Introduction to Functional Data Analysis and Its Role in the Age of AI

Jane-Ling Wang

University of California Davis

September 14, 2026
Connectomics, Non-Euclidean data, and Statistics
IMSI, Chicago

.

.

Acknowledgement: Research supported by NSF DMS-24-13924

Outline

1. Introduction to Functional Data Analysis
2. Mean and Covariance Estimation
3. Theory: Mean and Covariance Estimation

Outline

1. Introduction to Functional Data Analysis
2. Mean and Covariance Estimation
3. Theory: Mean and Covariance Estimation

What is Functional Data?

- Functional Data: A sample of random functions,
  with one function per subject.

  - These functions can be curves (1D), images (2D or 3D), or higher dimension object data.

What is Functional Data?

- Functional Data: A sample of random functions,
  with one function per subject.

  - These functions can be curves (1D), images (2D or 3D), or higher dimension object data.

- Characteristics of functional data:

  (i) The atom of functional data is a "function".
  (ii) They are $\infty-$dimensional data.

Example: Curve Data

- **Curve data**

  - real-valued functions defined on an interval $I \in \mathbb{R}$.

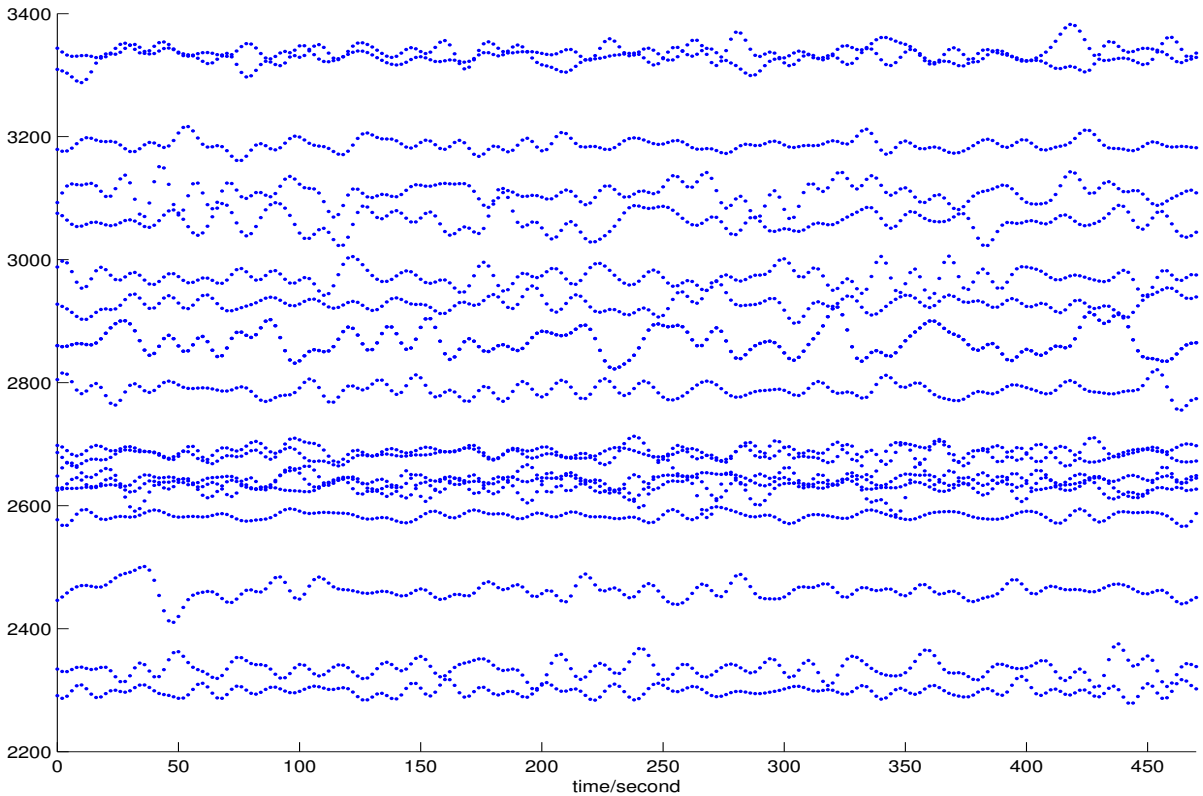
  - one curve per subject

  These curves are usually considered realizations of a stochastic process $X(t)$ in a **Hilbert space**, e.g.

  $L^2(I)$ or RKHS (reproducing kernel Hilbert space).

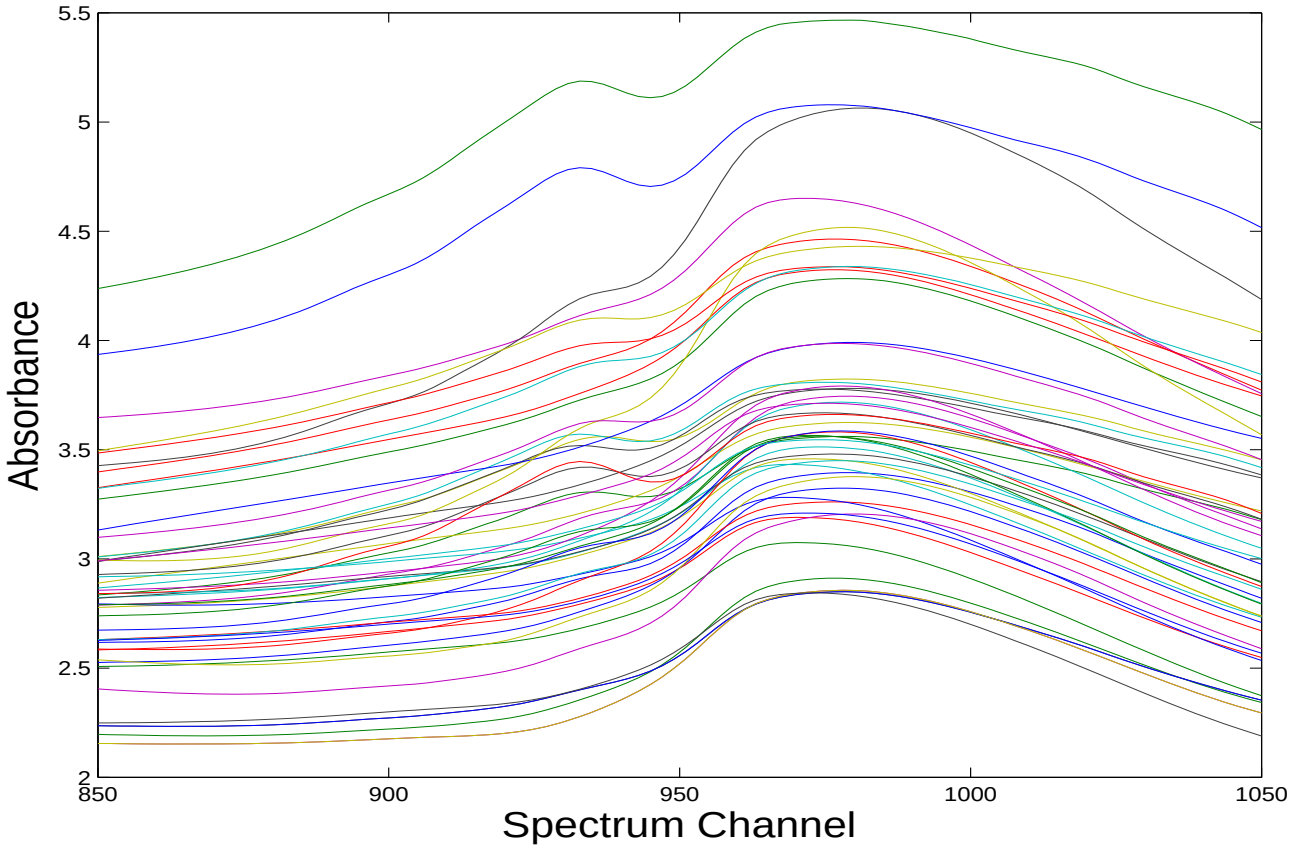
Examples of Functional Data

- fMRI data at a particular voxel for 20 subjects $\implies n = 20$.



Examples of Functional Data

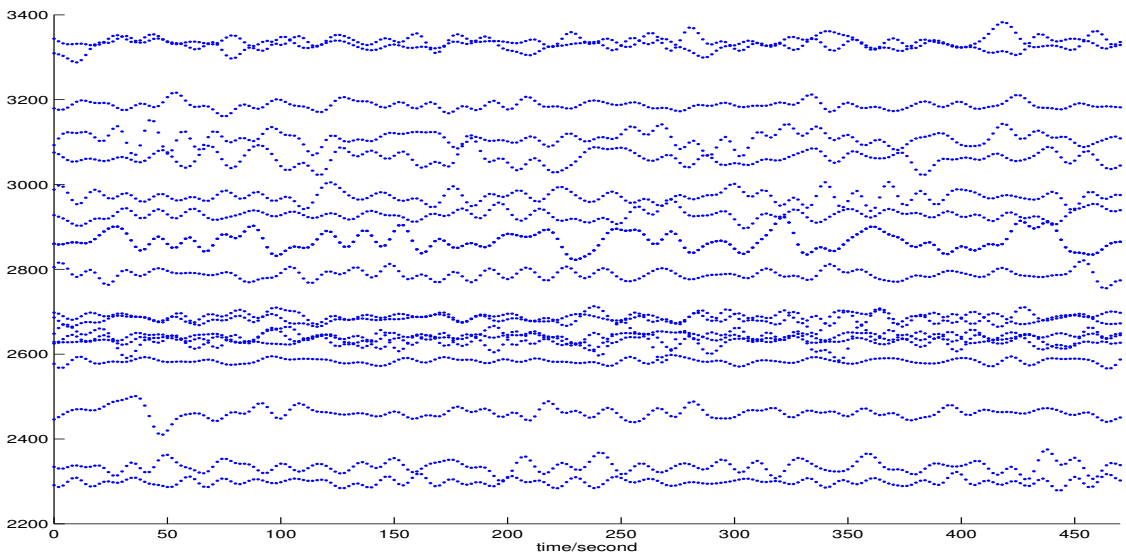
- Spectrum data for meat content - here the function is over the spectrum channels of n pieces of meat.



Real/Observed Functional Data

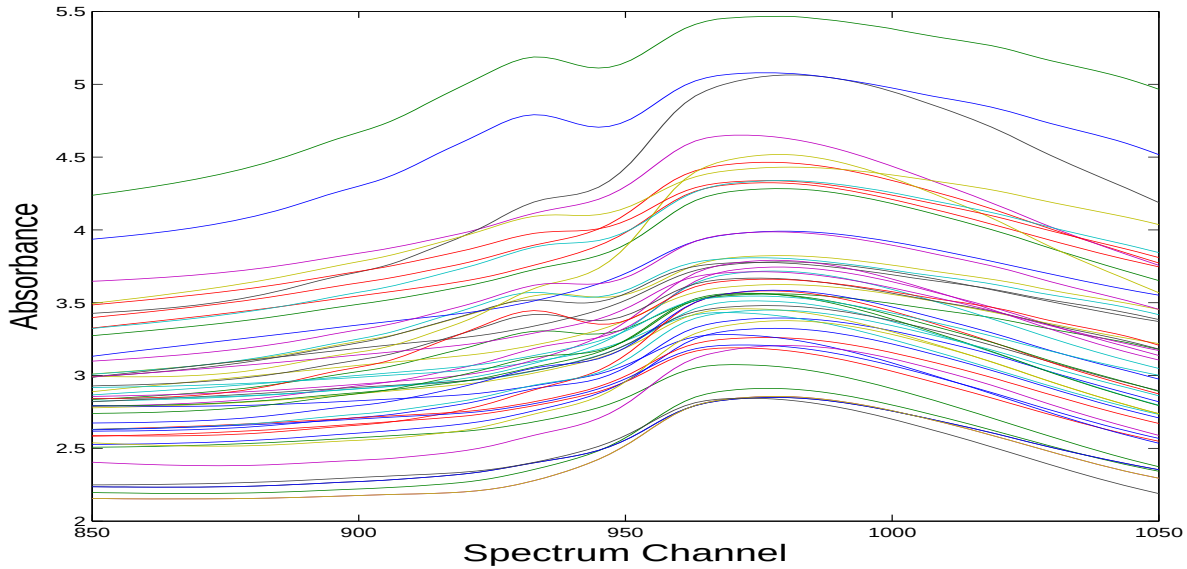
- In reality, functional data are recorded intensely on a time grid $\Longrightarrow$ high-dimensional data.

  * The fMRI data were recorded every two seconds for about 10 minutes (300 time points) $\Longrightarrow$ 300 dimensional data.



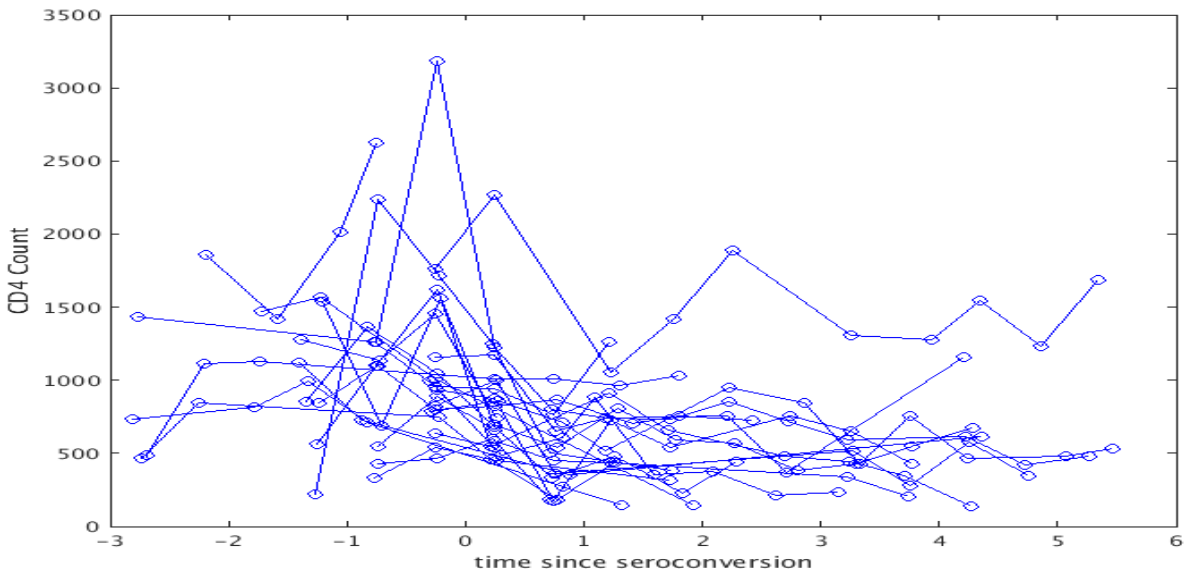
Real/Observed Functional Data

- The spectrum data were recorded at 100 frequency channels (hence 100-dim) and smoothed individually, i.e. pre-smoothed.



Longitudinal Data as Functional Data

- Longitudinal Data - Irregularly sampled functional data.
- They are often only a few measurements per subject, as in medical follow-up or social studies.



Longitudinal AIDS Data

- CD4 counts of 369 patients.

  $n_i = \#$ of repeated measurements for subject i,

  varies with subject.

  - An average of 6.44 measurement per subject.

Longitudinal AIDS Data

- CD4 counts of 369 patients.

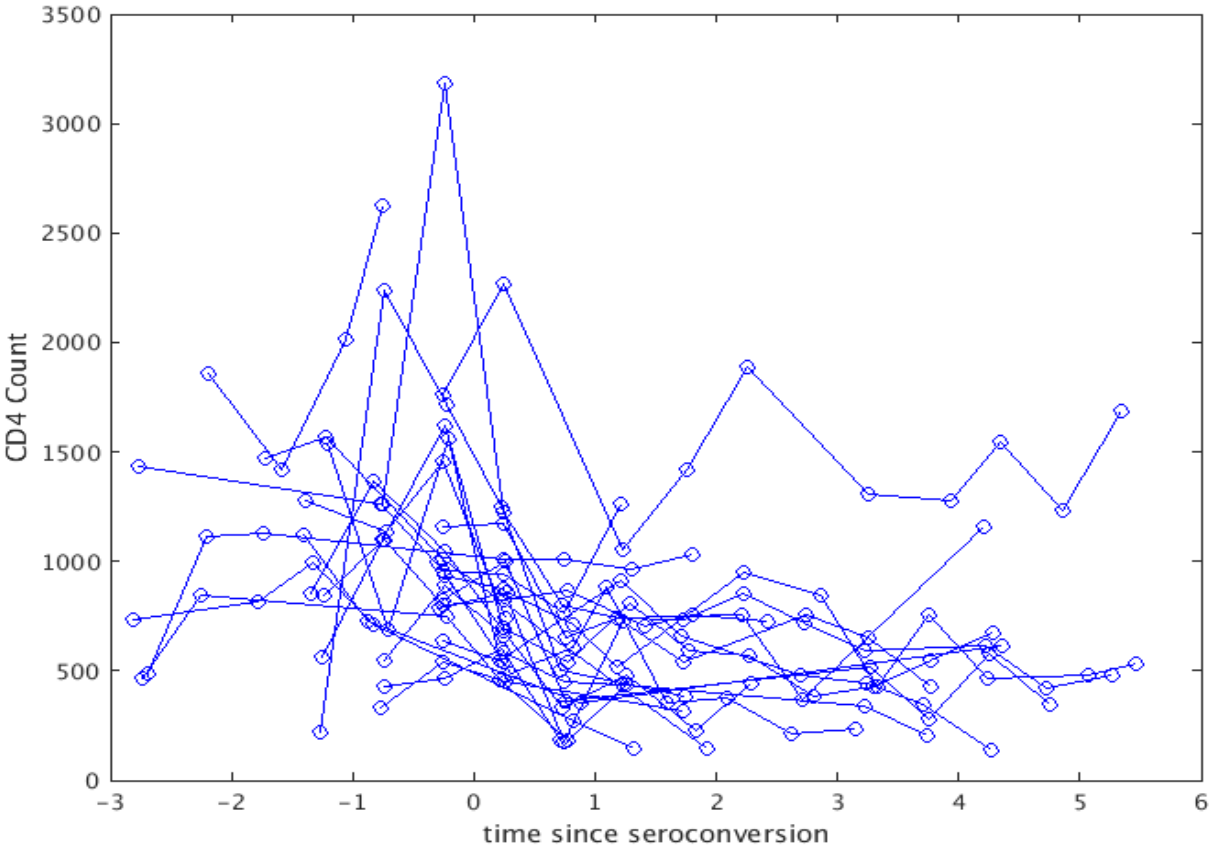
  $n_i = \#$ of repeated measurements for subject i,

  varies with subject.

  - An average of 6.44 measurement per subject.

- This results in longitudinal data with uneven # of measurements at irregular time-points.

CD4 Counts of First 25 Patients



Longitudinal vs Functional data

- Because of the intense recording, functional data have been customarily modeled with a nonparametric approach.

  - Often smoothness of the functions is assumed.

Longitudinal vs Functional data

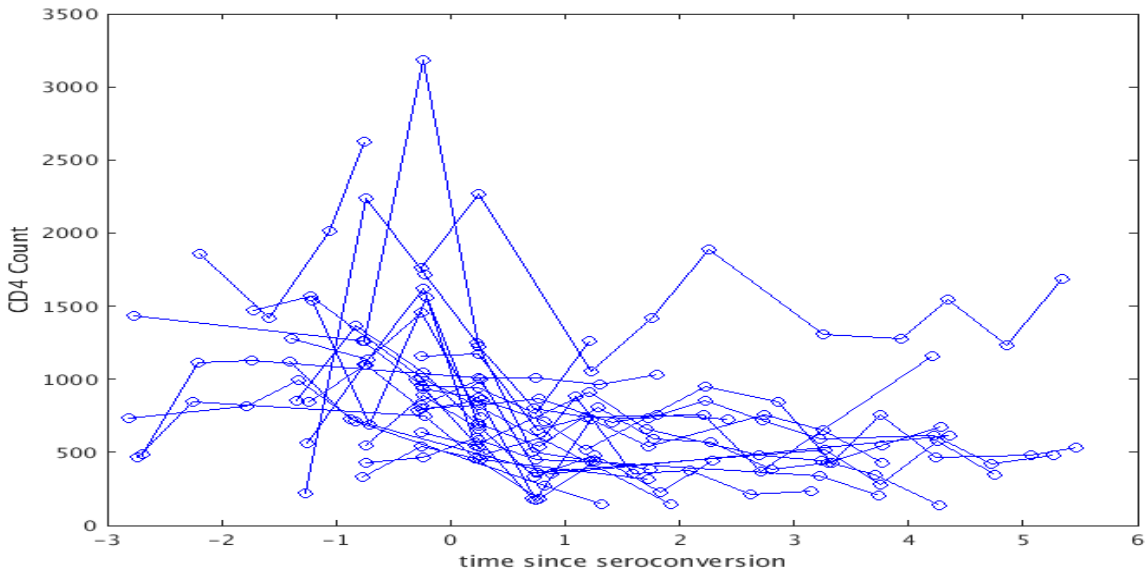
- Because of the intense recording, functional data have been customarily modeled with a nonparametric approach.

  - Often smoothness of the functions is assumed.

- Longitudinal data have traditionally been modeled by a parametric approach, such as a linear mixed-effects model.

  - However, it may not be easy to spot the pattern due to sparsity of and noise in the longitudinal data.

Longitudinal vs Functional data



- This motivates a data oriented nonparametric approach, which luckily is feasible under mild design conditions.

Longitudinal vs Functional data

- The analysis of functional data is termed "Functional Data Analysis" (FDA) by Ramsay (1982).

Longitudinal vs Functional data

- The analysis of functional data is termed "Functional Data Analysis" (FDA) by Ramsay (1982).

- Both longitudinal and functional data may be observed with noise (measurement errors).

  $\Longrightarrow$ the observed data for subject i might be

$$Y_{ij} = X_i(t_{ij}) + e_{ij}, \quad j = 1, \ldots, n_i,$$

  where $X_i(t)$ is a smooth random function,
  e_{ij} **are independent** $\forall\ i, j$.

Longitudinal vs Functional data

- The analysis of functional data is termed "Functional Data Analysis" (FDA) by Ramsay (1982).

- Both longitudinal and functional data may be observed with noise (measurement errors).

  $\Longrightarrow$ the observed data for subject i might be

$$Y_{ij} = X_i(t_{ij}) + e_{ij}, \quad j = 1, \ldots, n_i,$$

  where $X_i(t)$ is a smooth random function,
  e_{ij} **are independent** $\forall\ i, j$.

- A strength of the FDA approach is its ability to handle noise.

Summary

- There are the three types of functional data:

  (i) stochastic processes $\Longrightarrow \infty$-dimensional data
  (ii) intense/dense functional data $\Longrightarrow$ high dimensional data
  (iii) sparse functional/longitudinal data $\Longrightarrow$ irregular dim. data.

Summary

- There are the three types of functional data:

  (i) stochastic processes $\Longrightarrow \infty$-dimensional data
  (ii) intense/dense functional data $\Longrightarrow$ high dimensional data
  (iii) sparse functional/longitudinal data $\Longrightarrow$ irregular dim. data.

- The stochastic process $X_i(t)$ is assumed to be a continuous function.
  (But the observed data may contain noise, a.k.a. measurement error.)

Summary

- There are the three types of functional data:

  (i) stochastic processes $\Longrightarrow \infty$-dimensional data
  (ii) intense/dense functional data $\Longrightarrow$ high dimensional data
  (iii) sparse functional/longitudinal data $\Longrightarrow$ irregular dim. data.

- The stochastic process $X_i(t)$ is assumed to be a continuous function.
  (But the observed data may contain noise, a.k.a. measurement error.)

- Nonparametric approaches are typically employed to all three types of functional data.

Questions

- What are the pros and cons for each type of functional data?

  - Which one is the easiest or hardest to handle?
  (We'll discuss how to handle them in the next section.)

Questions

- What are the pros and cons for each type of functional data?

  - Which one is the easiest or hardest to handle?
  (We'll discuss how to handle them in the next section.)

  Answer: More is better!

Questions

- What are the pros and cons for each type of functional data?

  - Which one is the easiest or hardest to handle?
  (We'll discuss how to handle them in the next section.)

  Answer: More is better!

- How did we get away with the curse of high dimensionality?

Questions

- What are the pros and cons for each type of functional data?

  - Which one is the easiest or hardest to handle?
  (We'll discuss how to handle them in the next section.)

  Answer: More is better!

- How did we get away with the curse of high dimensionality?

  Answer: We borrow information by smoothing!

Questions

- What are the pros and cons for each type of functional data?

  - Which one is the easiest or hardest to handle?
  (We'll discuss how to handle them in the next section.)

  Answer: More is better!

- How did we get away with the curse of high dimensionality?

  Answer: We borrow information by smoothing!

- We can do so because we have a natural ordering of the data.

Questions

- What are the pros and cons for each type of functional data?

  - Which one is the easiest or hardest to handle?
  (We'll discuss how to handle them in the next section.)

  Answer: More is better!

- How did we get away with the curse of high dimensionality?

  Answer: We borrow information by smoothing!

- We can do so because we have a natural ordering of the data.

- Chen, Chen, M. and W. (2011) proposed a way to reorder multivariate data and convert high-dim data to functional data.

History of FDA

- The handling of longitudinal data as sparse functional data was the focus in Yao et al. (2005), but nonparametric approaches for longitudinal data had already been employed by Shi et al. (1996), Staniswalis and Lee (1998), James et al. (2000), and Rice and Wu (2001).

History of FDA

- The handling of longitudinal data as sparse functional data was the focus in Yao et al. (2005), but nonparametric approaches for longitudinal data had already been employed by Shi et al. (1996), Staniswalis and Lee (1998), James et al. (2000), and Rice and Wu (2001).

- First generation functional data typically consist of a random sample, $X_1(t), \ldots, X_n(t)$, of **independent real-valued functions**.

History of FDA

- The handling of longitudinal data as sparse functional data was the focus in Yao et al. (2005), but nonparametric approaches for longitudinal data had already been employed by Shi et al. (1996), Staniswalis and Lee (1998), James et al. (2000), and Rice and Wu (2001).

- First generation functional data typically consist of a random sample, $X_1(t), \ldots, X_n(t)$, of **independent real-valued functions**.

- Next generation functional data refer to functional data that are part of complex data objects,
possibly multivariate, correlated, or involve images and shapes.

History of FDA

- The handling of longitudinal data as sparse functional data was the focus in Yao et al. (2005), but nonparametric approaches for longitudinal data had already been employed by Shi et al. (1996), Staniswalis and Lee (1998), James et al. (2000), and Rice and Wu (2001).

- First generation functional data typically consist of a random sample, $X_1(t), \ldots, X_n(t)$, of **independent real-valued functions**.

- Next generation functional data refer to functional data that are part of complex data objects,
possibly multivariate, correlated, or involve images and shapes.

- Brain and neuroimaging data are examples of next generation functional data.



Outline

1 Introduction to Functional Data Analysis

2 Mean and Covariance Estimation

3 Theory: Mean and Covariance Estimation

Mean and Covariance Function

Data $\{X_1, \ldots, X_n\}$ are i.i.d. copies of a random function $X(t)$:

- Mean function: $\mu(t) = E(X(t))$
- Covariance function: $\Sigma(s,t) = \text{cov}(X(s), X(t))$, where s & $t \in$ interval I.

Mean and Covariance Function

Data $\{X_1, \ldots, X_n\}$ are i.i.d. copies of a random function $X(t)$:

- Mean function: $\mu(t) = E(X(t))$
- Covariance function: $\Sigma(s,t) = \text{cov}(X(s), X(t))$, where s & $t \in$ interval I.
- Regular functional data - All subjects are measured at the same time grid, $t_1, \ldots, t_m$, often equally spaced $\Longrightarrow$ multivariate data.

  Irregular functional data - The measurement schedule for subject i is $t_{i1}, \ldots, t_{in_i} \Longrightarrow$ longitudinal data.

Estimation of Mean and Covariance Functions: Dense and Regular Functional Data

- For regular or fully observed functional data, the cross-sectional (sample) mean at the observed time t can be used to estimate the mean function.

Estimation of Mean and Covariance Functions: Dense and Regular Functional Data

- For regular or fully observed functional data, the cross-sectional (sample) mean at the observed time t can be used to estimate the mean function.

  - Such a mean estimate provides a $\sqrt{n}$-consistent (pointwise) estimate of the mean $\mu(t)$ even in the presence of measurement errors. **(WHY?)**

Estimation of Mean and Covariance Functions: Dense and Regular Functional Data

- For regular or fully observed functional data, the cross-sectional (sample) mean at the observed time t can be used to estimate the mean function.

  - Such a mean estimate provides a $\sqrt{n}$-consistent (pointwise) estimate of the mean $\mu(t)$ even in the presence of measurement errors. **(WHY?)**

- The cross-section mean can further be smoothed slightly to obtain a smooth mean estimate (This requires a dense/intense measurement schedule).

  - **Think how you should smooth to retain the $\sqrt{n}$-consistent of the smooth estimate** .

Estimation of Mean and Covariance Functions: Dense and Regular Functional Data

- Likewise, the sample covariance matrix is also $\sqrt{n}$-consistent and can be slightly smoothed to retain the $\sqrt{n}$-rate of consistency.

Estimation of Mean and Covariance Functions: Irregular Functional Data

- The cross-sectional approaches no longer work for irregular data.

Estimation of Mean and Covariance Functions: Irregular Functional Data

- The cross-sectional approaches no longer work for irregular data.

  - Smoothing is needed to estimate the mean function.

Estimation of Mean and Covariance Functions: Irregular Functional Data

- The cross-sectional approaches no longer work for irregular data.

  - Smoothing is needed to estimate the mean function.

  - The mean estimate could be $\sqrt{n}$-consistent (pointwise) for "**dense**" (or intense) functional data but will have nonparametric rates otherwise.
  ("Dense" functional data $\Longleftrightarrow$ $\sqrt{n}$-rate of convergence is feasible.)

Estimation of Mean and Covariance Functions: Irregular Functional Data

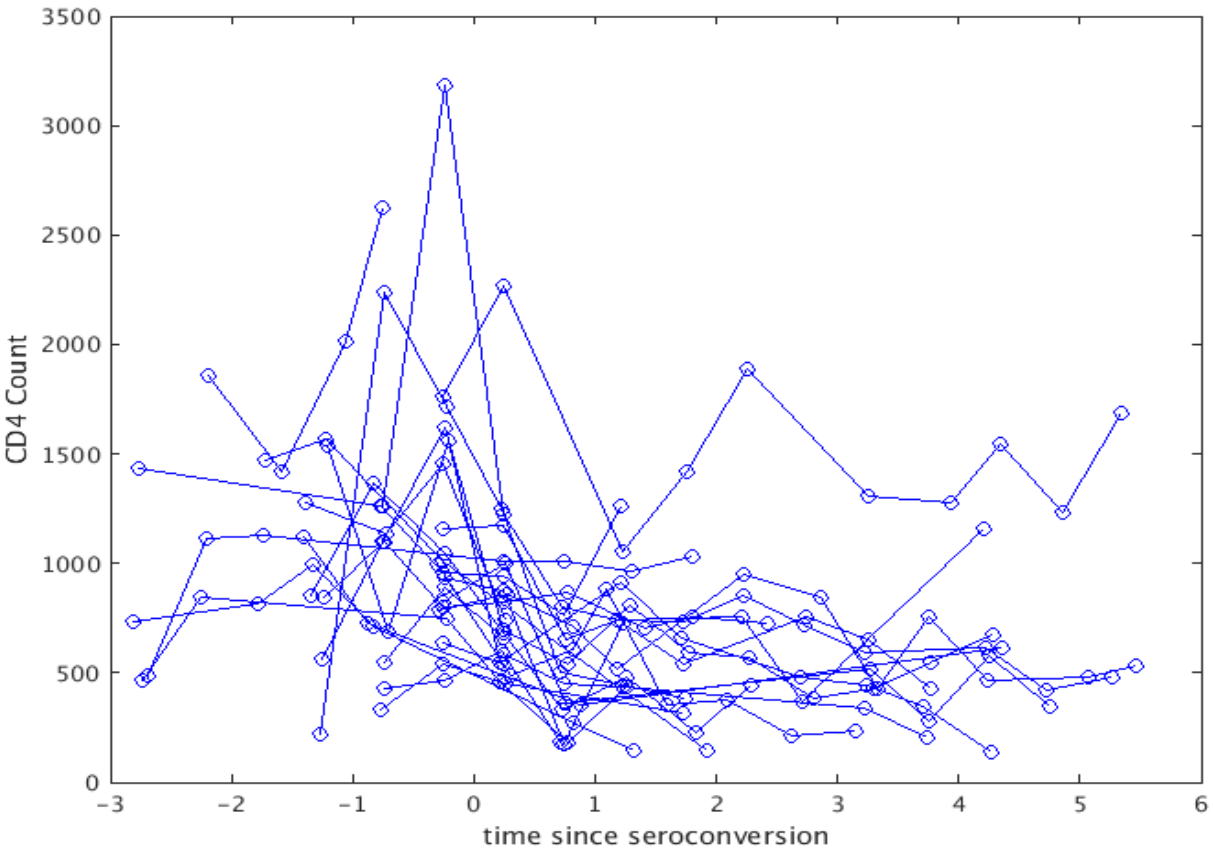
- The cross-sectional approaches no longer work for irregular data.

  - Smoothing is needed to estimate the mean function.

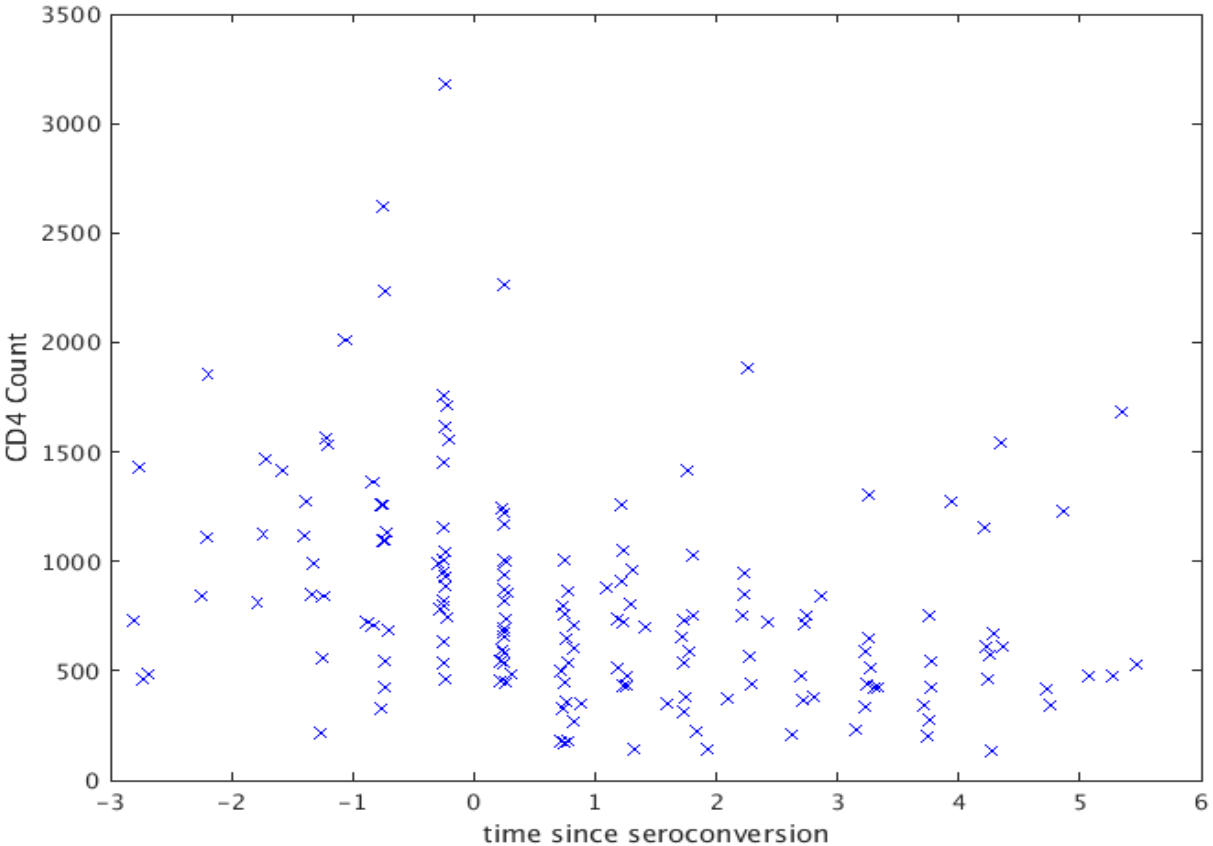
  - The mean estimate could be $\sqrt{n}$-consistent (pointwise) for "**dense**" (or intense) functional data but will have nonparametric rates otherwise.
  ("Dense" functional data $\Longleftrightarrow$ $\sqrt{n}$-rate of convergence is feasible.)

- Likewise, 2D smoothing is needed to estimate the covariance function. and $\sqrt{n}$-consistency can be achieved for **dense** functional data.

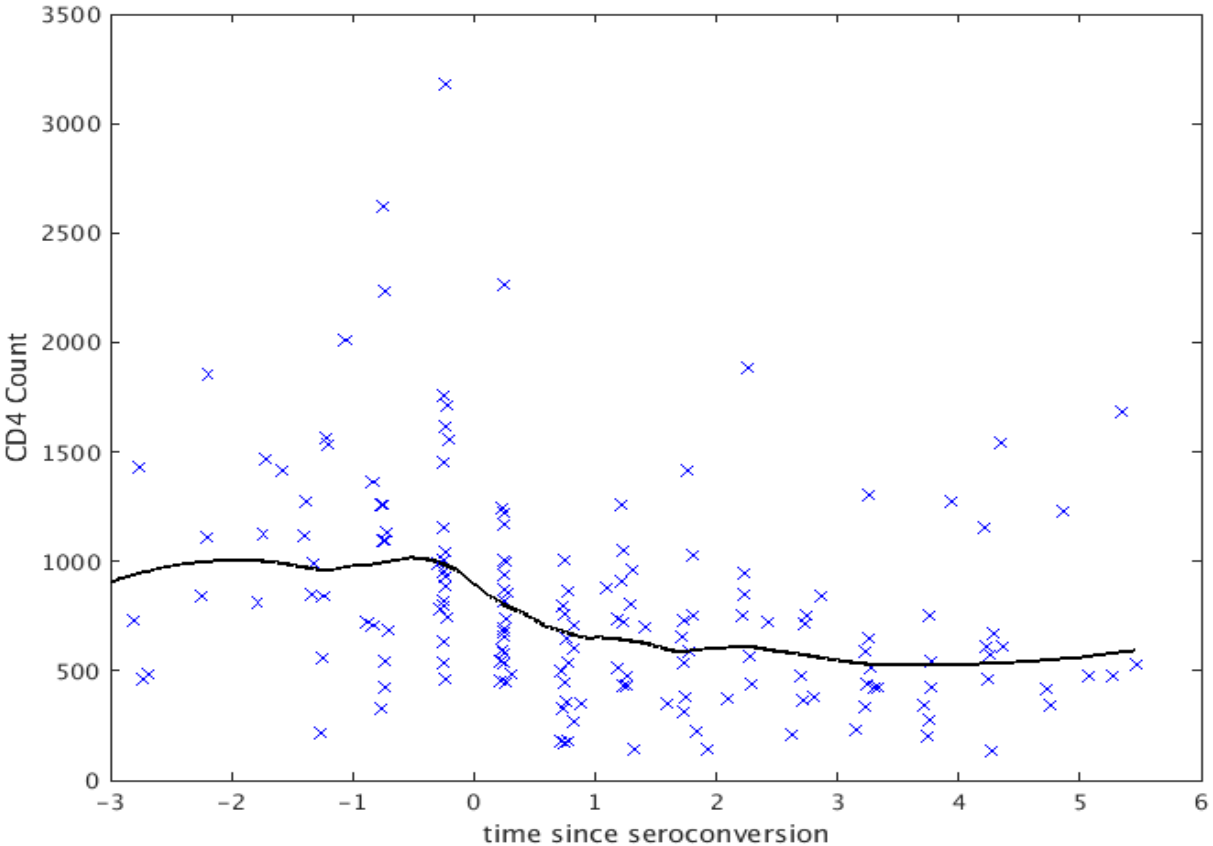
CD4 Counts of First 25 Patients



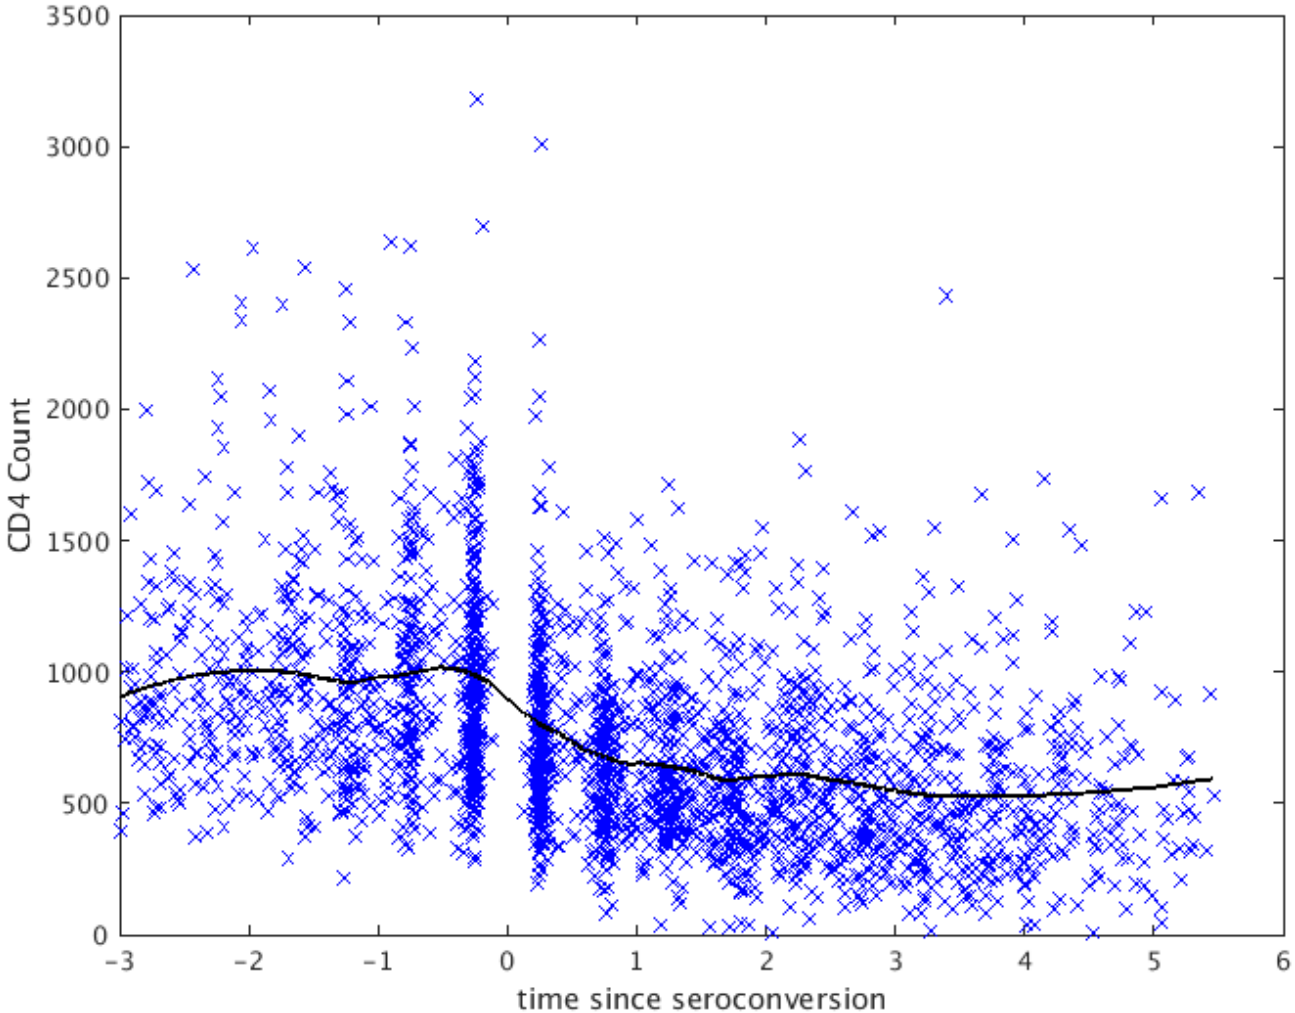
CD4 Counts of first 25 Patients



Mean Curve: CD4 Counts of first 25 Patients



Mean Curve: CD4 counts of all patients



Estimation of Mean Function

If we employ the local linear smoother, the estimate for the mean function is:

$$\hat{\mu}(t) = \hat{\beta}_0, \quad \text{where}$$

$$(\hat{\beta}_0, \hat{\beta}_1) = \underset{\beta_0, \beta_1}{\operatorname{argmin}} \sum_{i=1}^{n} \sum_{j=1}^{n_i} \Big[Y_{ij} - \beta_0 - \beta_1 (T_{ij} - t) \Big]^2 K_{h_\mu}(T_{ij} - t).$$

Remarks

- In addition to the local linear (or polynomial) smoother, any smoothing method, such as penalized splines, B-splines, Wavelets, and Fourier filtering, can be employed.

Remarks

- In addition to the local linear (or polynomial) smoother, any smoothing method, such as penalized splines, B-splines, Wavelets, and Fourier filtering, can be employed.

- These smoothing methods (scatter plot smoothers) can also be applied to dense data, whether regular or not, so a unified approach is feasible.

Remarks

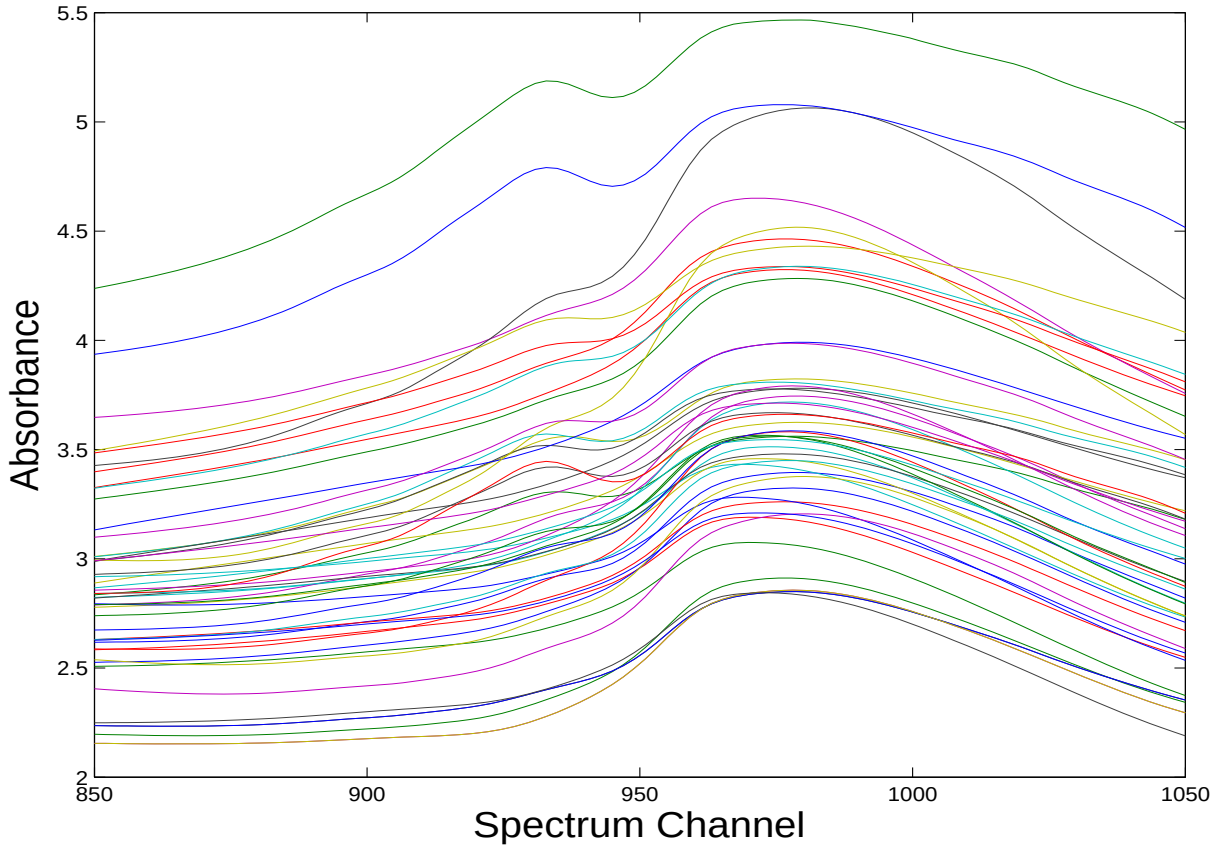
- In addition to the local linear (or polynomial) smoother, any smoothing method, such as penalized splines, B-splines, Wavelets, and Fourier filtering, can be employed.

- These smoothing methods (scatter plot smoothers) can also be applied to dense data, whether regular or not, so a unified approach is feasible.

- Another common approach for dense data is to pre-smooth the data from each subject separately, then take the cross-section mean at each time point t.

Remarks

This is the case for the spectrum data.



Remarks

- Such a pre-smoothing typically aims at:

  (i) getting rid of the noise in the data,
  (ii) revealing the latent smooth curve for each subject.

  Hall, M. W. (2006) and Zhang and Chen (2007).

Remarks

- Such a pre-smoothing typically aims at:

  (i) getting rid of the noise in the data,
  (ii) revealing the latent smooth curve for each subject.

  Hall, M. W. (2006) and Zhang and Chen (2007).

- However, both depend on the amount of smoothing and design of the dense data, so whether the mission has been accomplished is unclear.

- Because of this uncertainty and because pre-smoothing alters the data, we prefer not to adopt a pre-smoothing approach.



Alaska

Estimation of Covariance Function

- Our target is $\Sigma(s,t) = \text{cov}(X(s), X(t))$, but we do not observe X.

Estimation of Covariance Function

- Our target is $\Sigma(s,t) = \text{cov}(X(s), X(t))$, but we do not observe X.
- Observe $Y_{ij} = Y_i(t_{ij}) = X_{ij} + e_{ij}, \quad \text{where var}(e_{ij}) = \sigma^2(t_{ij})$.

  $\Longrightarrow \text{cov}(Y(s), Y(t)) = \text{cov}(X(s), X(t))$, when $s \neq t$,

  $$\text{var}(Y(t)) = \text{cov}(Y(t), Y(t)) = \text{cov}(X(t), X(t)) + \sigma^2(t).$$

Estimation of Covariance Function

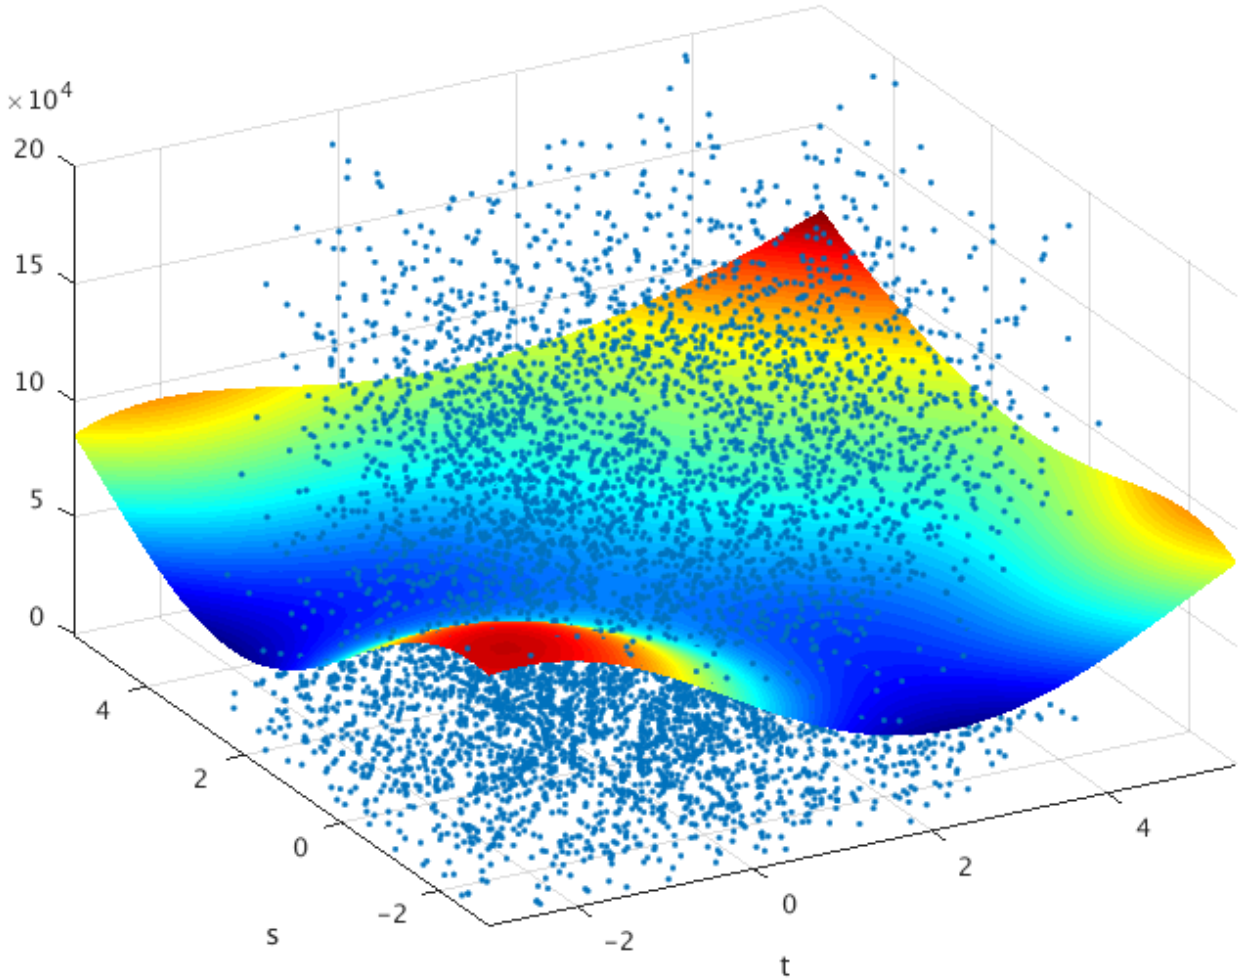
- Our target is $\Sigma(s,t) = \text{cov}(X(s), X(t))$, but we do not observe X.
- Observe $Y_{ij} = Y_i(t_{ij}) = X_{ij} + e_{ij}$, where $\text{var}(e_{ij}) = \sigma^2(t_{ij})$.

  $\Longrightarrow \text{cov}(Y(s), Y(t)) = \text{cov}(X(s), X(t))$, when $s \neq t$,

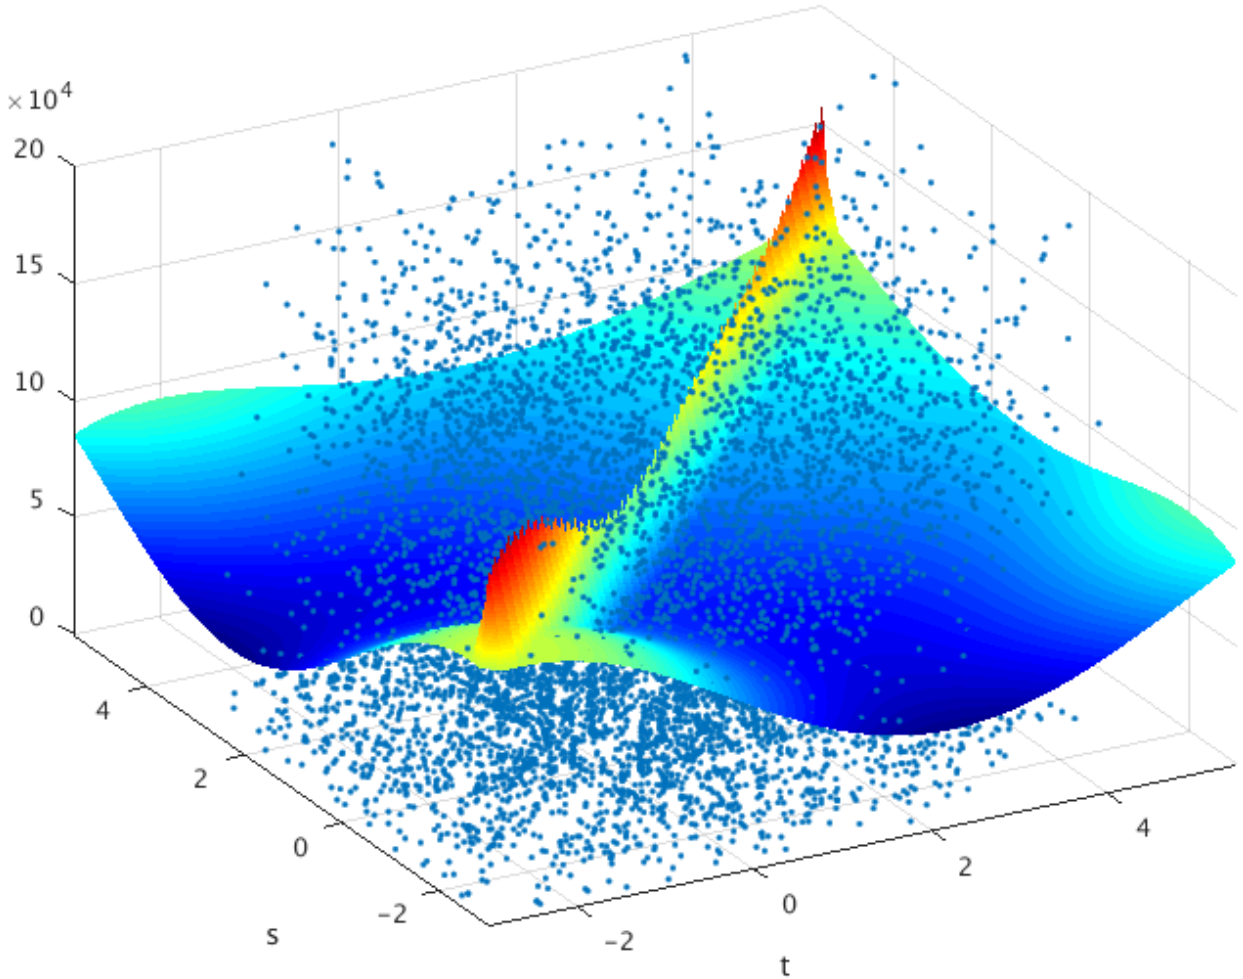
  $$\text{var}(Y(t)) = \text{cov}(Y(t), Y(t)) = \text{cov}(X(t), X(t)) + \sigma^2(t).$$

- This means we need to handle the diagonal part of $\text{cov}(Y)$, which corresponds to the variance of Y, differently from the rest of $\text{cov}(Y)$.

Covariance Surface of $X(t)$

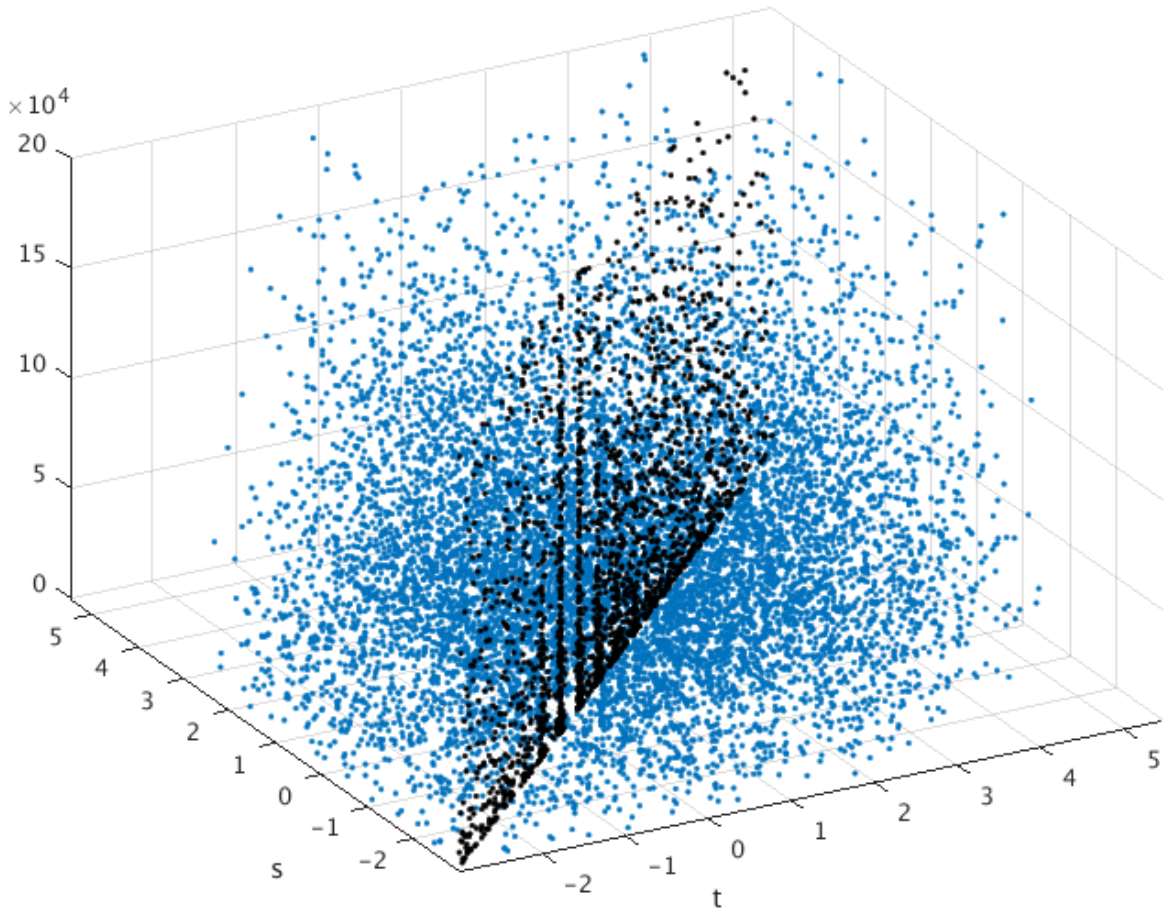


Covariance of $Y(t)$



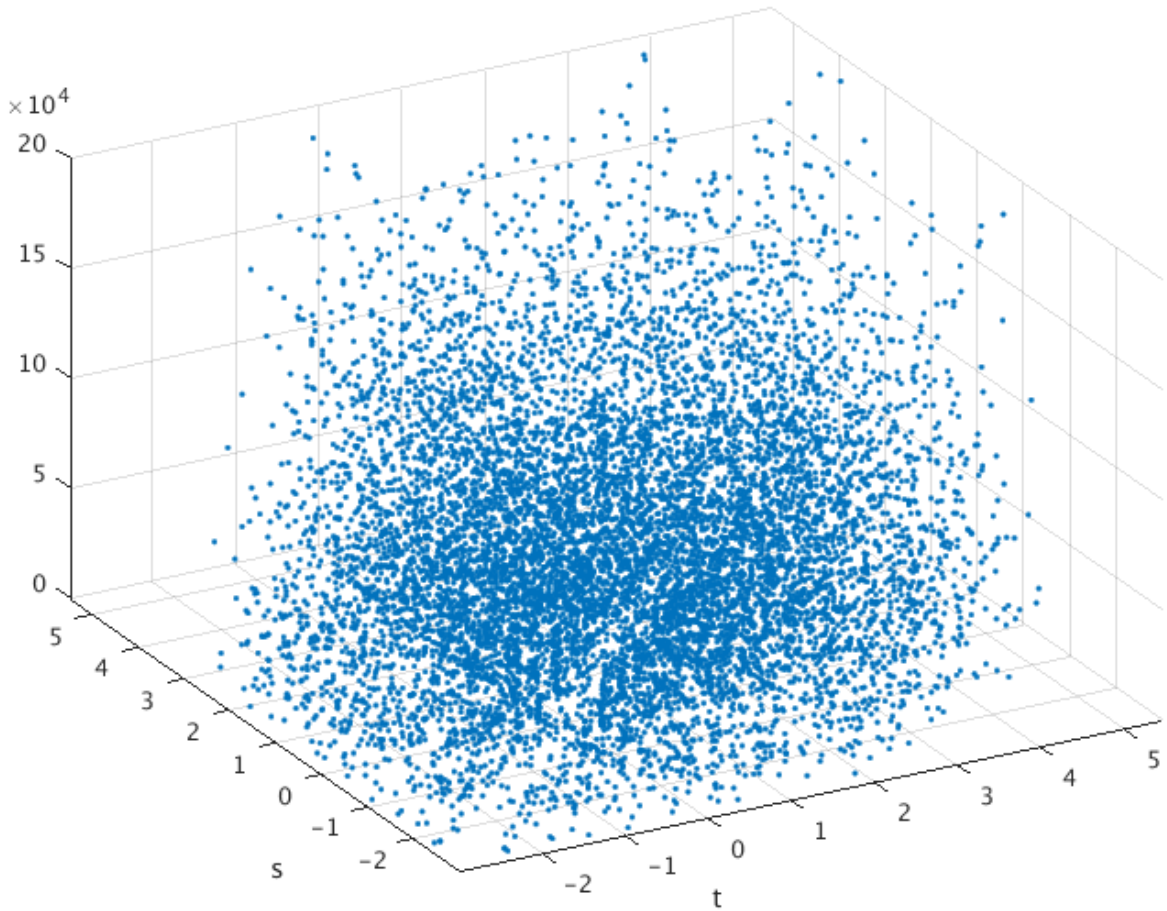
Raw Covariance Plot: Diagonal Data in Black

$$[Y(t_{ij}) - \hat{\mu}(t_{ij})][Y(t_{ik}) - \hat{\mu}(t_{ik})], \ \forall \ i, j, k$$

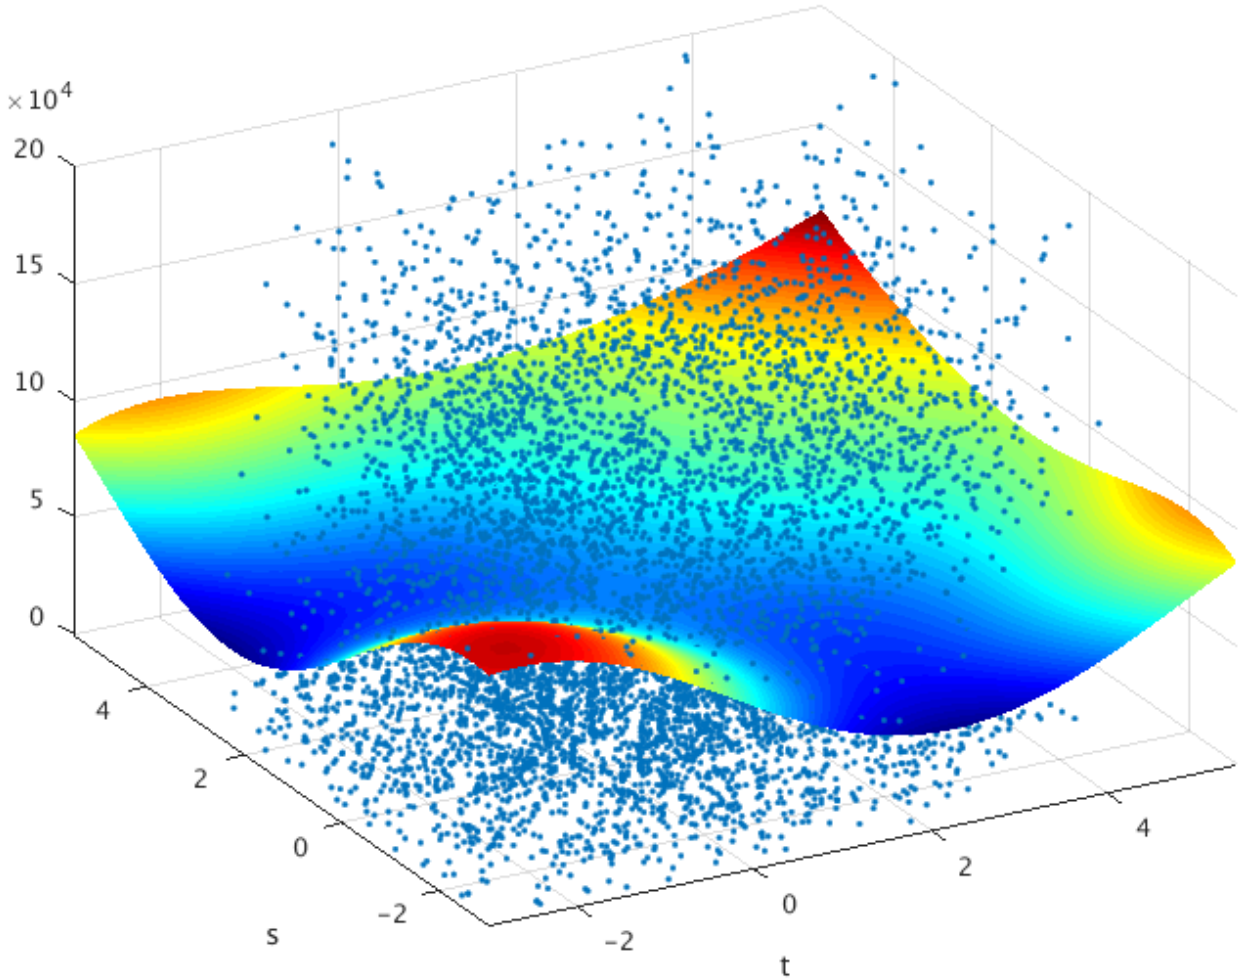


Raw Covariance with Diagonal Data Removed

$$[Y(t_{ij}) - \hat{\mu}(t_{ij)}][Y(t_{ik}) - \hat{\mu}(t_{ik})], \ \forall \ i, j \neq k$$

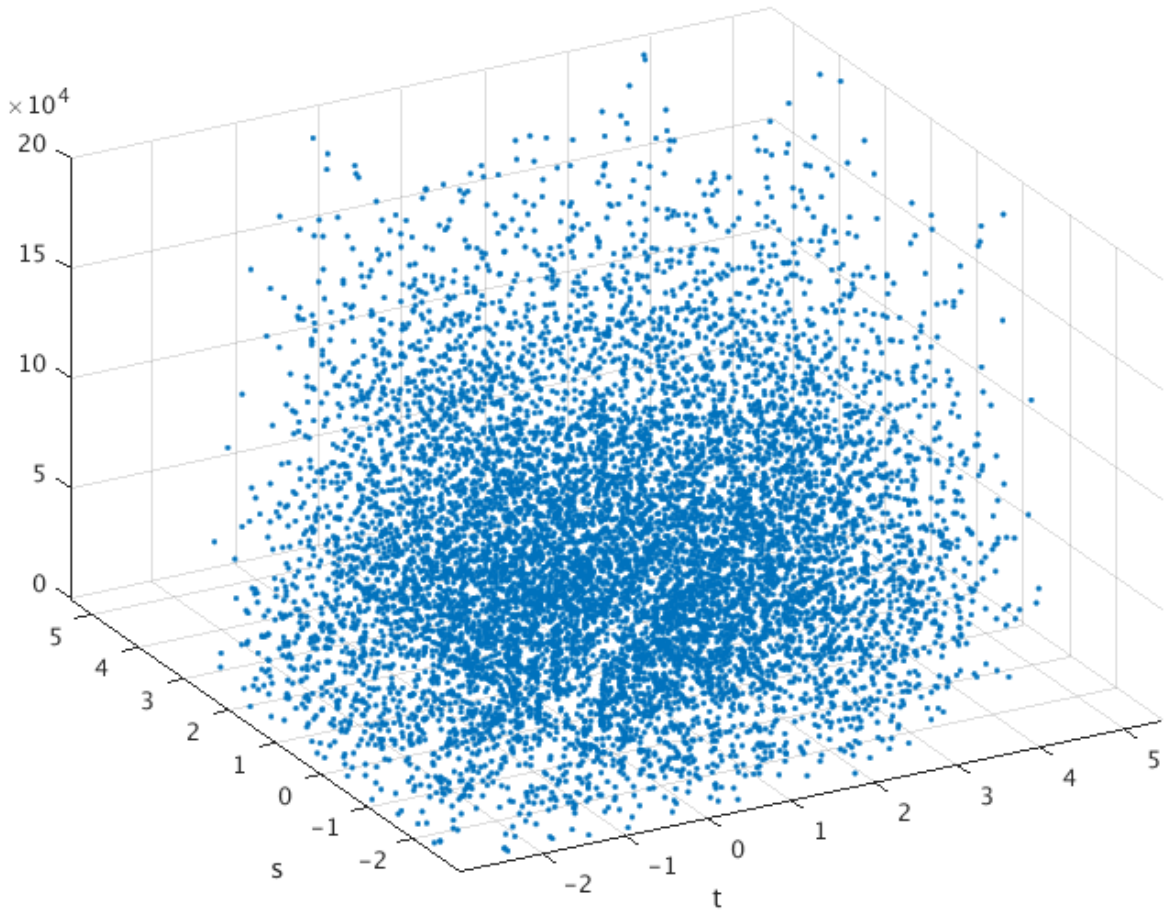


2D Smoother to Estimate the Covariance of $X(t)$

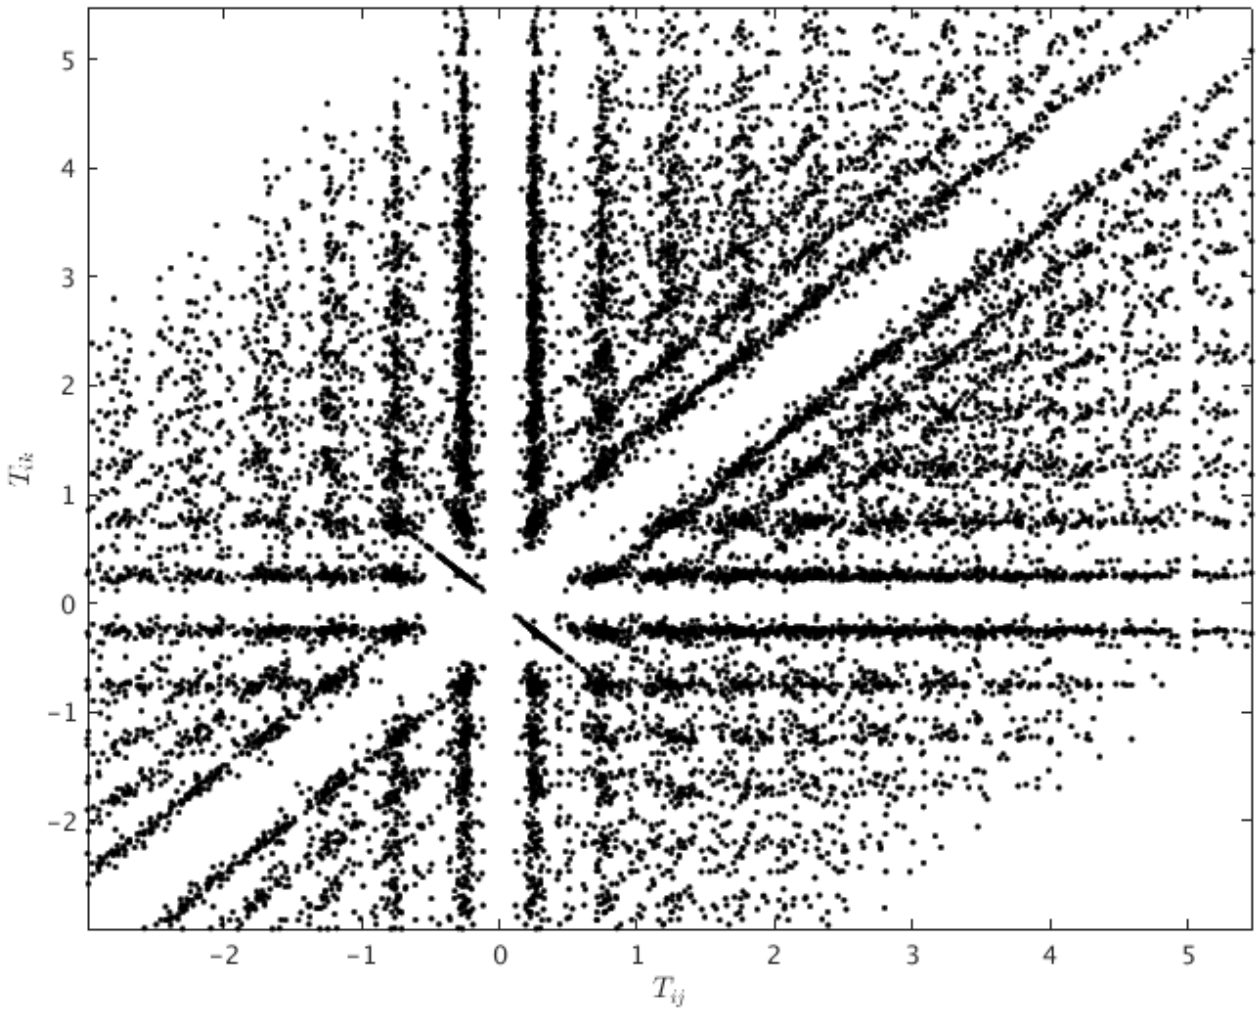


Raw Covariance with Diagonal Data Removed

$$[Y(t_{ij}) - \hat{\mu}(t_{ij)}][Y(t_{ik}) - \hat{\mu}(t_{ik})], \ \forall \ i, j \neq k$$



Design Plot for Covariance: $(t_{ij}, t_{ik}),\ \forall\ i, j \neq k$



Estimated Covariance Surface of $X(t)$

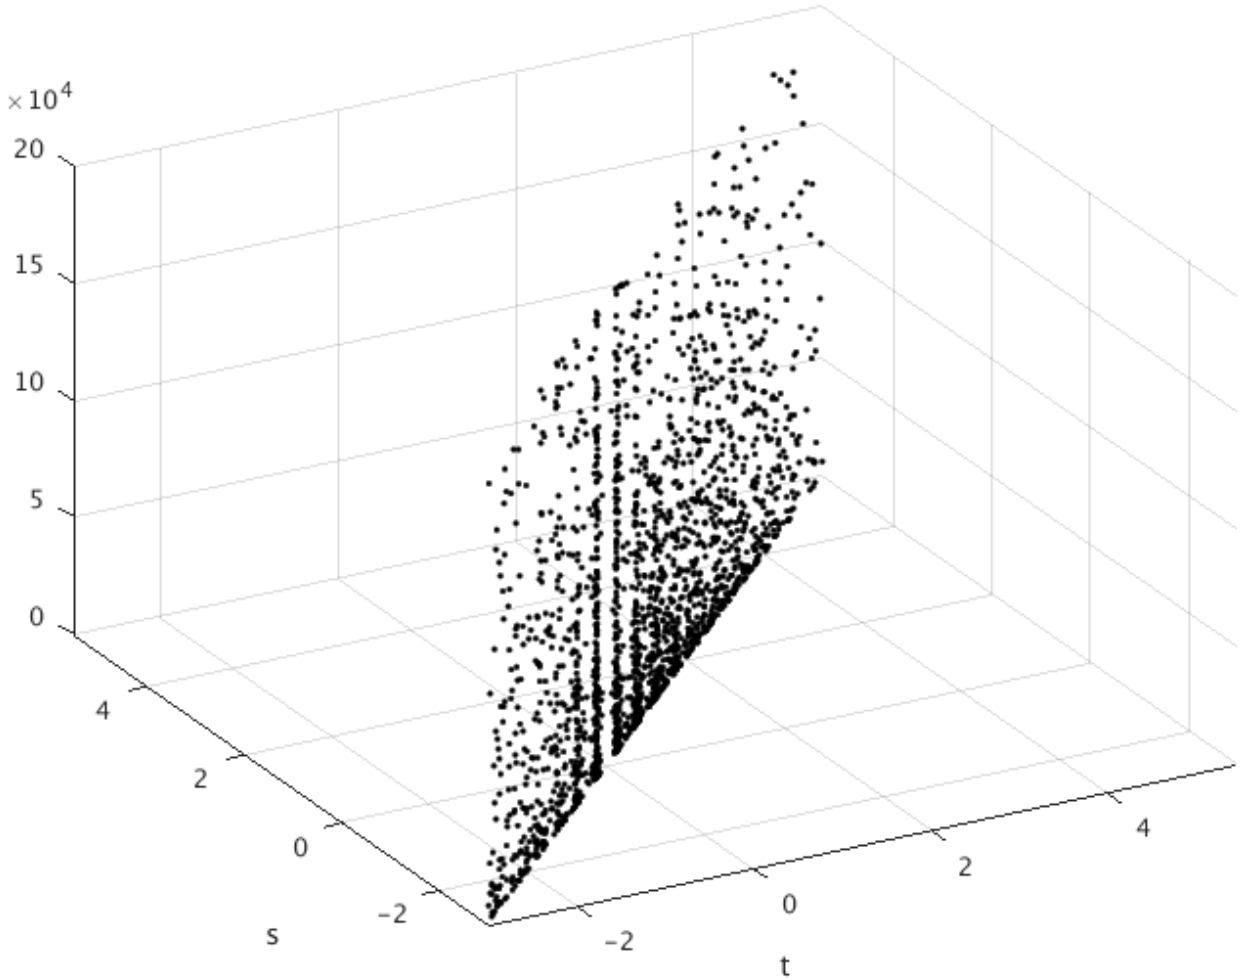
Let $D_{ijl} = (Y_{ij} - \hat{\mu}(t_{ij}))(Y_{il} - \hat{\mu}(t_{il}))$ be the raw covariances.

Employing a local linear smoother, the estimate for $\Sigma(s,t)$ is:

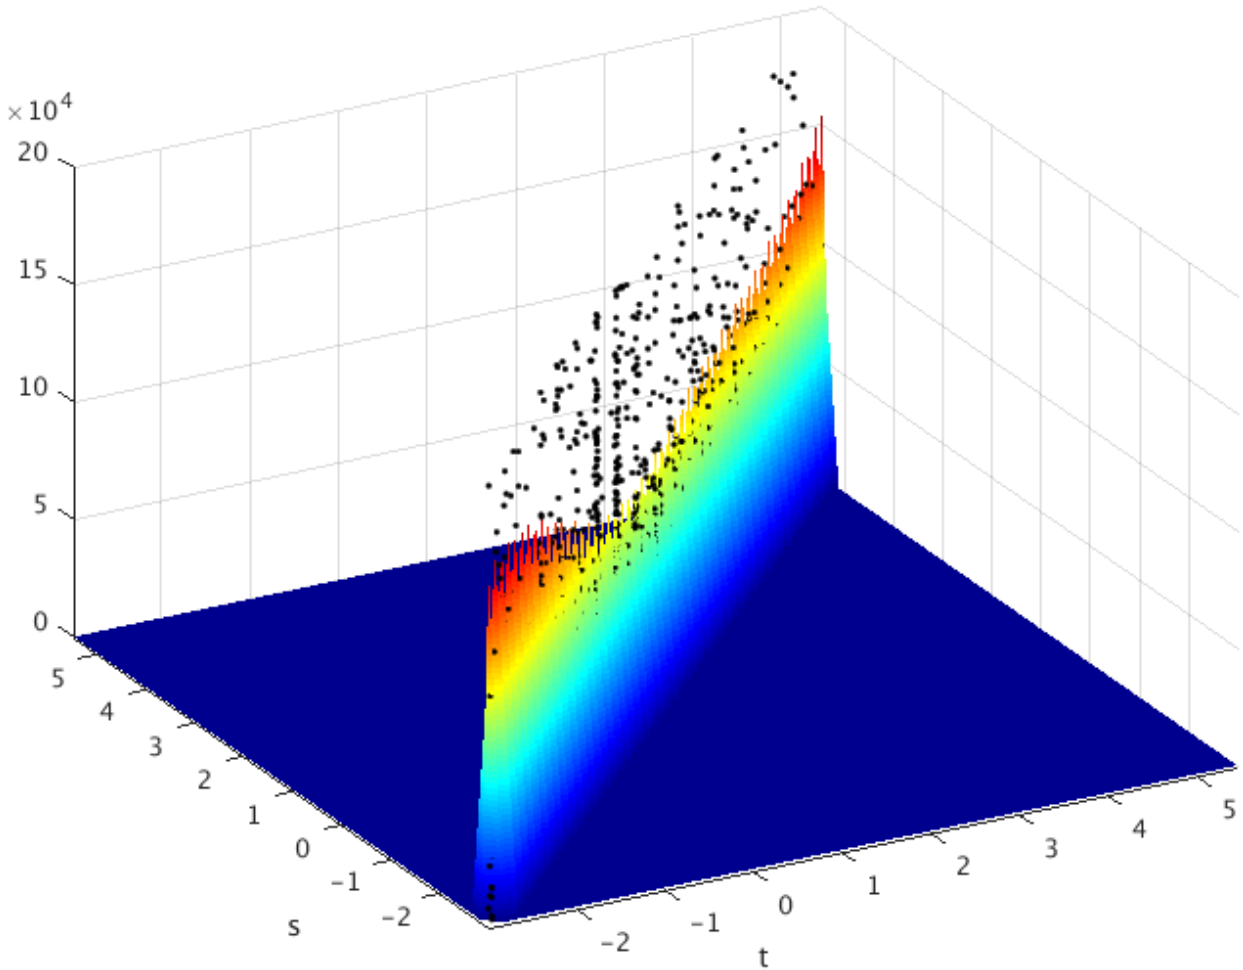
$\hat{\Sigma}(s,t) = \hat{\beta}_0$, where

$$(\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2) = \underset{\beta_0,\beta_1,\beta_2}{\operatorname{argmin}} \sum_{i=1}^{n} \sum_{1 \le j \ne l \le n_i} \Big[D_{ijl} - \beta_0 - \beta_1(t_{ij} - s) - \beta_2(t_{il} - t) \Big]^2 K_{h_\Sigma}(t_{ij} - s) K_{h_\Sigma}(t_{il} - t).$$

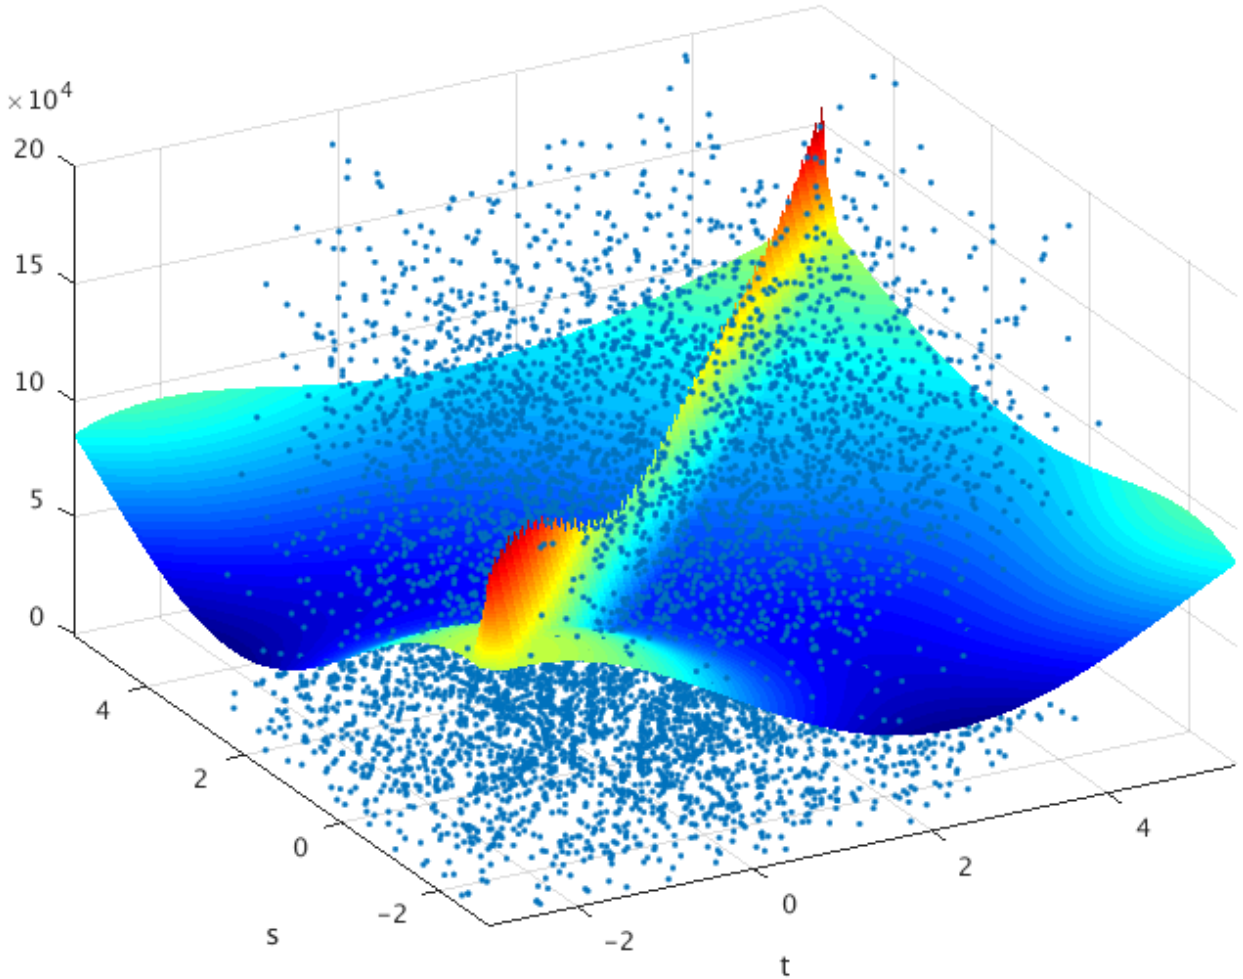
Raw Variance Plot:$[Y(t_{ij}) - \mu(t_{ij})]^2,\ \forall\ i, j$



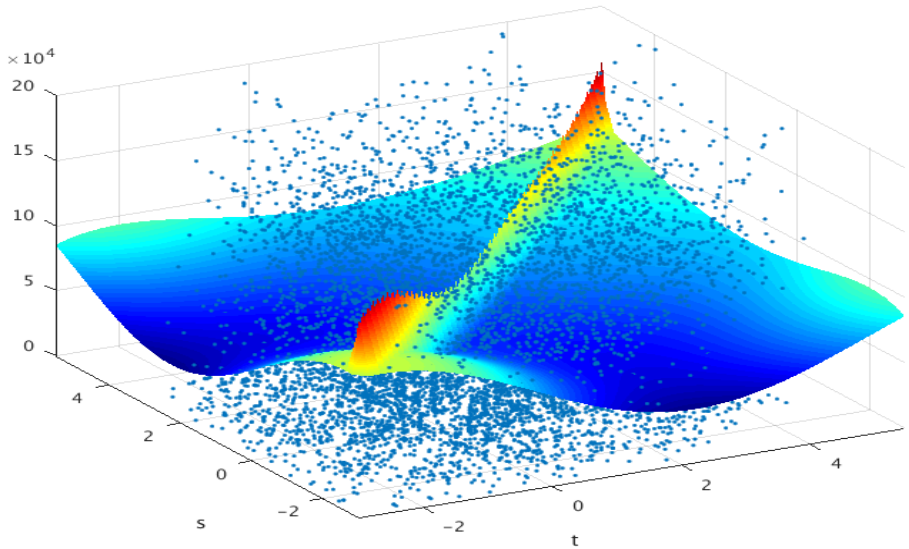
Estimated Variance function of $Y(t)$



Estimated Covariance & Variance of $Y(t)$

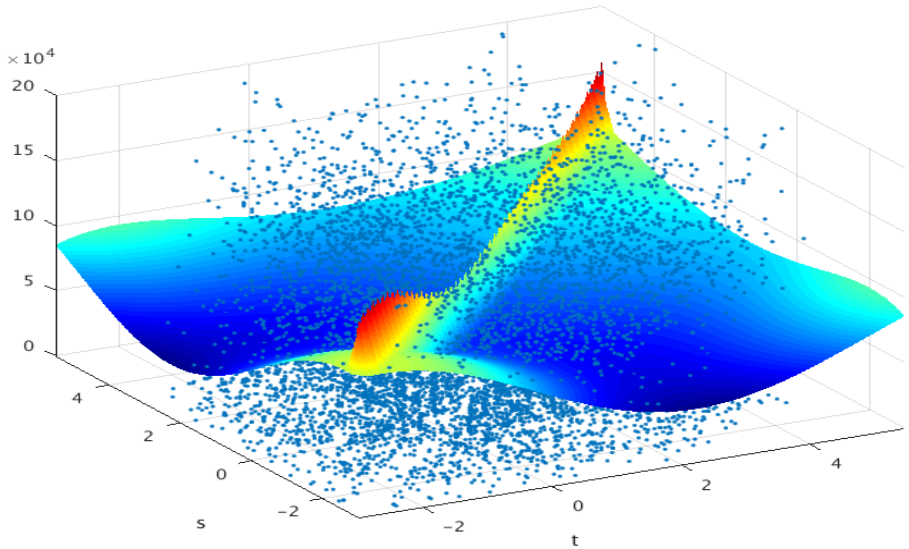


Estimates of $\sigma^2(t)$: Variance of Measurement Errors



$$\sigma^2(t) = \mathsf{var}(Y(t)) - \mathsf{var}(X(t)).$$

Estimates of $\sigma^2(t)$: Variance of Measurement Errors

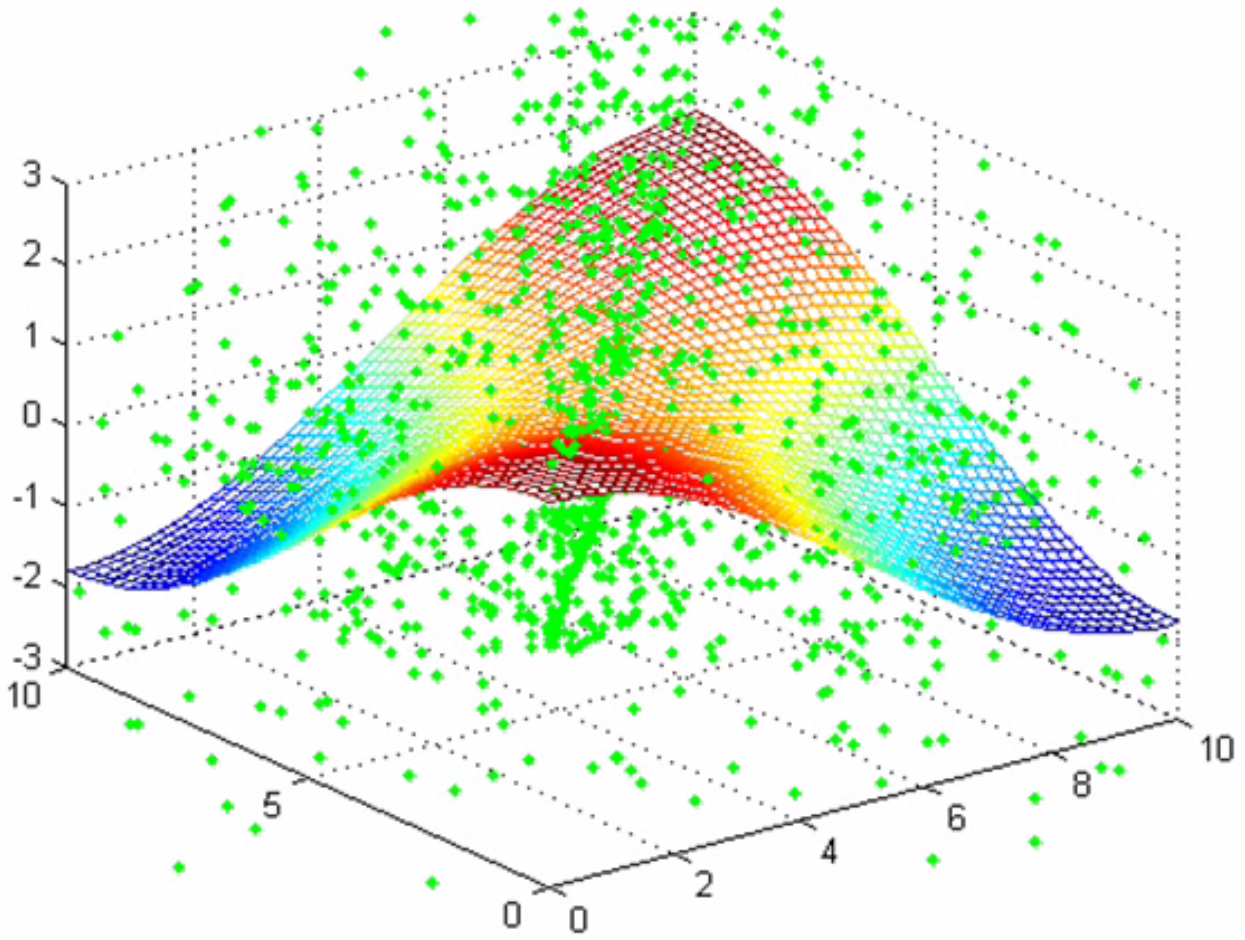


$\sigma^2(t) = \text{var}(Y(t)) - \text{var}(X(t))$.

When $\sigma^2(t) = \sigma^2$ for all t, one can estimate σ^2 by $\int_I \hat{\sigma}^2(t)dt$.

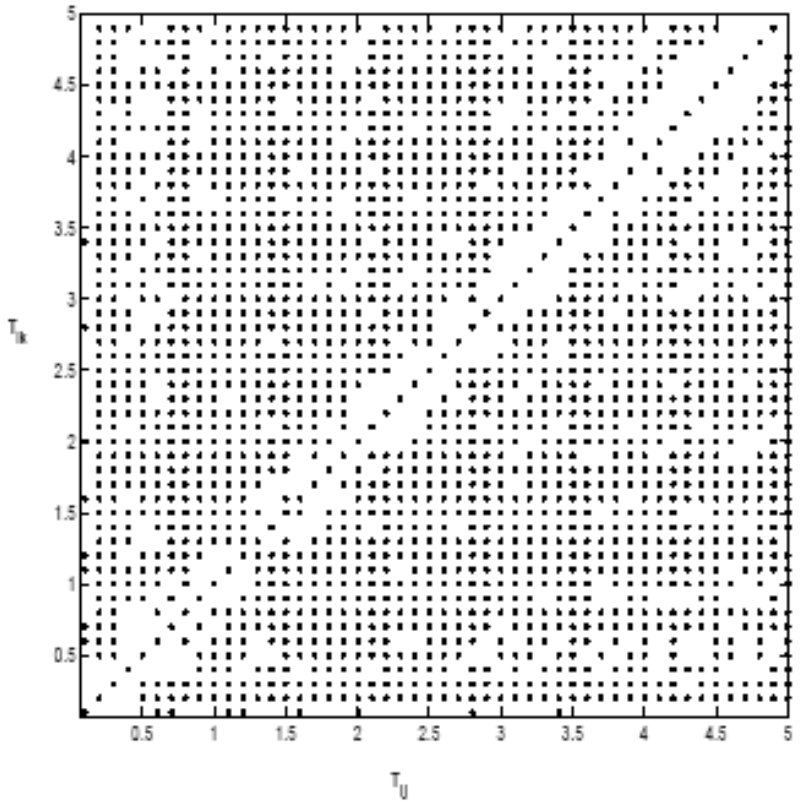
Due to boundary effects, PACE replaces I by a sub-interval.

Design Check



Raw Covariance Plot:

$[Y(t_ij) - \mu(t_ij)][Y(t_ik) - \mu(t_ik)], \forall j, k$



$Y(t) = X(t) + e(t)$
$Cov(Y(s), Y(t)) =$
$Cov(X(s), X(t))$, if $s \neq t$, but
$Var(Y(t)) = Var(X(t)) + \sigma^2$



Sodona, Arizona

Functional Principal Component Analysis: FPCA

- Since functional data is intrinsically $\infty-$dimensional, dimension reduction is needed to implement it.

- The most popular approach is via principal component analysis, termed FPCA.
  - This accounts for a large fraction of the FDA literature.

Functional Principal Component Analysis: FPCA

- Since functional data is intrinsically $\infty-$dimensional, dimension reduction is needed to implement it.

- The most popular approach is via principal component analysis, termed FPCA.
  - This accounts for a large fraction of the FDA literature.

- Once we have an estimate of the Covariance function, we can proceed to perform FPCA on a discrete time grid,
  i.e. perform multivariate eigen-analysis of the estimated covariance surface on the (intense) discrete time grid.

  Hall and Housseini-Nasab (2006) -Functional Data
  Hall, Müller and W. (2006) - Longitudinal Data
  Zhou , Wei and Yao (2025)

Outline

1 Introduction to Functional Data Analysis

2 Mean and Covariance Estimation

3 Theory: Mean and Covariance Estimation

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- So far, we have used "**scatter plot smoothers**" to estimate the mean and covariance functions.

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- So far, we have used "**scatter plot smoothers**" to estimate the mean and covariance functions.

  - This is the method of PACE (a Matlab package for functional data), first proposed in Yao, Müller and W. (2005).
  (An R package, fdapace, has been released recently.)

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- So far, we have used "**scatter plot smoothers**" to estimate the mean and covariance functions.

  - This is the method of PACE (a Matlab package for functional data), first proposed in Yao, Müller and W. (2005).
  (An R package, fdapace, has been released recently.)

  - It works for any type of sampling plan: regular, dense, or not.

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- So far, we have used "**scatter plot smoothers**" to estimate the mean and covariance functions.

  - This is the method of PACE (a Matlab package for functional data), first proposed in Yao, Müller and W. (2005).
  (An R package, fdapace, has been released recently.)

  - It works for any type of sampling plan: regular, dense, or not.

  - It assigns the same weight to each observation, so a subject with more measurements (large n_i) receives a larger total weight.

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- So far, we have used "**scatter plot smoothers**" to estimate the mean and covariance functions.

  - This is the method of PACE (a Matlab package for functional data), first proposed in Yao, Müller and W. (2005).
  (An R package, fdapace, has been released recently.)

  - It works for any type of sampling plan: regular, dense, or not.

  - It assigns the same weight to each observation, so a subject with more measurements (large n_i) receives a larger total weight.

- Another way to assign weights is to give equal weights, $\frac{1}{n}$, to all subjects as proposed in Li and Hsing (2010).

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- **Which weight assignment is better?**

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- **Which weight assignment is better?**

- Zhang and W. (2016) presented asymptotic properties of these estimators and showed that the equal weight per subject scheme is better for dense design but the equal weight per observation scheme is better for sparse data.

Theory: Mean and Covariance Estimation (Zhang and W., 2016)

- **Which weight assignment is better?**
- Zhang and W. (2016) presented asymptotic properties of these estimators and showed that the equal weight per subject scheme is better for dense design but the equal weight per observation scheme is better for sparse data.
- In a follow-up work Zhang and W. (2018), they further provided an optimal weighing scheme for mean and covariance estimation.

References for Mean and Covariance Estimation

- Yao Müeller and W. (2005, JASA)
- Li and Hsing (2010, AoS)
- Zhang and W. (2016, Ann. Stat.)

  Zhang and W. (2018, Statistics and Probability Letters)

- **Additional References:**

  Stainiswalis and Lee (1998, JASA),

  Yao (2005, J. Mult. Analy.),

  Review paper: Wang, Chiou and Müller (2016, Annual Review of Statistics and Its Applicaitons).

End of Introduction to Functional Data Analysis



Moraine Lake, Banff National Park