

Lecture 22

Computational Methods for GPs

Colin Rundel

04/12/2017

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 $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$:

The problem with GPs

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

Want to sample \mathbf{y} ?

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

The problem with GPs

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

Want to sample \mathbf{y} ?

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

Evaluate the (log) likelihood?

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

The problem with GPs

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

Want to sample \mathbf{y} ?

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

Evaluate the (log) likelihood?

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

Update covariance parameter?

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

$\mathcal{O}(n)$ - Linear complexity

$\mathcal{O}(n)$ - Linear complexity - Go for it

$\mathcal{O}(n)$ - Linear complexity - Go for it

$\mathcal{O}(n^2)$ - Quadratic complexity

$\mathcal{O}(n)$ - Linear complexity - Go for it

$\mathcal{O}(n^2)$ - Quadratic complexity - Pray

$\mathcal{O}(n)$ - Linear complexity - Go for it

$\mathcal{O}(n^2)$ - Quadratic complexity - Pray


$\mathcal{O}(n^3)$ - Cubic complexity

$\mathcal{O}(n)$ - Linear complexity - Go for it

$\mathcal{O}(n^2)$ - Quadratic complexity - Pray

$\mathcal{O}(n^3)$ - 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

$$y_p \sim \mathcal{N}(\mu_p + \Sigma_{po}\Sigma_o^{-1}(y_o - \mu_o), \Sigma_p - \Sigma_{po}\Sigma_o^{-1}\Sigma_{op})$$

Practice - Migratory Model Prediction


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


$$y_p \sim \mathcal{N}(\mu_p + \Sigma_{po}\Sigma_o^{-1}(y_o - \mu_o), \Sigma_p - \Sigma_{po}\Sigma_o^{-1}\Sigma_{op})$$

Step	CPU (secs)	CPU+GPU (secs)	Rel. Performance
1. Calc. Σ_p, Σ_{po}	1.080	0.046	23.0
2. Calc. $\text{chol}(\Sigma_p - \Sigma_{po}\Sigma_o^{-1}\Sigma_{op})$	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: CPU (28.9 min), CPU+GPU (7.8 min)

Cholesky CPU vs GPU (P100)






Relative Performance




Aside - Matrix Multiplication



An even bigger hammer

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,

$$M = U \text{diag}(S) V^t$$

$n \times m$ $n \times n$ $n \times m$ $m \times m$

where U are called 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.

Low rank approximations in general

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

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

where U are called 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.

It turns out (Eckart–Young theorem) that we can approximate M as having rank r by defining \tilde{S} to only have the r largest singular values (others set to zero).

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

Example

$$M = \begin{pmatrix} 1.000 & 0.500 & 0.333 & 0.250 \\ 0.500 & 0.333 & 0.250 & 0.200 \\ 0.333 & 0.250 & 0.200 & 0.167 \\ 0.250 & 0.200 & 0.167 & 0.143 \end{pmatrix} = U \operatorname{diag}(S) V^t$$

$$U = V = \begin{pmatrix} -0.79 & 0.58 & -0.18 & -0.03 \\ -0.45 & -0.37 & 0.74 & 0.33 \\ -0.32 & -0.51 & -0.10 & -0.79 \\ -0.25 & -0.51 & -0.64 & 0.51 \end{pmatrix}$$

$$S = \begin{pmatrix} 1.50 & 0.17 & 0.01 & 0.00 \end{pmatrix}$$

Example

$$M = \begin{pmatrix} 1.000 & 0.500 & 0.333 & 0.250 \\ 0.500 & 0.333 & 0.250 & 0.200 \\ 0.333 & 0.250 & 0.200 & 0.167 \\ 0.250 & 0.200 & 0.167 & 0.143 \end{pmatrix} = U \operatorname{diag}(S) V^t$$

$$U = V = \begin{pmatrix} -0.79 & 0.58 & -0.18 & -0.03 \\ -0.45 & -0.37 & 0.74 & 0.33 \\ -0.32 & -0.51 & -0.10 & -0.79 \\ -0.25 & -0.51 & -0.64 & 0.51 \end{pmatrix}$$

$$S = \begin{pmatrix} 1.50 & 0.17 & 0.01 & 0.00 \end{pmatrix}$$

Rank 2 approximation:

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

Approximation Error

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

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


Approximation Error

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


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

Strong dependence (large eff. range):

SVD




Low Rank SVD




Weak dependence (short eff. range):

SVD



Low Rank SVD



How does this help? (Sherman-Morrison-Woodbury)

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

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

How does this help? (Sherman-Morrison-Woodbury)

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

$$\begin{aligned}\tilde{M}_{n \times m}^{-1} &= \begin{pmatrix} A & U & S & V^t \\ n \times m & n \times k & k \times k & k \times m \end{pmatrix}^{-1} \\ &= A^{-1} - A^{-1}U(S^{-1} + V^tA^{-1}U)^{-1}V^tA^{-1}.\end{aligned}$$

How does this help?

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

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*.

- For a full rank Cholesky decomposition we get the determinant for “free”.

$$|M| = |LL^t| = \prod_{i=1}^n (\text{diag}(L)_i)^2$$

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*.

- For a full rank Cholesky decomposition we get the determinant for “free”.

$$|M| = |LL^t| = \prod_{i=1}^n (\text{diag}(L)_i)^2$$

- For a low rank approximation the Sherman-Morrison-Woodbury / Determinant lemma gives us,

$$\begin{aligned} \det(\tilde{M}) &= \det(A + USV^t) \\ &= \det(S^{-1} + V^t A^{-1} U) \det(S) \det(A) \end{aligned}$$

Low rank approximations for GPs

For a standard spatial random effects model,

$$y(\mathbf{s}) = x(\mathbf{s}) \boldsymbol{\beta} + w(\mathbf{s}) + \epsilon, \quad \epsilon \sim N(0, \tau^2 I)$$
$$w(\mathbf{s}) \sim \mathcal{N}(0, \boldsymbol{\Sigma}(\mathbf{s})), \quad \boldsymbol{\Sigma}(\mathbf{s}, \mathbf{s}') = \sigma \rho(\mathbf{s}, \mathbf{s}' | \theta)$$

if we can replace $\boldsymbol{\Sigma}(\mathbf{s})$ with a low rank approximation of the form

- $\boldsymbol{\Sigma}(\mathbf{s}) \approx \mathbf{U} \mathbf{S} \mathbf{V}^t$ where
- \mathbf{U} and \mathbf{V} are $n \times k$,
- \mathbf{S} is $k \times k$, and
- $\mathbf{A} = \tau^2 I$ or a similar diagonal matrix

Predictive Processes

For a rank k approximation,

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

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

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



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

Example

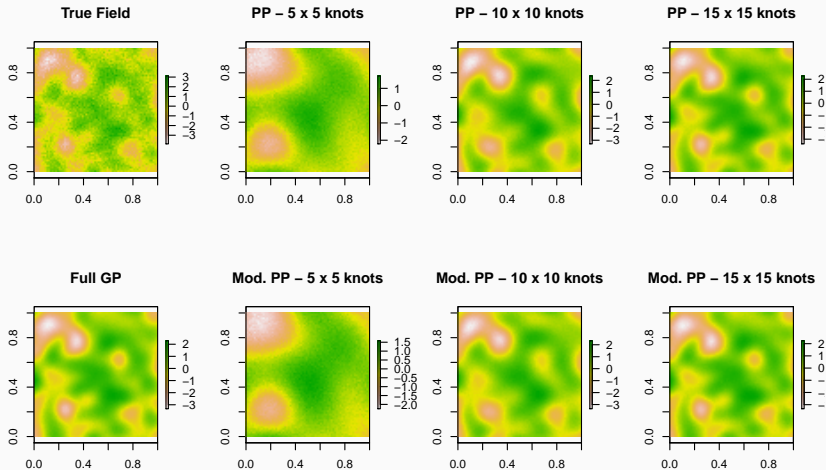
Below we have a surface generate from a squared exponential Gaussian Process where

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


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




Predictive Process Model Results



Performance



Parameter Estimates



Random Projections

Low Rank Approximations via Random Projections

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

Resulting approximation has a bounded expected error,

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

Halko, Martinsson, Tropp (2011)

Random Matrix Low Rank Approximations and GPs

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

Once again we have a bound on the error,

$$\mathbb{E} \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}'\mathbf{A}\mathbf{Q})\mathbf{Q}'\| = \mathbb{E} \|\mathbf{A} - \mathbf{U}\mathbf{S}\mathbf{U}'\| \lesssim c \cdot \sigma_{k+1}.$$

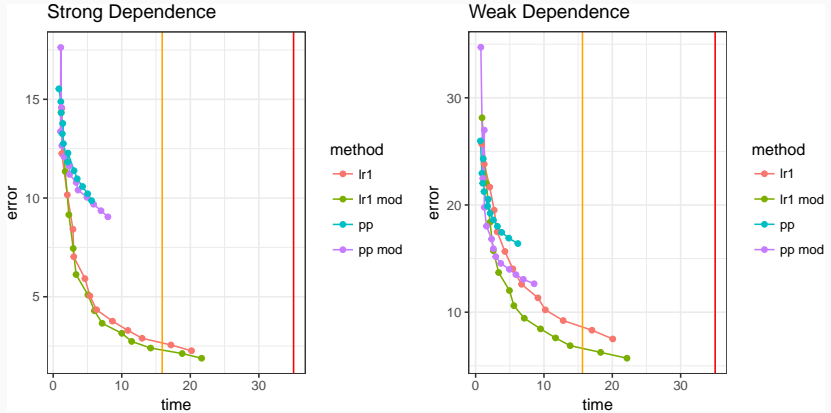
Halko, Martinsson, Tropp (2011), Banerjee, Dunson, Tokdar (2012)

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 ($\mathbf{A}\mathbf{\Omega}$) and several small matrix decompositions (QR, eigen).

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



Rand. Projection LR Decompositions for Prediction

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


$$\begin{aligned}y &\sim \mathcal{N}(0, \Sigma) \\ \Sigma &\approx \mathbf{U}\mathbf{S}\mathbf{U}^t = (\mathbf{U}\mathbf{S}^{1/2}\mathbf{U}^t)(\mathbf{U}\mathbf{S}^{1/2}\mathbf{U}^t)^t\end{aligned}$$

then

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

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

Dehdari, Deutsch (2012)



$n = 1000, \quad p = 10000$