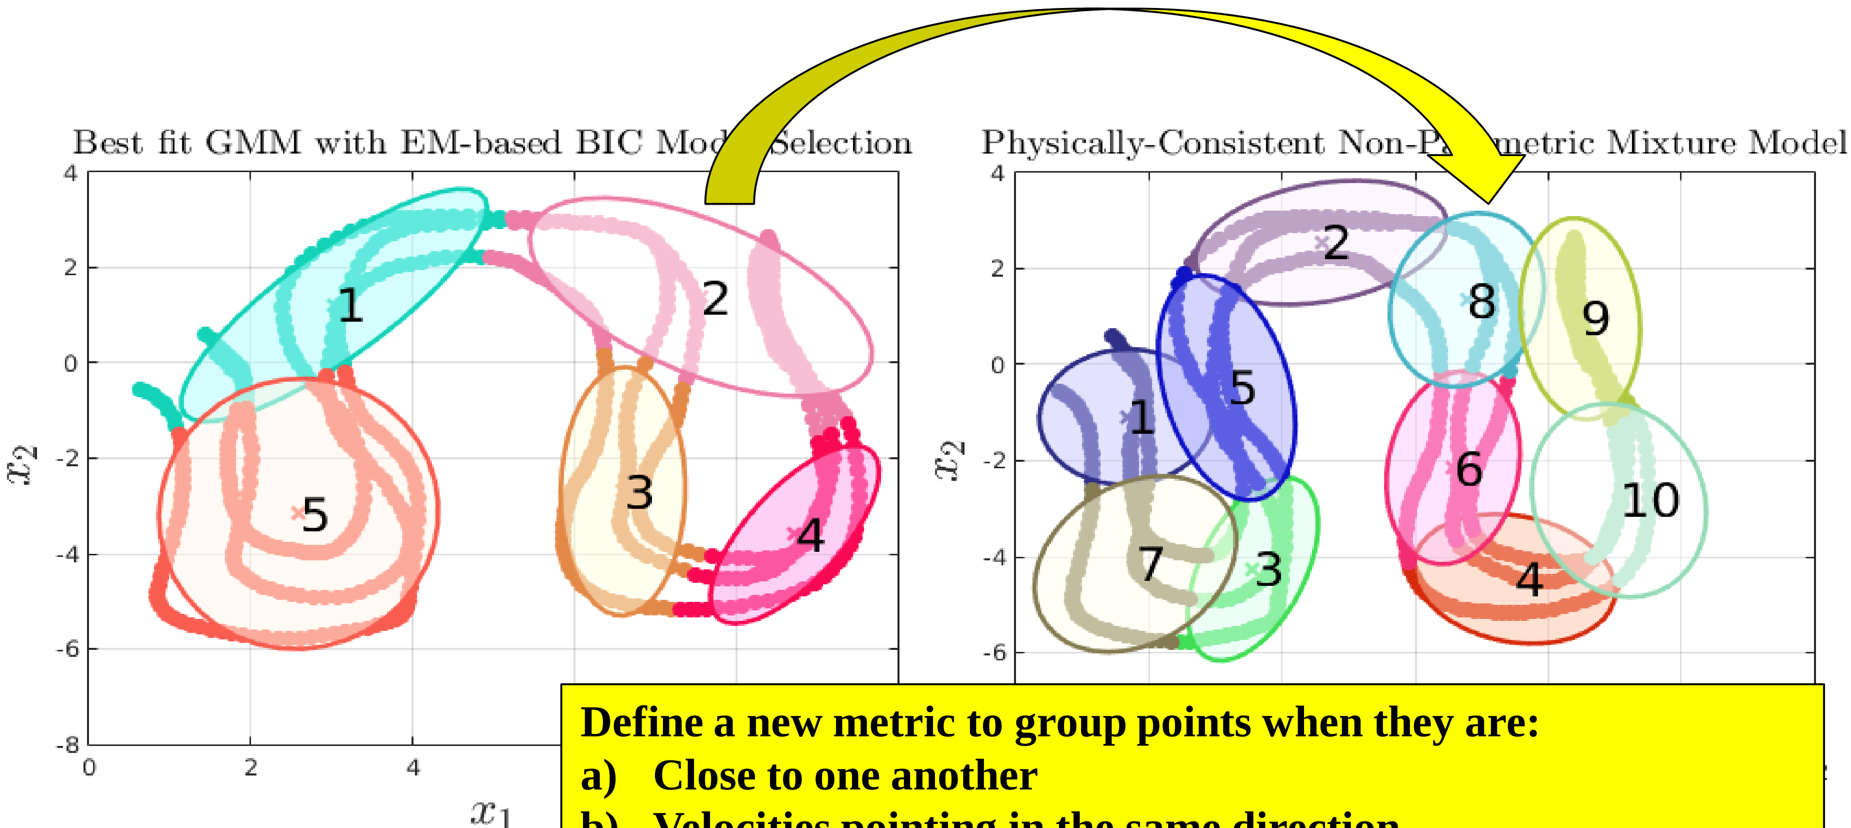

Define a new metric to group points when they are:

a) Close to one another

b) Velocities pointing in the same direction

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Physically-Consistent GMM

Introduce a new metric

$$\Delta_{ij}(x^i, x^j, \dot{x}^i, \dot{x}^j) = \underbrace{\left(1 + \frac{(\dot{x}^i)^T \dot{x}^j}{||\dot{x}^i|| \, ||\dot{x}^j||}\right)}_{\text{Directionality}} \underbrace{\exp\left(-l||x^i - x^j||^2\right)}_{\text{Locality}}.$$

$>> 0$

≈ 0

$>> 0$

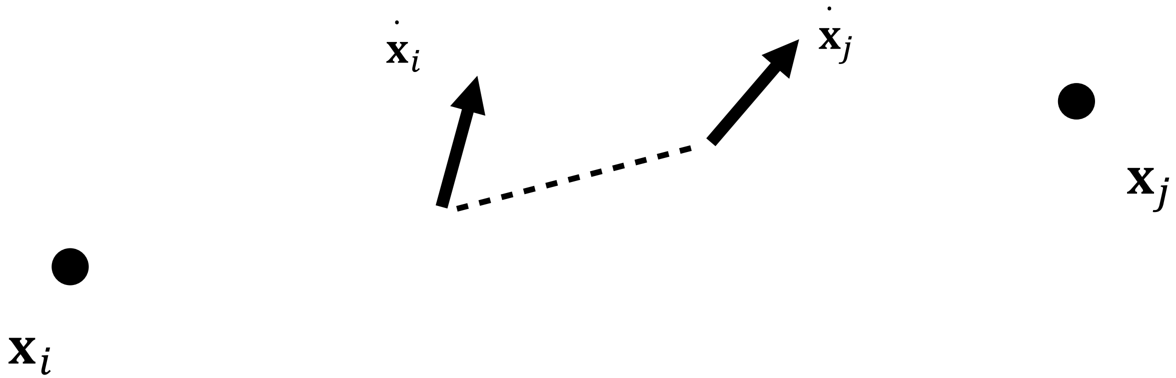

Use this metric to assign datapoints to a Gauss function.

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Recall: GMM Clustering Assignment

(see Applied Machine Learning course on clustering with GMM)

Likelihood of the mixture of K Gaussians: $L\left(\Theta=\left\{\pi_k,\mu^k,\Sigma^k\right\}_{k=1}^{K};x\right)=\sum_{k=1}^{K}\pi_k\cdot p\left(x;\mu^k,\Sigma^k\right)$

The mixing Coefficients are normalized. $\sum_{k=1}^{K}\pi_k=1$

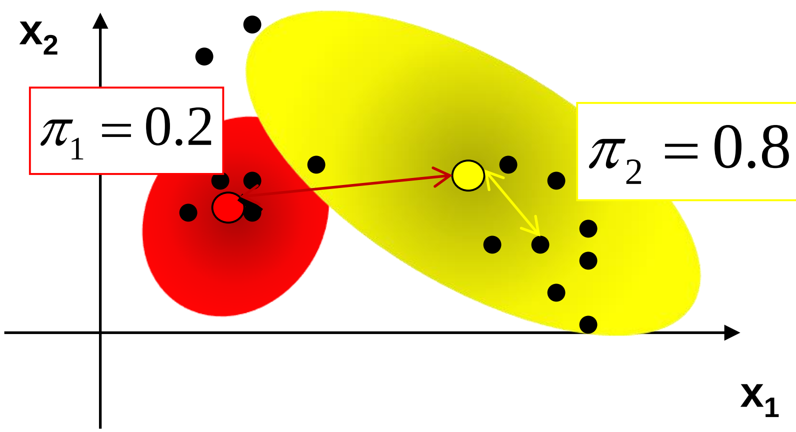

The number of clusters K is a hyperparameter, sometimes difficult to determine.

Bayesian Nonparametric Mixture Model

See Annexes B.3.2-3.3 for details

- **Bayesian**: Bayesian treatment of GMM training

→ No need to fix number of Gauss functions.

→ It learns both the GMM parameters and the number of these parameters required for an optimal fit of the data.

- ***Non-parametric*****:** Does NOT mean methods with “no parameters”, rather models whose complexity (# of states, # Gaussians) is inferred from the data.
  - Number of parameters grows with sample size.
  - Infinite-dimensional parameter space!

N. Figueroa and A. Billard. *. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.*

Bayesian Nonparametric Mixture Model

See Annexes B.3.2-3.3 for details

GMM is a hierarchical model, where each *k*th mixture component is viewed as a cluster, represented by a Gaussian distribution.

Each datapoint x^i is assigned to a cluster k via cluster assignment indicator variable Z= $\left\{ z_1, \ldots, z_M \right\}$.

$$z_i \in \{1, \ldots, K\}$$

$$p(z_i = k) = \pi_k$$

$$x_i | z_i = k \sim \mathcal{N}(\theta_\gamma^k).$$

Bayesian Nonparametric Mixture Model

See Annexes B.3.2-3.3 for details

1: Set priors on model parameters

Normal-inverse-Wishart distribution

See supplement on moodle on conjugate Bayesian analysis of the Gaussian distribution

$$\Theta = \left\{ \pi_k, \mu^k, \Sigma^k \right\}_{k=1}^{K}$$

Dirichlet Prior

The number of Gauss function is unknown and infinite,

$\Rightarrow K \to \infty$

The Dirichlet Process is used as a non-parametric prior on the mixing coefficients.

It removes the need to specify K.

2: Use *Bayesian inference* to estimate the parameters.

Dirichlet Process: Definition

A Dirichlet Process (DP) is a stochastic process that generates at each draw a probability distributions.

For K draws, we can write $D \sim p_1, p_2, \ldots p_K$

The range of realizations is a set of probability distributions.

This can be used to encapsulate prior knowledge on the distribution of random variables.

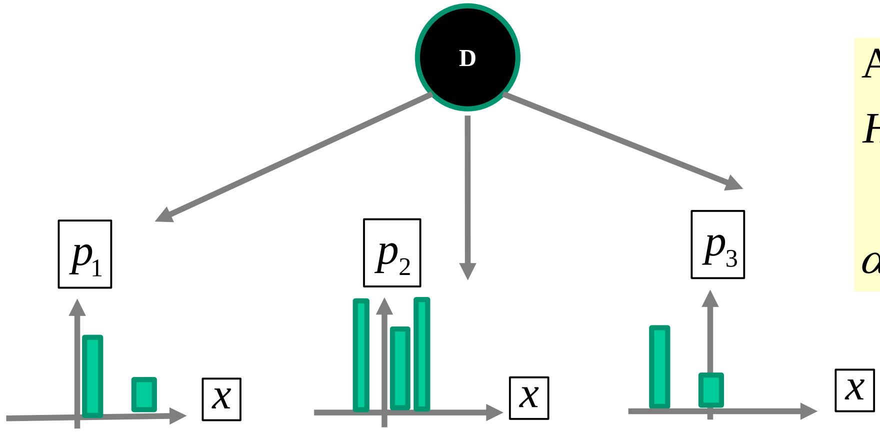

A DP is specified by two parameters

H a distribution called the base distribution.

This is the expected value of the process (the mean)

$\alpha \in \mathbb{R}$ the concentration parameter.

Dirichlet Process: Properties

K_n = # clusters in sample of size M

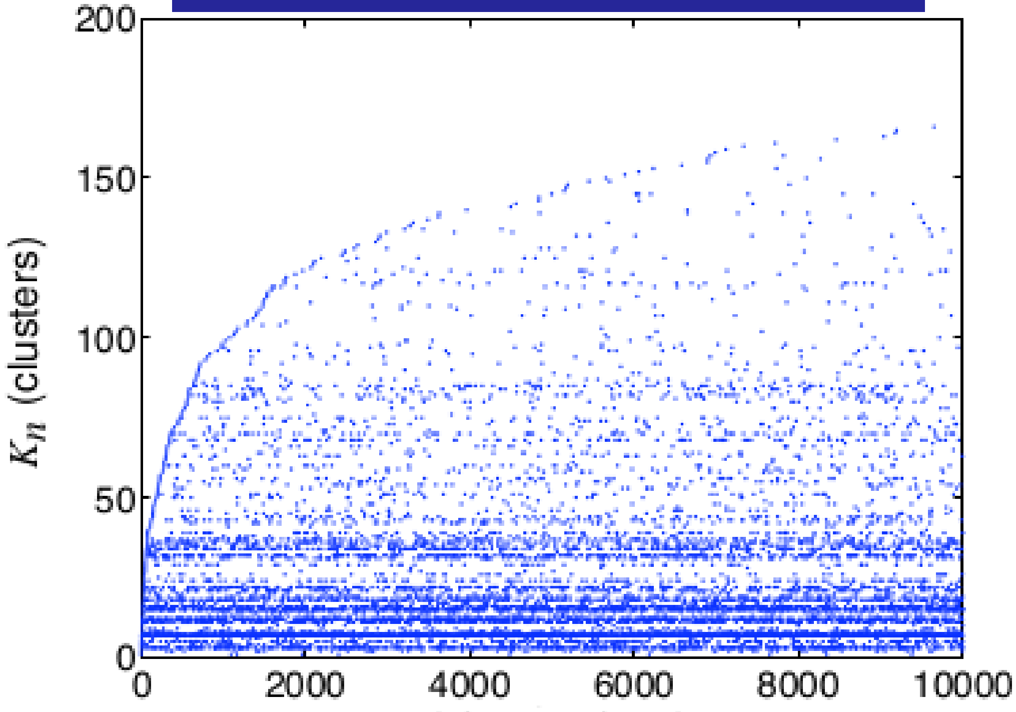

M (Observations)

The number of components needed to model M observations no longer depends on the open K parameters, but is ~$O(\alpha \log M)$.

As $K \rightarrow \infty$, the mixture model remains contained within $O(\alpha \log M)$ and is hence referred to as an infinite mixture model.

Bayesian Nonparametric Mixture Model

See Annexes B.3.2-3.3 for details

Dirichlet prior on mixing coefficients

$$\pi \sim \mathrm{Dir}\left(\frac{\alpha}{K}, \ldots, \frac{\alpha}{K}\right)$$

The vector of mixing coefficients is now considered as a categorical or multinomial distribution, which when sampled gives the probability of $p(z_i = k)$

$$z_i | \pi = \mathrm{Cat}(\pi)$$

$$x_i | z_i = k \sim \mathscr{F}(\theta_k).$$

The probability density of the GMM is then given by:

$$p\left(x; \Theta = \left\{\pi_k, \mu^k, \Sigma^k\right\}_{k=1}^{\infty}\right) = \sum_{k=1}^{\infty} p\left(z_i = k\right) \cdot p\left(x; \mu^k, \Sigma^k\right)$$

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Bayesian Nonparametric Mixture Model

See Annexes B.3.2-3.3 for details

Graphical model of the Bayesian GMModel.

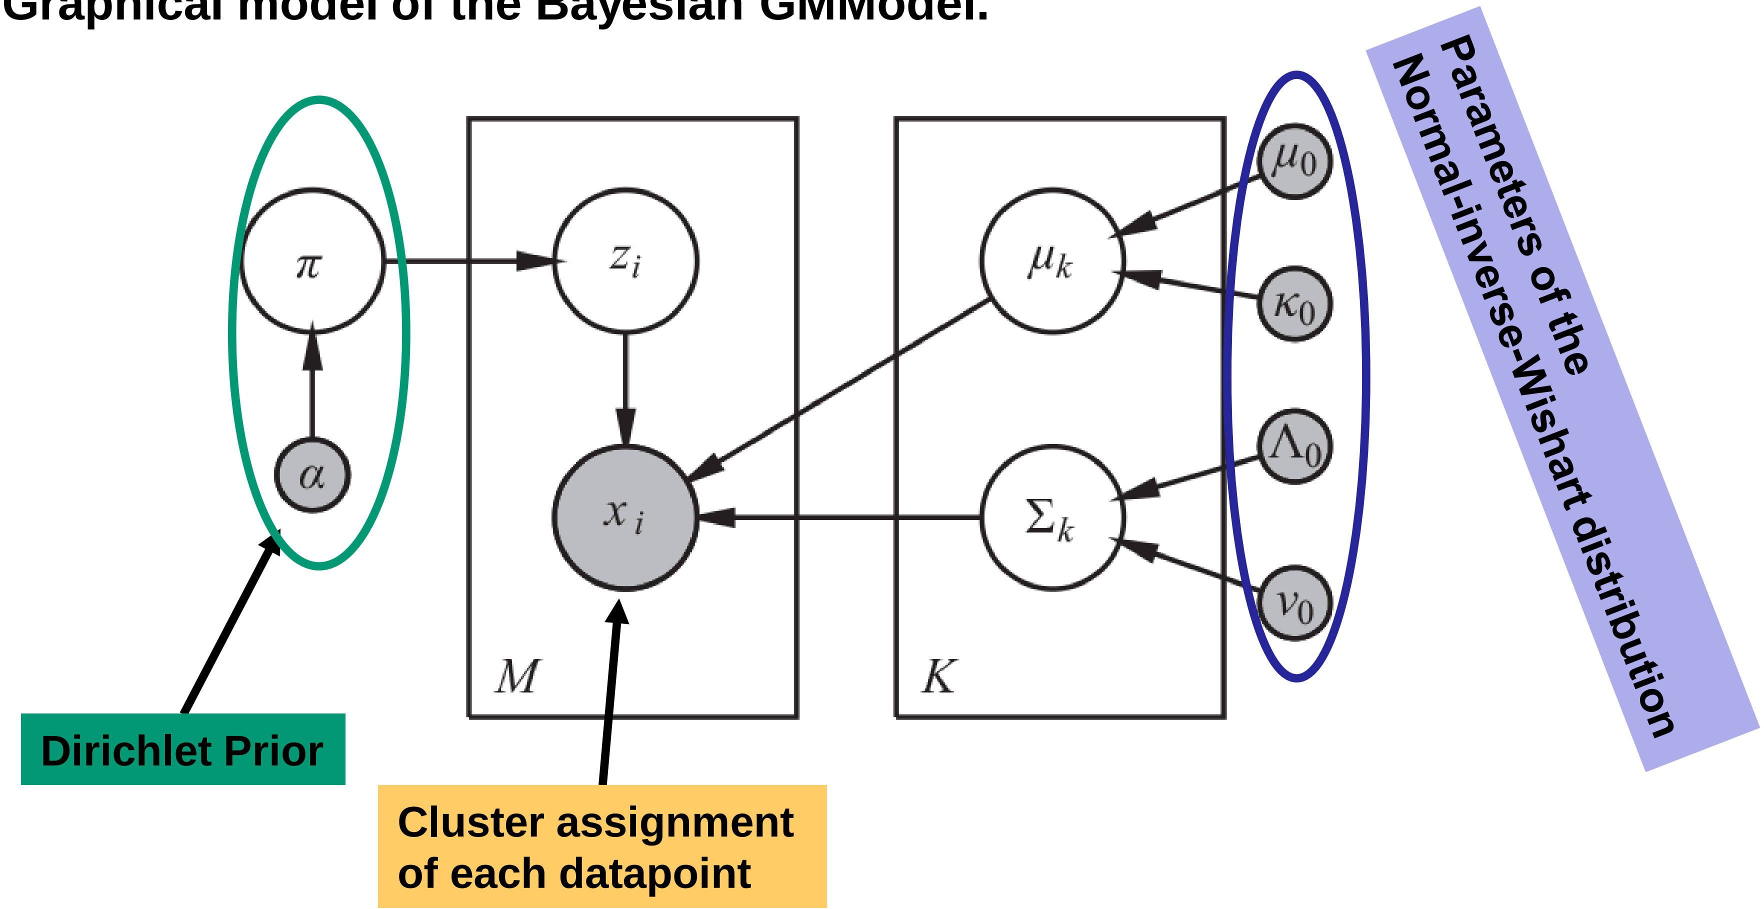

Algorithm: Physically-Consistent GMM

Uses a Bayesian Nonparametric Mixture Model combined with new metric to cluster data according to « physical consistency » (velocity alignment)

$$c_i \sim PC\!-\!CRP(\Delta,\alpha)$$

Physically Consistent Chinese Restaurant Process to assign datapoints to a Gauss function

Draw “seat” (cluster) assignment according to how close they are under this metric

$$p(c_i = j \mid C_{-i}, \mathbf{X}, \Delta, \alpha, \lambda) \propto \underbrace{p(c_i = j \mid \Delta, \alpha)}_{\text{Similarities in scaled velocity space}} \; \underbrace{p(\mathbf{X} \mid \mathbf{Z}(c_i = j \cup C_{-i}), \lambda)}_{\text{Observations in position space}},$$

$$p(C \mid \Delta, \alpha) = \prod_{i=1}^{M} p(c_i = j \mid \Delta, \alpha), \quad \text{where } p(c_i = j \mid \Delta, \alpha) = \begin{cases} \frac{\Delta_{ij}(\cdot)}{\sum_{j=1}^{M} \Delta_{ij}(\cdot) + \alpha} & \text{if} \quad i \neq j \\ \frac{\alpha}{M+\alpha} & \text{if} \quad i = j, \end{cases}$$

N. Figueroa and A. Billard. . *A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.*

Algorithm: Physically-Consistent GMM

$$c_i \sim PC\text{--}CRP(\Delta, \alpha)$$

Physically Consistent Chinese Restaurant Process to assign datapoints to a Gauss function

$$z_i = \mathbf{Z}(c_i)$$

Assign cluster label

$$\theta_\gamma^k \sim \mathcal{N}\mathrm{IW}(\lambda_0)$$

Draw mean and covariance from the inverse-Wishart distribution

$$x^i | z_i = k \sim \mathcal{N}(\theta_\gamma^k).$$

Compute distribution of datapoint from the Gauss distribution

N. Figueroa and A. Billard. . *A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.*

Examples: Physically-Consistent GMM

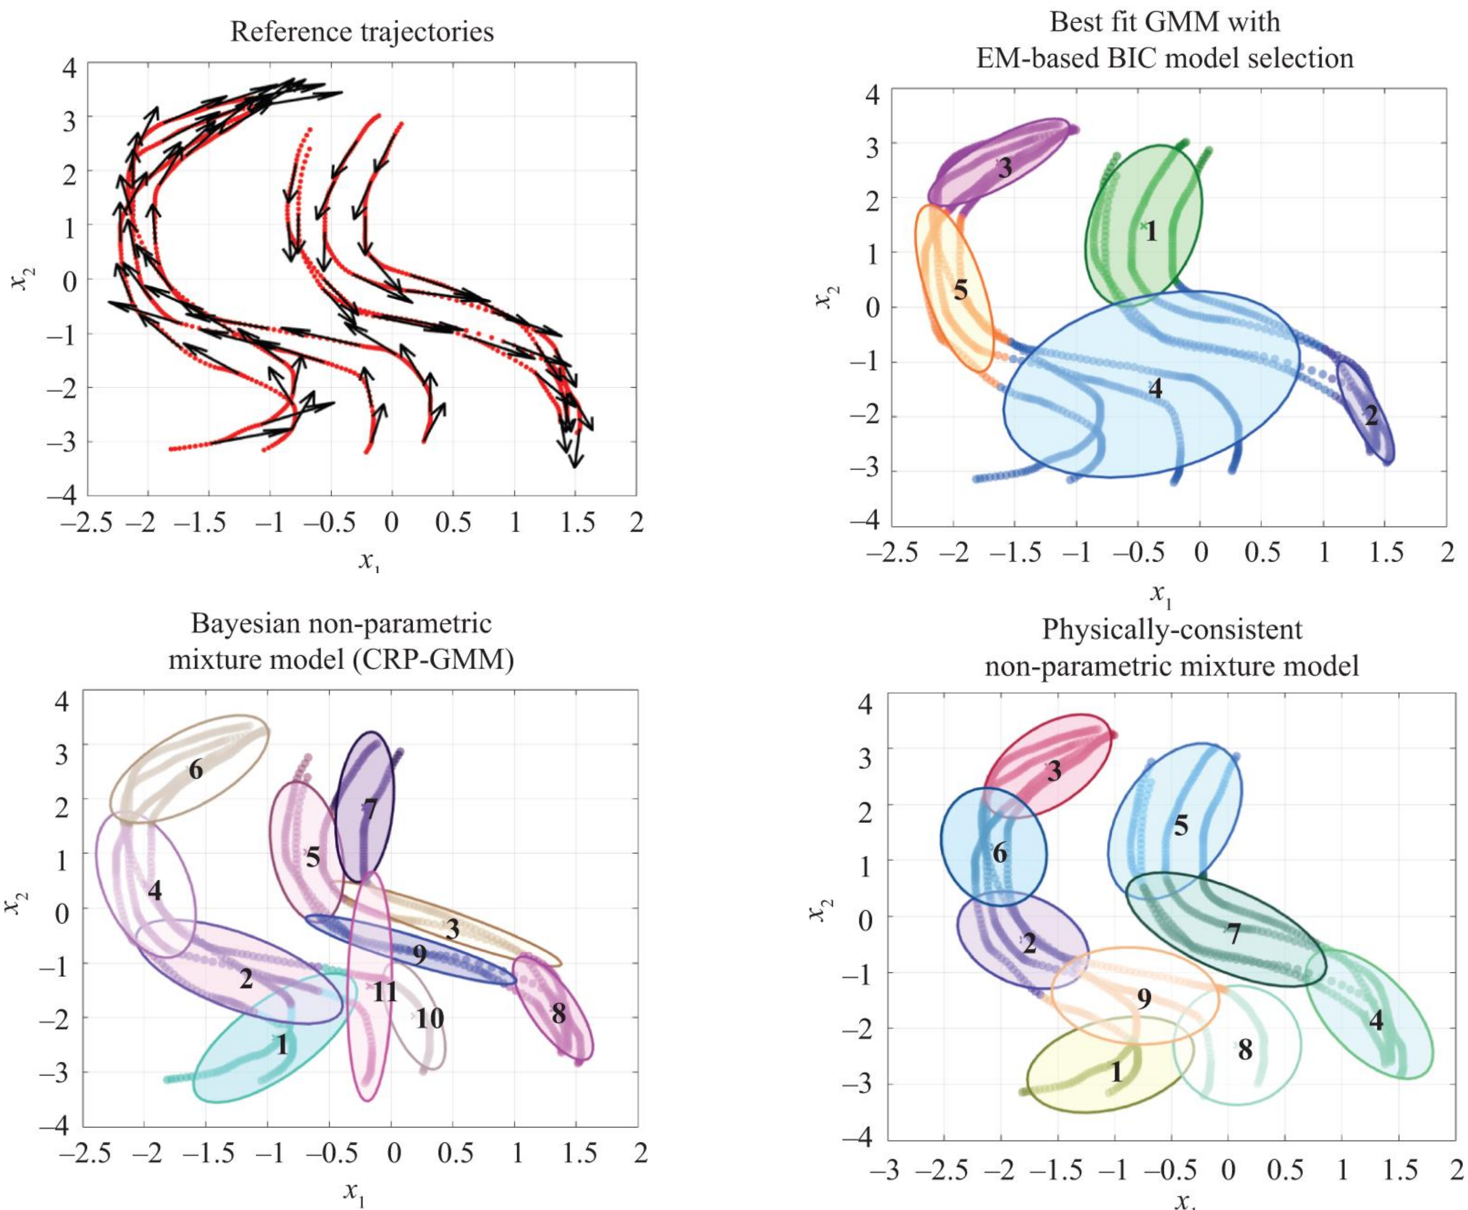

Examples: Physically-Consistent GMM

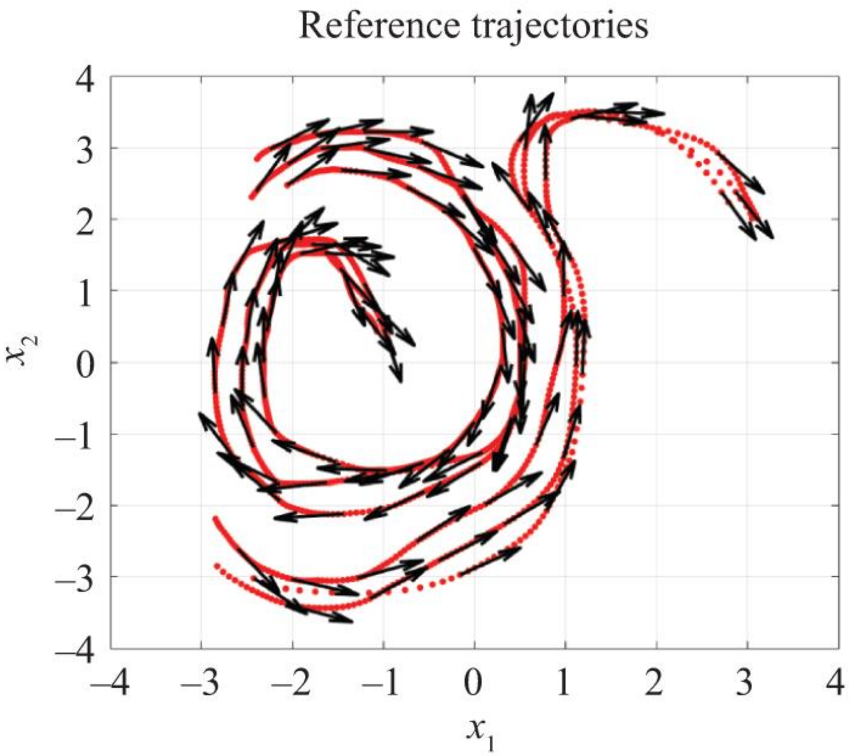

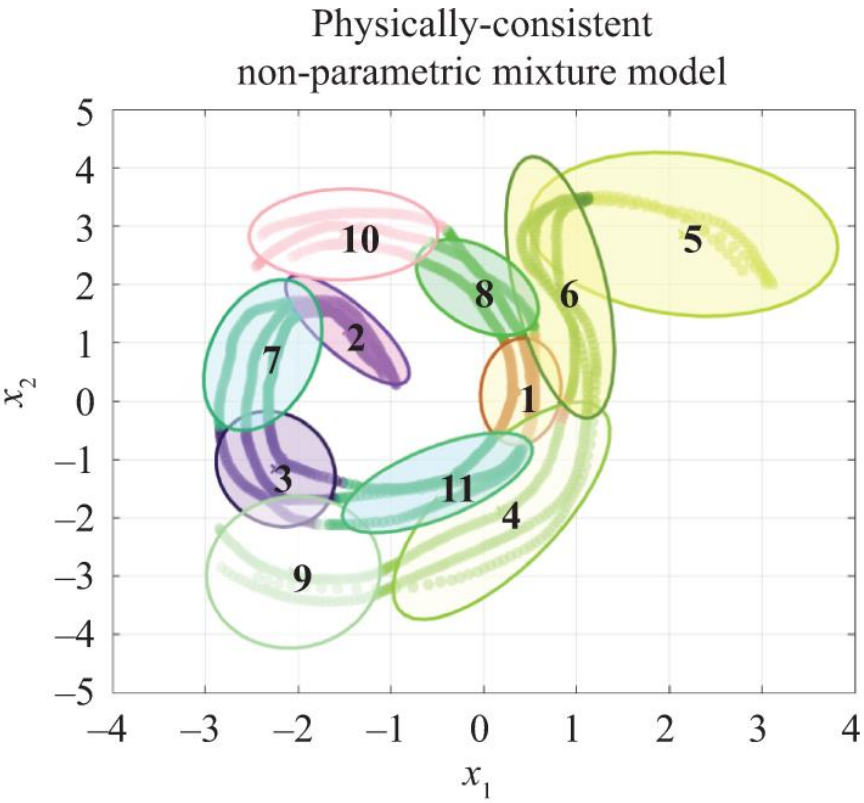

Examples: Physically-Consistent GMM

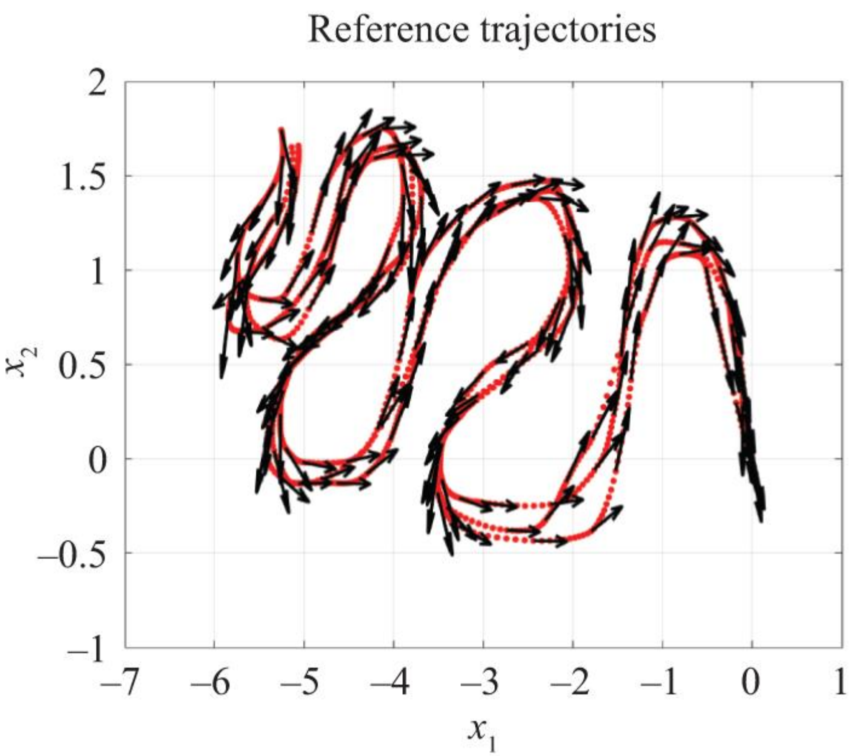

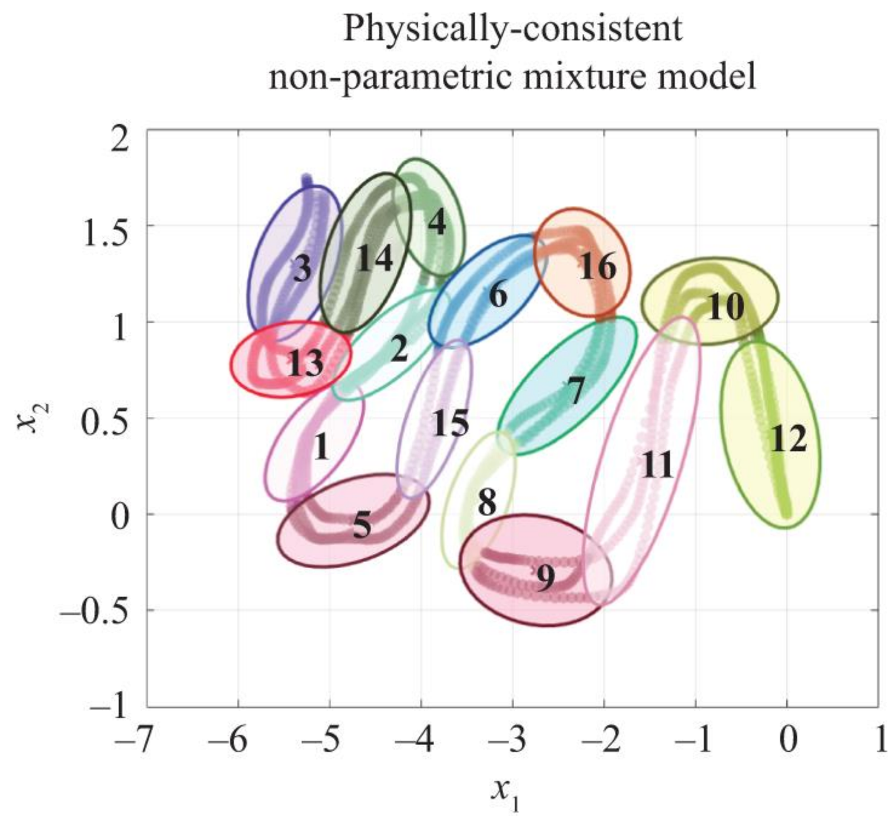

LPV-DS final optimization

Once the GMM parameters have been estimated with PC-GMM, we are left with satisfying the set of constraints for stability.

This leads to a non-convex but solvable optimization (see Section 3.4.3 of the book for details).

SEDS like

$$\min_{\Theta_f} J(\Theta_f) \quad \text{subject to}$$

P-QLF

$$(O1)\left\{(A^k)^T + A^k \prec 0, b^k = -A^k x^* \;\; \forall k = 1, \ldots, K\right.$$

$$(O2)\left\{(A^k)^T P + PA^k \prec 0, \;\; b^k = \mathbf{0} \;\; \forall k = 1, \ldots, K; \;\; P = P^T \succ 0\right.$$

$$(O3)\left\{(A^k)^T P + PA^k \prec Q^k, \;\; Q^k = (Q^k)^T \prec 0, \;\; b^k = -A^k x^* \;\; \forall k = 1, \ldots, K.\right.$$

LPV-DS final optimization

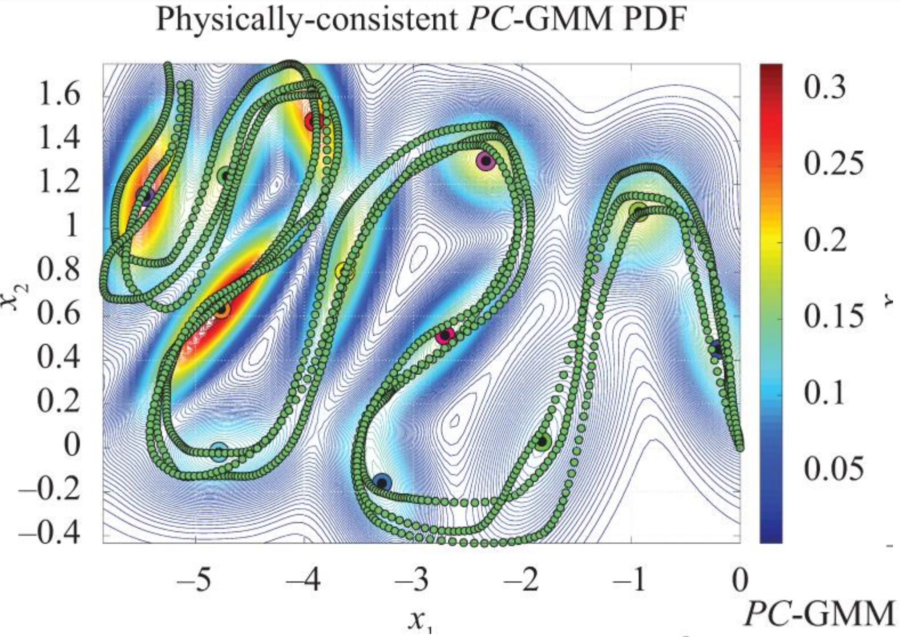

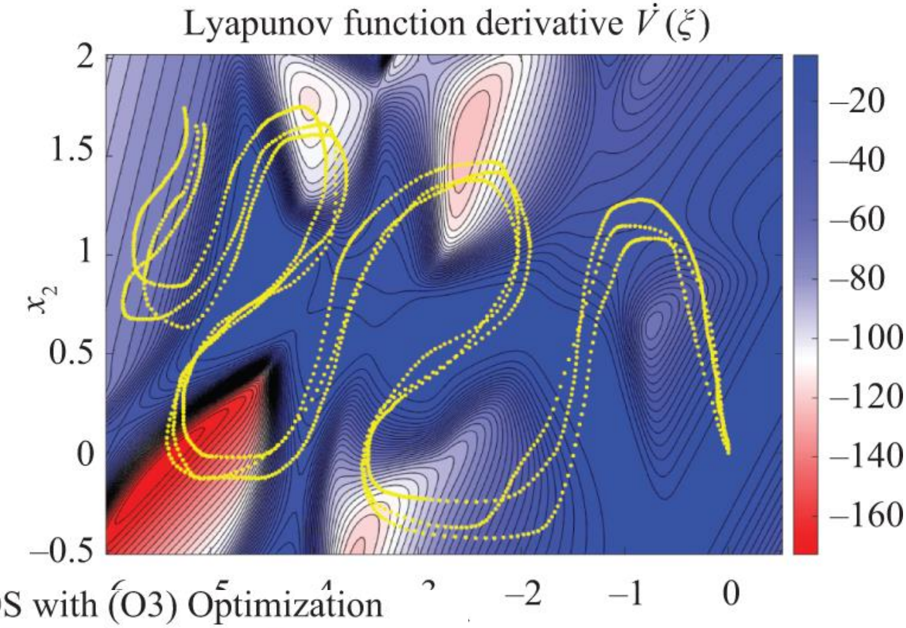

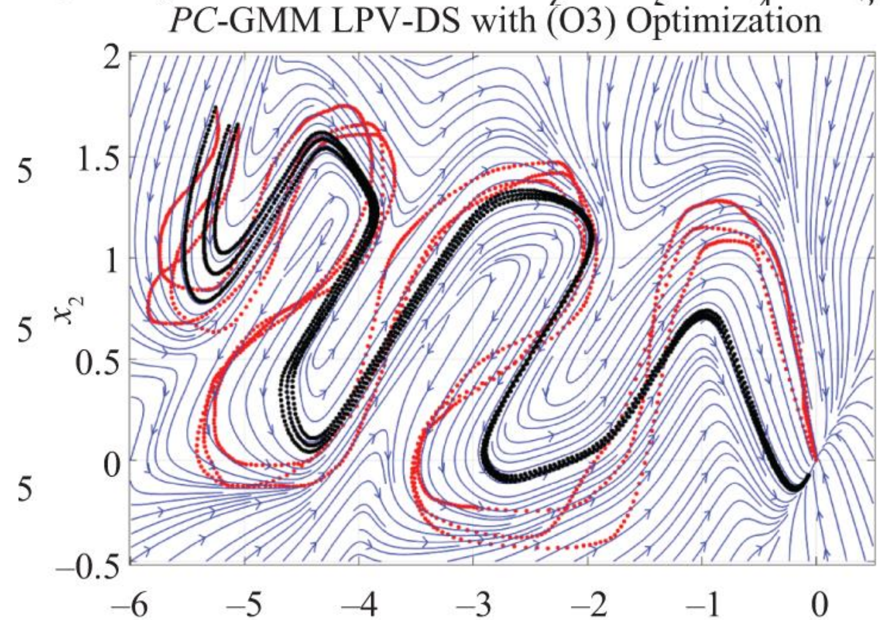

Result on previous example

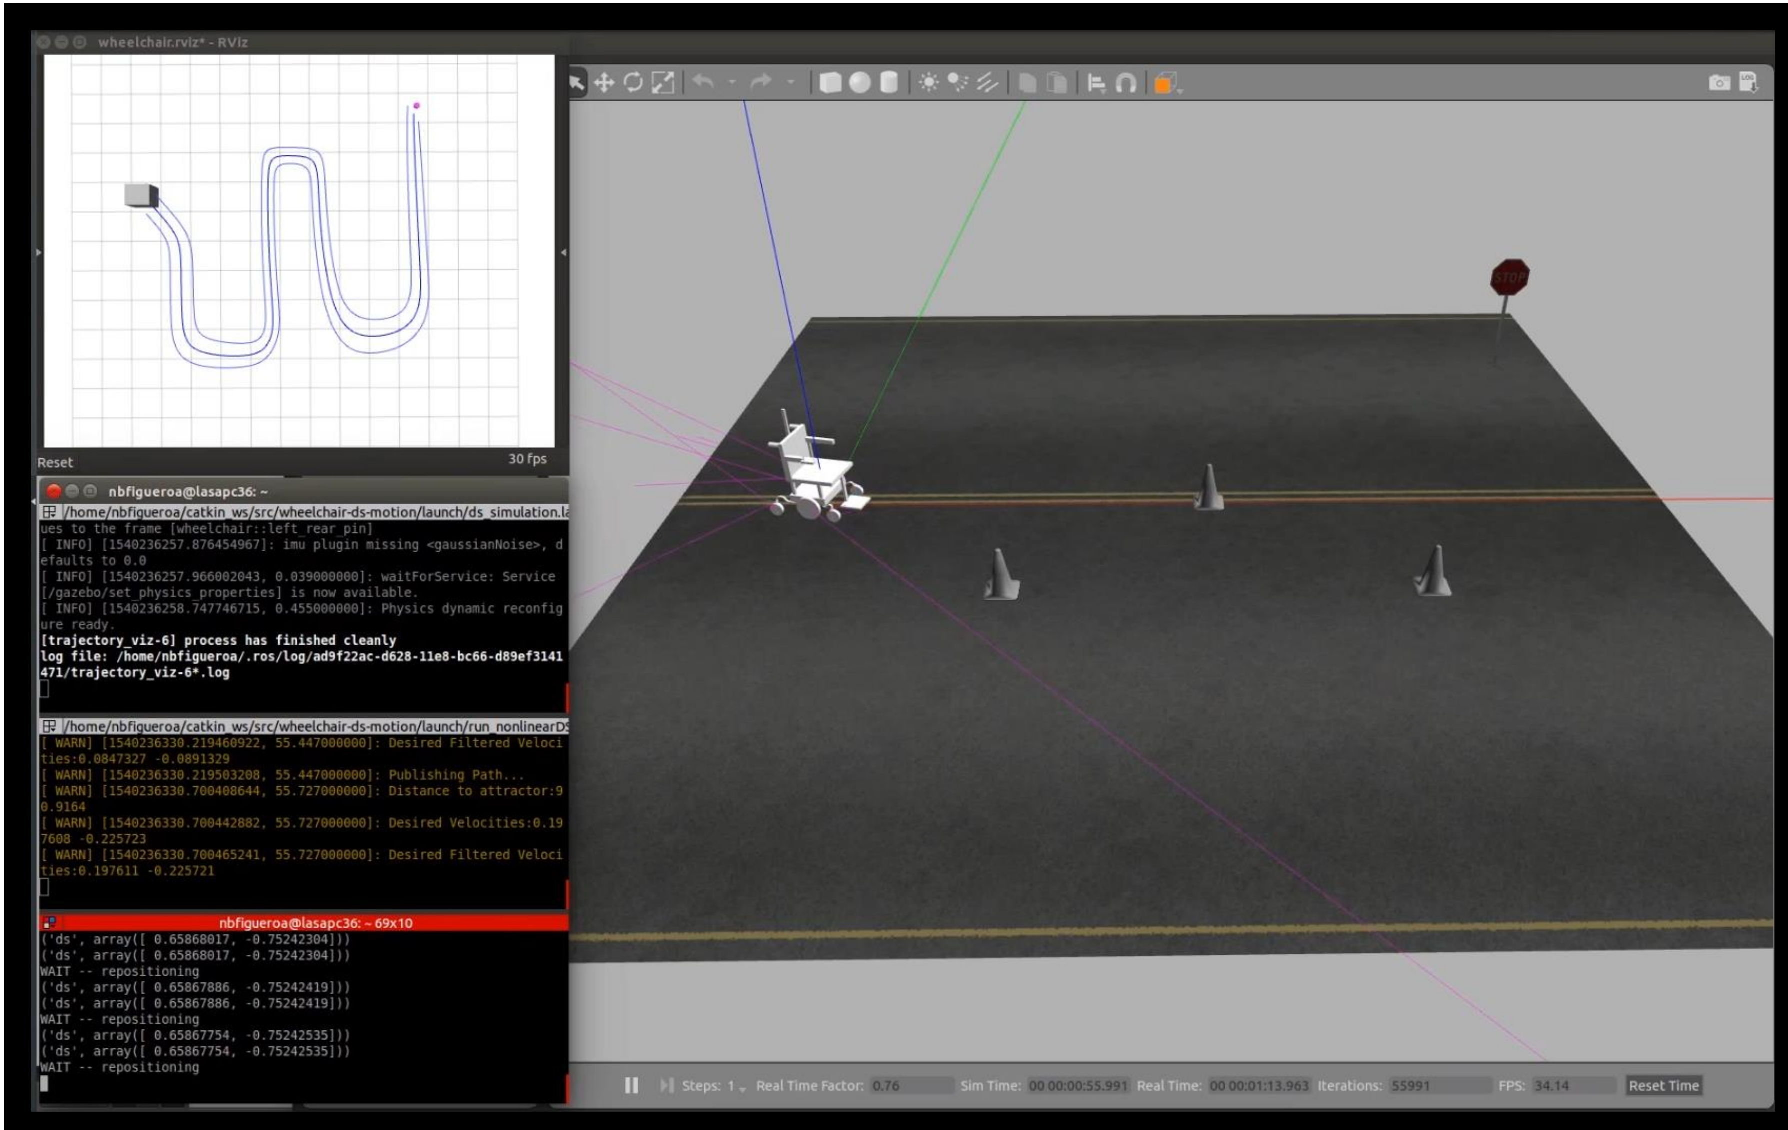

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Learning Physically-Consistent Gaussian Mixture Model

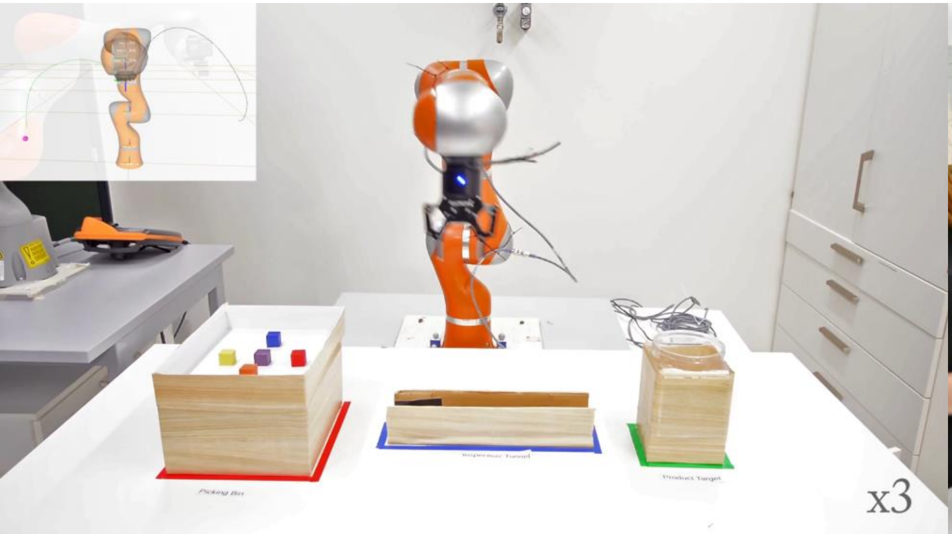

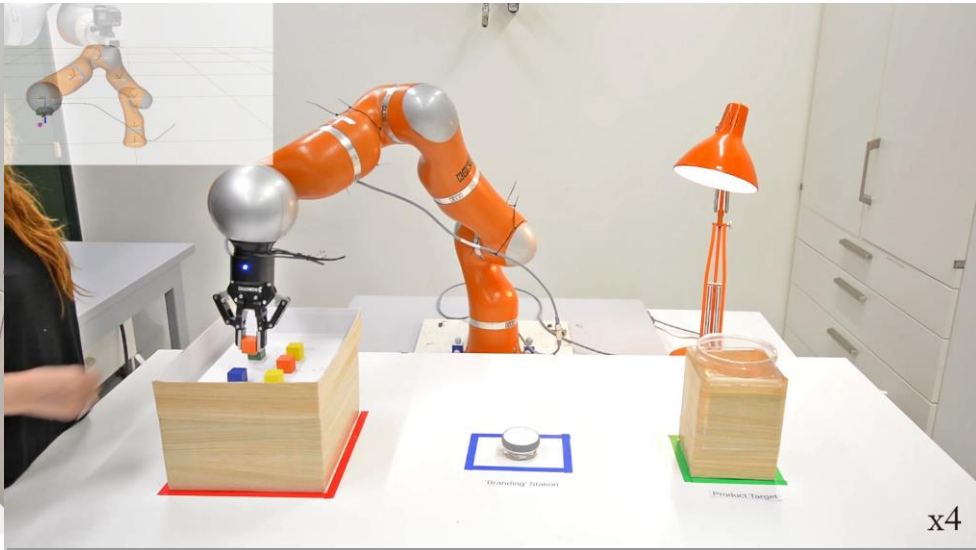

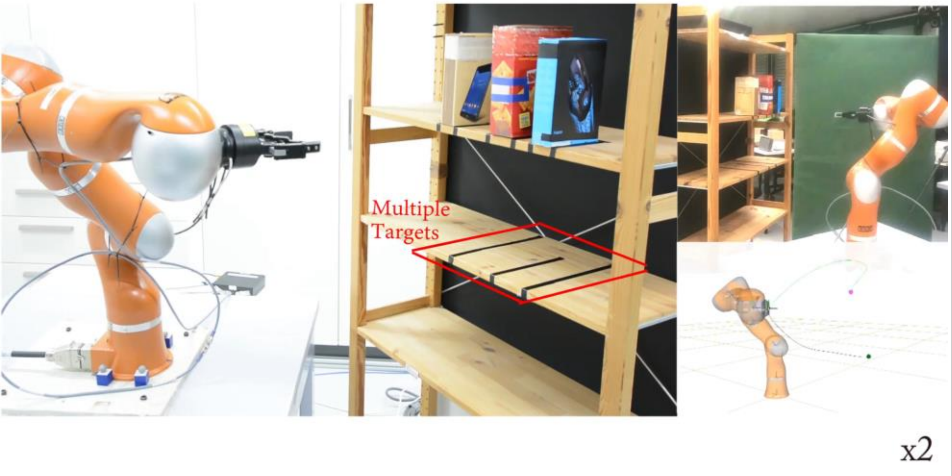

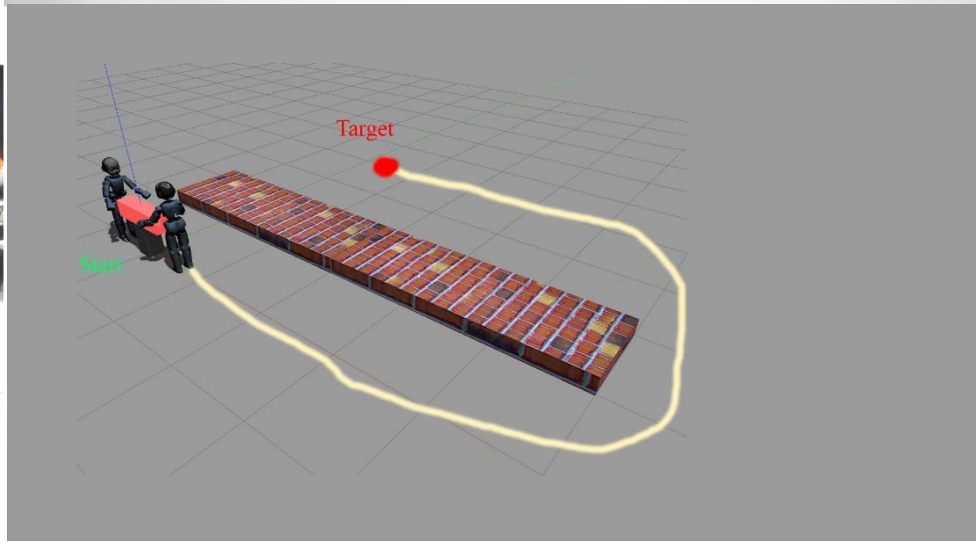

N. Figueroa and A. Billard.. A physically-consistent Bayesian non-parametric Mixture Model for dynamical system learning. CoRL. Oct 2018.

Summary LPV-DS

LPV-DS was offered as an alternative to SEDS to enable learning of more complex, and nonlinear DS from demonstrations.

SEDS	LPV - DS
Fix by hand number of Gaussians	**Learns automatically number of Gaussians**
Conservative stability constraints → Cannot learn highly non-linear trajectories	**Less conservative stability constraints → Can embed large non-linearities**