



Met Office
Hadley Centre



u^b

UNIVERSITÄT
BERN

OESCHGER CENTRE
CLIMATE CHANGE RESEARCH



Royal Netherlands
Meteorological Institute
Ministry of Infrastructure and the
Environment

Challenges and solutions for multiscale global temperature reconstruction

Finn Lindgren (finn.lindgren@ed.ac.uk)

The University of Edinburgh, Scotland

with Colin Morice, John Kennedy, Christopher Merchant, and the EUSTACE team



THE UNIVERSITY of EDINBURGH


EUSTACE has received funding from the European Union's Horizon 2020 Programme for Research and Innovation, under Grant Agreement no 640171



EUSTACE ANALYSIS

Combines in-situ and satellite data sources to derive daily air temperatures across the globe with quantified uncertainties.

- Daily mean air temperature (2 m) estimates from the mid-late 19th century at $\frac{1}{4}$ degree resolution.
- Observational dataset for use in climate monitoring, services and research.
 - Quantify bias and uncertainty arising from observational sampling (in space and time);
 - Quantify uncertainty from instrumental effects/network changes.
- Higher resolution daily gridded analyses for regional climate
 - Combine in situ and remote sensing data to support high resolution analysis.
 - Absolute temperature rather than anomaly product.



OBSERVATIONS


In situ air temperature:

- EUSTACE station dataset (UBERN) (GHCN-D, ECA&D, ISTI, DECADE, ERA-CLIM)
- HadNMAT-2 ship air temperatures (NOCS/Met Office)


Satellite skin temperature derived air temperature:

- Marine: ATSR (ESA CCI SST)
- Land: MODIS (USGS/NASA via ESA GlobTemperature)
- Ice: AVHRR (NOAA/FP7 NAACLIM)


Assimilated Observations 01/01/1880




Assimilated Observations 01/01/1955




Assimilated Observations 01/01/1930




Assimilated Observations 01/01/1980







Assimilated Observations 01/01/1855



Assimilated Observations 01/01/1930



Assimilated Observations 01/01/2005



Statistical model and method building blocks

Basic system components

- ▶ Temperature *processes on different spatial and temporal scales*
 - ▶ Seasonal
 - ▶ Slow climate processes
 - ▶ Medium-scale variability
 - ▶ Daily
- ▶ *Vast model size* ($\sim 10^{11}$ unknowns); need computationally efficient tools
- ▶ Hierarchical statistical model structure based on Gaussian processes
 - ▶ Stochastic PDEs translates to sparse precisions in *Gaussian Markov random fields*
- ▶ *Propagated uncertainty* via a Bayesian approach
 - ▶ Dependence structure parameters
 - ▶ Spatio-temporal process priors
 - ▶ Observation models; Multiple *observation sources*, with complex error *uncertainty structure*
- ▶ Goals:
 - ▶ a *best estimate*,
 - ▶ a *collection of samples*, and
 - ▶ more precise (and accurate) *uncertainty estimates*.

Example model: atérn driven heat equation on the sphere

The iterated heat equation is a simple non-separable space-time SPDE family:

$$(\kappa^2 - \Delta)^{\gamma/2} \left[\phi \frac{\partial}{\partial t} + (\kappa^2 - \Delta)^{\alpha/2} \right]^{\beta} x(\mathbf{s}, t) = \mathcal{W}(\mathbf{s}, t) / \tau$$

For constant parameters, $x(\mathbf{s}, t)$ has spatial Matérn covariance (for each t).

Discrete domain Gaussian Markov random fields (GMRFs)

$\mathbf{x} = (x_1, \dots, x_n) \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}^{-1})$ is Markov with respect to a neighbourhood structure $\{\mathcal{N}_i, i = 1, \dots, n\}$ if $Q_{ij} = 0$ whenever $j \notin \mathcal{N}_i \cup i$.

- ▶ Project the SPDE solution space onto local basis functions: random Markov dependent basis weights (Lindgren et al, 2011).

A finite element approximation has structure

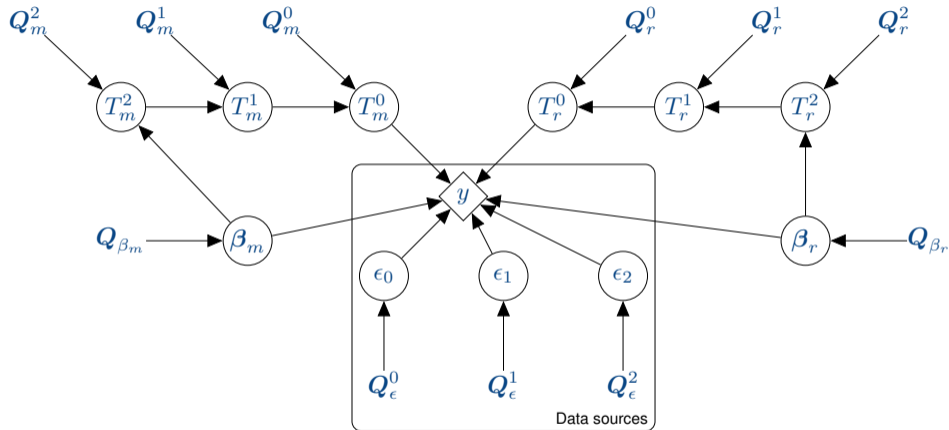
$$x(\mathbf{s}, t) = \sum_{i,j} \psi_i^{[s]}(\mathbf{s}) \psi_j^{[t]}(t) x_{ij}, \quad \mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{-1}), \quad \mathbf{Q} = \sum_{k=0}^{\alpha+\beta+\gamma} \mathbf{M}_k^{[t]} \otimes \mathbf{M}_k^{[s]}$$

even, e.g., if the spatial scale parameter κ is spatially varying.



Partial hierarchical representation

Observations of *mean, max, min*. Model *mean and range*.



Conditional specifications, e.g.

$$(T_m^0 | T_m^1, Q_m^0) \sim \mathcal{N}(T_m^1, Q_m^0)^{-1}$$

$$T_r^0 = \exp(T_r^1) G^{-1}[U_r^0(\mathbf{s}, t)], \quad U_r^0 \sim \mathcal{N}(\mathbf{0}, Q_r^0)^{-1}$$

Standardised observation uncertainty models

- ▶ Each data source may have complicated dependence structure
- ▶ To facilitate information blending, use a common error term structure

Common satellite derived data error model framework

The observational&calibration errors are modelled as three error components:

- ▶ independent (ϵ_0),
- ▶ spatially and/or temporally correlated (ϵ_1), and
- ▶ systematic (ϵ_2),

with distributions determined by the uncertainty information from satellite calibration models.

$$\text{E.g., } y_i = T_m(\mathbf{s}_i, t_i) + \epsilon_0(\mathbf{s}_i, t_i) + \epsilon_1(\mathbf{s}_i, t_i) + \epsilon_2(\mathbf{s}_i, t_i)$$

In practice, each data source might have several different components of each type; independent components can be merged, but not necessarily correlated or systematic components.

Station observation & homogenisation model

Daily means

For station k at day t_i ,

$$y_m^{k,i} = T_m(\mathbf{s}_k, t_i) + \sum_{j=1}^{J_k} H_j^k(t_i) e_m^{k,j} + \epsilon_m^{k,i},$$

where $H_j^k(t)$ are temporal step functions, $e_m^{k,j}$ are latent bias variables, and $\epsilon_m^{k,i}$ are independent measurement and discretisation errors.

Daily mean/max/min

For station k at day t_i , $y_m^{k,i} = T_m(\mathbf{s}_k, t_i) + \tilde{H}_m^k(t_i) + \epsilon_m^{k,i}$,

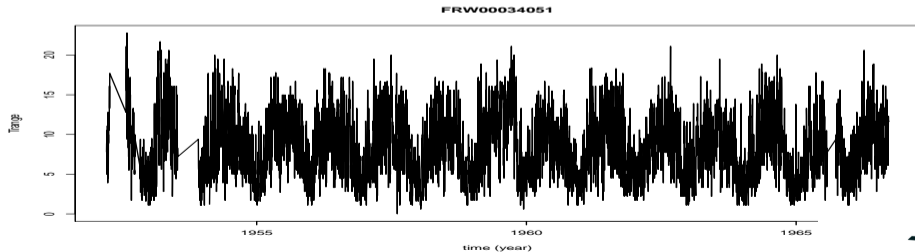
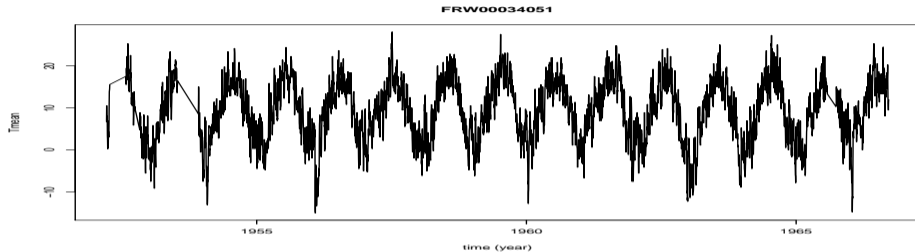
$$y_x^{k,i} = T_m(\mathbf{s}_k, t_i) + \tilde{H}_{r,m}^k(t_i) + \frac{\tilde{H}_{r,r}^k(t_i)}{2} T_r(\mathbf{s}_k, t_i) + \epsilon_x^{k,i},$$

$$y_n^{k,i} = T_m(\mathbf{s}_k, t_i) + \tilde{H}_{r,m}^k(t_i) - \frac{\tilde{H}_{r,r}^k(t_i)}{2} T_r(\mathbf{s}_k, t_i) + \epsilon_n^{k,i},$$


where \tilde{H}_\cdot are the total bias correction variables for each observation.

Observed data

Observed daily T_{mean} and T_{range} for station FRW00034051





Multiscale model component samples







Combined model samples for T_m and T_r

(Proof of concept; no actual data was involved in this figure)



Estimates of median & scale for T_m and T_r



February climatology

(Preliminary estimates, using only in-situ land station data)

Linearised inference

All Spatio-temporal latent random processes combined into $\mathbf{x} = (\mathbf{u}, \boldsymbol{\beta}, \mathbf{b})$, with joint expectation $\boldsymbol{\mu}_x$ and precision \mathbf{Q}_x :

$$(\mathbf{x} \mid \boldsymbol{\theta}) \sim \mathcal{N}(\boldsymbol{\mu}_x, \mathbf{Q}_x^{-1}) \quad (\text{Prior})$$

$$(\mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta}) \sim \mathcal{N}(h(\mathbf{x}), \mathbf{Q}_{y|x}^{-1}) \quad (\text{Observations})$$

$$p(\mathbf{x} \mid \mathbf{y}, \boldsymbol{\theta}) \propto p(\mathbf{x} \mid \boldsymbol{\theta}) p(\mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta}) \quad (\text{Conditional posterior})$$

Non-linear and/or non-Gaussian observations

For a non-linear $h(\mathbf{x})$ with Jacobian \mathbf{J} at $\mathbf{x} = \tilde{\boldsymbol{\mu}}$, iterate:

$$(\mathbf{x} \mid \mathbf{y}, \boldsymbol{\theta}) \stackrel{\text{approx}}{\sim} \mathcal{N}(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{Q}}^{-1}) \quad (\text{Approximate conditional posterior})$$

$$\tilde{\mathbf{Q}} = \mathbf{Q}_x + \mathbf{J}^\top \mathbf{Q}_{y|x} \mathbf{J}$$

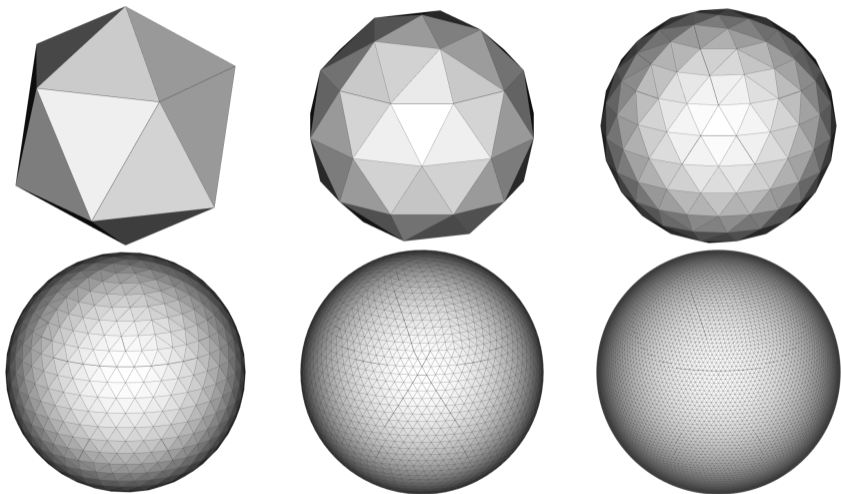
$$\tilde{\boldsymbol{\mu}}' = \tilde{\boldsymbol{\mu}} + a \tilde{\mathbf{Q}}^{-1} \left\{ \mathbf{J}^\top \mathbf{Q}_{y|x} [\mathbf{y} - h(\tilde{\boldsymbol{\mu}})] - \mathbf{Q}_x (\tilde{\boldsymbol{\mu}} - \boldsymbol{\mu}_x) \right\}$$

for some $a > 0$ chosen by line-search.

Iterative solutions for $\sim 10^{11}$ latent variables

- ▶ Full non-linear models solver components:
 - ▶ Nonlinear Newton iteration with robust line-search
 - ▶ Preconditioned conjugate gradient (PCG) iteration for $Q(\mu - \hat{\mu}) = r = b - Q\hat{\mu}$
 - ▶ Local and multiscale approximations for preconditioning: $M^{-1}Q \approx I$
 - ▶ Sampling with PCG: $Q(x - \hat{\mu}) = Lw$
Requires only a rectangular pseudo-Cholesky factorisation $LL^T = Q$.
Possible due to the kronecker product sum precision structure.
- ▶ Pragmatic implementation:
 - ▶ Daily mean temperature only
 - ▶ $\sim 60,000$ conditionally independent days (on the fine temporal scale):
embarrassingly parallel daily direct solves
 - ▶ Multiscale component grouped into three superblocks
 - ▶ Reduced spatial resolution

Triangulations for all corners of Earth



MULTI-SCALE ANALYSIS MODEL

Statistical model for temperature variations and different scales (space and time):

- **Climatological variation:** local seasonal cycle with effects of latitude, altitude and coastal influence.
- **Large-scale variation:** Slowly varying climatological mean temperature field. Station homogenisation.
- **Daily Local:** daily variability associated with weather. Satellite retrieval biases.

Simultaneously estimates observational biases of known bias structures:

- e.g. satellite biases, station homogenisation.

Processed on STFC's LOTUS cluster www.jasmin.ac.uk:

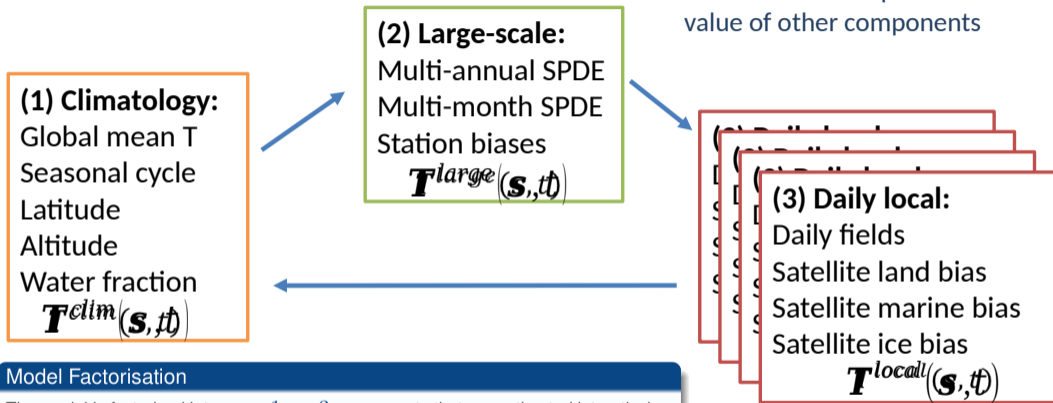
- Largest solves processed on 20 core/256GB RAM node.
- Highly parallel observation pre-processing.

| Element | Resolution | N Variables |
|-------------|---------------------|-------------|
| Seasonal | Bimonthly x 1° SPDE | 245,772 |
| Slow-scale* | 5 year x 5° SPDE | 107,604 |
| Latitude | 0.5° latitude SPDE | 721 |
| Altitude | (0.25° grid) | 1 |
| Coastal | (0.25° grid) | 1 |
| Grand mean | Analysis mean | 1 |

| Element | Resolution | N Variables |
|--------------|---------------------|-------------|
| Large-scale | 3 monthly x 5° SPDE | 1,752,408 |
| Station bias | NA | 82,072 |

| Element | Resolution | N Variables per day |
|-------------------------|--------------------------------|---------------------|
| Daily local | ~0.5 degree SPDE | 162,842 |
| Satellite bias (marine) | Global | 1 |
| Satellite bias (land) | Global + 2.5 degree SPDE | 1 + 40,962 |
| Satellite bias (ice) | Hemispheric + 2.5 degree SPDE* | 2 + 40,962 |

ITERATIVE SOLUTION



Model Factorisation

The model is factorised into $m = 1, \dots, 3$ components that are estimated iteratively, substituting \tilde{y}_m for y :

$$\tilde{y}_m = y - \sum_{n \neq m} J_n \mu_{x_n} \tilde{y}_n$$





MULTI-SCALE ANALYSIS MODEL

Statistical model for temperature variations and different scales (space and time):

- **Climatological variation**: local seasonal cycle with effects of latitude, altitude and coastal influence.
- **Large-scale variation**: Slowly varying climatological mean temperature field.
- **Daily Local**: daily variability associated with weather.


Simultaneously estimates observational biases of known bias structures:

- e.g. satellite biases, station homogenisation.



SATELLITE BIAS MODELS


- Simplified model of known error structures in satellite air temperature retrievals:
 - Global/hemispheric systematic bias covariates.
 - Daily estimates of spatially varying bias as a spatial random field.
- Estimated jointly with daily temperature variability.




ENSEMBLE ANALYSIS

- Samples drawn from joint posterior distribution of temperature and bias variables.
- Temperature model samples projected onto analysis grid.
- Spatial/temporal correlation in analysis errors is encoded into the ensemble.
- Summary statistics can be derived from the ensemble. Expected value, total uncertainty and observation constraint information also available.

EUSTACE Ensemble 04/08/2003-13/08/2003





EUSTACE Ensemble 30/07/2010-05/08/2010





ENSEMBLE ANALYSIS

- Samples drawn from joint posterior distribution of temperature and bias variables.
- Temperature model samples projected onto analysis grid.
- Spatial/temporal correlation in analysis errors is encoded into the ensemble.
- Summary statistics can be derived from the ensemble. Expected value, total uncertainty and observation constraint information also available.

EUSTACE Ensemble 01/01/2006-14/01/2006



EUSTACE Ensemble 30/07/2010-05/08/2010



Summary

Not covered in this talk:

- ▶ Pure conditional block updates risk getting stuck; need for convergence acceleration
- ▶ Overlapping space-time blocks for preconditioning
- ▶ Non-stationary random field parameter estimation
- ▶ Direct & iterative variance calculations to eliminate or reduce Monte Carlo error in the reconstruction uncertainties
- ▶ Fast approximate handling of correlated error components

Summary:

- ▶ Challenging statistical problem, in both size and complexity
- ▶ Approximate calculation techniques allows some of the complexity to be handled with reasonable computational resources
 - ▶ SPDEs and Gaussian Markov random fields
 - ▶ Fast local sparse solves
 - ▶ Global multiscale block iteration
- ▶ Close collaboration between climate scientists, statisticians, and software engineers is essential



Modelling non-Gaussian quantities

Power tail quantile (POQ) model

The quantile function $F_{\theta}^{-1}(p)$, $p \in [0, 1]$, is defined through a quantile blend of left- and right-tailed generalised Pareto distributions:

$$f_{\theta}^{-}(p) = \begin{cases} \frac{1-(2p)^{-\theta}}{2\theta}, & \theta \neq 0, \\ \frac{1}{2} \log(2p), & \theta = 0, \end{cases}$$

$$f_{\theta}^{+}(p) = -f_{\theta}^{-}(1-p) = \begin{cases} \frac{(2(1-p))^{-\theta}-1}{2\theta}, & \theta \neq 0, \\ -\frac{1}{2} \log(2(1-p)), & \theta = 0. \end{cases}$$

$$F_{\theta}^{-1}(p) = \theta_0 + \frac{\tau}{2} [(1-\gamma)f_{\theta_3}^{-}(p) + (1+\gamma)f_{\theta_4}^{+}(p)].$$



The parameters $\theta = (\theta_0, \theta_1 = \log \tau, \theta_2 = \text{logit}[(\gamma+1)/2], \theta_3, \theta_4)$ control the median, spread/scale, skewness, and the left and right tail shape.

This model is also known as the *five parameter lambda model* (Gilchrist, 2000).







Copula transformation: $G^{-1}[u(\mathbf{s}, t)] = F_{\theta(\mathbf{s}, t)}^{-1}\{\Phi[u(\mathbf{s}, t)]\}$

Diurnal range distributions

RSM00025594 (BUHTA PROVIDENJA)



SP000060040 (LANZAROTE/AEROPUERT)





For these stations, POQ does a slightly better job than a Gamma distribution.









Diurnal range distributions

ASN00005008 (MARDIE)



ASN00023738 (MYPONGA)



For these stations only POQ comes close to representing the distributions.

Note: Some shapes may be due to unmodeled station inhomogeneities.




Overlapping blocks and multigrid

Overlapping block preconditioning

Let D_k^\top be a restriction matrix to subdomain Ω_k , and let W_k be a diagonal weight matrix. Then an additive Schwarz preconditioner is

$$M^{-1}x = \sum_{k=1}^K W_k D_k (D_k^\top Q D_k)^{-1} D_k^\top W_k x$$

Multigrid and/or approximate multiscale Schur complements



The hierarchy of scales and preconditioning ($\mathbf{x}_0 = \mathbf{B}\mathbf{x}_1 + \text{fine scale variability}$):

Multiscale Schur complement approximation

Solving $\mathbf{Q}_{x|y}\mathbf{x} = \mathbf{b}$ can be formulated using two solves with the upper (fine) block $\mathbf{Q}_0 + \mathbf{A}^\top \mathbf{Q}_\epsilon \mathbf{A}$, and one solve with the *Schur complement*

$$\mathbf{Q}_1 + \mathbf{B}^\top \mathbf{Q}_0 \mathbf{B} - \mathbf{B}^\top \mathbf{Q}_0 \left(\mathbf{Q}_0 + \mathbf{A}^\top \mathbf{Q}_\epsilon \mathbf{A} \right)^{-1} \mathbf{Q}_0$$

By mapping the fine scale model onto the coarse basis used for the coarse model, we get an *approximate* (and sparse) Schur solve via

$$\begin{bmatrix} \tilde{\mathbf{Q}}_B + \mathbf{B}^\top \mathbf{A}^\top \mathbf{Q}_\epsilon \mathbf{A} \mathbf{B} & -\tilde{\mathbf{Q}}_B \\ -\tilde{\mathbf{Q}}_B & \mathbf{Q}_1 + \tilde{\mathbf{Q}}_B \end{bmatrix} \begin{bmatrix} \text{ignored} \\ \mathbf{x}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \tilde{\mathbf{b}} \end{bmatrix}$$

where $\tilde{\mathbf{Q}}_B = \mathbf{B}^\top \mathbf{Q}_0 \mathbf{B}$.

The block matrix can be interpreted as the precision of a bivariate field on a common, coarse spatio-temporal scale, and the same technique applied to this system, with $\mathbf{x}_{1,1} = \mathbf{B}_{1|2}\mathbf{x}_{1,2} + \text{finer scale variability}$.



Also applies to the station data bias homogenisation coefficients.



Variance calculations

Sparse partial inverse: Takahashi recursions postprocesses Cholesky

Takahashi recursions compute \mathbf{S} such that $\mathbf{S}_{ij} = (\mathbf{Q}^{-1})_{ij}$ for all $Q_{ij} \neq 0$. Postprocessing of the (sparse) Cholesky factor.

Basic Rao-Blackwellisation of sample estimators

Let $\mathbf{x}^{(j)}$ be samples from a Gaussian posterior and let $\mathbf{a}^\top \mathbf{x}$ be a linear combination of interest. Then, for any subdomain $\Omega_k \subset \Omega$,

$$\mathbb{E}(\mathbf{a}^\top \mathbf{x}) = \mathbb{E} [\mathbb{E}(\mathbf{a}^\top \mathbf{x} \mid \mathbf{x}_{\Omega_k^*})] \approx \frac{1}{J} \sum_{j=1}^J \mathbb{E}(\mathbf{a}^\top \mathbf{x} \mid \mathbf{x}_{\Omega_k^*}^{(j)})$$

$$\begin{aligned} \text{Var}(\mathbf{a}^\top \mathbf{x}) &= \mathbb{E} [\text{Var}(\mathbf{a}^\top \mathbf{x} \mid \mathbf{x}_{\Omega_k^*})] + \text{Var} [\mathbb{E}(\mathbf{a}^\top \mathbf{x} \mid \mathbf{x}_{\Omega_k^*})] \\ &\approx \text{Var}(\mathbf{a}^\top \mathbf{x} \mid \mathbf{x}_{\Omega_k^*}^j) + \frac{1}{J} \sum_{j=1}^J \left[\mathbb{E}(\mathbf{a}^\top \mathbf{x} \mid \mathbf{x}_{\Omega_k^*}^{(j)}) - \mathbb{E}(\mathbf{a}^\top \mathbf{x}) \right]^2 \end{aligned}$$



Efficient if $\mathbf{a}\mathbf{a}^\top$ sparsity matches \mathbf{S} for each subdomain.

