The Completion of Covariance Kernels

Victor M. Panaretos

(joint with Kartik Waghmare)

victor.panaretos@epfl.ch

EPFL

Based on joint work[1] with Kartik Waghmare

[1] Waghmare & Panaretos (2021). *The Completion of Covariance Kernels.* `arXiv:2107.07350`

In a nutshell:

- Let Ω be a closed, connected, and symmetric union (possibly uncountable) of subsquares in $[0,1]^2$ covering the diagonal:

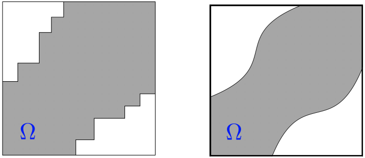

- Let $K_\Omega(s,t) : \Omega \to \mathbb{R}$ be a partial covariance kernel[2] on Ω.

[2] i.e. $\forall\ I \times I \subset \Omega$, the restriction $K_\Omega|_{I\times I}$ is a covariance kernel

In a nutshell:

- Let Ω be a closed, connected, and symmetric union (possibly uncountable) of subsquares in $[0,1]^2$ covering the diagonal:

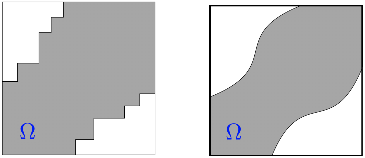

- Let $K_\Omega(s,t) : \Omega \to \mathbb{R}$ be a partial covariance kernel[2] on Ω.

We consider the following problem:

How can $K_\Omega(s,t)$ be completed to a covariance kernel $K(s,t)$ on $[0,1]^2$?

[2] i.e. $\forall\ I \times I \subset \Omega$, the restriction $K_\Omega|_{I\times I}$ is a covariance kernel

- Do there always exist completions? How many?
- Is there canonical choice among them? Is it constructible?
- Is a unique completion necessarily canonical?
- Can we find necessary and sufficient conditions for unique completion?
- Can we constructively characterise all completions?
- How do completions vary when we perturb K_Ω? (estimation)
- How do these questions relate to a process $\{X(t) : t \in [0,1]\}$ such that

$$\mathrm{Cov}\{X(u), X(v)\} = K_\Omega(u,v), \qquad (u,v) \in \Omega.$$

Answers can depend on the form of the domain Ω.

- We primarily consider **serrated domains** 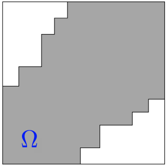

- ... and discuss extensions to **nearly serrated domains** 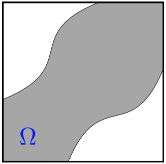

Enough to cover motivating problems.

1. Probability/Analysis: continuation of positive definite *functions*

   is a p.d. function ϕ determined by its restriction on $(-\delta, \delta)$?

   Equivalent to our problem in stationary case,

   $$K_\Omega(u,v) = \phi(u-v), \quad \Omega = \{|u-v| < \delta\}$$

   - Related to moment problem and continuation of characteristic functions (e.g. Gnedenko, Esseen)
   - Major results by Krein and co-workers.

2. Statistics:
   - Matrix case & Multivariate Analysis: e.g. Gohberg, Johnson, Dempster.
   - Functional Data Analysis: Descary & Panaretos (2019), Delaigle et al. (2020), Lin et al (2020), Kneip & Leibl (2020)...

Recovering Covariance from Sample Path Fragments/Snippets

Can we estimate $K = \mathrm{Cov}\{X(u), X(v)\}$ on $[0,1]^2$ when only observing copies of X on subintervals of $[0,1]$ of length $\delta < 1$?

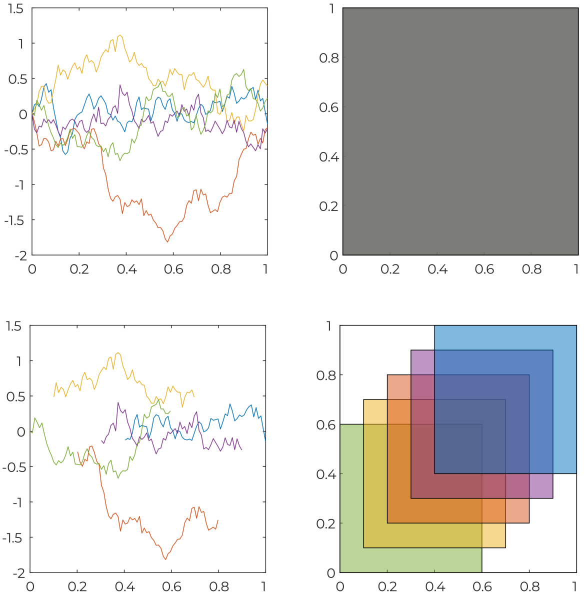
1.5
1
0.5
0
-0.5
-1
-1.5
-2
0
0.2
0.4
0.6
0.8
1

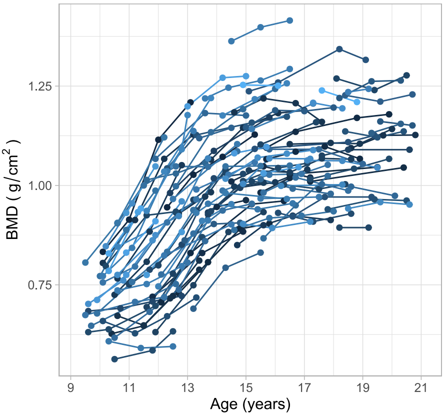

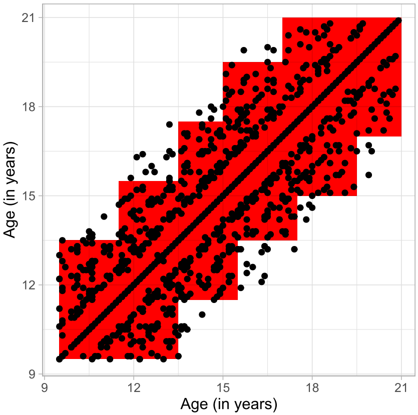

BMD measurements for 117 females taken between the ages of 9.5 and 21 years

Define the set of completions as

$$\mathcal{C}(K_\Omega) = \{K \succeq 0 \text{ on } [0,1]^2 : K|_\Omega = K_\Omega\}.$$

- Previous work focusses on sufficient conditions for $|\mathcal{C}(K_\Omega)| = 1$.
- We wish to comprehensively understand the set $\mathcal{C}(K_\Omega)$
- A priori, it is unclear if $\mathcal{C}(K_\Omega)$ is empty or not – did not define K_Ω as a *restriction* of a covariance
- $\mathcal{C}(K_\Omega)$ is convex & bounded (when K_Ω bounded), though.

Therefore:

$\mathcal{C}(K_\Omega)$ can either be empty, a singleton, or uncountably infinite.

Define the set of completions as

$$\mathcal{C}(K_\Omega) = \{K \succeq 0 \text{ on } [0,1]^2 : K|_\Omega = K_\Omega\}.$$

- Previous work focusses on sufficient conditions for $|\mathcal{C}(K_\Omega)| = 1$.
- We wish to comprehensively understand the set $\mathcal{C}(K_\Omega)$
- A priori, it is unclear if $\mathcal{C}(K_\Omega)$ is empty or not – did not define K_Ω as a *restriction* of a covariance
- $\mathcal{C}(K_\Omega)$ is convex & bounded (when K_Ω bounded), though.

Therefore:

$\mathcal{C}(K_\Omega)$ can either be empty, a singleton, or uncountably infinite.

It turns out that:

Theorem (Waghmare & Panaretos, 2021)

The set of completions $\mathcal{C}(K_\Omega)$ over a serrated domain is always non-empty. In particular, it always includes an explicitly constructible element.

We start with an easy case: the **2-serrated** case.

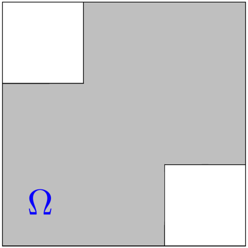

$$\Omega = (I_1 \times I_1) \cup (I_2 \times I_2) \quad \text{with} \quad I_1 = [0, b],\ I_2 = [a, 1]\, a \leq b.$$

For notational ease, we write

$$K_A = K_\Omega|_{A \times A}.$$

for any product set $A \times A \subset \Omega$.

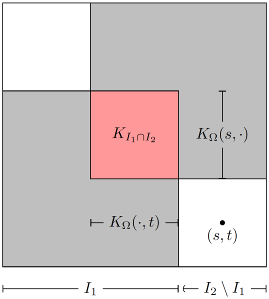

Define $K_\star : [0,1]^2 \to \mathbb{R}$ as

$$K_\star(s,t) = \begin{cases} K_\Omega(s,t), & (s,t) \in \Omega \\ \langle K_\Omega(s,\cdot), K_\Omega(\cdot,t) \rangle_{\mathcal{H}(K_{I_1 \cap I_2})}, & (s,t) \notin \Omega \end{cases}$$

where $\mathcal{H}(C)$ denotes the RKHS of a covariance C.

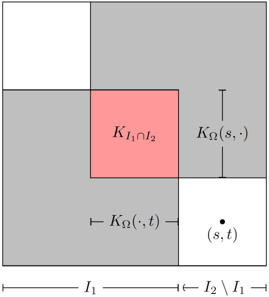

Define $K_\star : [0,1]^2 \to \mathbb{R}$ as

$$K_\star(s,t) = \begin{cases} K_\Omega(s,t), & (s,t) \in \Omega \\ \langle K_\Omega(s,\cdot), K_\Omega(\cdot,t) \rangle_{\mathcal{H}(K_{I_1 \cap I_2})}, & (s,t) \notin \Omega \end{cases}$$

where $\mathcal{H}(C)$ denotes the RKHS of a covariance C.

Proposition (Waghmare & Panaretos, 2021)

$K_\star$ is a bona fide covariance and $K_\star \in \mathcal{C}(K_\Omega)$.

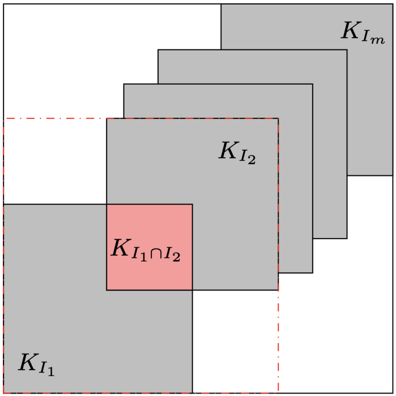
K_{I_m}
K_{I_2}
$K_{I_1 \cap I_2}$
K_{I_1}

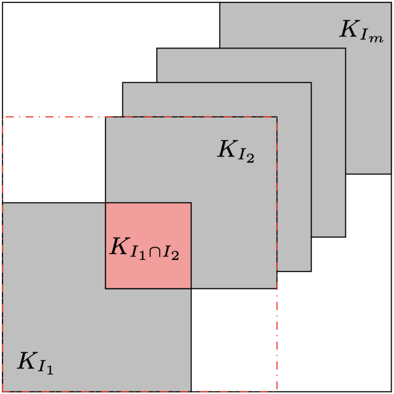

Theorem (Waghmare & Panaretos, 2021)

Recursive application of the 2-serrated formula yields a valid completion $K^* \in \mathcal{C}(K_\Omega)$, indeed the same completion irrespective of the order it is applied in.

As an example, let $I = [0,1]$

$$K_\Omega(s,t) = s \wedge t, \quad (s,t) \in \Omega = \underbrace{([0,2/3] \times [0,2/3])}_{I_1} \cup \underbrace{([1/3,1] \times [1/3,1])}_{I_2}.$$

Clearly, this can be completed to the covariance of standard Brownian motion,

$$K(s,t) = s \wedge t, \qquad (s,t) \in [0,1]^2.$$

As an example, let $I = [0,1]$

$$K_\Omega(s,t) = s \wedge t, \quad (s,t) \in \Omega = \underbrace{([0,2/3] \times [0,2/3])}_{I_1} \cup \underbrace{([1/3,1] \times [1/3,1])}_{I_2}.$$

Clearly, this can be completed to the covariance of standard Brownian motion,

$$K(s,t) = s \wedge t, \qquad (s,t) \in [0,1]^2.$$

Let's check if this is what our construction yields:

- For $f, g : \underbrace{[1/3, 2/3]}_{I_1 \cap I_2} \to \mathbb{R}$, $\langle f, g \rangle_{\mathcal{H}(K_{[1/3,2/3]})} = \frac{1}{(1/3)} \int_{1/3}^{2/3} f'(u) g'(u) du$.
- Thus, for $s \in (2/3, 1]$ and $t \in [0, 1/3)$,

$$K^*(s,t) = \frac{1}{(1/3)} \int_{1/3}^{2/3} \underbrace{\frac{\partial}{\partial u} K_\Omega(s,u)}_{=1} \underbrace{\frac{\partial}{\partial u} K_\Omega(u,t)}_{=t} du = t = s \wedge t, \qquad \text{since } t < s.$$

Iterating, we see that the extension method yields Brownian motion for any serrated domain.

Theorem (Waghmare & Panaretos, 2021)

The covariance K^* is the only completion of K_Ω such that the associated Gaussian process forms an undirected graphical model w.r.t. $G = ([0,1], \Omega)$

Theorem (Waghmare & Panaretos, 2021)

The covariance K^* is the only completion of K_Ω such that the associated Gaussian process forms an undirected graphical model w.r.t. $G = ([0,1], \Omega)$

What does it mean to say

"*the Gaussian process X forms an undirected graphical model w.r.t.* $([0,1], \Omega)$"

Define an (uncountable) graph $G = ([0,1], \Omega)$, i.e. $s \leftrightarrow t$ whenever $(s,t) \in \Omega$.

Theorem (Waghmare & Panaretos, 2021)

The covariance K^* is the only completion of K_Ω such that the associated Gaussian process forms an undirected graphical model w.r.t. $G = ([0,1], \Omega)$

What does it mean to say

"*the Gaussian process X forms an undirected graphical model w.r.t.* $([0,1], \Omega)$"

Define an (uncountable) graph $G = ([0,1], \Omega)$, i.e. $s \leftrightarrow t$ whenever $(s,t) \in \Omega$.

Allows a usual conditional independence interpretation of the graph:

$$\underbrace{\{X(t) : t \in I\} \perp\!\!\!\perp \{X(t) : t \in J\} \,\Big|\, \{X(t) : t \in S\}}$$

$$\Updownarrow$$

$$\overbrace{S \text{ separates } I \text{ from } J \text{ w.r.t. } \Omega}$$

Theorem (Waghmare & Panaretos, 2021)

The covariance K^* is the only completion of K_Ω such that the associated Gaussian process forms an undirected graphical model w.r.t. $G = ([0,1], \Omega)$

What does it mean to say

"*the Gaussian process X forms an undirected graphical model w.r.t.* $([0,1], \Omega)$"

Define an (uncountable) graph $G = ([0,1], \Omega)$, i.e. $s \leftrightarrow t$ whenever $(s,t) \in \Omega$.

Allows a usual conditional independence interpretation of the graph:

$$\underbrace{\{X(t) : t \in I\} \perp\!\!\!\perp \{X(t) : t \in J\} \,\big|\, \{X(t) : t \in S\}}_{\Updownarrow}$$

$$\overbrace{S \text{ separates } I \text{ from } J \text{ w.r.t. } \Omega}$$

Say that $S \subset [0,1]$ separates $I \subset [0,1]$ from $J \subset [0,1]$ w.r.t. $G = ([0,1], \Omega)$ if,

1. $S^2 \subseteq \Omega$
2. for any path $I \ni t_1 < t_2 < \ldots < t_n \in J$ with $(t_j, t_{j+1}) \in \Omega$ for $j \in \{2, \ldots, n-1\}$, there exists an $m \in \{2, \ldots, n-1\}$ such that $t_m \in S$

- K^* has the global Markov property w.r.t. edge set Ω
- Intuitively, relies exclusively on correlations intrinsic to Ω — propagates only "observed" correlations via the Markov property, without introducing arbitrary unseen correlations.
- It is unique in doing so among all possible completions

For all these reasons:

We call the completion K^* the *canonical completion*.

- Interestingly, best linear prediction of fragments based on $K_\star$ is equivalent to optimal predictors introduced by Kneip & Leibl (2020).

Based on the method of constructing $K_\star$, we can go backwards and prove that:

Theorem (Characterisation of Graphical Models)

Let $\{X_t : t \in I\}$ be a Gaussian process with covariance K. Then, X forms an undirected graphical model with respect to a serrated Ω if and only if $K \in \mathcal{G}_\Omega$, where

$$\mathcal{G}_\Omega = \Big\{K \in \mathcal{C} : K(s,t) = \big\langle K|_\Omega(s,\cdot), K|_\Omega(\cdot,t)\big\rangle_{\mathcal{H}(K_J)} \text{ for all } J \subset I \text{ separating } s,t \in I \text{ in } \Omega\Big\}.$$

So our previous result can now be interpreted as saying:

Completions and Graphical Models

$$\mathcal{C}(K_\Omega) \cap \mathcal{G}_\Omega = \{K_\star\}$$

For $A \subset B \subset \Omega$, let K_B/K_A be the Schur complement of K_B w.r.t. K_A,

$$(K_B/K_A)(s,t) = K_B(s,t) - \big\langle K_B(s,\cdot), K_B(\cdot,t)\big\rangle_{\mathcal{H}(K_A)}$$

i.e. the covariance of the *residuals* $\{X_t - \Pi(X_t|X_A) : t \in B \setminus A\}$.

For $A \subset B \subset \Omega$, let K_B/K_A be the Schur complement of K_B w.r.t. K_A,

$$(K_B/K_A)(s,t) = K_B(s,t) - \langle K_B(s,\cdot), K_B(\cdot,t)\rangle_{\mathcal{H}(K_A)}$$

i.e. the covariance of the *residuals* $\{X_t - \Pi(X_t|X_A) : t \in B \setminus A\}$.

Theorem (Waghmare & Panaretos, 2021)

Let K_Ω be a partial covariance kernel on a serrated domain Ω of m intervals.

The following two statements are equivalent:

1. K_Ω admits a unique completion on $[0,1]^2$, i.e. $\mathcal{C}(K_\Omega)$ is a singleton.
2. there exists an $r \in \{1,\dots,m\}$, such that

$$K_{I_p}/K_{I_p \cap I_{p+1}} = 0, \text{ for } 1 \le p < r \quad \text{and} \quad K_{I_{q+1}}/K_{I_q \cap I_{q+1}} = 0, \text{ for } r \le q < m.$$

- Condition (2) is strictly weaker than any previous sufficient condition (so those were not necessary)
- It implies that $X(t) = \Pi[X(t)|\{X(s) : t \in I_r\}]$ for one of the intervals I_r defining the serrated domain.
- So when unique completion is possible, the process $\{X(t) : t \in [0,1]\}$ is a **deterministic linear transformation** of its restriction $\{X(t) : t \in I_r\}$ to one of the intervals I_r defining the serrated domain.
- In any case, when a unique completion exists, it must be the canonical one.
- Condition is checkable at the level of K_Ω, i.e. at the level of observables
- Notice that *identifiability* of K from $K|_\Omega$ does not *require* unique completion conditions on $K|_\Omega$ – can assume, for example that $K \in \mathcal{G}_\Omega$ (a strictly weaker assumption)

K is a completion of K_Ω if and only if

$$K = K_\star + C$$

where C is a valid cross-covariance between

$$X_1 \sim N(0, K_{I_1}/K_{I_1\cap I_2}) \quad \text{and} \quad X_2 \sim N(0, K_{I_2}/K_{I_1\cap I_2})$$

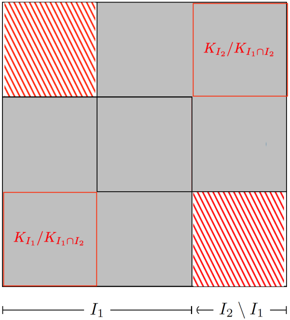

- Valid C are easily but arbitrarily obtained:
  - *Any* coupling of $X_1 \sim N(0, K_{I_1}/K_{I_1\cap I_2})$ and $X_2 \sim N(0, K_{I_2}/K_{I_1\cap I_2})$ will yield valid cross-covariance $C(s,t) = \text{cov}\{X_1(s), X_2(t)\}$
  - Like assigning a correlation to two variances – think of 3×3 matrices

$$\begin{pmatrix} \sigma_1^2 & * & ? \\ * & * & * \\ ? & * & \sigma_2^2 \end{pmatrix}$$

- Can characterise in operator notation – choose $\|\Psi\| = 1$ arbitrarily, then

$$\mathbf{K}f = \mathbf{K}_\star f + \underbrace{\left(\begin{array}{c|c|c} 0 & 0 & (\mathbf{L}_1^{1/2}\Psi\mathbf{L}_2^{1/2})^* \\ \hline 0 & 0 & 0 \\ \hline \mathbf{L}_1^{1/2}\Psi\mathbf{L}_2^{1/2} & 0 & 0 \end{array}\right)}_{\mathbf{C}} \begin{pmatrix} f|_{I_1\setminus I_2} \\ f|_{I_1\cap I_2} \\ f|_{I_2\setminus I_1} \end{pmatrix}$$

- Any completion other than canonical one introduces arbitrary correlations

Everything in black depends only on K_Ω (equiv. on its canonical extension $K_\star$):

Theorem (Waghmare & Panaretos, 2021)

Let K_Ω be a continuous partial covariance on a serrated domain Ω of m intervals. Then K is a completion of K_Ω if and only if its operator $f \mapsto \mathcal{K}f$ has the form

$$\mathcal{K}f(t) = \sum_{j:t\in I_j} \mathcal{K}_j f_{I_j}(t) + \sum_{p:t\in S_p} \mathcal{R}_p f_{D_p}(t) + \sum_{p:t\in D_p} \mathcal{R}_p^* f_{S_p}(t) - \sum_{p:t\in I_p\cap I_{p+1}} \mathcal{J}_p f_{J_p}(t) \text{ a.e.}$$

where for $1 \le p < m$,

$$\mathcal{R}_p = \underbrace{\left[\mathcal{J}_p^{-1/2}\mathcal{S}_p^*\right]^*\left[\mathcal{J}_p^{-1/2}\mathcal{D}_p\right]}_{\text{w/ kernel } K_\star|_{R_p}\text{, step } p \text{ of algorithm}} + \mathcal{U}_p^{1/2}\Psi_p\mathcal{V}_p^{1/2}$$

$$\mathcal{U}_p = \mathcal{K}_{S_p} - \left[\mathcal{J}_p^{-1/2}\mathcal{S}_p^*\right]^*\left[\mathcal{J}_p^{-1/2}\mathcal{S}_p^*\right], \qquad \mathcal{V}_p = \mathcal{K}_{D_p} - \left[\mathcal{J}_p^{-1/2}\mathcal{D}_p^*\right]^*\left[\mathcal{J}_p^{-1/2}\mathcal{D}_p^*\right]$$

and $\Psi_p : L^2(D_p) \to L^2(S_p)$ are bounded linear maps with $\|\Psi_p\| \le 1$.

Furthermore, taking $\Psi_1 = \Psi_2 = \ldots = \Psi_m = 0$ yields the canonical completion.

The Picture that Illustrates the Formula

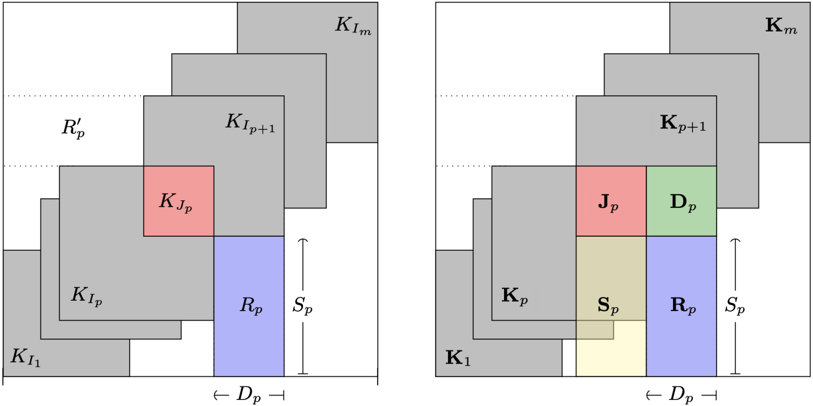

Makes sense to choose canonical completion as target of estimation:

- When completion is unique, it will be canonical
- When completion non-unique, canonical completion is least presumptuous

$\Longrightarrow$ It is always an identifiable and interpretable target of estimation

Makes sense to choose canonical completion as target of estimation:

- When completion is unique, it will be canonical
- When completion non-unique, canonical completion is least presumptuous

$\Longrightarrow$ It is always an identifiable and interpretable target of estimation

Estimating specifically the canonical completion is qualitatively different under non-uniqueness than all previous approaches (which focussed on uniqueness)

1. If we impose uniqueness by way of assumption (a very strong assumption), then one can use, for example, series estimators or matrix completion.
2. However such estimators will yield arbitrary (almost certainly non-canonical) completions if uniqueness does not actually hold.
3. To guarantee canonicity, we need to satisfy the system of operator equations on the previous slide – an inverse problem
4. Can be seen as an adaptive approach: will yield the unique completion when uniqueness holds, and a stable/canonical one otherwise.

Let $\widehat{K}_\Omega$ be an estimator of K_Ω.

Define $\widehat{K}_\star$ to be the estimator of $K_\star$ based on solving a regularised version of the linear operator system defining $K_\star$ (i.e. with all $\Psi_p = 0$).

Concretely, since

$$\mathcal{R}_p = \left[\mathcal{J}_p^{-1/2}\mathcal{S}_p^*\right]^* \left[\mathcal{J}_p^{-1/2}\mathcal{D}_p\right]$$

we start from $p = 1$ and recursively define the regularised empirical versions of $\mathcal{R}_p$:

$$\hat{\mathcal{R}}_p = \sum_{k=1}^{N_p} \frac{1}{\hat{\lambda}_{p,k}} \cdot \hat{\mathcal{S}}_p \, \hat{e}_{p,k} \otimes \hat{\mathcal{D}}_p^* \, \hat{e}_{p,k},$$

where:

- $\hat{\lambda}_{p,k}$ and $\hat{e}_{p,k}$ denote the kth eigenvalue and eigenfunction of $\hat{\mathcal{J}}_p$
- N_p is the truncation parameter
- $\hat{\mathcal{S}}_p$ has kernel $\hat{K}_\star|_{S_p \times J_p}$.

Let $A_{p,k}$ be the squared Hilbert-Schmidt error when approximating $\mathcal{R}_p = \left[\mathcal{J}_p{}^{-1/2}\mathcal{S}_p^*\right]^* \left[\mathcal{J}_p{}^{-1/2}\mathcal{D}_p\right]$ by replacing $\mathcal{J}_p$ with its rank-k truncation.

Theorem (Waghmare & Panaretos, 2021 (perturbation version))

Assume that for every $1 \leq p < m$, *we have*

- $\lambda_{p,k} \sim k^{-\alpha}$
- $A_{p,k} \sim k^{-\beta}$.

then

$$\|\widehat{K}_\star - K_\star\|_{L^2(I\times I)} \preceq \|\hat{K}_\Omega - K_\Omega\|_{L^2(\Omega)}^{\gamma_{m-1}}$$

where

$$\gamma_{m-1} = \frac{\beta}{4\alpha+\beta+3}\left[\frac{\beta}{2\alpha+\beta+1}\right]^{m-2}, \qquad m > 1,$$

provided the regularisation parameters are chosen to satisfy

$$N_p \sim \|\hat{K}_\Omega - K_\Omega\|_{L^2(\Omega)}^{-2\gamma_p/\beta}.$$

Theorem (Waghmare & Panaretos, 2021 (statistical version))

Assume that for every $1 \leq p < m$, *we have*

- $\lambda_{p,k} \sim k^{-\alpha}$
- $A_{p,k} \sim k^{-\beta}$

If

$$\|\hat{K}_\Omega - K_\Omega\|_{L^2(\Omega)} = O_{\mathbb{P}}(1/n^{\zeta})$$

then for every $\varepsilon > 0$,

$$\|\hat{K}_\star - K_\star\|_{L^2(I \times I)} = O_{\mathbb{P}}(1/n^{\zeta\gamma_{m-1}-\varepsilon})$$

provided the truncation parameters $\mathbf{N} = (N_p)_{p=1}^{m-1}$ *scale according to the rule*

$$N_p \sim n^{\gamma_p/\beta}$$

where $\gamma_{m-1} = \frac{\beta}{\beta+2\alpha+3/2}\left[\frac{\beta}{\beta+\alpha+1/2}\right]^{m-2}$.

$$\|\widehat{K}_\star - K_\star\|_{L^2(I\times I)} \preceq \|\hat{K}_\Omega - K_\Omega\|_{L^2(\Omega)}^{\gamma_{m-1}}$$

where $\gamma_{m-1} = \frac{\beta}{4\alpha+\beta+3}\left[\frac{\beta}{2\alpha+\beta+1}\right]^{m-2}$ for $m > 1$.

Remarks on the exponent γ_{m-1}:

- It strictly decreases as a function of the number of intervals m
- It can get arbitrarily close to 1 for a large enough rate of decay of approximation errors β.
- An increase in the rate of decay of eigenvalues α is accompanied by a decrease in the rate of convergence.
- If $K_\Omega \in C^r(\Omega)$ then the same applies to the kernels $K_\Omega|_{J_p\times J_p}$ of $\mathcal{J}_p$ implying $\lambda_{p,k}$ is $o(1/k^{r+1})$ for every $1 \le p < m$ and thus $\alpha = r+1$.
- All other things being equal, an increase in the smoothness of K_Ω also tends to a decrease in the rate of convergence

$$K_1(s,t) = \sum_{j=1}^{4} \frac{\phi_j(s)\phi_j(t)}{2^{j-1}} \qquad K_2(s,t) = s \wedge t \qquad K_3(s,t) = 10ste^{-10|s-t|^2}$$

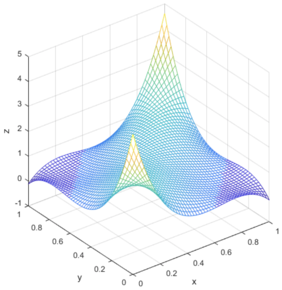

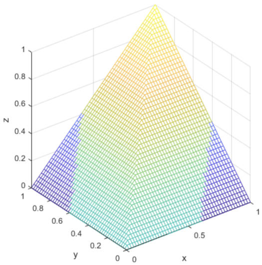

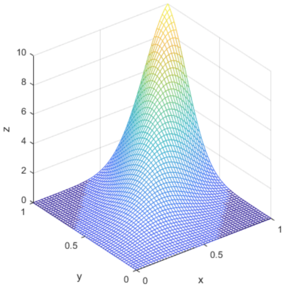

$$\mathrm{RRE} = \frac{\int_{\Omega^c} |\hat{K}_\star - K|^2 / \int_{\Omega^c} |K|^2}{\int_{\Omega} |\hat{K}_\Omega - K_\Omega|^2 / \int_{\Omega} |K|^2}.$$

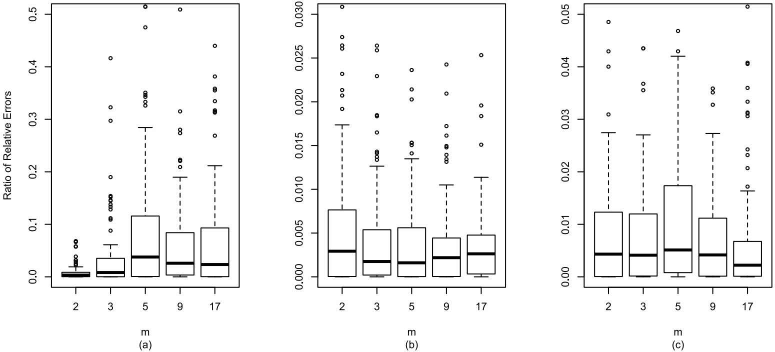

$\hat{K}_\Omega \to$ empirical covariance on Ω based on $n = 100$ fragments

Some Plots: $K_1(s,t) = \sum_{j=1}^{4} \frac{1}{2^{j-1}} \phi_j(s)\phi_j(t)$

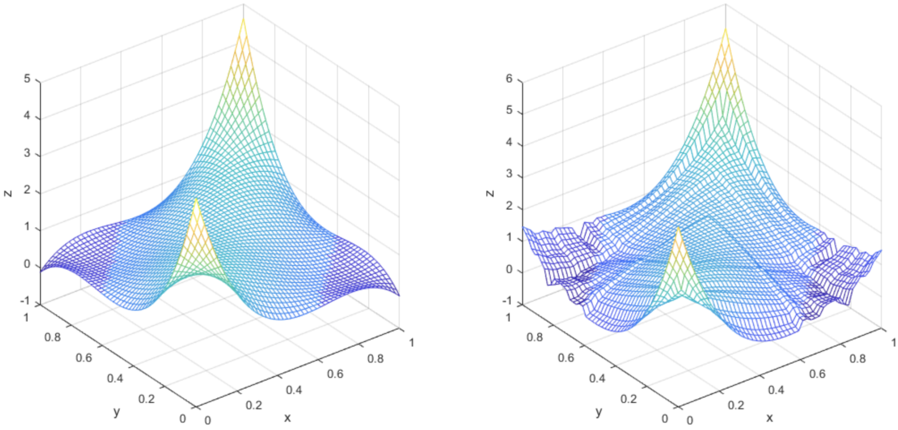

Some Plots: $K_2(s,t) = s \wedge t$

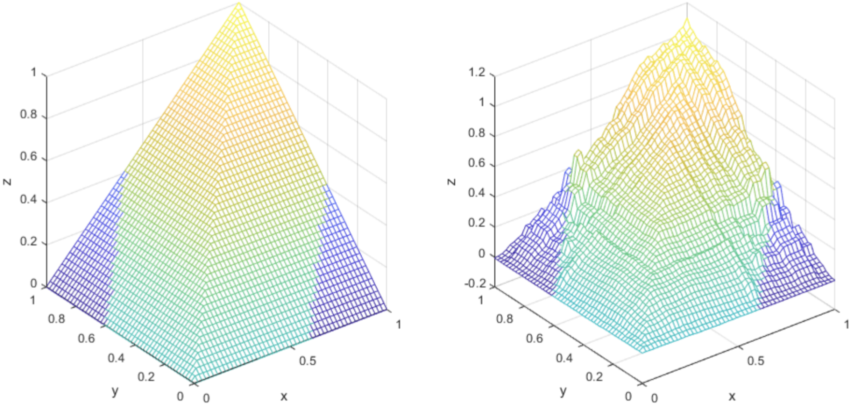

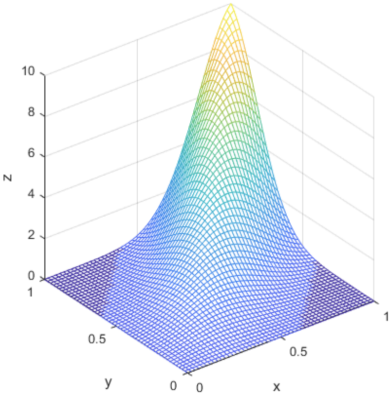
10
8
6
4
2
0
z
1
0.5
y
0
0
0.5
x
1

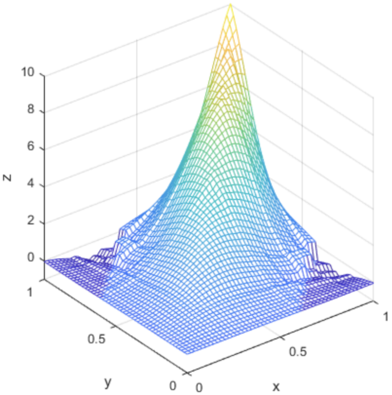
10
8
6
4
2
0
z
1
0.5
y
0
0
0.5
x
1

Definition

$\widetilde{\Omega} \subset I \times I$ is a nearly serrated domain if for every $\epsilon > 0$, there exist serrated domains $\Omega_\epsilon \subset \widetilde{\Omega} \subset \Omega^\epsilon$ such that $d_H(\widetilde{\Omega}, \Omega_\epsilon), d_H(\widetilde{\Omega}, \Omega^\epsilon) < \epsilon$

Of particular importance is a band $\widetilde{\Omega} = \{(s,t) \in I \times I : |s-t| \leq \delta\}$ which occurs when fragments are observable only over intervals of constant length.

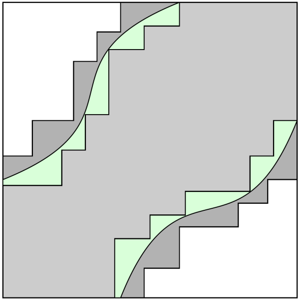

Main takeaway: can glean information via $\Omega_\epsilon \subset \widetilde{\Omega} \subset \Omega^\epsilon$

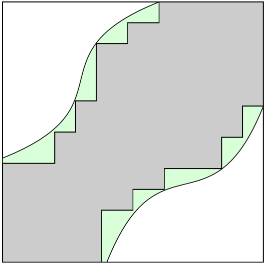

Proposition (Waghmare & Panaretos, 2021) – Checking uniqueness via serration

Let $K_{\widetilde{\Omega}}$ be a partial covariance on a nearly serrated domain $\widetilde{\Omega}$ and let $\Omega \subset \widetilde{\Omega}$ be a serrated domain. If the restriction $K_{\widetilde{\Omega}}\big|_{\Omega}$ admits a unique completion, so does $K_{\widetilde{\Omega}}$.

- The proposition, via our checkable necessary and sufficient conditions for uniqueness on serrated domains, gives can yield uniqueness under weaker conditions than previously known for banded domains.

Theorem (Waghmare & Panaretos, 2021) – Unique completions remain canonical

If $K_{\widetilde{\Omega}}$ on a nearly serrated $\widetilde{\Omega}$ completes uniquely, then the completion is canonical.

- So targeting a canonical completion remains a good strategy – under uniqueness, the unique completion is canonical.
- Canonical, in this case, means Ω-Markov.
- But how do we construct it?

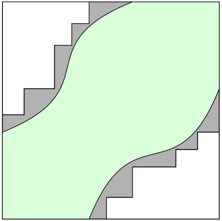

Theorem (Waghmare & Panaretos, 2021) – Constructibility of Canonical Completion

A covariance $K_\star$ on I can be recovered as the canonical completion of its restriction $K_\star|_\Omega$ on a serrated domain Ω if and only if it is the canonical completion of a partial covariance on some nearly serrated domain $\widetilde{\Omega} \subset \Omega$.

- In particular, if a unique completion of $K|_{\widetilde{\Omega}}$ exists then it equals the canonical completion of $K|_\Omega$ for a (in fact any) serrated $\Omega \supset \widetilde{\Omega}$
- Alternatively, if the process $X \sim N(0, K)$ is $\widetilde{\Omega}$-Markov for $\widetilde{\Omega}$ nearly serrated, then $K = (K|_\Omega)_\star$ for any serrated $\Omega \supset \widetilde{\Omega}$
- Consequential for inference from sample path fragments.

In practice: observable domain unclear a priori

Can consistently estimate any serrated restriction within $\Omega_\infty = \limsup_k I_k \times I_k$

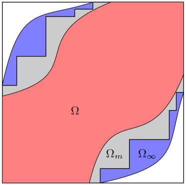

So need $K \in \mathcal{G}_\Omega$ with $\Omega \subset \Omega_\infty$ for identifiability, and can estimate from any

$$\Omega \subset \Omega_m \subset \Omega_\infty$$

Remarks:

- We can't know Ω_∞ and $\cup_{j \le n} I_j^2$ is a bad (overfitting) estimator thereof.
- "Well populated" regions are better proxies for Ω_∞
- Ω_m should represent a "well populated" nearly serrated region.
- Balance with choosing small m – large m introduces additional ill-posedness.
- Choosing Ω_m does not necessarily discard information, to the contrary it protects from boundary effects.

Example: Bone Mineral Density

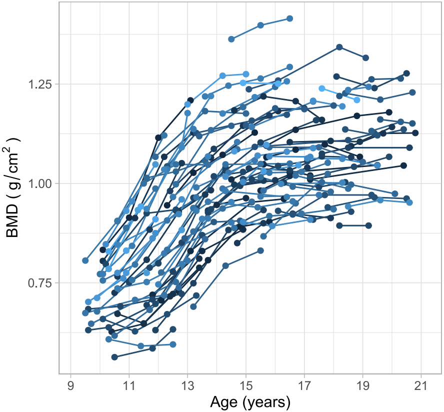

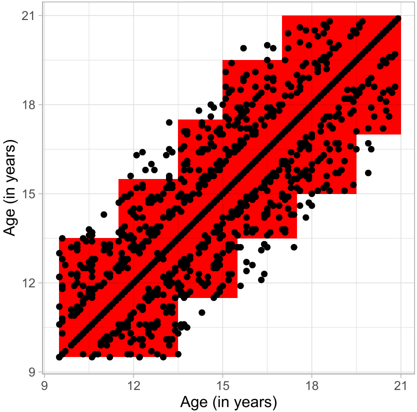

BMD measurements for 117 females taken between the ages of 9.5 and 21 years

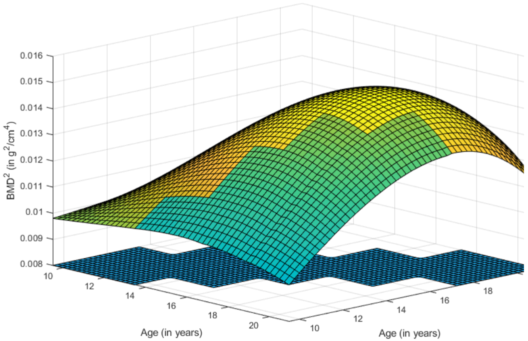

Figure: Completed covariance of the BMD data.

- Delaigle, Hall, Huang & Kneip (2020). *Estimating the covariance of fragmented and other related types of functional data.* Journal of the American Statistical Association.
- Descary & Panaretos (2019). *Recovering Covariance from Sample Path Fragments*. Biometrika.
- Kneip & Liebl (2020). *On the optimal reconstruction of partially observed functional data.* Annals of Statistics
- Lin, Wang, & Zhong (2020). *Basis Expansions for Functional Snippets*. Biometrika.
- Waghmare & Panaretos (2021). *The Completion of Covariance Kernels.* `arXiv:2107.07350`