


PART II

COUPLING AND MODULATING CONTROLLERS

Chapter 8

Adapting and Modulating an Existing Control Law


Adapting a controller



Adapting a controller

untouched

One should not re-estimate the entire system.




A modulation of the system may be necessary.

The modulation should be *local*.

Initial – Nominal DS


Start with an initial DS $\dot{x} = f(x)$ with attractor $\dot{x}^* = f(x^*) = 0$.



We will call this the *nominal* DS

DS after Local Modulation

Generate a *local* modulation



How can we modulate the DS while preserving the stability properties?



Modulation – Theoretical Approach

Idea:

Modulate the original DS in such a way that the new DS:

- remains a first order DS, ✓
- preserves asymptotic stability at its attractor x^* , ?
- has still a single attractor. ?

Which properties do we need for M to satisfy the other two requirements?

The modulated dynamics is given by:

$$\dot{x} = M(x) f(x), \quad M(x) \in \mathbb{R}^{N \times N}$$

Does this satisfy our three desiderata?

The matrix M above should not be confused with the metric M of contraction theory introduced in Lecture 3.

DS after Local Modulation

$$\dot{x} = M(x) f(x)$$

The modulation should be local.

Preserve stability at attractor.

$$\Rightarrow M(x^*) = I.$$

Original attractor.

Locality of the modulation

\exists a compact region $B(x, x^*)$ around x^* where $M(x) = I., \forall x \in B(x, x^*)$.

Uniqueness of attractor

$$\Rightarrow M(x) f(x) \neq 0, \forall x \neq x^*$$

M must be full rank.

DS after Local Modulation

$$\dot{x} = M(x) f(x)$$


$$M(x) = R(\phi(x)), \quad \phi(x) = e^{-\|x-x^o\|} \phi_o.$$

Rotation $R \in \mathbb{R}^{N \times N}$





Local Rotation

Active only locally.

$$\begin{aligned} R(x) &\neq I, \quad x \in \Omega \\ R(x) &= I, \quad \forall x \notin \Omega \end{aligned}$$




Modulation – Local Rotations - Examples



Modulation – Rotation Leading to a Limit Cycle


Modulations can be used to generate **new properties** for the DS

$$M = \begin{bmatrix} \cos(\gamma\theta) & -\sin(\gamma\theta) \\ \sin(\gamma\theta) & \cos(\gamma\theta) \end{bmatrix}, \quad \gamma(x, x^o) = e^{-\frac{1}{\sigma^2}(\|x\| - r)^2}, \quad \sigma > 0.$$




Diverging Systems

Invert flow and change stability of the nominal system (see Today's exercises).




Flow diverging from attractor and limit cycle



Modulation - Scaling


Multiplicative term in front of the modulation will **increase/decrease the speed** of the DS but *will not modify the type of dynamics*. $\dot{x} = \lambda M(x) f(x)$, $\lambda \in \mathbb{R}$.

Modulated DS with limit cycle



$\lambda = 1$

Modulated DS with limit cycle




$\lambda = 10$

Modulation – Other Properties

Modulations can be used to generate **new properties** for the DS

$$M = \begin{bmatrix} 1 - 2\gamma(x, x^o) & 0 \\ 0 & 1 - 2\gamma(x, x^o) \end{bmatrix}, \quad \gamma(x, x^o) = e^{-\frac{1}{\sigma^2} \|x - x^o\|}, \quad \sigma > 0.$$

ERRATUM: book, section 8.1.1



Parameterization of the modulation

$$\dot{x} = M(x) f(x) \quad M(x) = (1 + \kappa(x)) R(x)$$

Modulates speed

Rotation $R \in \mathbb{R}^{N \times N}$

Ω

Active only locally,
zero elsewhere

$$\kappa(x) > 0, \quad x \in \Omega$$

$$\kappa(x) = 0, \quad \forall x \notin \Omega$$

$$R(x) \neq I, \quad x \in \Omega$$

$$R(x) = I, \quad \forall x \notin \Omega$$

Parameterization of the modulation

The modulation function can rotate and speed up the dynamics locally.

$$M(x) = (1 + \kappa(x)) R(x) \leftarrow \text{Rotation}$$

Direction / axes can be learned from data

Modulates the speed

Speed and region of activation can be learned from data.

In 2D, a rotation is defined solely by an angle ϕ , $R(\phi) = \begin{bmatrix} \cos(\phi) & -\sin(\phi) \\ \sin(\phi) & \cos(\phi) \end{bmatrix}$.

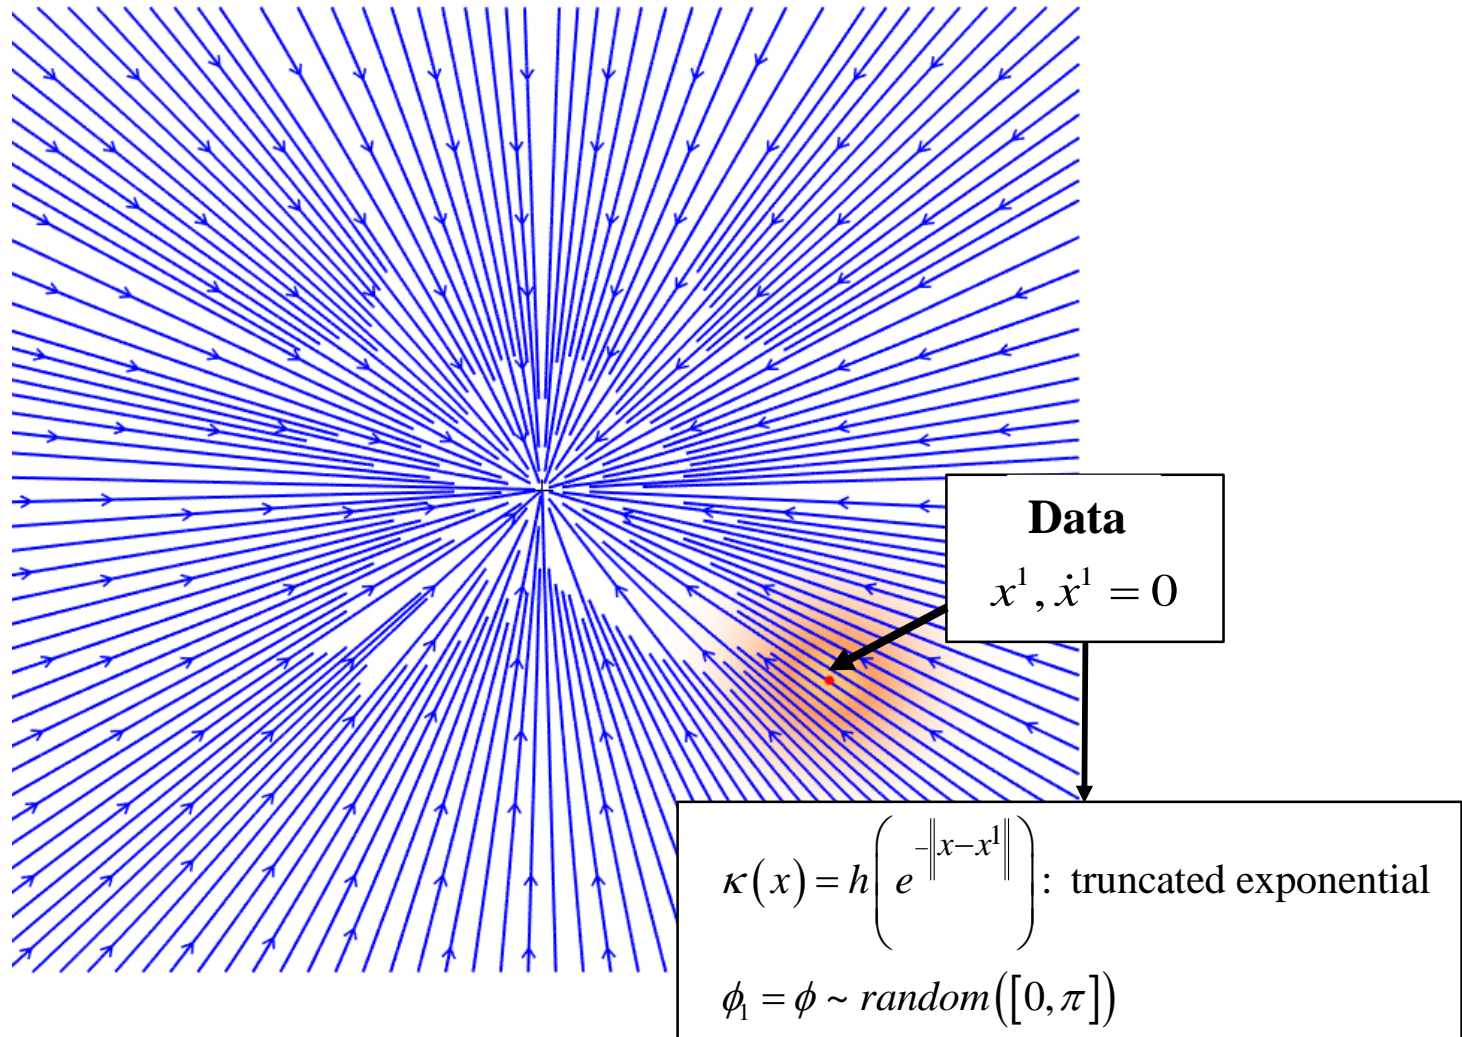
In 3D, a rotation $R(\phi, \mu_R)$ is defined by an angle ϕ around a vector μ_R .

The modulation function is hence parameterized as:


$$M(x; \Theta = \{\phi, \mu_R, \kappa\}) = (1 + \kappa(x; X, \dot{X})) R(\phi(x; X), \mu_R(x; X)), \quad X, \dot{X} : \text{data for learning}$$

Learning the modulation

Example – generated a rotation around a single demonstrated point




Example – learning a rotation around a single point




This can model only a single rotation

Use nonlinear regression techniques to model more complex modulations from full demonstrated trajectories.

Original linear dynamics: $\dot{x} = f(x)$



Original linear dynamics: $\dot{x} = f(x)$




Draw training datapoints

DATA: set of M trajectories

$$\{X, \dot{X}\} = \{x^m, \dot{x}^m\}_{m=1}^M$$

Reshaped dynamics: $\dot{x} = M(x) f(x)$



Learning a Modulation (2D example)

General formulation


$$\dot{\mathbf{x}} = \mathbf{g}(\mathbf{x}) = \mathbf{M}(\mathbf{x})\mathbf{f}(\mathbf{x})$$

\mathbf{x}	Robot's state
$\mathbf{R}(\mathbf{x})$	Rotation Matrix
$[\kappa(\mathbf{x})]$	Scaling factor

Modulation composed of a scaling and rotation:

$$\mathbf{M}(\mathbf{x}) = [1 + \kappa(\mathbf{x})]\mathbf{R}(\mathbf{x})$$

$$\mathbf{R}(\mathbf{x}) = \begin{bmatrix} \cos(\phi(\mathbf{x})) & -\sin(\phi(\mathbf{x})) \\ \sin(\phi(\mathbf{x})) & \cos(\phi(\mathbf{x})) \end{bmatrix}$$




Define the parameter vector:

$$\boldsymbol{\theta}(\mathbf{x}) = \begin{bmatrix} \kappa(\mathbf{x}) \\ \phi(\mathbf{x}) \end{bmatrix}$$

For dimension >2 , the parameter vector is expanded by parameters describing the rotation set.

From Trajectory Data to Reshaping Parameters

x	Robot's state
$R(x)$	Rotation Matrix
$[\kappa(x)]$	Scaling factor
x_m	Original data-points
x_m^o	New data-points



Trajectory data

$$\{x^m, \dot{x}^m\}_{m=1}^M$$

Nominal vel.

$${}^o\dot{x}^m = f(x^m)$$

Nominal pos.

$$\{{}^o x^m\}_{m=1}^M$$

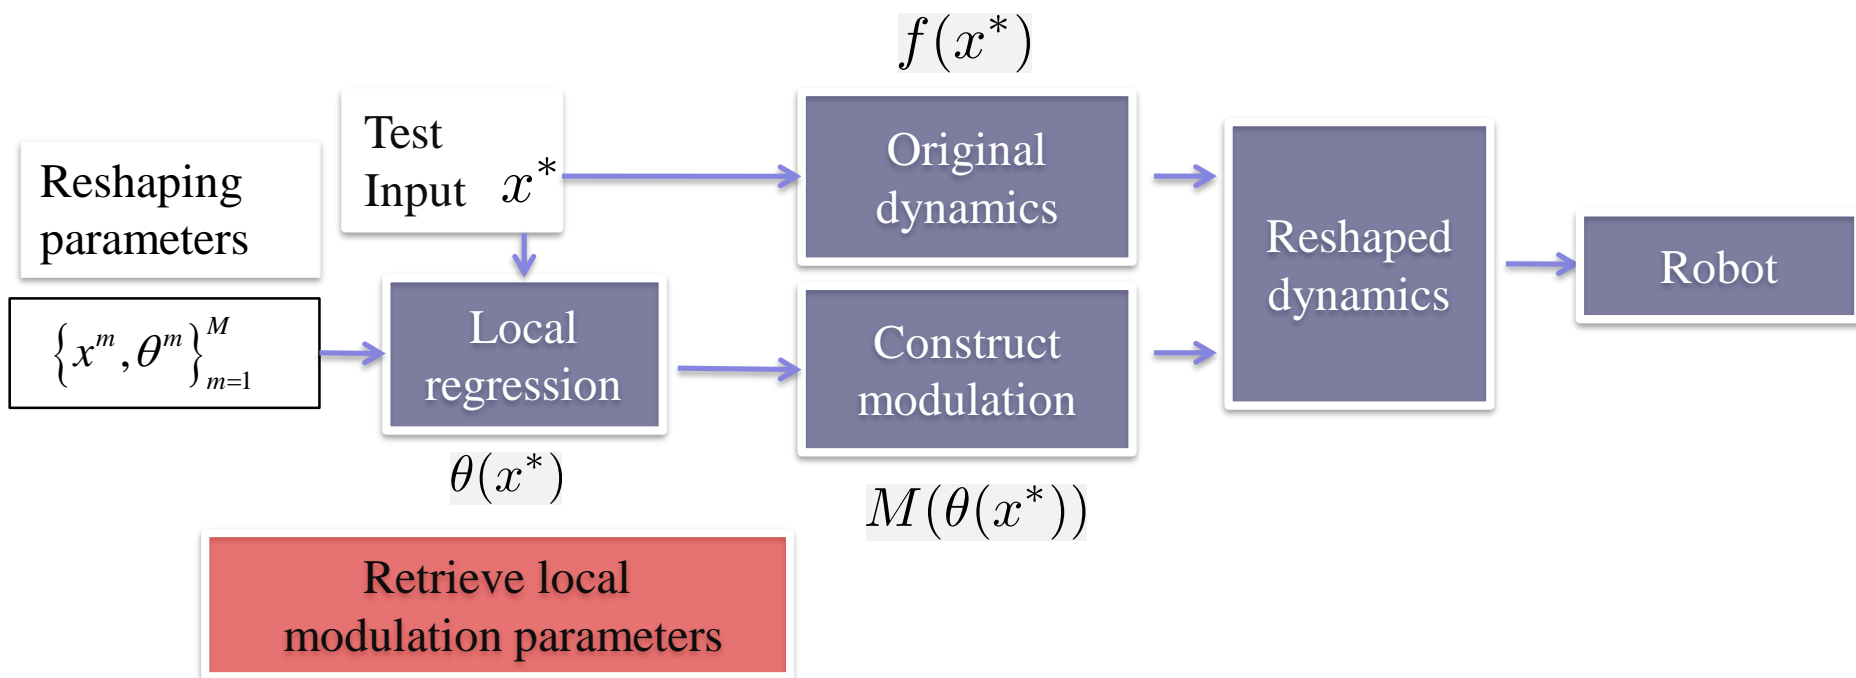
$$\kappa_m = \frac{\|\dot{x}^m\|}{\|{}^o\dot{x}^m\|} - 1$$

$$\phi_m = \frac{\arccos(\langle \dot{x}^m, {}^o\dot{x}^m \rangle)}{\|\dot{x}^m\| \|{}^o\dot{x}^m\|}$$

Reshaping parameters

$$\{x^m, \theta^m\}_{m=1}^M$$

Learning and Using the Reshaped Dynamics



Pseudo-Code for Converting Data into Modulation Parameters

Algorithm 8.1 Procedure for converting 2D or 3D trajectory data to modulation data.

Require: Trajectory data $\{x^m, \dot{x}^m\}_{m=1}^M$

1: **for** $m = 1 \rightarrow M$ **do**

2: Compute nominal velocity: ${}^o\dot{x}^m = f(x^m)$

3: Compute rotation vector (3D only): $\mu^m = \frac{\dot{x}^m \times {}^o\dot{x}^m}{\|\dot{x}^m\| \|{}^o\dot{x}^m\|}$

4: Compute rotation angle: $\phi^m = \arccos \frac{\dot{x}^{mT} {}^o\dot{x}^m}{\|\dot{x}^m\| \|{}^o\dot{x}^m\|}$

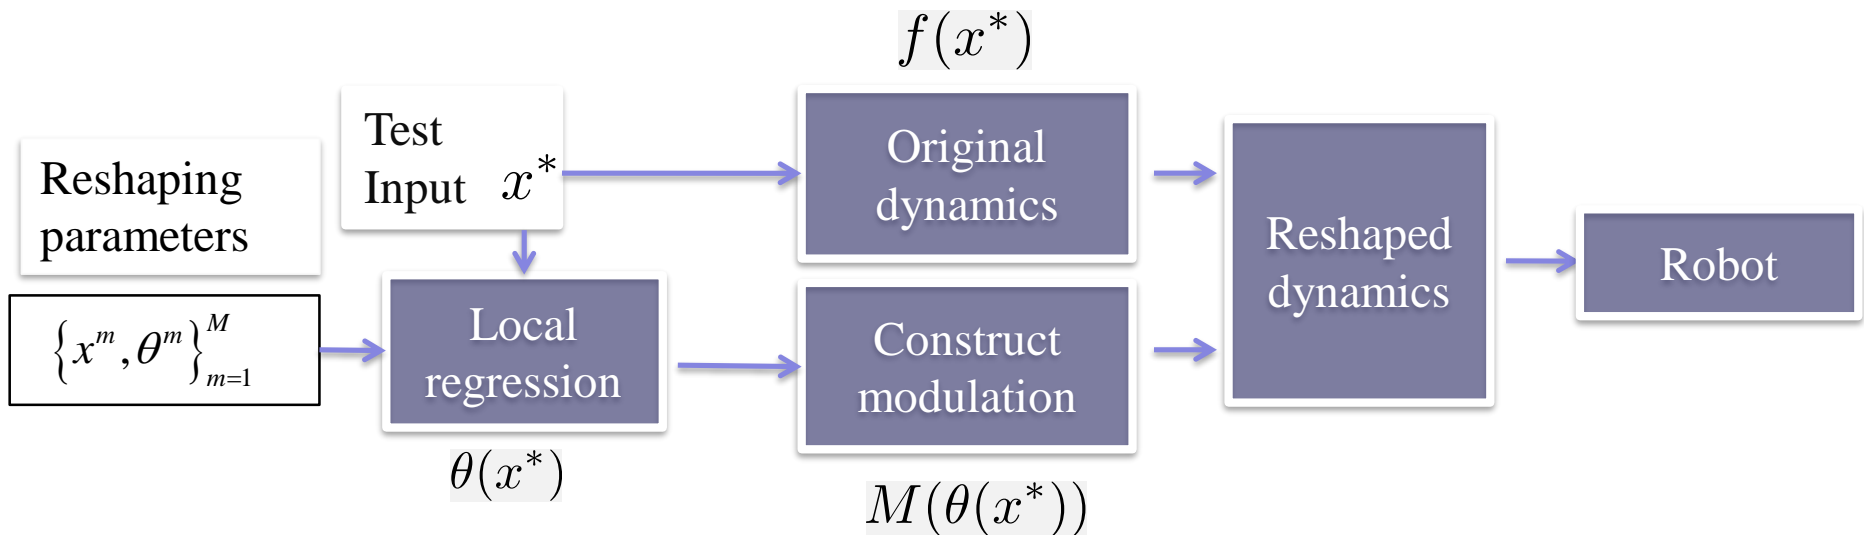
5: Compute scaling: $\kappa_m = \frac{\|\dot{x}^m\|}{\|{}^o\dot{x}^m\|} - 1$

6: 3D: $\theta^m = [\phi^m \mu^m, \kappa^m]$, 2D: $\theta^m = [\phi^m, \kappa^m]$

7: **end for**

8: **return** Modulation data $\{x^m, \theta^m\}_{m=1}^M$

Learning and Using the Reshaped Dynamics



Use Gaussian Process Regression (GPR)

$$M(x; \Theta = \{\phi, \mu_R, \kappa\}) = \left(1 + \kappa(x; X, \dot{X})\right) R(\phi(x; X), \mu_R(x; X))$$


Estimate the vector through more GPRs

Estimate the speed scaling through 1 GPR
Scalar function changing in space.

1 GPR function to estimate
state-dependent angle

Gaussian Process Regression (GPR) for Learning the Modulation

Original linear dynamics: $\dot{x} = f(x)$



DATA: set of M datapoints

$$\{X, \dot{X}\} = \{x^m, \dot{x}^m\}_{m=1}^M$$

Gaussian Process Regression (GPR) for Learning the Modulation

GPR model:


$$\dot{x} = \sum_{i=1}^M \alpha_i k(x, x^i)$$

with $\alpha = \left[K(X, X) + \sigma_n^2 I \right]^{-1} \begin{bmatrix} \dot{x}^1 \\ \dots \\ \dot{x}^M \end{bmatrix}$

Gram matrix: $K(X, X) = \begin{bmatrix} k(x^1, x^1) & k(x^1, x^2) & \dots & k(x^1, x^M) \\ \cdot & \cdot & \cdot & \cdot \\ k(x^M, x^1) & k(x^M, x^2) & \dots & \cdot \end{bmatrix}$

Kernel:

$$k(x, x') = \sigma_f e^{-\frac{\|x-x'\|}{2l}}$$



DATA: set of M datapoints

$$\{X, \dot{X}\} = \{x^m, \dot{x}^m\}_{m=1}^M$$


Hyperparameters

$l > 0$: lengthscale parameter \sim kernel width - controls tightness of the fit.

$\sigma_f > 0$: scales the kernel to the range of values of the output.

$\sigma_n > 0$: noise parameter - controls tolerance on fit imprecision.

Lengthscale parameter



Too small



Too large




Good fit

Kernel:


$$k(x, x') = \sigma_f e^{-\frac{\|x-x'\|}{2l}}$$

$l > 0$: controls tightness of the fit.


Noise parameter



Too small



Too large




Good fit

$$\dot{x} = \sum_{i=1}^M \alpha_i k(x, x^i), \text{ with } \alpha = \left[K(X, X) + \sigma_n^2 I \right]^{-1} \begin{bmatrix} \dot{x}^1 \\ \dots \\ \dot{x}^M \end{bmatrix}$$

$\sigma_n > 0$: noise parameter - controls tolerance on fit imprecision

Convergence away from data

Fit converges to zero – mean of GPR far from datapoints



GPR model:

$$\dot{x} = \sum_{i=1}^M \alpha_i k(x, x^i)$$

$k(x, x^i) \sim 0$ away from data.

Convergence away from data

Modulated dynamics:

$$\dot{x} = \left(1 + \underset{\rightarrow 0}{\kappa(x)}\right) \underset{\rightarrow 0}{R(\phi(x))} f(x)$$

$\rightarrow I$

Active only locally,
zero elsewhere


$$\kappa(x) > 0, x \in \Omega$$

$$\kappa(x) = 0, \forall x \notin \Omega$$

$$R(x) \neq I, x \in \Omega$$

$$R(x) = I, \forall x \notin \Omega$$

Hard to determine where the GPR function vanishes exactly.



Forces stability away from modulation

GPR model:


$$\dot{x} = \sum_{i=1}^M \alpha_i k(x^*, x^i), \quad x^* : \text{query point}$$

$$\text{with } \alpha = \left[K(X, X) + \sigma_n^2 I \right]^{-1} \begin{bmatrix} \dot{x}^1 \\ \dots \\ \dot{x}^M \end{bmatrix}$$

Rewrite as a function of query point

$$\dot{x} = \underbrace{\left[K(X, X) + \sigma_n^2 I \right]^{-1}}_{\alpha(x^*)} K_{Xx^*}$$


Truncate the GPR function at fixed bounds.



$$\alpha'(x^*) = \begin{cases} 0 & \alpha(x^*) < \underline{\alpha} \\ \frac{1}{2} \left(1 + \sin \left(\frac{2\pi(\alpha(x^*) - \underline{\alpha})}{2\rho} - \frac{\pi}{2} \right) \right) \alpha(x^*) & \underline{\alpha} \leq \alpha(x^*) \leq \underline{\alpha} + \rho. \\ \alpha(x^*) & \underline{\alpha} + \rho < \alpha(x^*) \end{cases}$$


Example: Modulation Learned by GPR

Reshaped dynamics: $\dot{x} = M(x) f(x)$




Example: Modulation Learned by GPR


Original dynamics: $\dot{x} = f(x)$



Example: Modulation Learned by GPR


Reshaped dynamics: $\dot{x} = M(x) f(x)$





Can be used to improve SEDS fit on nonlinear trajectories.


Example: 3D Modulation Learned by GPR



Training data



Training points used by GP-MDS



Influence region of the GP








Example trajectories from
original dynamical system



Example trajectories from
reshaped dynamical system

Example: 3D Modulation For Robot Control



Learning an **external** modulation

Learning an **external** modulation

Until now, the modulation depends solely on the current state of the system.

$$\dot{x} = M(x) f(x).$$

To provide more flexibility and allow the user to control for the activation of the modulation, we make it dependent on an external signal.

$$\dot{x} = M(x, s) f(x), \quad s \in \mathbb{R}^M : \text{external input}$$

s : could be measured by any external sensor

Stability conditions

- $M(x, s)$ locally active
- $M(x, s)$ is full rank
- $M(x, s)f(x)$ is stable at x^*

How should we design $M(x, s)$ to preserve properties of nominal DS?

Design of the modulation function

The modulation combines a **state and input-dependent** scaling and rotation.

$$M(x, s) = (1 + \kappa(x, s)) R(x, s)$$

Parameterized by $\theta(x, s) = h_s(s) [\phi(x) \mu_R, \kappa(x)]$

Rotation parameters $\phi(x): \mathbb{R}^N \rightarrow [-\pi; \pi]$

Design an state-dependent angle function: $\phi(x) = h_x(x) \phi_c,$
 $\phi_c \in [-\pi, \pi]$

Speed scaling

$$\kappa(x): \mathbb{R}^N \rightarrow \mathbb{R}^+$$

Design of the modulation function

The modulation combines a **state and input–dependent** scaling and rotation.

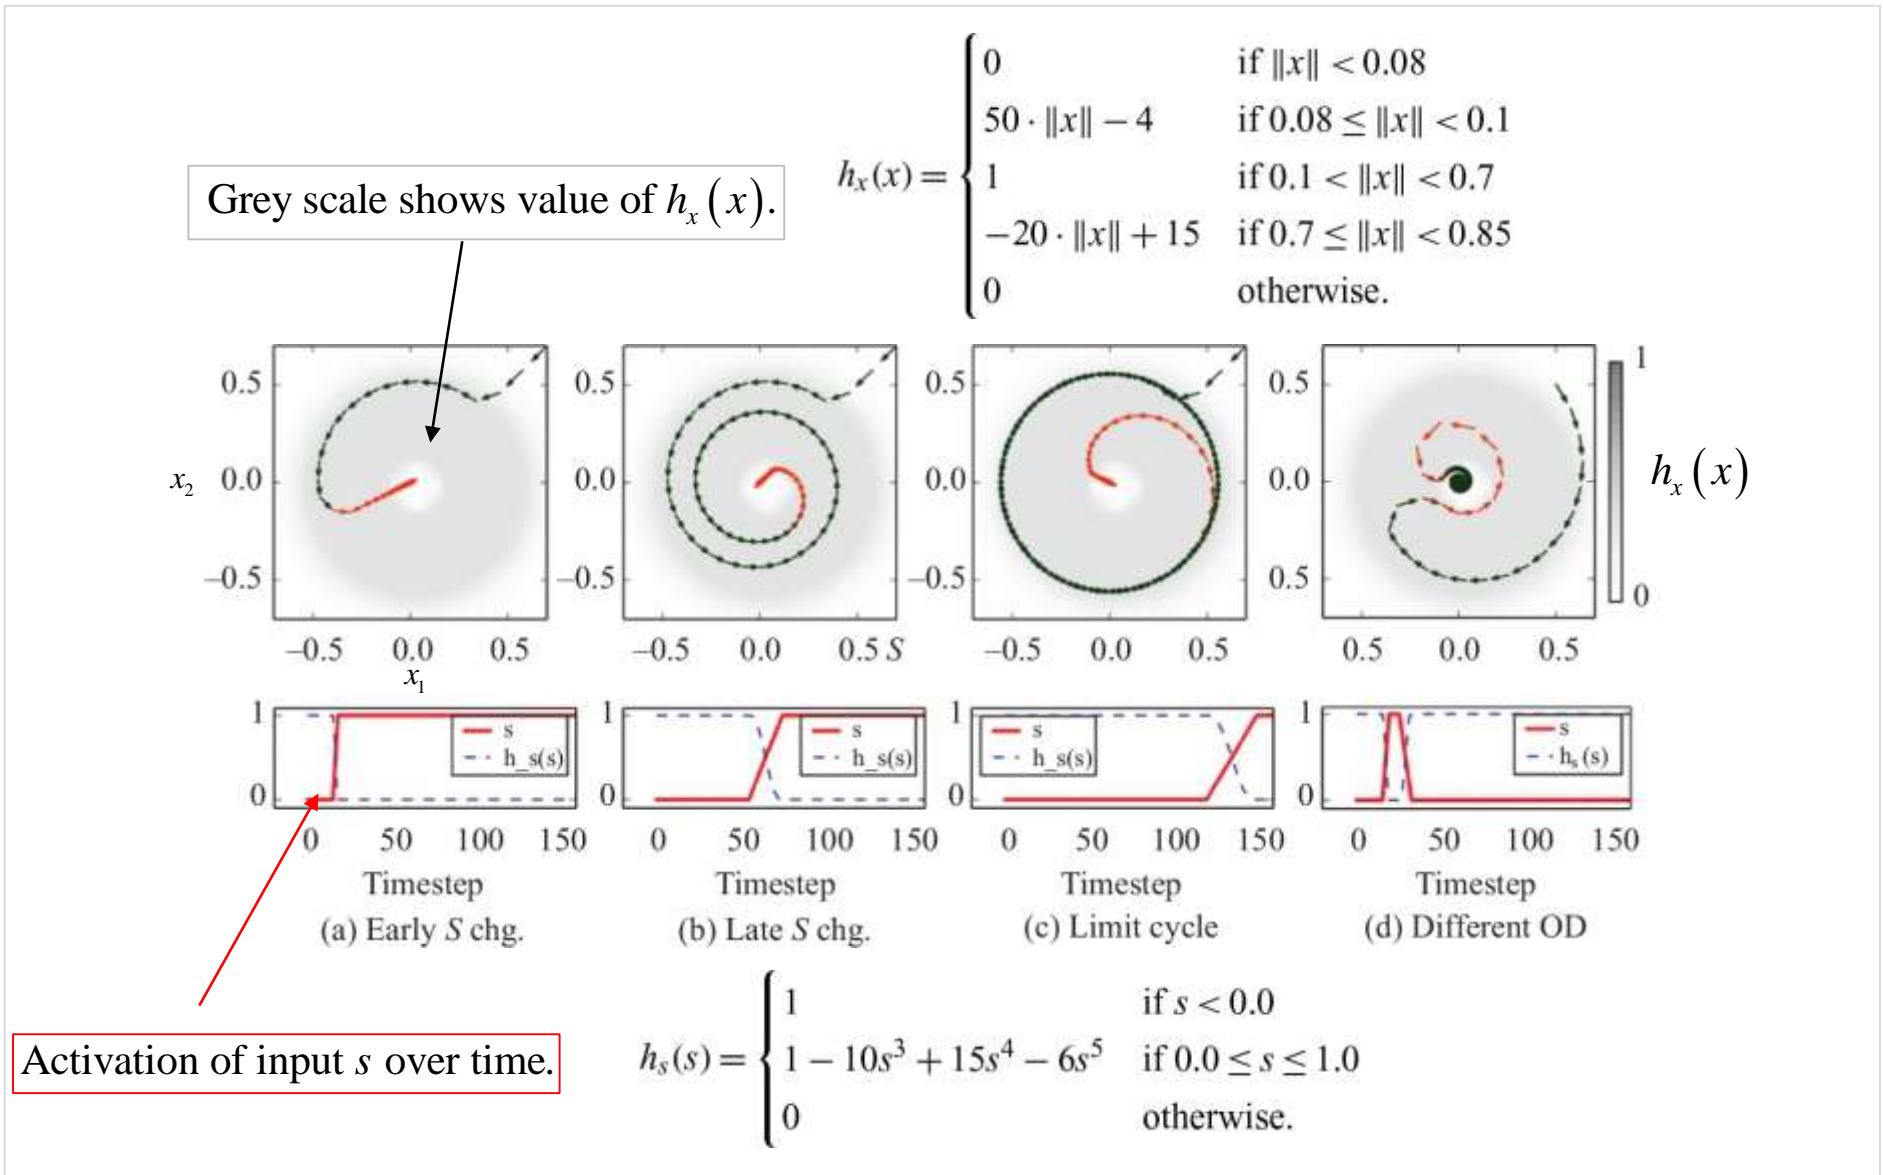
$$M(x, s) = (1 + \kappa(x, s)) R(x, s)$$

Parameterized by $\theta(x, s) = h_s(s) [\phi(x) \mu_R, \kappa(x)]$

$h_s(s): \mathbb{R}^M \rightarrow [0, 1]$ - determines when the modulation is active depending on external input


h_x determines the shape of the nonlinearity: $\phi(x) = h_x(x) \phi_c$.

Example of a modulation function




Application for tactile exploration

1. Original dynamics




2. Learn parameterization of modulation

$$\theta(x,s) = h_s(s) [\phi(x) \mu_R, \kappa(x)]$$



3. Learn activation function

$$h_s(s) : \mathbb{R}^M \rightarrow [0,1]$$



Application to tactile exploration


Scenario:

- “Blind” exploration of the space
- Each time make contact with object, localize the object and estimate its orientation
- Position the hand in the right orientation to grasp the object once certain of its location



Learn modulation function of the input

The more certain the algorithm is on position of object, the larger h_s .





s : variance of object's position estimate

Learn modulation function of the input

Use particle filter to determine the likelihood of the object's location



Spread of particles at initialization



Exp 1: example 1


Application to obstacle avoidance

- External input: collision detection
- Learn a modulation to adapt trajectory depending on collision information




Original dynamics

External signal: s
Time since last contact
+
Angle of last collision





Modulated dynamics

Resulting dynamics with different values of activation



Application to obstacle avoidance

Teaching the activation function




8 demonstrations



Learn activation function
(using Gaussian Process
Regression)

Application to obstacle avoidance

No modulation



With modulation



Summary

- Introduced a closed-form modulation of DS
 - Active locally only
 - Can be activated through external input
 - Can be designed to preserve stability properties of original DS
- The modulation can be learned from data
- The modulated system inherits DS properties:
 - enables highly reactive motion
- Modulating a DS is advantageous:
 - Enables to shape a simple nominal (linear DS) to make it nonlinear
 - Well suited for learning from demonstration