## Approximate regression based on a Reproducing

 kernel Hilbert spaces approach
## Damiano Varagnolo, Gianluigi Pillonetto, Luca Schenato

Department of Information Engineering - University of Padova (Italy)

$$
\text { May } 5^{\text {th }}, 2010
$$



## Problem statement

inputs: set of noisy measurements of a certain signal:

$$
y^{m}=f\left(x^{m}\right)+\nu^{m} \quad m=1, \ldots, M
$$

goal: estimate $f(x)$


## Parametric approach

## Parametric approach

assumption: known structure but unknown parameters
example: exponential:

$$
f(x)=\exp (-\theta x) \quad \theta, x \in \mathbb{R}^{+}
$$


goal: estimate $\theta$ starting from the data set $\left\{\left(x^{m}, y^{m}\right)\right\}$

## Parametric approach - interpretation

assume we don't know how the function is made: $f(\cdot)$ could be "almost everything"

$$
\Downarrow
$$

$f(\cdot)$ lives in an infinite dimensional space $\rightarrow$ there is infinite uncertainity

## Parametric approach - interpretation

assume we don't know how the function is made: $f(\cdot)$ could be "almost everything"

$$
\Downarrow
$$

$f(\cdot)$ lives in an infinite dimensional space $\rightarrow$ there is infinite uncertainity
parametric approach: restrict the function to live in a known and finite-dimensional space
$\Rightarrow$ it adds an infinite amount of prior information

## Parametric approach - order estimation

Quite important to estimate the order (e.g. for ARMA models)
Usual methods:

- Bayesian information criterion
- Akaike information criterion
- Mallow's $C_{p}$
general aim: find a trade-off between estimation error bias and estimation error variance


## Nonparametric approach

## Nonparametric approach

assumption: signal $f$ lives in a certain functions space:

$$
f \in \mathcal{H}_{K}
$$

goal: search the estimate $\widehat{f}$ directly inside this space, in general via:

$$
\widehat{f}=\arg \min _{\widetilde{f} \in \mathcal{H}_{K}}\left(\text { Loss function }\left(\widetilde{f},\left\{y^{m}\right\}\right)+\gamma\|\widetilde{f}\|_{\mathcal{H}_{K}}^{2}\right)
$$

motivations: functional structure of $f$ could be not easily managed with parametric structures

## Nonparametric approach - initial hypotheses

measurement model:

$$
y^{m}=L_{m}(f)+\nu^{m}
$$

where:

- functional $L_{m}(f)$ is linear and continuous in $f$
- measurement noise $\nu^{m}$ is:
- zero-mean Gaussian
- i.i.d.
- independent on $f$ and on $L_{m}(\cdot)$
- $f \in \mathcal{H}_{K}$
- $\mathcal{H}_{K}$ is an infinite-dimensional Hilbert space


## From infinite to finite dimensionality

Theorem (Representer theorem - hypothesis)
Given the cost-function minimization problem:

$$
\widehat{f}=\arg \min _{\tilde{f} \in \mathcal{H}_{K}} Q\left(L_{1}(\widetilde{f}), \ldots, L_{M}(\widetilde{f}), y^{1}, \ldots, y^{M},\|\widetilde{f}\|_{\mathcal{H}_{K}}^{2}\right)
$$

assume:

- $L_{m}(\widetilde{f})$ are linear and continuous in $\widetilde{f}$
- $Q(\cdot)$ is strictly increasing in $\|\widetilde{f}\|_{\mathcal{H}_{\kappa}}$
- there exists a solution to

$$
\arg \min _{\widetilde{f} \in \mathcal{H}_{K}} Q(\cdot)
$$

## From infinite to finite dimensionality

Theorem (Representer theorem - conclusion)
... then the solution is on the form

$$
\widehat{f}(\cdot)=\sum_{m=1}^{M} c^{m} g_{m}(\cdot)
$$

with:

- (using Riesz' representation theorem)

$$
L_{m}(f)=\left\langle g_{m}, f\right\rangle_{\mathcal{H}_{K}}
$$

- $\operatorname{span}\left\langle g_{1}, \ldots, g_{M}\right\rangle$ is at most M-dimensional
- weights $c^{m}$ depend on $Q(\cdot)$ (will be derived later)


## Usual cost functions

with quadratic losses:

$$
Q(\widetilde{f})=\sum_{m=1}^{M} \frac{\left(\widetilde{f}\left(x^{m}\right)-y^{m}\right)^{2}}{\sigma^{2}}+\gamma\|\widetilde{f}\|_{\mathcal{H}_{K}}^{2}
$$

with Vapnik's $\epsilon$-insensitive losses:

$$
Q(\widetilde{f})=\sum_{m=1}^{M} V\left(\widetilde{f}\left(x^{m}\right), y^{m}\right)+\gamma\|\widetilde{f}\|_{\mathcal{H}_{K}}^{2}
$$

where:

$$
V\left(\widetilde{f}\left(x^{m}\right), y^{m}\right):= \begin{cases}0 & \text { if }\left|\widetilde{f}\left(x^{m}\right)-y^{m}\right| \leq \epsilon \\ \left|\widetilde{f}\left(x^{m}\right)-y^{m}\right|-\epsilon & \text { otherwise }\end{cases}
$$

## Reproducing kernel Hilbert spaces

Definition
An Hilbert space $\mathcal{H}_{K}$ is said to have a reproducing kernel if there exists:

$$
K(\cdot, \cdot): \mathcal{D} \times \mathcal{D} \rightarrow \mathcal{M}
$$

such that:

$$
f(x)=\langle f(\cdot), K(x, \cdot)\rangle_{\mathcal{H}_{K}}
$$

(called the reproducing property)

Theorem
If the reproducing kernel $K(\cdot, \cdot)$ exists then it is unique

## How to compute the optimal estimate

$$
\text { Representer theorem } \Rightarrow \widehat{f}(\cdot)=\sum_{m=1}^{M} c^{m} g_{m}(\cdot)
$$

Reproducing kernel property $\quad \Rightarrow \quad g_{m}(\cdot)=K\left(x^{m}, \cdot\right)$

Together $\Rightarrow \quad \widehat{f}(\cdot)=\sum_{m=1}^{M} c^{m} K\left(x^{m}, \cdot\right)$

## Numerical solution with quadratic loss functions

If:

$$
\widehat{f}=\arg \min _{\widetilde{f} \in \mathcal{H}_{K}}\left(\sum_{m=1}^{M} \frac{\left(\widetilde{f}\left(x^{m}\right)-y^{m}\right)^{2}}{\sigma^{2}}+\gamma\|\widetilde{f}\|_{\mathcal{H}_{K}}^{2}\right)
$$

then:

$$
\left[\begin{array}{c}
c^{1} \\
\vdots \\
c^{M}
\end{array}\right]=\left(\left[\begin{array}{ccc}
K\left(x^{1}, x^{1}\right) & \cdots & K\left(x^{1}, x^{M}\right) \\
\vdots & & \vdots \\
K\left(x^{M}, x^{1}\right) & \cdots & K\left(x^{M}, x^{M}\right)
\end{array}\right]+\gamma I_{M}\right)^{-1}\left[\begin{array}{c}
y^{1} \\
\vdots \\
y^{M}
\end{array}\right]
$$

## Numerical solution in Bayesian frameworks

first hypothesis: $f$ is a realization of a zero-mean Gaussian process with covariance $K$ :

$$
\operatorname{cov}\left(f\left(x^{m}\right), f\left(x^{n}\right)^{T}\right)=K\left(x^{m}, x^{n}\right)
$$

second hypothesis: $f$ is independent on the measurement noise


It is equal to the quadratic cost-function based estimator

## Drawbacks

$$
\text { Optimal estimate: } \quad \widehat{f}(\cdot)=\sum_{m=1}^{M} c^{m} K\left(x^{m}, \cdot\right)
$$

$1^{\circ}$ feature: must invert $\left(K+\gamma I_{M}\right)^{-1}$
$2^{\circ}$ feature: must store $\left[c^{1}, \ldots, c^{M}\right]$

Possible problems: if $M$ is big then it could be:

- computationally hard to find (invert an $M \times M$ matrix)
- hard to store or communicate (representation can be quite big)


## Approximated regression

## Approximated non parametric regression -

 introduction need for reduction in computational complexity, i.e.- need estimation algorithms with an $O(\cdot)$ smaller than $O\left(M^{3}\right)$
- need representations using less than $M$ scalars

$$
\Downarrow
$$

must find:

- an $E$-dimensional model with $E \ll M$ such that:

$$
M:=\left[\phi_{1}(\cdot), \ldots, \phi_{E}(\cdot)\right] \mathbb{R}^{E} \quad M \subseteq \mathcal{H}_{K}
$$

- how to map the data set $\{\mathcal{X}, \mathcal{Y}\}$ into $M$


## Notation

Extension of finite linear algebra operations:

$$
\begin{gathered}
f^{T} g:=\int f(x)^{T} g(x) d x \\
A f\left(x^{\prime}\right):=\int A\left(x^{\prime}, x\right) f(x) d x \\
f^{T} A g:=\iint f\left(x^{\prime}\right)^{T} A\left(x^{\prime}, x\right) g(x) d x^{\prime} d x
\end{gathered}
$$

## How to map data sets into the estimation model

assume basis $\Phi:=\left[\phi_{1}(\cdot), \ldots, \phi_{E}(\cdot)\right]$ is given

If the inner product $P$ of $\mathcal{H}_{K}$ is given then:

- the projection operator $\mathcal{P}$ is:

$$
\mathcal{P}=\Phi\left(\Phi^{T} P \Phi\right)^{-1} \Phi^{T} P
$$

- the remainder operator $\mathcal{R}$ is given by:

$$
\mathcal{R}=I-\mathcal{P}
$$

- $\mathcal{P}$ and $\mathcal{R}$ are such that:

$$
\|f\|_{\mathcal{H}_{K}}^{2}=\|\mathcal{P} f\|_{\mathcal{H}_{K}}^{2}+\|\mathcal{R} f\|_{\mathcal{H}_{K}}^{2} \quad \forall f \in \mathcal{H}_{K}
$$

## How to map data sets into the estimation model

Given the projection operator $\mathcal{P}$,
if optimal estimate in $\mathcal{H}_{K}: \quad \widehat{f}(\cdot)=\sum_{m=1}^{M} c^{m} K\left(x^{m}, \cdot\right)$
then optimal estimate in $M: \quad \widehat{\mathcal{P}}(\cdot)$
drawback: still requires the explicit computation of the optimal $\widehat{f}$ conceptual advantage: the optimal basis $\Phi$ is the one that maximizes $\mathbb{E}\left[\|\mathcal{P} \widehat{f}\|_{\mathcal{H}_{K}}^{2}\right] \rightarrow$ gives the idea of how to find the optimal basis

## How to find the optimal estimation model

Imposition of additional hypotheses:

- $K(\cdot, \cdot)$ is a Mercer Kernel:
- continuous
- symmetric
- definite positive ${ }^{\star}$
- the input locations domain $\mathcal{D}$ is compact


## How to find the optimal estimation model - first implications

1: $K(\cdot, \cdot)$ defines a compact linear positive definite integral operator:

$$
\left(L_{K} f\right)\left(x^{\prime}\right):=\int_{\mathcal{D}} K\left(x^{\prime}, x\right) f(x) d x=K f\left(x^{\prime}\right)
$$

2: there are at most a numerable set of eigenfunctions $\phi(\cdot)$ :

$$
K \phi_{k}(\cdot)=\lambda_{k} \phi_{k}(\cdot) \quad k=1,2, \ldots
$$

## How to find the optimal estimation model - second

 implicationsTheorem (Mercer's)
with the previous hypotheses:

- $\left\{\lambda_{k}\right\}$ are real and non-negative: $\lambda_{1} \geq \lambda_{2} \geq \ldots \geq 0$
- $\left\{\phi_{k}(\cdot)\right\}$ is an orthonormal basis for the space

$$
\mathcal{H}_{K}=\left\{f \in \mathcal{L}^{2} \text { s.t. } f=\sum_{k=1}^{\infty} a_{k} \phi_{k} \left\lvert\, \sum_{k=1}^{\infty} \frac{a_{k} \cdot a_{k}}{\lambda_{k}}<+\infty\right.\right\}
$$

- $f_{1}=\sum_{k=1}^{\infty} a_{k} \phi_{k} \quad f_{2}=\sum_{k=1}^{\infty} b_{k} \phi_{k} \Rightarrow\left\langle f_{1}, f_{2}\right\rangle_{\mathcal{H}_{k}}=\sum_{k=1}^{\infty} \frac{a_{k} \cdot b_{k}}{\lambda_{k}}$


## How to find the optimal estimation model

## use the PCA idea to find the optimal basis $\Phi$

$\Rightarrow$ optimal $\Phi$ is the set the first $E$ eigenfunctions
note: $\mathbb{E}\left[\|\widehat{f}\|_{\mathcal{H}_{K}}^{2}\right]=\sum_{k=1}^{\infty} \lambda_{k} \Rightarrow\left\{\begin{array}{l}\mathbb{E}\left[\|\mathcal{P} \widehat{f}\|_{\mathcal{H}_{K}}^{2}\right]=\sum_{k=1}^{E} \lambda_{k} \\ \mathbb{E}\left[\|\mathcal{R} \widehat{f}\|_{\mathcal{H}_{K}}^{2}\right]=\sum_{k=E+1}^{\infty} \lambda_{k}\end{array}\right.$
how to choose $E$ : approximation error effect $\sum_{k=E+1}^{\infty} \lambda_{k}$ should be comparable to the measurement noise

## Desired qualities of the approximated regression algorithms

We are looking for an estimate living in a E-dimensional space spanned by eigenfunctions $\phi_{1}(\cdot), \ldots, \phi_{E}(\cdot)$, i.e.: $\widehat{f}=\sum_{k=1}^{E} a_{k} \phi_{k}$

$$
\text { Question: how to compute } a_{1}, \ldots, a_{E} \text { ? }
$$

Constraints:

- we don't want to compute the optimal estimate $\sum_{m=1}^{M} c^{m} K\left(x^{m}, \cdot\right)$
- we don't want to use the projection operator $\mathcal{P}$


## New notation

## measurement model:

$$
y^{m}=\sum_{k=1}^{+\infty} a_{k} \phi_{k}\left(x^{m}\right)+\nu^{m} \quad \rightarrow \quad \mathcal{Y}=C \mathbf{a}+\mathbf{e}+\mathcal{V}
$$

definitions:

$$
\begin{gathered}
\mathcal{Y}:=\left[\begin{array}{c}
y^{1} \\
\vdots \\
y^{M}
\end{array}\right] \quad C:=\left[\begin{array}{ccc}
\phi_{1}\left(x^{1}\right) & \ldots & \phi_{E}\left(x^{1}\right) \\
\vdots & & \vdots \\
\phi_{1}\left(x^{M}\right) & \ldots & \phi_{E}\left(x^{M}\right)
\end{array}\right] \\
\mathbf{a}:=\left[\begin{array}{c}
a_{1} \\
\vdots \\
a_{E}
\end{array}\right] \quad \mathbf{e}:=\left[\begin{array}{c}
\sum_{k=E+1}^{+\infty} a_{k} \phi_{k}\left(x^{1}\right) \\
\vdots \\
\sum_{k=E+1}^{+\infty} a_{k} \phi_{k}\left(x^{M}\right)
\end{array}\right] \quad \mathcal{V}:=\left[\begin{array}{c}
\nu^{1} \\
\vdots \\
\nu^{E}
\end{array}\right]
\end{gathered}
$$

## Approximated learning - kind of approaches

cost-function:

- data fitting $\rightarrow$ loss functions
- not overfit $\rightarrow$ Tikhonov regularizer

$$
\widehat{f}=\arg \min _{\widetilde{f} \in \mathcal{H}_{K}^{E}}\left(\sum_{m=1}^{M} \frac{\left(\widetilde{f}\left(x^{m}\right)-y^{m}\right)^{2}}{\sigma^{2}}+\gamma\|\widetilde{f}\|_{\mathcal{H}_{K}^{E}}^{2}\right)
$$

Bayesian:

- put a prior on the eigenfunctions weights $a_{k}$
- find the best linear unbiased estimator:

$$
\widehat{\mathbf{a}}=\operatorname{cov}(\mathbf{a}, \mathcal{Y}) \operatorname{var}(\mathcal{Y})^{-1} \mathcal{Y}
$$

## Approximated learning - cost-function approach

$$
\begin{gathered}
\widehat{f}=\arg \min _{\widetilde{f} \in \mathcal{H}_{K}^{E}}\left(\sum_{m=1}^{M} \frac{\left(\tilde{f}\left(x^{m}\right)-y^{m}\right)^{2}}{\sigma^{2}}+\gamma\|\widetilde{f}\|_{\mathcal{H}_{K}^{E}}^{2}\right) \\
\Downarrow \\
\widehat{\mathbf{a}}=\left(\sigma^{2} \Sigma_{\mathbf{a}} C^{T} C+\gamma I_{E}\right)^{-1} \Sigma_{\mathbf{a}} C^{T} \mathcal{Y}
\end{gathered}
$$

$$
\left(\text { with } \quad \Sigma_{\mathbf{a}}:=\mathbb{E}\left[\mathbf{a a}^{T}\right]=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{E}\right)\right)
$$

computations load: $O\left(E^{3}+E^{2} M+E M^{2}\right)$ operations representations size: $E$ scalars

## Approximated learning - Bayesian approach

$$
\begin{gathered}
\text { prior: } a_{k} \sim \mathcal{N}\left(0, \lambda_{k}\right) \\
\widehat{\mathbf{a}}=\operatorname{cov}(\mathbf{a}, \mathcal{Y}) \operatorname{var}(\mathcal{Y})^{-1} \mathcal{Y} \\
\Downarrow \\
\widehat{\mathbf{a}}=\Sigma_{\mathbf{a}} C^{T}\left(C \Sigma_{\mathbf{a}} C^{T}+\Sigma_{\mathbf{e}}+\sigma^{2} I_{M}\right)^{-1} \mathcal{Y} \\
\left(\text { with } \Sigma_{\mathbf{e}}:=\mathbb{E}\left[\mathbf{e e}^{T}\right]\right)
\end{gathered}
$$

computations load: $O\left(M^{3}\right)$ operations
representations size: $E$ scalars

## Approximated learning - comparisons of the

 numerical solutionscost-function approach:

$$
\widehat{\mathbf{a}}=\left(\sigma^{2} \Sigma_{\mathbf{a}} C^{T} C+\gamma I_{E}\right)^{-1} \Sigma_{\mathbf{a}} C^{T} \mathcal{Y} \quad \rightarrow \quad O\left(E^{3}+E^{2} M+E M^{2}\right)
$$

Bayesian approach:

$$
\begin{gathered}
\widehat{\mathbf{a}}=\Sigma_{\mathbf{a}} C^{T}\left(C \Sigma_{\mathbf{a}} C^{T}+\Sigma_{\mathrm{e}}+\sigma^{2} I_{M}\right)^{-1} \mathcal{Y} \quad \rightarrow \quad O\left(M^{3}\right) \\
\Downarrow \\
\text { not equivalent! }
\end{gathered}
$$

## Eigenfunctions estimation

## Estimation of the eigenfunctions - introduction

## Questions:

- how to obtain the eigenfunctions $\phi_{k}(\cdot)$ given the kernel $K(\cdot, \cdot)$ ?
- how to obtain the eigenfunctions $\phi_{k}(\cdot)$ if we don't know even the kernel $K(\cdot, \cdot)$ ?


## Estimation of the eigenfunctions - introduction

## Questions:

- how to obtain the eigenfunctions $\phi_{k}(\cdot)$ given the kernel $K(\cdot, \cdot)$ ?
- how to obtain the eigenfunctions $\phi_{k}(\cdot)$ if we don't know even the kernel $K(\cdot, \cdot)$ ?

Remark: we work in a subspace of $\mathcal{L}^{2}$ :

- $K(\cdot, \cdot)$ is continuous (already given since it is Mercer)
- $\phi_{k}(\cdot)$ is a continuous function (already given by Mercer's theorem)


## Estimation of the eigenfunctions given the kernel

Suppose $K(\cdot, \cdot)$ is given. Then if $\phi(\cdot)$ is eigenfunction and $\lambda$ is its eigenvalue:

$$
\int_{\mathcal{D}} K\left(x, x^{\prime}\right) \phi\left(x^{\prime}\right) d x^{\prime}=\lambda \phi(x)
$$

we can approximate:

$$
\int_{\mathcal{D}} K\left(x, x^{\prime}\right) \phi\left(x^{\prime}\right) d x^{\prime} \approx \sum_{j=1}^{Q} K\left(x^{i}, x^{j}\right) \phi\left(x^{j}\right) w_{j}
$$

Linear system from which to estimate $\phi(\cdot)$ and $\lambda$ :

$$
\sum_{j=1}^{Q} K\left(x^{i}, x^{j}\right) \phi\left(x^{j}\right) w_{j}=\lambda \phi\left(x^{i}\right) \quad i=1, \ldots, Q
$$

## Estimation of the eigenfunctions given the kernel

$$
\begin{gathered}
\sum_{j=1}^{Q} K\left(x^{i}, x^{j}\right) \phi\left(x^{j}\right) w_{j}=\lambda \phi\left(x^{i}\right) \quad i=1, \ldots, Q \\
{\left[\begin{array}{ccc}
K\left(x^{1}, x^{1}\right) w_{1} & \cdots & K\left(x^{1}, x^{Q}\right) w_{Q} \\
\vdots & \vdots \\
K\left(x^{Q}, x^{1}\right) w_{1} & \cdots & K\left(x^{Q}, x^{Q}\right) w_{Q}
\end{array}\right]\left[\begin{array}{c}
\phi\left(x^{1}\right) \\
\vdots \\
\phi\left(x^{Q}\right)
\end{array}\right]=\lambda\left[\begin{array}{c}
\phi\left(x^{1}\right) \\
\vdots \\
\phi\left(x^{Q}\right)
\end{array}\right]}
\end{gathered}
$$

solve an eigenvalue-eigenvector problem

Note: choice of $\left\{x^{i}\right\}$ and $\left\{w_{i}\right\}$ can be critical

## Estimation of the eigenfunctions without knowing

 the kernelIf $K(\cdot, \cdot)$ is unknown then:
(1) estimate the covariance of the stochastic process and obtain $\widehat{C}$
(2) assume the kernel is the estimated covariance, i.e. $K(\cdot, \cdot)=\widehat{C}$
(0) proceed as before

Note: choice of $\left\{x^{i}\right\}$ and $\left\{w_{i}\right\}$ is less critical than then the estimation of $\widehat{C}$

## Example of eigenfunctions

Kernel for BIBO stable linear time-invariant systems:

$$
K\left(x, x^{\prime} ; \beta\right)= \begin{cases}\frac{\exp (-2 \beta x)}{2}\left(\exp \left(-\beta x^{\prime}\right)-\frac{\exp (-\beta x)}{3}\right) & \text { if } x \leq x^{\prime} \\ \frac{\exp \left(-2 \beta x^{\prime}\right)}{2}\left(\exp (-\beta x)-\frac{\exp \left(-\beta x^{\prime}\right)}{3}\right) & \text { if } x \geq x^{\prime}\end{cases}
$$



## Drawbacks

$\phi_{k}(\cdot)$ cannot be computed from $\phi_{k-1}(\cdot), \ldots, \phi_{1}(\cdot)$

can be computationally expensive if eigenfunctions have to be estimated "on-the-fly"

## Distributed estimation

## Distributed approximated regression - Introduction

## Our framework:

- there is a zero-mean Gaussian process $\mathcal{F}$ of which we know the covariance-kernel:

$$
\operatorname{cov}\left(\mathcal{F}(x, t), \mathcal{F}(x, t)^{T}\right)
$$

(e.g.: wind blowing on a wind farm: $x=[$ lat. lon. height $]$ )

- there are $S$ sensors that sample the same realization $f$ drawn from $\mathcal{F}$ :

$$
y_{s}^{m}=f\left(x_{s}^{m}, t_{s}^{m},\right)+\nu_{s}^{m}
$$

## Distributed approximated regression - Introduction

"our goal": distributely estimate the realization $f$ our constraint: sensors can exchange a limited amount of information


## Distributed approximated regression - Introduction

"our goal": distributely estimate the realization $f$ our constraint: sensors can exchange a limited amount of information

our actual goal: find distributed algorithms and characterize their performances (variance of the estimation error)

## Distributed estimation: first algorithm

First step: think to an effective estimator simplificative hypothesis: sensors measure the same realization




Appreciable characteristics:

- no common sampling grid
- unknown time delays


## Distributed estimation with known delays

If we know the delays between the various functions we can:
(1) (locally) shift the various data sets
(2) (locally) compute the eigenfunctions weights $a_{k}^{s}$
(0) (distributely) make average consensus on the weights $a_{k}^{s}$

- (locally) shift back the representation


$y_{3}$



## Distributed estimation with known delays

If we know the delays between the various functions we can:
(1) (locally) shift the various data sets
(2) (locally) compute the eigenfunctions weights $a_{k}^{s}$
(3) (distributely) make average consensus on the weights $a_{k}^{s}$
( (locally) shift back the representation




## Distributed estimation with known delays

If we know the delays between the various functions we can:
(1) (locally) shift the various data sets
(2) (locally) compute the eigenfunctions weights $a_{k}^{s}$
(3) (distributely) make average consensus on the weights $a_{k}^{s}$
( (locally) shift back the representation




## Distributed estimation with known delays

If we know the delays between the various functions we can:
(1) (locally) shift the various data sets
(2) (locally) compute the eigenfunctions weights $a_{k}^{s}$
(3) (distributely) make average consensus on the weights $a_{k}^{s}$
( (locally) shift back the representation




## Distributed estimation with known delays

If we know the delays between the various functions we can:
(1) (locally) shift the various data sets
(2) (locally) compute the eigenfunctions weights $a_{k}^{s}$
(0) (distributely) make average consensus on the weights $a_{k}^{s}$
( (locally) shift back the representation


$y_{3}$


## Distributed estimation with known delays

If we know the delays between the various functions we can:
(1) (locally) shift the various data sets
(2) (locally) compute the eigenfunctions weights $a_{k}^{s}$
(3) (distributely) make average consensus on the weights $a_{k}^{s}$
( (locally) shift back the representation



results in general not equivalent to centralized estimate!

## Distributed estimation with unknown delays

## And if we do not know the delays?

first formulate a centralized optimization problem with a cost-function based regularization:

$$
-\ln p\left(x_{1}^{1}, y_{1}^{1}, \ldots, x_{S}^{M}, y_{S}^{M} \mid \tau_{1}, \ldots, \tau_{S}, a_{1}, \ldots, a_{E}\right)+\gamma \sum_{k=1}^{E} \frac{a_{k}^{2}}{\lambda_{k}}
$$

then distributely solve it
Note: both minimizations use 2 -steps gradient descents:
(1) keep delays $\tau_{s}$ fixed and update the weights $a_{k}$
(2) keep the weights $a_{k}$ fixed and update the delays $\tau_{s}$

## Gradient descents steps: intuition

How do the gradient descent steps work?

Weights $a_{k}$ update: ( $\tau_{s}$ fixed)

Time delays $\tau_{s}$ update: ( $a_{k}$ fixed)

## Gradient descents steps: intuition

How do the gradient descent steps work?

Weights $a_{k}$ update: ( $\tau_{s}$ fixed)
(1) join all the shifted data sets


Time delays $\tau_{s}$ update: ( $a_{k}$ fixed)

## Gradient descents steps: intuition

How do the gradient descent steps work?

Weights $a_{k}$ update: ( $\tau_{s}$ fixed)
(1) join all the shifted data sets
(2) compute $\widehat{f}$ as before


Time delays $\tau_{s}$ update: ( $a_{k}$ fixed)

## Gradient descents steps: intuition

How do the gradient descent steps work?

Weights $a_{k}$ update: ( $\tau_{s}$ fixed)
(1) join all the shifted data sets
(2) compute $\widehat{f}$ as before


Time delays $\tau_{s}$ update: ( $a_{k}$ fixed)
(1) shift optimally each data set


## Gradient descents steps: intuition

How do the gradient descent steps work?

Weights $a_{k}$ update: ( $\tau_{s}$ fixed)
(1) join all the shifted data sets
(2) compute $\widehat{f}$ as before


Time delays $\tau_{s}$ update: ( $a_{k}$ fixed)
(1) shift optimally each data set


## Simulations - distributed function estimation



## Simulations - distributed function estimation


wind strength (sensor B)

process realization
$x \times \times \times \times \times \times \times \times$ measurements

## Simulations - distributed function estimation



\[

\]

## Simulations - distributed function estimation



## Characterization of the distributed algorithms

these algorithms can be effective $\Rightarrow$ worthy to be characterized
let's start with the simplest case:
(1) each sensor knows exactly $S$ ( $\mathrm{n}^{\circ}$ of sensors)
(2) no time-delay between measured signals
(3) common input-locations grid among sensors


## Simplest case: optimal distributed algorithm

 there exists a distributed strategy equivalent to the centralized one:(1) (locally) make initial estimations:

$$
\hat{\mathbf{a}}_{s}=\Sigma_{\mathbf{a}} C^{T}\left(C \Sigma_{\mathbf{a}} C^{T}+\Sigma_{\mathbf{e}}+\frac{\sigma^{2}}{S} I_{M}\right)^{-1} \mathcal{Y}_{s}
$$

(2) (distributely) make an average consensus on the various $\widehat{\mathbf{a}}_{s}$

Difference from pure local estimators: how to weight the measurement noise:

$$
\widehat{\mathbf{a}}_{s}^{\text {loc }}=\Sigma_{\mathbf{a}} C^{T}\left(C \Sigma_{\mathbf{a}} C^{T}+\Sigma_{\mathrm{e}}+\sigma^{2} I_{M}\right)^{-1} \mathcal{Y}_{s}
$$

## Guessed distributed strategy

hypothesis removal: sensors do not know $S$ ( $\mathrm{n}^{\circ}$ of sensors)

$$
\text { all sensors make the same guess: } \left.S_{g} \quad \text { ("g" }=\text { guess }\right)
$$

## how distributed estimator changes?

distributed strategy:
(1) (locally) make initial estimations:

$$
\hat{\mathbf{a}}_{s}\left(S_{g}\right)=\Sigma_{\mathrm{a}} C^{T}\left(C \Sigma_{\mathrm{a}} C^{T}+\Sigma_{\mathrm{e}}+\frac{\sigma^{2}}{S_{g}} I_{M}\right)^{-1} \mathcal{Y}_{s}
$$

(2) (distributely) make an average consensus on the various $\widehat{\mathbf{a}}_{s}\left(S_{g}\right) \equiv$

## Comparisons between estimators performances

## performance " $=$ " estimation error variance

centralized vs local: centralized is always better than local
centralized vs guessed distributed: centralized is always better than guessed distributed (equal iff $S=S_{g}$, (guess is correct))
guessed distributed vs local: depends!!
Proposition
If $S_{g} \in[1,2(S-1)]$ then guessed distributed strategy is better than local independently of the kernel, noise power, number of measurements, etc.

## Current research on performances characterization

remove the common grid hypothesis and perform similar comparative analyses between different algorithms of increasing complexity:

- simple average consensus of locally optimal estimates
- average consensus of local estimates with weighted measurement noise covariance
- local construction of pseudo-measurements on a common grid, then use the pseudo-measurements as before


## Other research directions

distributed number of sensors statistical estimation: • (locally) generate $y_{s}$ from a known probability distribution

- (distributely) combine these $y_{s}$ using a known function $f(\cdot)$
- (locally) use ML, MMSE or MAP strategies to estimate the actual number of sensors
distributed fault detection: (with faults on the measurements)
- make a distributed estimation
- make also a local estimation
- compare the local and the distributed estimations
- use statistical decision theory to locally say if there are problems on the measurements


## Appendix

## Bias vs. Variance tradeoff

$$
\begin{aligned}
\mathbb{E}_{\text {data set }}\left[(y-f(x))^{2}\right] & =\mathbb{E}_{x}\left[\mathbb{E}_{y}\left[(y-\mathbb{E}[y \mid x])^{2} \mid x\right]\right] \\
& \left.+\mathbb{E}_{x}\left[\mathbb{E}_{y}\left|(f(x)-\mathbb{E}[f(x)])^{2}\right| x\right]\right] \\
& +\mathbb{E}_{x}\left[\mathbb{E}_{y}\left[(\mathbb{E}[y \mid x]-\mathbb{E}[f(x)])^{2} \mid x\right]\right] \\
& =\mathbb{E}_{x}[\operatorname{lvar}(y \mid x)] \\
& +\mathbb{E}_{x}[\operatorname{var}(f(x))] \\
& +\mathbb{E}_{x}\left[(\operatorname{bias}(f(x)))^{2}\right]
\end{aligned}
$$

## Riesz' representation theorem

Definition (dual of an Hilbert space)
If $\mathcal{H}_{K}$ is a Hilbert space, then the space of the continuous linear functionals $L: \mathcal{H}_{K} \rightarrow \mathbb{R}$ is called its dual and indicated with $\mathcal{H}_{K}^{*}$

Theorem (Riesz' representation theorem)
If $\mathcal{H}_{K}$ is a Hilbert space and $\mathcal{H}_{K}^{*}$ is its dual, then

$$
\forall L \in \mathcal{H}_{K}^{*} \exists!g \in \mathcal{H}_{K} \text { s.t. } L(f)=\langle g, f\rangle \quad \forall f \in \mathcal{H}_{K}
$$

