

Computational Methods for GPs

Lecture 23

Dr. Colin Rundel

GPs and Computational Complexity

The problem with GPs

Unless you are lucky (or clever), Gaussian process models are difficult to scale to large problems. For a Gaussian process $\mathbf{y} \sim (\boldsymbol{\mu}, \boldsymbol{\Sigma})$:

Want to sample \mathbf{y} ?

$$\boldsymbol{\mu} + \text{Chol}(\boldsymbol{\Sigma}) \times \mathbf{Z} \text{ with } Z_i \sim (0, 1) \quad (\mathbf{n}^3)$$

Evaluate the (log) likelihood?

$$-\frac{1}{2} \log |\boldsymbol{\Sigma}| - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) - \frac{n}{2} \log 2\pi \quad (\mathbf{n}^3)$$

Update covariance parameter?

$$\{\boldsymbol{\Sigma}\}_{ij} = \sigma^2 \exp(-\{d\}_{ij} \phi) + \sigma_n^2 \mathbf{1}_{i=j} \quad (\mathbf{n}^2)$$

A simple guide to computational complexity

(n) - Linear complexity

Go for it!


(n^2) - Quadratic complexity

Pray

(n^3) - Cubic complexity

Give up

How bad is the problem?



Practice - Migratory Model Prediction

After fitting the GP need to sample from the posterior predictive distribution at ~ 3000 locations

$$\mathbf{y}_p \sim \left(\boldsymbol{\mu}_p + \boldsymbol{\Sigma}_{p0} \boldsymbol{\Sigma}_0^{-1} (\mathbf{y}_0 - \boldsymbol{\mu}_0), \boldsymbol{\Sigma}_p - \boldsymbol{\Sigma}_{p0} \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\Sigma}_{0p} \right)$$

Step	CPU (secs)
1. Calc $\boldsymbol{\Sigma}_p, \boldsymbol{\Sigma}_{p0}, \boldsymbol{\Sigma}_0$	1.080
2. Calc $\text{chol}(\boldsymbol{\Sigma}_p - \boldsymbol{\Sigma}_{p0} \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\Sigma}_{0p})$	0.467
3. Calc $\boldsymbol{\mu}_{p 0} + \text{chol}(\boldsymbol{\Sigma}_{p 0}) \times \mathbf{Z}$	0.049
4. Calc Allele Prob	0.129
Total	1.732

Total run time for 1000 posterior predictive draws:

- CPU (28.9 min)

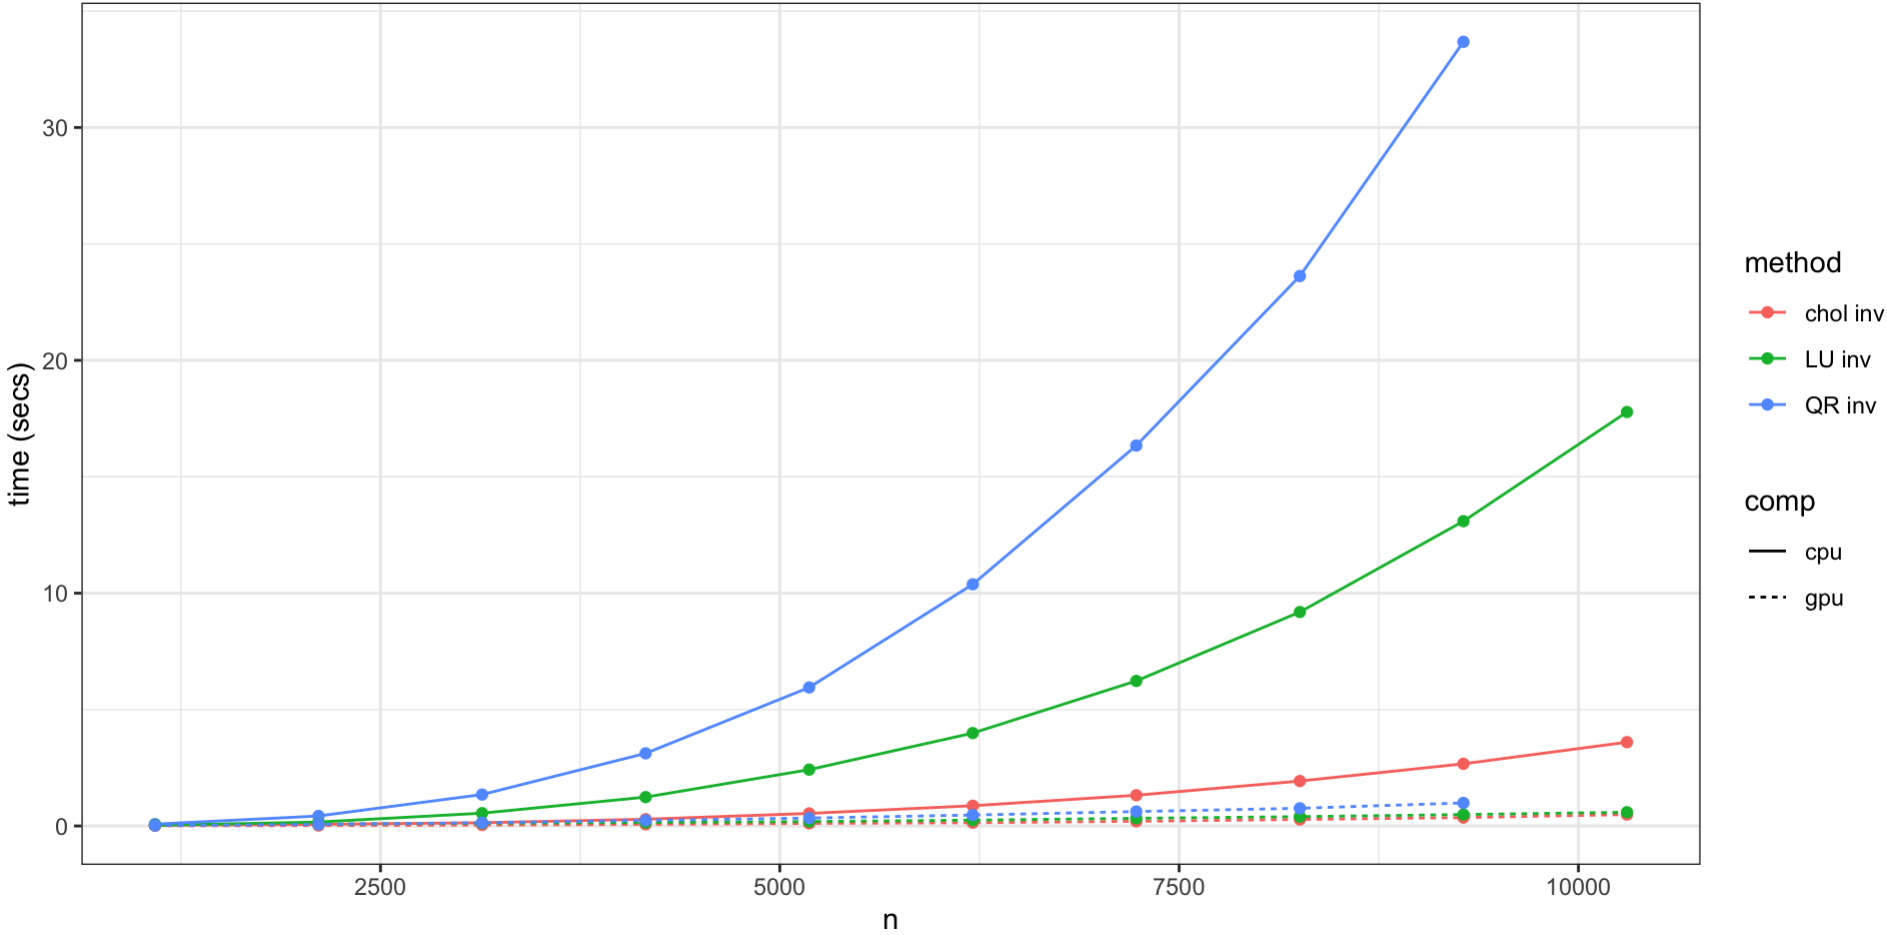
A bigger hammer?


Step	CPU (secs)	CPU+GPU (secs)	Rel. Perf
1. Calc. $\Sigma_p, \Sigma_{p_0}, \Sigma_p$	1.080	0.046	23.0
2. Calc. $\text{chol}(\Sigma_p - \Sigma_{p_0}\Sigma_o^{-1}\Sigma_{op})$	0.467	0.208	2.3
3. Calc. $\mu_{p o} + \text{chol}(\Sigma_{p o}) \times Z$	0.049	0.052	0.9
4. Calc. Allele Prob	0.129	0.127	1.0
Total	1.732	0.465	3.7

Total run time for 1000 posterior predictive draws:


- CPU (28.9 min)
- CPU+GPU (7.8 min)

Cholesky CPU vs GPU (P100)







Relative Performance



Aside (1) - Matrix Multiplication (P100)

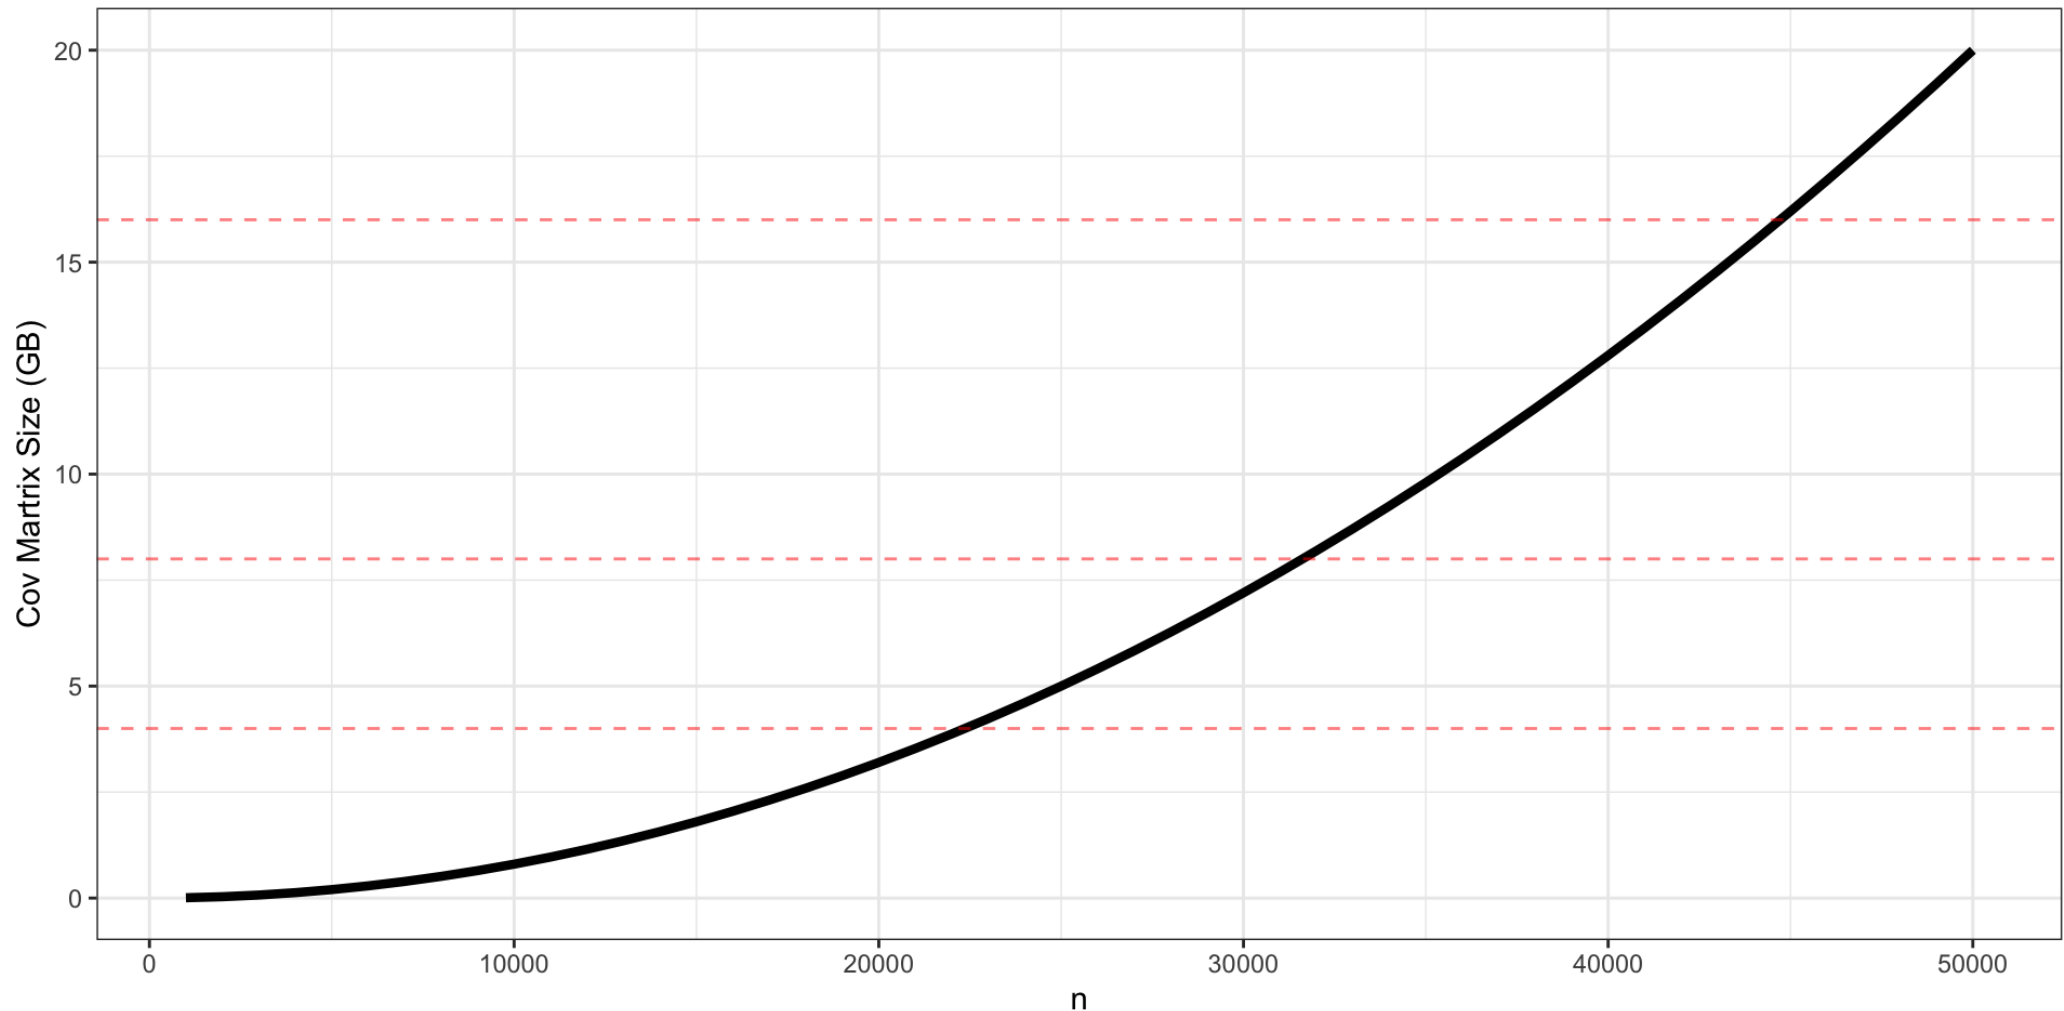


Matrix Multiplication - Relative Performance



Aside (2) - Memory Limitations

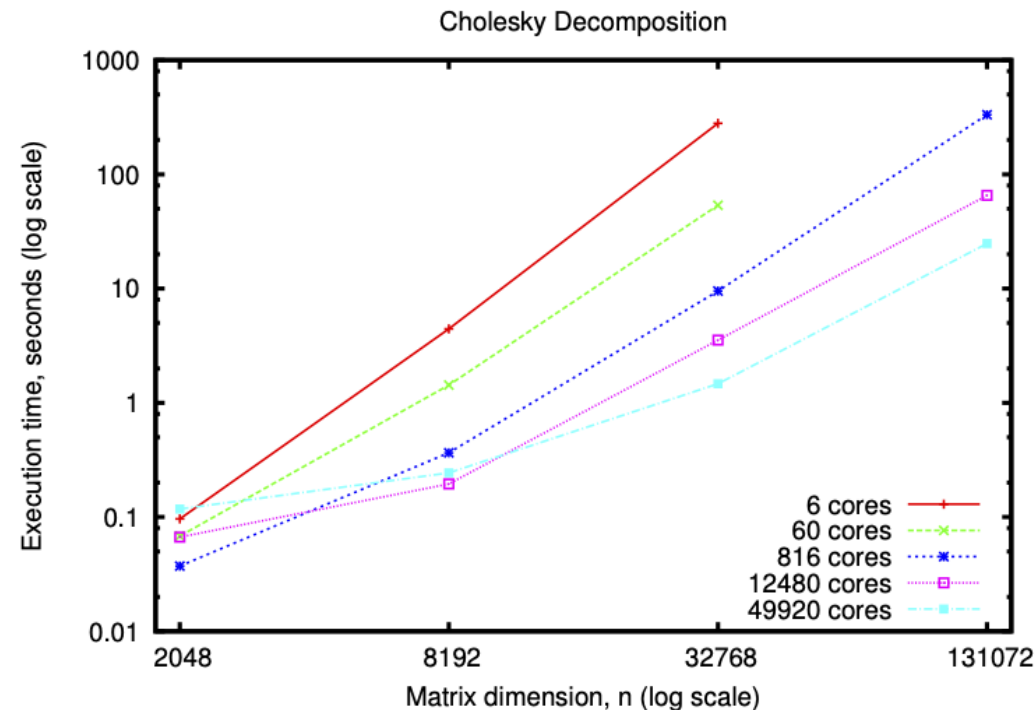
A general covariance is a dense $n \times n$ matrix, meaning it will require $n^2 \times 64$ -bits to store.



Other big hammers

`bigGP` is an R package written by Chris Paciorek (UC Berkeley), et al.

- Specialized distributed implementation of linear algebra operation for GPs
- Designed to run on large super computer clusters
- Uses both shared and distributed memory
- Able to fit models on the order of $n = 65k$ (32 GB Cov. matrix)



More scalable solutions?

- Spectral domain / basis functions
- Covariance tapering
- GMRF approximations
- Low-rank approximations
- Nearest-neighbor models

Low Rank Approximations

Low rank approximations in general

Lets look at the example of the singular value decomposition of a matrix,

$$\underset{n \times m}{M} = \underset{n \times n}{U} \underset{n \times m}{\text{diag}(S)} \underset{m \times m}{V^t}$$

where U are the left singular vectors, V the right singular vectors, and S the singular values. Usually the singular values and vectors are ordered such that the singular values are in descending order.

The Eckart–Young theorem states that we can construct an approximation of M with rank k by setting S to contain only the k largest singular values and all other values set to zero.

$$\begin{aligned} \underset{n \times m}{M} &= \underset{n \times n}{U} \underset{n \times m}{\text{diag}(S)} \underset{m \times m}{V^t} \\ &= \underset{n \times k}{\tilde{U}} \underset{k \times k}{\tilde{\text{diag}}(S)} \underset{k \times m}{\tilde{V}^t} \end{aligned}$$

Example

$$M = \begin{pmatrix} 1.000 & 0.500 & 0.333 & 0.250 \\ 0.500 & 0.333 & 0.250 & 0.200 \\ 0.333 & 0.250 & 0.200 & 0.167 \\ 0.250 & 0.200 & 0.167 & 0.143 \end{pmatrix} = U \text{diag}(S) V^t$$
$$U = V = \begin{pmatrix} -0.79 & 0.58 & -0.18 & -0.03 \\ -0.45 & -0.37 & 0.74 & 0.33 \\ -0.32 & -0.51 & -0.10 & -0.79 \\ -0.25 & -0.51 & -0.64 & 0.51 \end{pmatrix}$$
$$S = (1.50 \quad 0.17 \quad 0.01 \quad 0.00)$$

Rank 2 approximation

$$M = \begin{pmatrix} -0.79 & 0.58 \\ -0.45 & -0.37 \\ -0.32 & -0.51 \\ -0.25 & -0.51 \end{pmatrix} \begin{pmatrix} 1.50 & 0.00 \\ 0.00 & 0.17 \end{pmatrix} \begin{pmatrix} -0.79 & -0.45 & -0.32 & -0.25 \\ 0.58 & -0.37 & -0.51 & -0.51 \end{pmatrix}$$
$$= \begin{pmatrix} 1.000 & 0.501 & 0.333 & 0.249 \\ 0.501 & 0.330 & 0.251 & 0.203 \\ 0.333 & 0.251 & 0.200 & 0.166 \\ 0.249 & 0.203 & 0.166 & 0.140 \end{pmatrix}$$

$$M = \begin{pmatrix} 1.000 & 0.500 & 0.333 & 0.250 \\ 0.500 & 0.333 & 0.250 & 0.200 \\ 0.333 & 0.250 & 0.200 & 0.167 \\ 0.250 & 0.200 & 0.167 & 0.143 \end{pmatrix}$$

Approximation Error

We can measure the error of the approximation using the Frobenius norm,

$$\|M - \tilde{M}\|_F = \left(\sum_{i=1}^m \sum_{j=1}^n (M_{ij} - \tilde{M}_{ij})^2 \right)^{1/2}$$


$$M - \tilde{M} = \begin{pmatrix} 0.00022 & -0.00090 & 0.00012 & 0.00077 \\ -0.00090 & 0.00372 & -0.00053 & -0.00317 \\ 0.00012 & -0.00053 & 0.00013 & 0.00039 \\ 0.00077 & -0.00317 & 0.00039 & 0.00277 \end{pmatrix}$$

$$\|M - \tilde{M}\|_F = 0.00674$$


Strong dependence

For a covariance matrix with a large effective range,

SVD




Low Rank SVD




Weak dependence

For a covariance matrix with a large effective range,

SVD



Low Rank SVD



How does this help?

There is an immensely useful linear algebra identity, the Sherman-Morrison-*Woodbury* formula, for the inverse (and determinant) of a decomposed matrix,

$$\begin{aligned} \mathbf{M}_{n \times m}^{-1} &= \left(\begin{array}{ccc} \mathbf{A} & \mathbf{U} & \mathbf{S} \\ \mathbf{V}^t & & \end{array} \right)^{-1} \\ &= \mathbf{A}^{-1} - \mathbf{A}^{-1} \mathbf{U} (\mathbf{S}^{-1} + \mathbf{V}^t \mathbf{A}^{-1} \mathbf{U})^{-1} \mathbf{V}^t \mathbf{A}^{-1}. \end{aligned}$$

How does this help?

- Imagine that $\mathbf{A} = \text{diag}(\mathbf{A})$, then it is trivial to find \mathbf{A}^{-1} .
- \mathbf{S}^{-1} is $k \times k$ which is hopefully small, or even better $\mathbf{S} = \text{diag}(\mathbf{S})$.
- $(\mathbf{S}^{-1} + \mathbf{V}^t \mathbf{A}^{-1} \mathbf{U})$ is $k \times k$ which is also small.

Aside - Determinant

Remember for any MVN distribution when evaluating the likelihood

$$-\frac{1}{2} \log |\Sigma| - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) - \frac{n}{2} \log 2\pi$$

we need the inverse of Σ as well as its *determinant*.

Low rank approximations for GPs

For a standard spatial random effects model,

$$y(\mathbf{s}) = \mathbf{x}(\mathbf{s}) \boldsymbol{\beta} + w(\mathbf{s}) + \epsilon, \quad \epsilon \sim \mathbf{N}(0, \tau^2 \mathbf{I})$$

$$w(\mathbf{s}) \sim (0, \boldsymbol{\Sigma}(\mathbf{s})), \quad \boldsymbol{\Sigma}(\mathbf{s}, \mathbf{s}') = \sigma^2 \rho(\mathbf{s}, \mathbf{s}' | \theta)$$

if we can replace $\boldsymbol{\Sigma}(\mathbf{s})$ with a low rank approximation of the form $\boldsymbol{\Sigma}(\mathbf{s}) \approx \mathbf{U} \mathbf{S} \mathbf{U}^t$ where

- \mathbf{U} is $n \times k$,
- \mathbf{S} is $k \times k$, and
- $\mathbf{A} = \tau^2 \mathbf{I}$ or a similar diagonal matrix

Predictive Processes

Gaussian Predictive Processes

For a rank k approximation,

- Pick k knot locations s^*
- Calculate knot covariance, $\Sigma(s^*)$, and knot cross-covariance, $\Sigma(s, s^*)$
- Approximate full covariance using

$$\Sigma(s) \approx \underbrace{\Sigma(s, s^*)}_{n \times k} \underbrace{\Sigma(s^*)^{-1}}_{k \times k} \underbrace{\Sigma(s^*, s)}_{k \times n}.$$

- PPs systematically underestimates variance (σ^2) and inflate τ^2 ,
Modified predictive process corrects this using

$$\begin{aligned} \boldsymbol{\Sigma}(\mathbf{s}) \approx & \boldsymbol{\Sigma}(\mathbf{s}, \mathbf{s}^*) \boldsymbol{\Sigma}(\mathbf{s}^*)^{-1} \boldsymbol{\Sigma}(\mathbf{s}^*, \mathbf{s}) \\ & + \text{diag}\left(\boldsymbol{\Sigma}(\mathbf{s}) - \boldsymbol{\Sigma}(\mathbf{s}, \mathbf{s}^*) \boldsymbol{\Sigma}(\mathbf{s}^*)^{-1} \boldsymbol{\Sigma}(\mathbf{s}^*, \mathbf{s})\right). \end{aligned}$$

Example


Below we have a surface generate from a squared exponential Gaussian Process where

$$\{\Sigma\}_{ij} = \sigma^2 \exp(-(\phi d)^2) + \tau^2 I$$


$$\sigma^2 = 1 \quad \phi = 9 \quad \tau^2 = 0.1$$

Predictive Process Model Results


True Field




PP - 5 x 5 knots



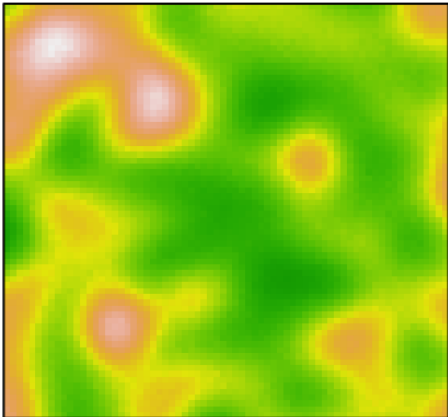
PP - 10 x 10 knots




PP - 15 x 15 knots




Full GP




Mod. PP - 5 x 5 knots




Mod. PP - 10 x 10 knots




Mod. PP - 15 x 15 knots



Performance



Parameter Estimates



Random Projections

Low Rank via Random Projections

1. Starting with an matrix \mathbf{A} .
 $m \times n$
2. Draw a Gaussian random matrix $\mathbf{\Omega}$.
 $n \times k+p$
3. Form $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$ and compute its QR factorization $\mathbf{Y} = \mathbf{Q} \mathbf{R}$
4. Form $\mathbf{B} = \mathbf{Q}' \mathbf{A}$.
5. Compute the SVD of $\mathbf{B} = \hat{\mathbf{U}} \mathbf{S} \mathbf{V}'$.
6. Form the matrix $\mathbf{U} = \mathbf{Q} \hat{\mathbf{U}}$.
7. Form $\tilde{\mathbf{A}} = \mathbf{U} \mathbf{S} \mathbf{V}'$

Resulting approximation has a bounded expected error,

$$\mathbb{E} \| \mathbf{A} - \mathbf{USV}' \|_{\text{F}} \leq \left[1 + \frac{4\sqrt{k+p}}{p-1} \sqrt{\min(m, n)} \right] \sigma_{k+1}.$$

Random Matrix Low Rank Approxs and GPs

The preceding algorithm can be modified slightly to take advantage of the positive definite structure of a covariance matrix.

1. Starting with an $n \times n$ covariance matrix A .
2. Draw Gaussian random matrix $\mathbf{\Omega}$ $\begin{matrix} n \times k+p \end{matrix}$.
3. Form $Y = A \mathbf{\Omega}$ and compute its QR factorization $Y = Q R$.
4. Form the $B = Q' A Q$.
5. Compute the eigen decomposition of $B = \hat{U} S \hat{U}'$.
6. Form the matrix $U = Q \hat{U}$.

Once again we have a bound on the error,

$$E \|A - USU'\|_F \approx c \cdot \sigma_{k+1}.$$

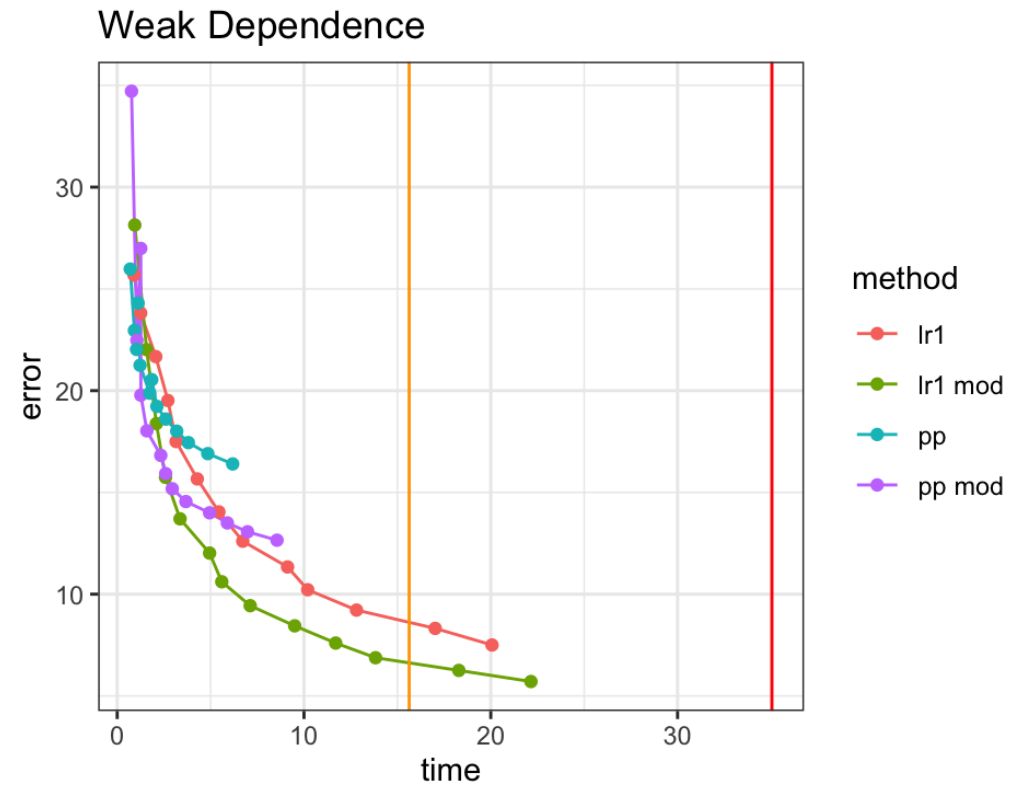
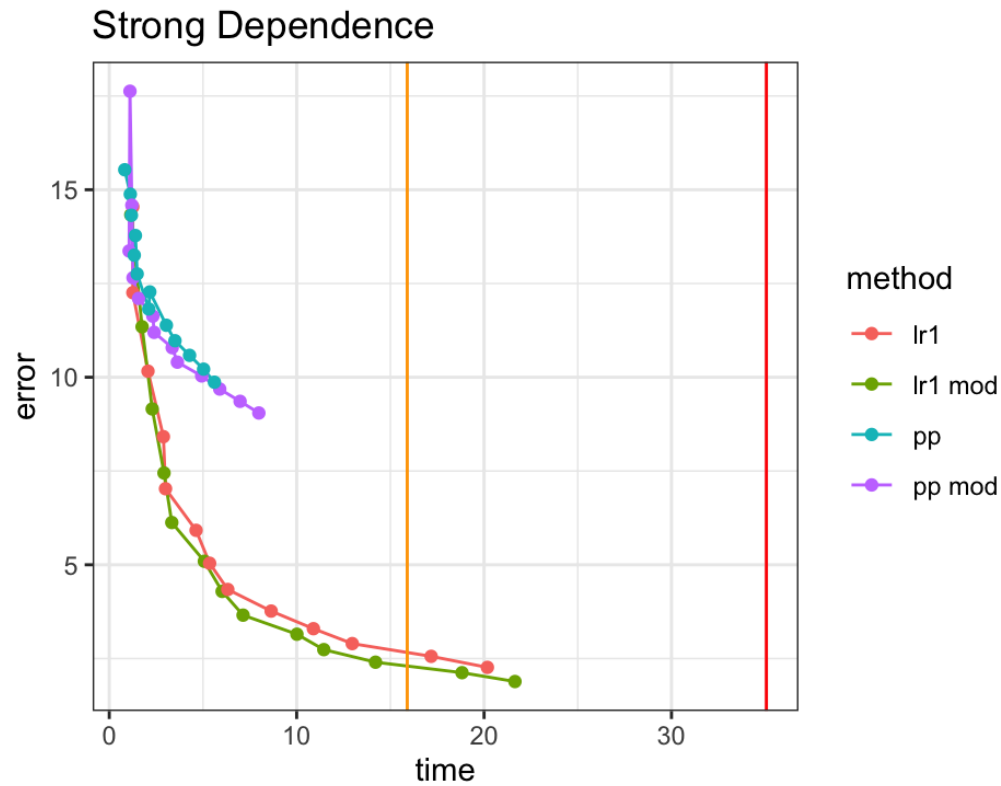
Low Rank Approximations and GPUs

Both predictive process and random matrix low rank approximations are good candidates for acceleration using GPUs.

- Both use Sherman-Woodbury-Morrison to calculate the inverse (involves matrix multiplication, addition, and a small matrix inverse).
- Predictive processes involves several covariance matrix calculations (knots and cross-covariance) and a small matrix inverse.
- Random matrix low rank approximations involves a large matrix multiplication ($A \Omega$) and several small matrix decompositions (QR, eigen).

Comparison

$n = 15,000$, $k = \{100, \dots, 4900\}$



Rand. Projection LR Depositions for Prediction

This approach can also be used for prediction, if we want to sample


$$\mathbf{y} \sim (0, \mathbf{\Sigma})$$

$$\mathbf{\Sigma} \approx \mathbf{USU}^t = (\mathbf{US}^{1/2} \mathbf{U}^t)(\mathbf{US}^{1/2} \mathbf{U}^t)^t$$

then

$$\mathbf{y}_{\text{pred}} = (\mathbf{U} \mathbf{S}^{1/2} \mathbf{U}^t) \times \mathbf{Z} \text{ where } Z_i \sim (0, 1)$$

because $\mathbf{U}^t \mathbf{U} = \mathbf{I}$ since \mathbf{U} is an orthogonal matrix.



$n = 1000, p = 10000$
 Sta 344 - Fall 2022