Stability and localization in the ensemble square root filter

Sarah Dance[1], Ruth Petrie[1], Nancy Nichols[1], Mark Butala[2]

1. University of Reading, UK
2. University of Illinois, Urbana-Champaign, USA

Adjoint workshop 2009

University of Reading

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Outline

1. What is an ensemble square root filter?
2. What are the known problems with under-sampling?
3. What implications does under-sampling have for filter stability?
4. Localization and the Schur product
5. The Khatri-Rao product
6. Conclusions

The ensemble filter idea

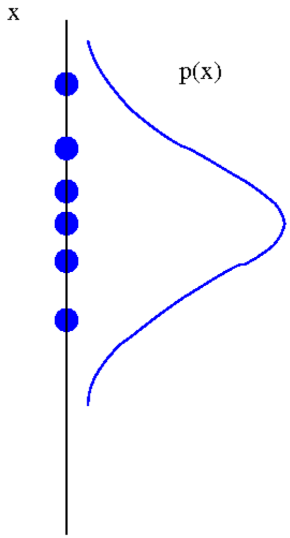

The ensemble filter idea

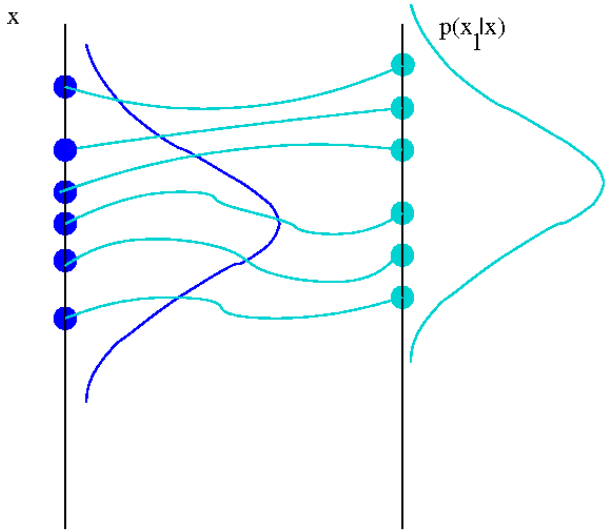

The ensemble filter idea

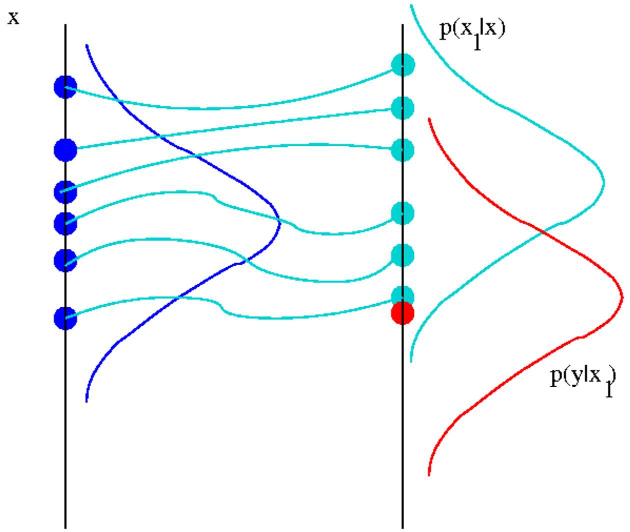

The ensemble filter idea

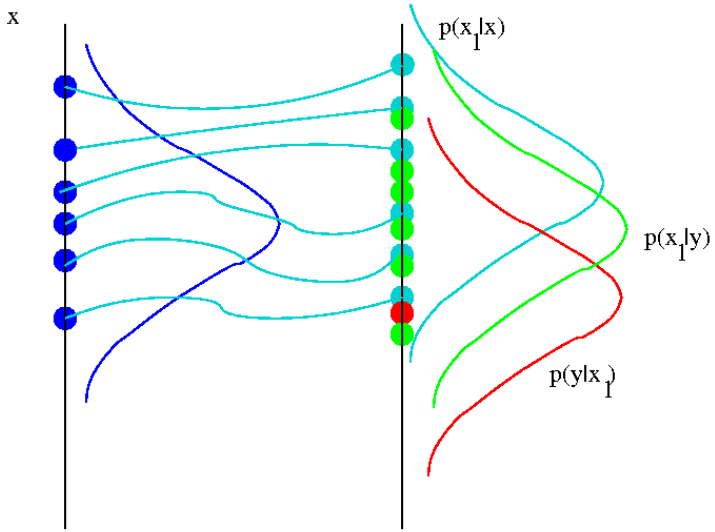

The ensemble square root filter (Tippett et al. 2003)

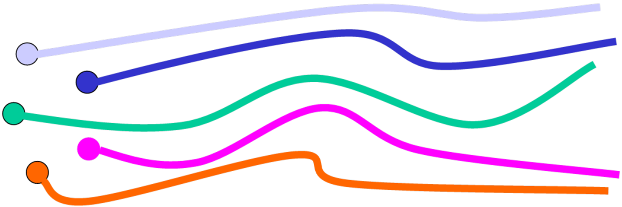

Ensemble : $\{\mathbf{x}_i : \mathbf{x}_i \in \mathbb{R}^n\}_{i=1,\ldots,m}$

The ensemble square root filter (Tippett et al. 2003)

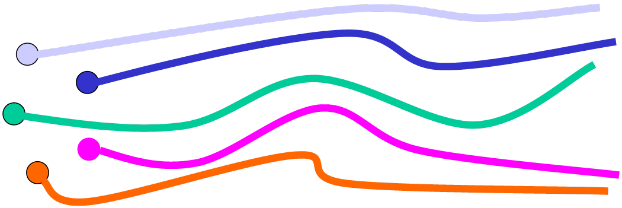

Ensemble : $\{\mathbf{x}_i : \mathbf{x}_i \in \mathbb{R}^n\}_{i=1,\dots,m}$

Ensemble mean: $\bar{\mathbf{x}} = \frac{1}{m}\sum_{i=1}^{m} \mathbf{x}_i$

The ensemble square root filter (Tippett et al. 2003)

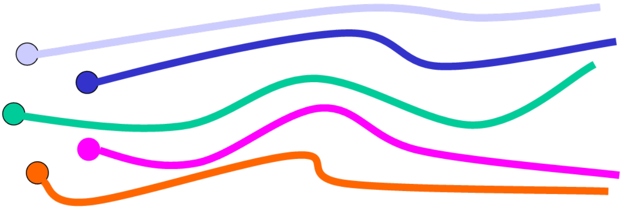

Ensemble : $\{\mathbf{x}_i : \mathbf{x}_i \in \mathbb{R}^n\}_{i=1,\ldots,m}$

Ensemble mean: $\bar{\mathbf{x}} = \frac{1}{m}\sum_{i=1}^{m} \mathbf{x}_i$

Perturbation matrix: $\mathbf{X} = \frac{1}{\sqrt{m-1}} \begin{pmatrix} \mathbf{x}_1 - \bar{\mathbf{x}} & \mathbf{x}_2 - \bar{\mathbf{x}} & \ldots & \mathbf{x}_m - \bar{\mathbf{x}} \end{pmatrix}.$

size $n \times m$, low rank

The ensemble square root filter (Tippett et al. 2003)

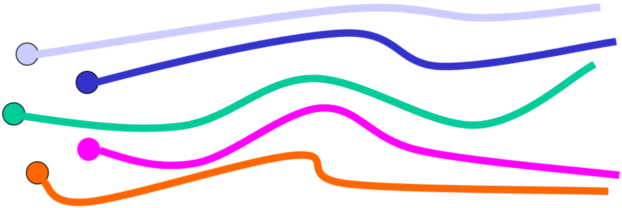

Ensemble : $\{\mathbf{x}_i : \mathbf{x}_i \in \mathbb{R}^n\}_{i=1,\ldots,m}$

Ensemble mean: $\bar{\mathbf{x}} = \frac{1}{m}\sum_{i=1}^{m} \mathbf{x}_i$

Perturbation matrix: $\mathbf{X} = \frac{1}{\sqrt{m-1}} \left(\mathbf{x}_1 - \bar{\mathbf{x}} \quad \mathbf{x}_2 - \bar{\mathbf{x}} \quad \ldots \quad \mathbf{x}_m - \bar{\mathbf{x}} \right).$

Covariance: $\mathbf{P} = \mathbf{X}\mathbf{X}^T = \frac{1}{m-1}\sum_{i=1}^{m}(\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T.$

size $n \times n$, low rank

Ensemble forecast

Each member is forecast using a nonlinear model

$$\mathbf{x}(t_{k+1}) = \mathbf{f}(\mathbf{x}(t_k)),$$

to the time of an observation $\mathbf{y}_k \in \mathbb{R}^p$.

$$\mathbf{y}_k = H(\mathbf{x}(t_k)) + \epsilon_k,$$

where

- H is an observation operator such that $H : \mathbb{R}^n \to \mathbb{R}^p$
- ϵ_k is a stochastic variable with mean zero and covariance $\mathbf{R}_k$.

Observation update

Update mean: $\overline{\mathbf{x}^a} = \overline{\mathbf{x}^f} + \mathbf{K}(\mathbf{y} - \overline{\mathbf{y}^f})$

Update perturbations: $\mathbf{X}^a = \mathbf{X}^f \mathbf{T}$

where

Forecast obs ensemble $\mathbf{y}_i^f = H(\mathbf{x}_i^f)$

Perturbation matrix $\mathbf{Y}^f = \frac{1}{\sqrt{m-1}}\left(\mathbf{y}_i^f - \overline{\mathbf{y}^f}\right)$.

as columns

Gain matrix

$$\mathbf{K} = \mathbf{X}^f (\mathbf{Y}^f)^T \mathbf{D}^{-1}$$
$$\mathbf{D} = \mathbf{Y}^f (\mathbf{Y}^f)^T + \mathbf{R}.$$
$$\mathbf{T}\mathbf{T}^T = \mathbf{I} - (\mathbf{Y}^f)^T \mathbf{D}^{-1} \mathbf{Y}^f.$$

We never have to compute $\mathbf{P} = \mathbf{X}\mathbf{X}^T$!

Known problems with under-sampling (1)

- Under-estimation of variance (e.g. Anderson and Anderson, 1999; Hamill et al., 2001; Furrer and Bengtsson, 2007)
- Background is too strongly weighted compared with the observation.

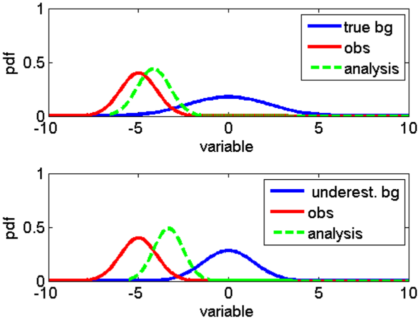

- Covariance inflation

Known problems with under-sampling (2)

Spurious long range correlations

$$\overline{\mathbf{x}^a} = \overline{\mathbf{x}^f} + \mathbf{P}^f\mathbf{H}^T(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}(\mathbf{y} - \overline{\mathbf{y}^f})$$

$$= \overline{\mathbf{x}^f} +$$

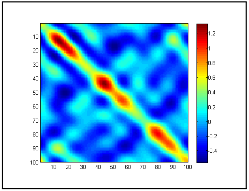

$$\mathbf{H}^T(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}(\mathbf{y} - \overline{\mathbf{y}^f})$$

$\Rightarrow$ unphysical analysis increments far from the location of the observation

- Covariance localization

Feedback form

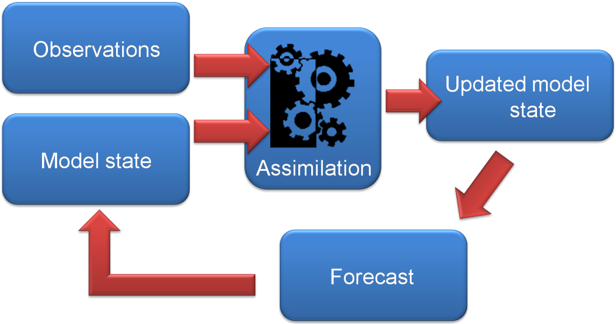

Combine the forecast and observation update steps to write the filter in feedback form:

$$\overline{\mathbf{x}^a(t_{k+1})} = \overline{\mathbf{f}(\mathbf{x}^a(t_k))} + \mathbf{K}_{k+1}(\mathbf{y}_{k+1} - \overline{H(\mathbf{f}(\mathbf{x}^a(t_k)))}).$$

Error equation

Let $\mathbf{e}_i(t_k) = \mathbf{x}_i^a(t_k) - \mathbf{x}^t(t_k)$, i.e., the error in i-th analysis ensemble member at time t_k.

Seek an approximate error evolution equation by linearizing about the true state . . .

Error equation

Let $\mathbf{e}_i(t_k) = \mathbf{x}_i^a(t_k) - \mathbf{x}^t(t_k)$, i.e., the error in i-th analysis ensemble member at time t_k.

Seek an approximate error evolution equation by linearizing about the true state ...

Error equation

$$\overline{\mathbf{e}(t_{k+1})} = (\mathbf{I} - \mathbf{K}_{k+1}\mathbf{H})\,\mathbf{F}\overline{\mathbf{e}(t_k)} + \text{higher order terms.}$$

Here $\mathbf{F}$ and $\mathbf{H}$ are the Jacobians of $\mathbf{f}$ and H respectively, evaluated at $\mathbf{x}^t$

- $\overline{\mathbf{x}_k^a}$ will converge to the true state if $|\overline{\mathbf{e}_k}| \to 0$ as $t_k \to \infty$.
- We expect that the eigenvalues of $(\mathbf{I} - \mathbf{K}_{k+1}\mathbf{H})\,\mathbf{F}$ lie within the unit circle.

What can go wrong? An example

Suppose $\mathbf{F} = \alpha\mathbf{I}$ and $p = n$, $\mathbf{H} = \mathbf{I}$.

The gain matrix becomes

$$\mathbf{K} = \mathbf{X}^f(\mathbf{X}^f)^T\mathbf{D}^{-1},$$

an $n \times n$, square matrix with rank at most $m - 1$.

What can go wrong? An example

Suppose $\mathbf{F} = \alpha\mathbf{I}$ and $p = n$, $\mathbf{H} = \mathbf{I}$.

The gain matrix becomes

$$\mathbf{K} = \mathbf{X}^f(\mathbf{X}^f)^T\mathbf{D}^{-1},$$

an $n \times n$, square matrix with rank at most $m - 1$.

Writing $\mathbf{K}$ in its Jordan normal form,

$$\mathbf{K} = \mathbf{E}\begin{pmatrix} \mathbf{J}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}\mathbf{E}^{-1},$$

we have

$$(\mathbf{I} - \mathbf{KH})\,\mathbf{F} = \alpha\mathbf{E}\begin{pmatrix} \mathbf{I} - \mathbf{J}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}\mathbf{E}^{-1}.$$

Thus we have a set of eigenvalues equal to α that do not lie within the unit circle. Errors in the analysis will grow over time!

What can we do about it?

- The problem in this example arose because $\mathbf{P}_k^f$ was not full rank.
- Clearly it would be desirable if we could modify the algorithms so that the approximation to this matrix is full rank.
- Although it would not necessarily guarantee that the algorithms are not unstable for some other reason!

The Schur product

The Schur product (Schur, 1911) is defined as an elementwise product between two matrices of the same size, thus, if $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{k \times l}$ then the i, j-th element of their Schur product may be written

$$[\mathbf{A} \circ \mathbf{B}]_{ij} = [\mathbf{A}]_{ij}[\mathbf{B}]_{ij}.$$

Example

$$\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \circ \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix} = \begin{pmatrix} 1 \times 5 & 2 \times 6 \\ 3 \times 7 & 4 \times 8 \end{pmatrix}.$$

Properties of the Schur product (1)

Theorem (Schur's theorems)

1. *If* $\mathbf{A}$ *is strictly positive definite and* $\mathbf{B}$ *is positive semi-definite with all its main diagonal diagonal entries positive, then* $\mathbf{A} \circ \mathbf{B}$ *is postive definite.*

Properties of the Schur product (1)

Theorem (Schur's theorems)

1. *If* $\mathbf{A}$ *is strictly positive definite and* $\mathbf{B}$ *is positive semi-definite with all its main diagonal diagonal entries positive, then* $\mathbf{A} \circ \mathbf{B}$ *is postive definite.*
2. *The Schur product of two covariance matrices is a covariance matrix.*

Properties of the Schur product (1)

Theorem (Schur's theorems)

1. *If* $\mathbf{A}$ *is strictly positive definite and* $\mathbf{B}$ *is positive semi-definite with all its main diagonal diagonal entries positive, then* $\mathbf{A} \circ \mathbf{B}$ *is postive definite.*
2. *The Schur product of two covariance matrices is a covariance matrix.*

Now $\mathbf{P} = \mathbf{X}\mathbf{X}^T$ so, for non-degenerate cases, its main diagonal entries will be strictly positive.

Properties of the Schur product (1)

Theorem (Schur's theorems)

1. *If* $\mathbf{A}$ *is strictly positive definite and* $\mathbf{B}$ *is positive semi-definite with all its main diagonal diagonal entries positive, then* $\mathbf{A} \circ \mathbf{B}$ *is postive definite.*
2. *The Schur product of two covariance matrices is a covariance matrix.*

Now $\mathbf{P} = \mathbf{X}\mathbf{X}^T$ so, for non-degenerate cases, its main diagonal entries will be strictly positive.

Hence, if we choose a positive definite covariance matrix ρ,

$$\rho \circ \mathbf{P}$$

is a full rank covariance matrix!

Localization

If we pick ρ to be the matrix to be a positive definite banded covariance matrix, then we can also use ρ to remove spurious correlations (e.g. Hamill et al, 2001).

$$\rho \circ \mathbf{P} \quad =$$

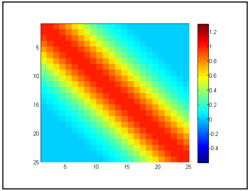

$\circ$

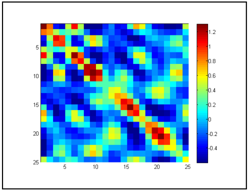

$=$

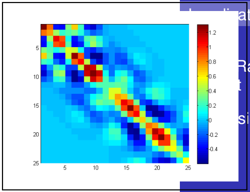

Localization in a square root filter

- For a SRF, it is not obvious how to apply the localization technique consistently using the Schur product.
- Firstly, the calculation of $\rho \circ \mathbf{P}$ in the Kalman gain requires the calculation of $\mathbf{P} = \mathbf{X}\mathbf{X}^T$
- How can you compute a consistent ensemble perturbation update?

The Khatri-Rao product (Rao and Rao, 1998)

The Khatri-Rao product provides us with a factorization of the Schur product equation:

Theorem (Factorization of the Schur product)

If $\mathbf{P}$ *is an* $n \times n$ *matrix such that* $\mathbf{P} = \mathbf{X}\mathbf{X}^T$ *with* $\mathbf{X}$ *of size* $n \times m$*, and* ρ *is also an* $n \times n$ *matrix, such that* $\rho = \mathbf{C}\mathbf{C}^T$ *with* $\mathbf{C}$ *of size* $n \times k$*, then*

$$\rho \circ \mathbf{P} = (\mathbf{C}^T \odot \mathbf{X}^T)^T (\mathbf{C}^T \odot \mathbf{X}^T).$$

But what is $\odot$?

The Khatri-Rao product

Let **A** and **B** be a matrices with r columns. Let α_i be the i-th column of **A** and β_i be the i-th column of **B**. The Khatri-Rao product of **A** and **B** is the partitioned matrix

$$\mathbf{A} \odot \mathbf{B} = \left(\alpha_1 \otimes \beta_1 | \alpha_2 \otimes \beta_2 | \ldots | \alpha_r \otimes \beta_r\right),$$

where $\otimes$ indicates the well-known Kronecker product.

Example

$$\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} \odot \begin{pmatrix} 5 & 6 \\ 7 & 8 \\ 9 & 10 \end{pmatrix} = \left(\begin{array}{cc} 1 \times \begin{bmatrix} 5 \\ 7 \\ 9 \end{bmatrix} & 2 \times \begin{bmatrix} 6 \\ 8 \\ 10 \end{bmatrix} \\ 3 \times \begin{bmatrix} 5 \\ 7 \\ 9 \end{bmatrix} & 4 \times \begin{bmatrix} 6 \\ 8 \\ 10 \end{bmatrix} \end{array} \right).$$

Applying the product to the SRF

- We would like to replace $\mathbf{X}$ with $(\mathbf{C}^T \odot \mathbf{X}^T)^T$ in the SRF algorithm.
- But $\mathbf{X}$ is $n \times m$ and $(\mathbf{C}^T \odot \mathbf{X}^T)^T$ is $n \times mn$, so we need to reduce the size of this matrix again at the end of the analysis step
- Bishop and Hodyss (2009) have done something similar for ETKF, using a different approach to reduce the ensemble size.
- Future work to compare these approaches, and apply the K-R product to different SRFs

Conclusions

- The ensemble SRF suffers from problems with under-sampling
- It is well known that this causes under-estimation of variance and spurious correlations
- New examples show that low rank may lead to filter divergence
- Rank problem can be corrected using Schur product, or Khatri-Rao product for an SRF
- We are working on implementing the K-R product for a number of filters