On the efficiency and consitency of covariance localisation in the EnKF

Alban Farchi and Marc Bocquet
CEREA, laboratoire commun École des Ponts ParisTech and EDF R&D, Université Paris-Est, Champs-sur-Marne, France

Tuesday, 4 June 2018

International EnKF workshop

Contents

The ensemble Kalman filter analysis

- We focus on the Kalman filter analysis step:

$$\mathbf{K} = \mathbf{B}\mathbf{H}^{\mathrm{T}} \left(\mathbf{R} + \mathbf{H}\mathbf{B}\mathbf{H}^{\mathrm{T}}\right)^{-1}, \tag{1}$$

$$\mathbf{x}_{\mathrm{a}} = \mathbf{x}_{\mathrm{b}} + \mathbf{K}\left(\mathbf{y} - \mathbf{H}\mathbf{x}_{\mathrm{b}}\right), \tag{2}$$

$$\mathbf{P}_{\mathrm{a}} = \left(\mathbf{I} - \mathbf{K}\mathbf{H}\right)\mathbf{B}, \tag{3}$$

- In the EnKF, the statistics are carried through by the *ensemble* $\left\{\mathbf{x}^i,\ i = 1 \dots N_{\mathrm{e}}\right\}$:

$$\mathbf{x}_{\mathrm{b}} = \overline{\mathbf{x}}, \tag{4}$$

$$\mathbf{B} = \mathbf{X}\mathbf{X}^{\mathrm{T}}, \tag{5}$$

where $\overline{\mathbf{x}}$ is the *ensemble mean* and $\mathbf{X}$ is the *normalised anomaly matrix*.

Rank deficiency of the EnKF

- The matrix $\mathbf{X}\mathbf{X}^{\mathrm{T}}$ has rank limited by $N_{\mathrm{e}} - 1$, too small to accurately represent $\mathbf{B}$ in a high-dimensional system ($N_{\mathrm{x}} \gg N_{\mathrm{e}}$).
- When N_{e} is too small, $\mathbf{X}\mathbf{X}^{\mathrm{T}}$ is characterised by large sampling errors, which often take the form of *spurious correlations* at long distance.
- A common solution is to use localisation: either make the analysis local (domain localisation) or use a localised $\mathbf{B}$ (covariance localisation).

domain localisation	**covariance localisation**
relies on a collection of local analyses	relies on a localised background error $\mathbf{B}$
embarrassingly parallel	no obvious parallelisation of the perturbation update
ad hoc treatment of non-local observations	rigourous treatment of non-local observations

Covariance localisation in the deterministic EnKF

The (exact) LEnSRF

Regularise $\mathbf{X}\mathbf{X}^{\mathrm{T}}$ with a localisation matrix $\boldsymbol{\rho}$:

$$\mathbf{B} = \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^{\mathrm{T}}\right). \tag{6}$$

Update the perturbation as:

$$\mathbf{T} = \mathbf{I} + \mathbf{B}\mathbf{H}^{\mathrm{T}}\mathbf{R}^{-1}\mathbf{H}, \tag{7}$$

$$\mathbf{X}_{\mathrm{a}} = \mathbf{T}^{-1/2}\mathbf{X}. \tag{8}$$

- First focus: *efficient implementation* (accuracy/speed) of the perturbation update.
- Second focus: *consistency* (how well does $\mathbf{X}_{\mathrm{a}}\mathbf{X}_{\mathrm{a}}^{\mathrm{T}}$ match $\mathbf{P}_{\mathrm{a}}$) of the perturbation update.

Contents

Using an augmented ensemble to represent $\mathbf{B}$

- In the analysis step, we choose to compute $\widehat{N}_{\mathrm{e}} > N_{\mathrm{e}}$ perturbations $\widehat{\mathbf{X}}$ such that $\mathbf{B} = \boldsymbol{\rho} \circ \mathbf{X}\mathbf{X}^{\mathrm{T}} \approx \widehat{\mathbf{X}}\widehat{\mathbf{X}}^{\mathrm{T}}$. There are two methods.

The modulation method

Suppose that there is a matrix $\mathbf{W}$ with N_{m} columns such that $\boldsymbol{\rho} \approx \mathbf{W}\mathbf{W}^{\mathrm{T}}$. Let $\widehat{\mathbf{X}}$ be the matrix with $N_{\mathrm{m}} N_{\mathrm{e}}$ columns:

$$\left[\widehat{\mathbf{X}}\right]_n^{jN_{\mathrm{e}}+i} = [\mathbf{W}]_n^j \, [\mathbf{X}]_n^i \, . \tag{9}$$

The random svd method

Compute the truncated svd $\mathbf{B} = \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^{\mathrm{T}}\right) \approx \mathbf{U}\boldsymbol{\Sigma}\mathbf{U}^{\mathrm{T}}$ with N_{m} columns, and form the matrix $\widehat{\mathbf{X}} = \mathbf{U}\boldsymbol{\Sigma}^{1/2}$.

The modulation method

1: Compute the modulation product $\widehat{\mathbf{X}} = \mathbf{W}\boldsymbol{\Delta}\mathbf{X}$

$$\left[\widehat{\mathbf{X}}\right]_n^{jN_\mathrm{e}+i} = [\mathbf{W}]_n^j \, [\mathbf{X}]_n^i \,. \tag{10}$$

- This is a mix between a Schur product (for the state variable index n) and a tensor product (for the ensemble indices i and j). [Buehner, 2005]
- The modulation product is based on a factorisation property shown by [Lorenc, 2003] and is currently used for covariance localisation [e.g., Bishop et al., 2017], including in operational centers [e.g., Arbogast et al., 2017].

The random svd method

1: Compute the truncated svd $\mathbf{B} = \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^{\mathrm{T}}\right) \approx \mathbf{U}\boldsymbol{\Sigma}\mathbf{U}^{\mathrm{T}}$ with N_{m} columns.
2: Form the matrix $\widehat{\mathbf{X}} = \mathbf{U}\boldsymbol{\Sigma}^{1/2}$.

- The matrix $\mathbf{B}$ is sparse, which means that efficient (and parallelisable) methods can be used to compute the truncated svd, e.g., the *random svd* algorithm of [Halko et al., 2011].

[Farchi and Bocquet, 2019]

Compute the updated perturbations

- How to obtain N_e updated members using the $\widehat{N}_\mathrm{e}$ analysis perturbations $\widehat{\mathbf{X}}_\mathrm{a}$ of the augmented ensemble?
- A solution is to use the augmented ensemble to compute an approximate *left-transform update* of the (non-augmented) ensemble as follows:

$$\mathbf{X}_\mathrm{a} = \left(\mathbf{I} + \mathbf{B}\mathbf{H}^\mathrm{T}\mathbf{R}^{-1}\mathbf{H}\right)^{-1/2}\mathbf{X}, \tag{11}$$

$$\mathbf{X}_\mathrm{a} = \left(\mathbf{I} + \widehat{\mathbf{X}}\widehat{\mathbf{X}}^\mathrm{T}\mathbf{H}^\mathrm{T}\mathbf{R}^{-1}\mathbf{H}\right)^{-1/2}\mathbf{X}, \tag{12}$$

$$\mathbf{X}_\mathrm{a} = \left\{\mathbf{I} - \widehat{\mathbf{X}}\left(\mathbf{I} + \widehat{\mathbf{Y}}^\mathrm{T}\mathbf{R}^{-1}\widehat{\mathbf{Y}} + \left(\mathbf{I} + \widehat{\mathbf{Y}}^\mathrm{T}\mathbf{R}^{-1}\widehat{\mathbf{Y}}\right)^{1/2}\right)^{-1}\widehat{\mathbf{Y}}^\mathrm{T}\mathbf{R}^{-1}\widehat{\mathbf{H}}\right\}\mathbf{X}, \tag{13}$$

where $\widehat{\mathbf{Y}} = \mathbf{H}\widehat{\mathbf{X}}$.

- The linear algebra is performed in the augmented ensemble space (i.e., using $\widehat{N}_\mathrm{e} \times \widehat{N}_\mathrm{e}$ matrices) using the formula derived by [Bocquet, 2016], later used by [Bishop et al., 2017] under the name *gain form of the ETKF*.

The Lorenz 1996 model

- We use the *L96 model* with $N_\mathrm{x} = 400$ variables:

$$\frac{\mathrm{d}x_n}{\mathrm{d}t} = (x_{n+1} - x_{n-2})\, x_{n-1} - x_n + 8, \quad n = 1 \ldots N_\mathrm{x}. \tag{14}$$

- The observations are given every $\Delta t = 0.05$ by

$$\mathbf{y} = \mathbf{x} + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(\mathbf{0},\ \mathbf{I}). \tag{15}$$

- The localisation matrix is constructed using the *Gaspari–Cohn* function, assuming that the grid points are equally distributed in space.
- All algorithms use an ensemble of $N_\mathrm{e} = 10$ members.
- The runs are $2 \times 10^4 \Delta t$ long and our criterion is the time-average *analysis RMSE*.

Results with the L96 model

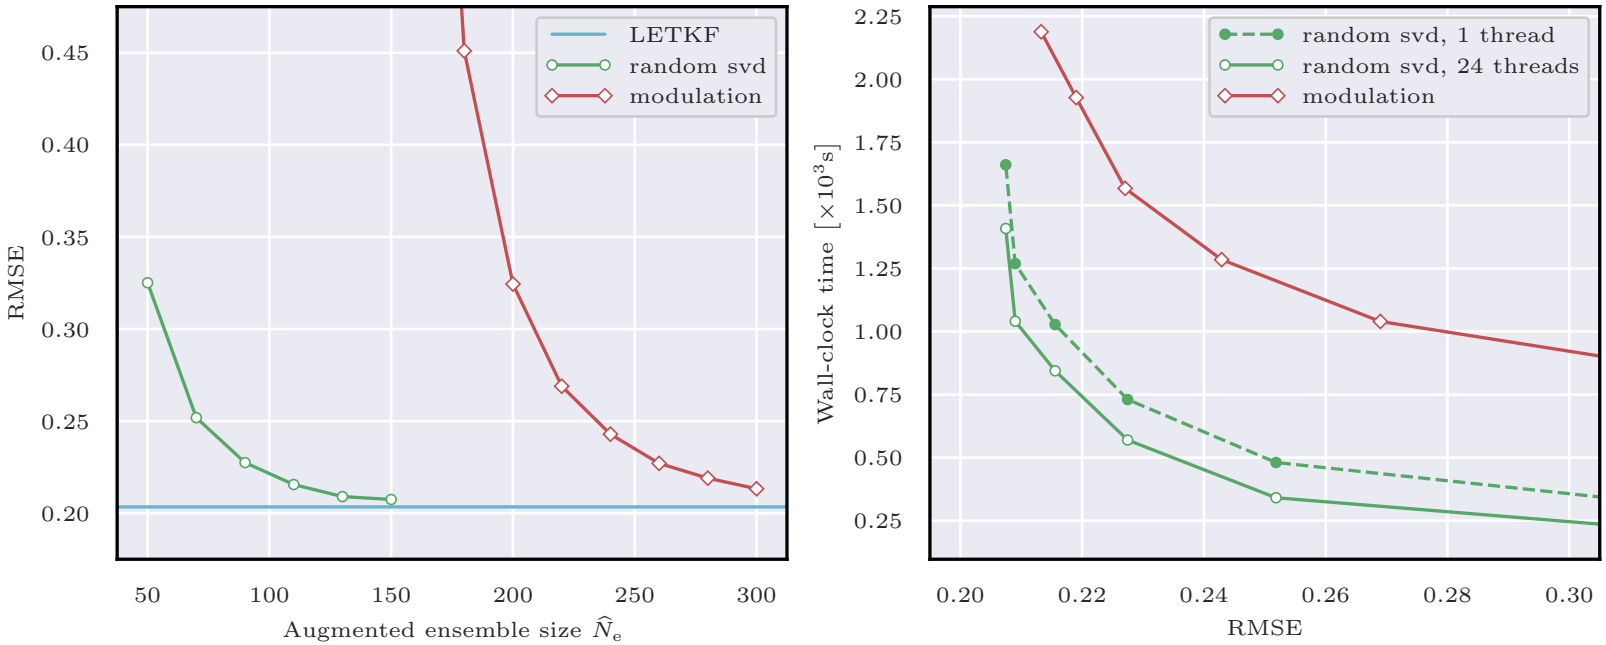

- Both methods can yield similar RMSE scores as the LETKF.
- The modulation method requires a larger augmented ensemble size to yield similar RMSE scores as the random svd method.
- For a given level of RMSE score, the random svd method is faster than the modulation method.

[Farchi and Bocquet, 2019]

A multilayer extension of the L96 model

- We introduce the *mL96 model*, that consists of $P_z = 32$ coupled layers of the L96 model with $P_\mathrm{h} = 40$ variables:

$$\begin{aligned}\frac{\mathrm{d}x_{(z,\,h)}}{\mathrm{d}t} = &\left(x_{(z,\,h+1)} - x_{(z,\,h-2)}\right) x_{(z,\,h-1)} - x_{(z,\,h)} + F_z \\ &+ \delta_{\{z>0\}} \left(x_{(z-1,\,h)} - x_{(z,\,h)}\right) \\ &+ \delta_{\{z<P_z\}} \left(x_{(z+1,\,h)} - x_{(z,\,h)}\right). \end{aligned} \tag{16}$$

- The forcing term linearly decreases from $F_1 = 8$ to $F_{32} = 4$.

Satellite observations for the mL96 model

- Each column is observed independently via:

$$y_{c,\,h} = \sum_{z=1}^{P_z} [\mathbf{\Omega}]_{c,\,z}\, x_{z,\,h} + v_{c,\,h}, \qquad v_{c,\,h} \sim \mathcal{N}(0,1), \quad (17)$$

where $\mathbf{\Omega}$ is a weighting matrix with $N_\mathrm{c} = 8$ channels that is designed to mimic *satellite radiances*.

- The 8×40 observations are available every $\Delta t = 0.05$.
- The runs are $10^4 \Delta t$ long.
- All algorithms use an ensemble of $N_\mathrm{e} = 8$ members.

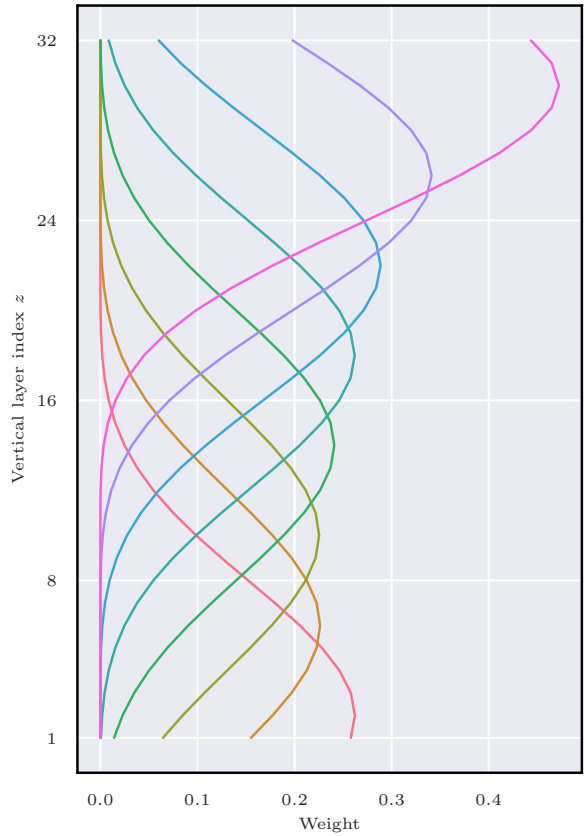

Covariance localisation (with augmented ensembles) is used *only in the vertical direction*. Domain localisation (LETKF-like) is used in the horizontal direction.

Results with the mL96 model

- Using covariance localisation in the vertical direction yields better RMSE scores than the LETKF.
- The modulation method requires a larger augmented ensemble size to yield similar RMSE scores as the random svd method.
- Both methods benefit from the parallelisation of the local analyses, but the parallelisation potential of the random svd method is not fully exploited because of our limited computational platform.

[Farchi and Bocquet, 2019]

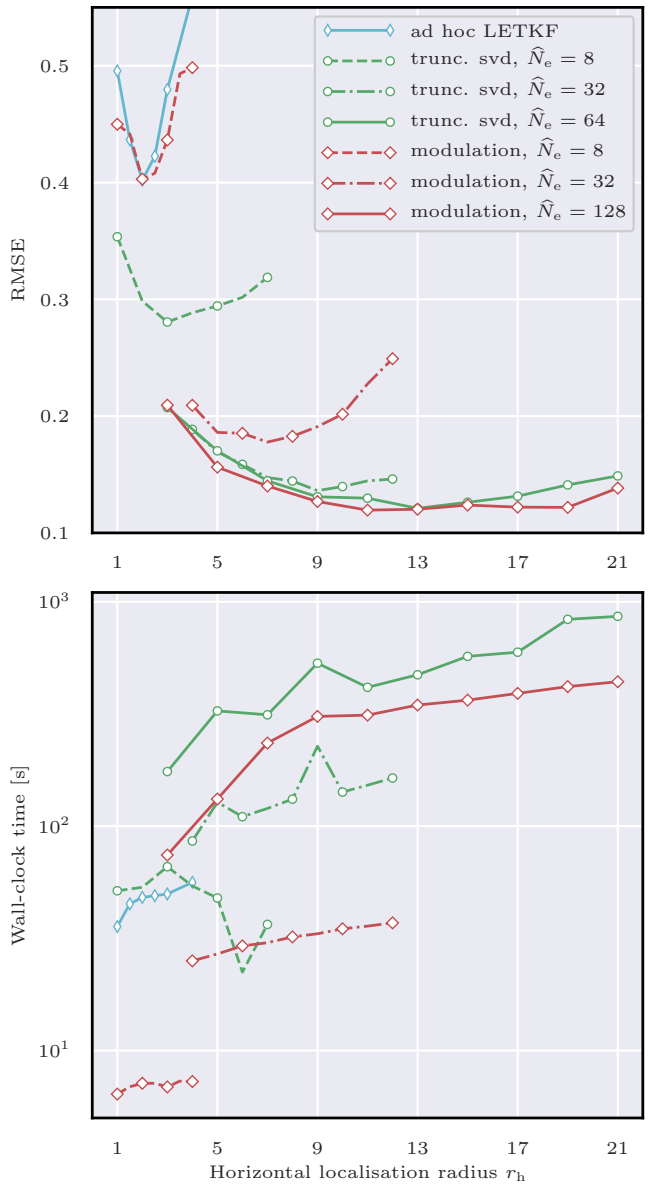

Conclusions related to the augmented ensembles

- We have shown how to use *augmented ensembles* to implement covariance localisation in the deterministic EnKF.
- We have proposed an alternative to the modulation method to construct the augmented ensemble, based on the *random svd* algorithm.
- We have compared both methods using the L96 model and a multilayer extension of the L96 model with satellite-like observations.
- We have shown that the augmented ensemble size need to be smaller with the random svd method than with the modulation method.
- We have seen that using *domain localisation* in the horizontal and *covariance localisation* in the vertical seems to be an adequate approach to assimilate *satellite radiances* in a spacially extended model.

Contents

Failed deterministic sampling

- The perturbation update of the LEnSRF:

$$\mathbf{T} = \mathbf{I} + \mathbf{B}\mathbf{H}^{\mathrm{T}}\mathbf{R}^{-1}\mathbf{H}, \tag{18}$$

$$\mathbf{X}_{\mathrm{a}} = \mathbf{T}^{-1/2}\mathbf{X}. \tag{19}$$

- A first source of inconsistency lies in the fact that $\mathbf{B} \neq \mathbf{X}\mathbf{X}^{\mathrm{T}}$.
- A potential fix: apply the left-transform to $\widetilde{\mathbf{X}}$, defined as the N_{e} dominant modes of $\mathbf{B} = \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^{\mathrm{T}}\right)$, for which $\mathbf{B} \approx \widetilde{\mathbf{X}}\widetilde{\mathbf{X}}^{\mathrm{T}}$.
- Empirically, this fix systematically makes the EnKF diverge. Intuitively, this comes from a *double* application of localisation.

A new approach

- For the LEnSRF update, instead of using the left-transform it could be more consistent with how the perturbations are defined to look for a low-rank perturbation matrix $\mathbf{X}_\mathrm{a}$ such that

$$\mathbf{P}_\mathrm{a} \approx \boldsymbol{\rho} \circ \left(\mathbf{X}_\mathrm{a}\mathbf{X}_\mathrm{a}^\mathrm{T}\right). \tag{20}$$

A solution of Eq. (20) trades the accuracy of the representation of the long range covariances (which may eventually be discarded at the next cycle) for a potentially better accuracy of the short range covariances.

- In this context, the goal is to solve the optimisation problem

$$\mathbf{X}_\mathrm{a} = \underset{\mathrm{rank}(\mathbf{X}) \leq N_\mathrm{e}-1}{\arg\min} \mathcal{L}\left(\mathbf{X}\right), \quad \text{with} \quad \mathcal{L}\left(\mathbf{X}\right) = \ln \left\| \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^\mathrm{T}\right) - \mathbf{P}_\mathrm{a} \right\|_\mathrm{F}, \tag{21}$$

where $\|*\|_\mathrm{F}$ is the Frobenius matrix norm.

[Bocquet and Farchi, 2019]

A new approach

- This problem is similar to a *weighted low-rank approximation* (WLRA) problem. We expect that it has *no tractable solution* as opposed to the unweighted case for which the *Eckart-Young* solution holds.
- However, the gradient of the cost function can be computed:

$$\nabla \mathcal{L}(\mathbf{X}) = 2 \left\| \boldsymbol{\Delta} \right\|_{\mathrm{F}}^{-2} (\boldsymbol{\rho} \circ \boldsymbol{\Delta}) \mathbf{X}, \tag{22}$$

with $\boldsymbol{\Delta} = \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^{\mathrm{T}} \right) - \mathbf{P}_{\mathrm{a}}$.

- Therefore, a local solution can be found using a *gradient-based minimisation algorithm*.

Computing the gradient and the cost function

- We use a *mode expansion* with $\widehat{N}_\mathrm{e}$ modes for the analysis error covariance matrix: $\mathbf{P}_\mathrm{a} \approx \widehat{\mathbf{X}}_\mathrm{a}\widehat{\mathbf{X}}_\mathrm{a}^\mathrm{T}$.
- To compute $\mathcal{L}$ and $\nabla\mathcal{L}$, we use the (classical) formula

$$\boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^\mathrm{T}\right) \cdot \mathbf{v} = \sum_{i=1}^{N_\mathrm{e}} \mathbf{X}^{(i)} \circ \left[\boldsymbol{\rho} \cdot \left(\mathbf{X}^{(i)} \circ \mathbf{v}\right)\right], \tag{23}$$

for any matrix $\mathbf{X}$ of size $N_\mathrm{x} \times N_\mathrm{e}$ and any vector $\mathbf{v}$ of size N_x, and where $\mathbf{X}^{(i)}$ is the i-th column of $\mathbf{X}$.

- Therefore, the numerical cost of computing $\mathcal{L}$ and $\nabla\mathcal{L}$ is:
  - $\mathcal{O}\left(N_\mathrm{e}\widehat{N}_\mathrm{e}N_\mathrm{x}N_\mathrm{b}\right)$ if $\boldsymbol{\rho}$ is banded with bandwidth N_b;
  - $\mathcal{O}\left(N_\mathrm{e}\widehat{N}_\mathrm{e}N_\mathrm{x}\ln N_\mathrm{x}\right)$ if $\boldsymbol{\rho}$ is homogeneous.
- The numerical cost can be reduced by a factor N_e or $\widehat{N}_\mathrm{e}$ if Eq. (23) is computed in parallel.

The Lorenz 1996 and the Kuramoto–Sivashinsky models

- We use the *L96 model* with $N_{\mathrm{x}} = 40$ variables, in the same configuration as in the previous test series.
- As a complement, we use the *KS model*:

$$\partial_t u = -u\partial_x u - \partial_x^2 u - \partial_x^4 u. \tag{24}$$

Equation (24) is solved over the domain $x \in [0, 32\pi]$ using a pseudo-spectral integration with $N_{\mathrm{x}} = 128$. For this model, the observations are given every $\Delta t = 1$ by

$$\mathbf{y} = \mathbf{x} + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \tag{25}$$

- For both models, the localisation matrix is constructed using the *Gaspari–Cohn* function, assuming that the grid points are equally distributed in space.
- The runs are $2 \times 10^4 \Delta t$ long and our criterion is the time-average *analysis RMSE*.
- We compare:
  - the LETKF algorithm;
  - the LEnSRF algorithm (exact left transform update);
  - the new LEnSRF algorithm (new perturbation update scheme).

Accuracy for the L96 model and KS models

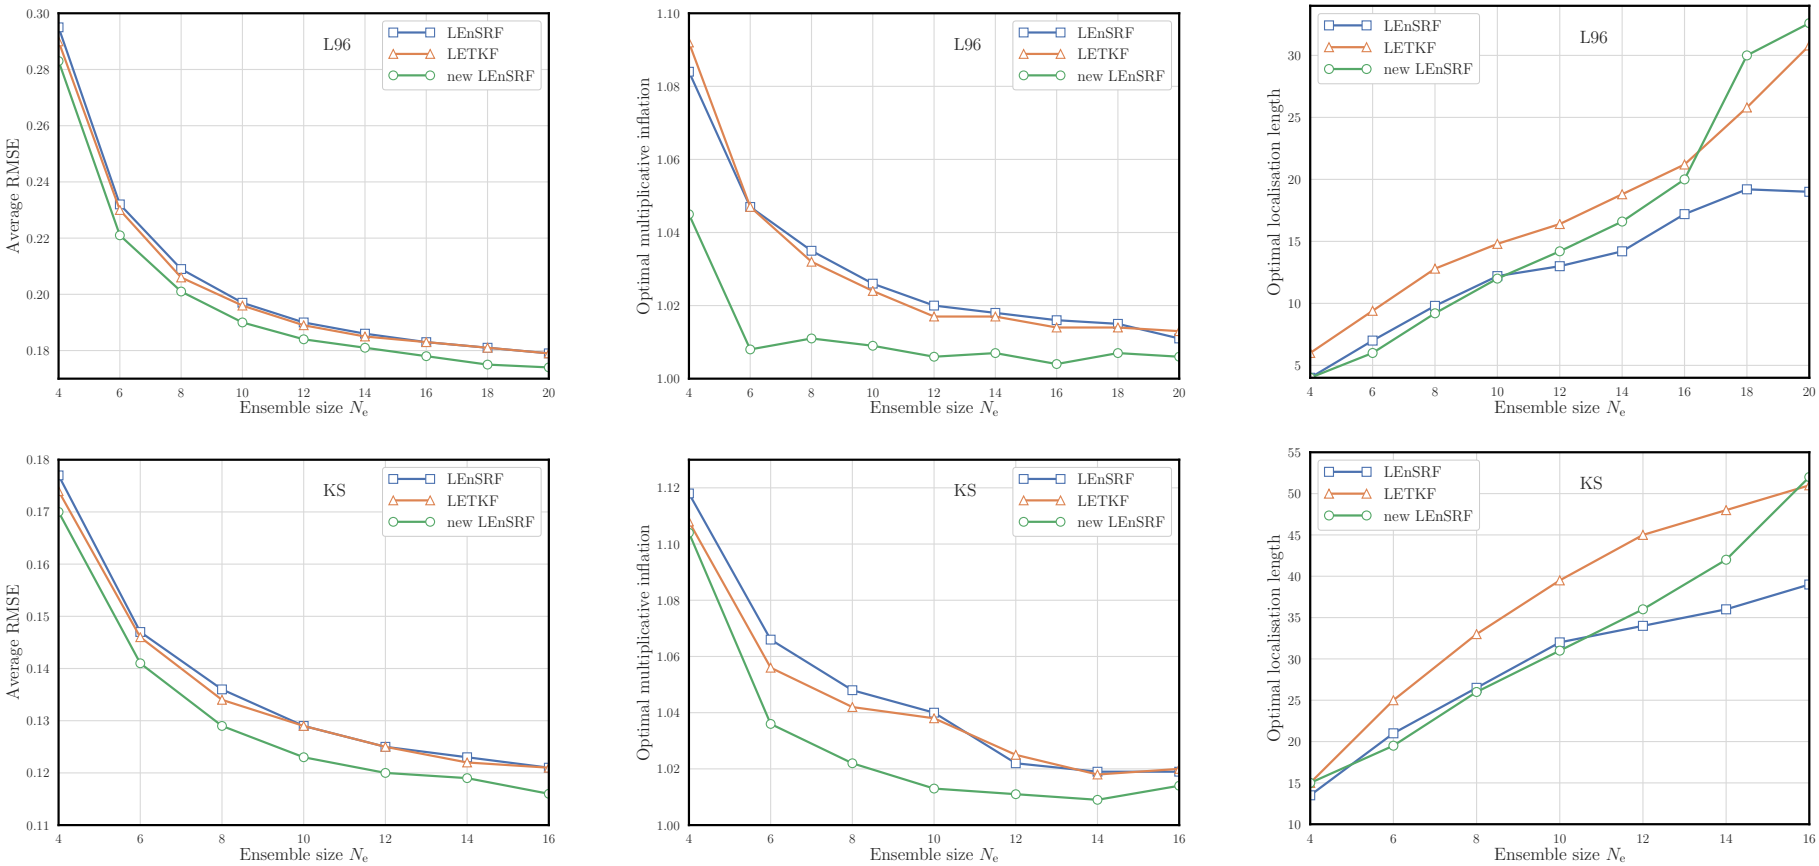

Comparison of the LETKF, the LEnSRF, and the LEnSRF with the new update scheme, applied to L96 (top) and to KS (bottom). The RMSE, optimal localisation and optimal inflation are plotted as functions of the ensemble size N_e.

[Bocquet and Farchi, 2019]

Robustness of the new scheme

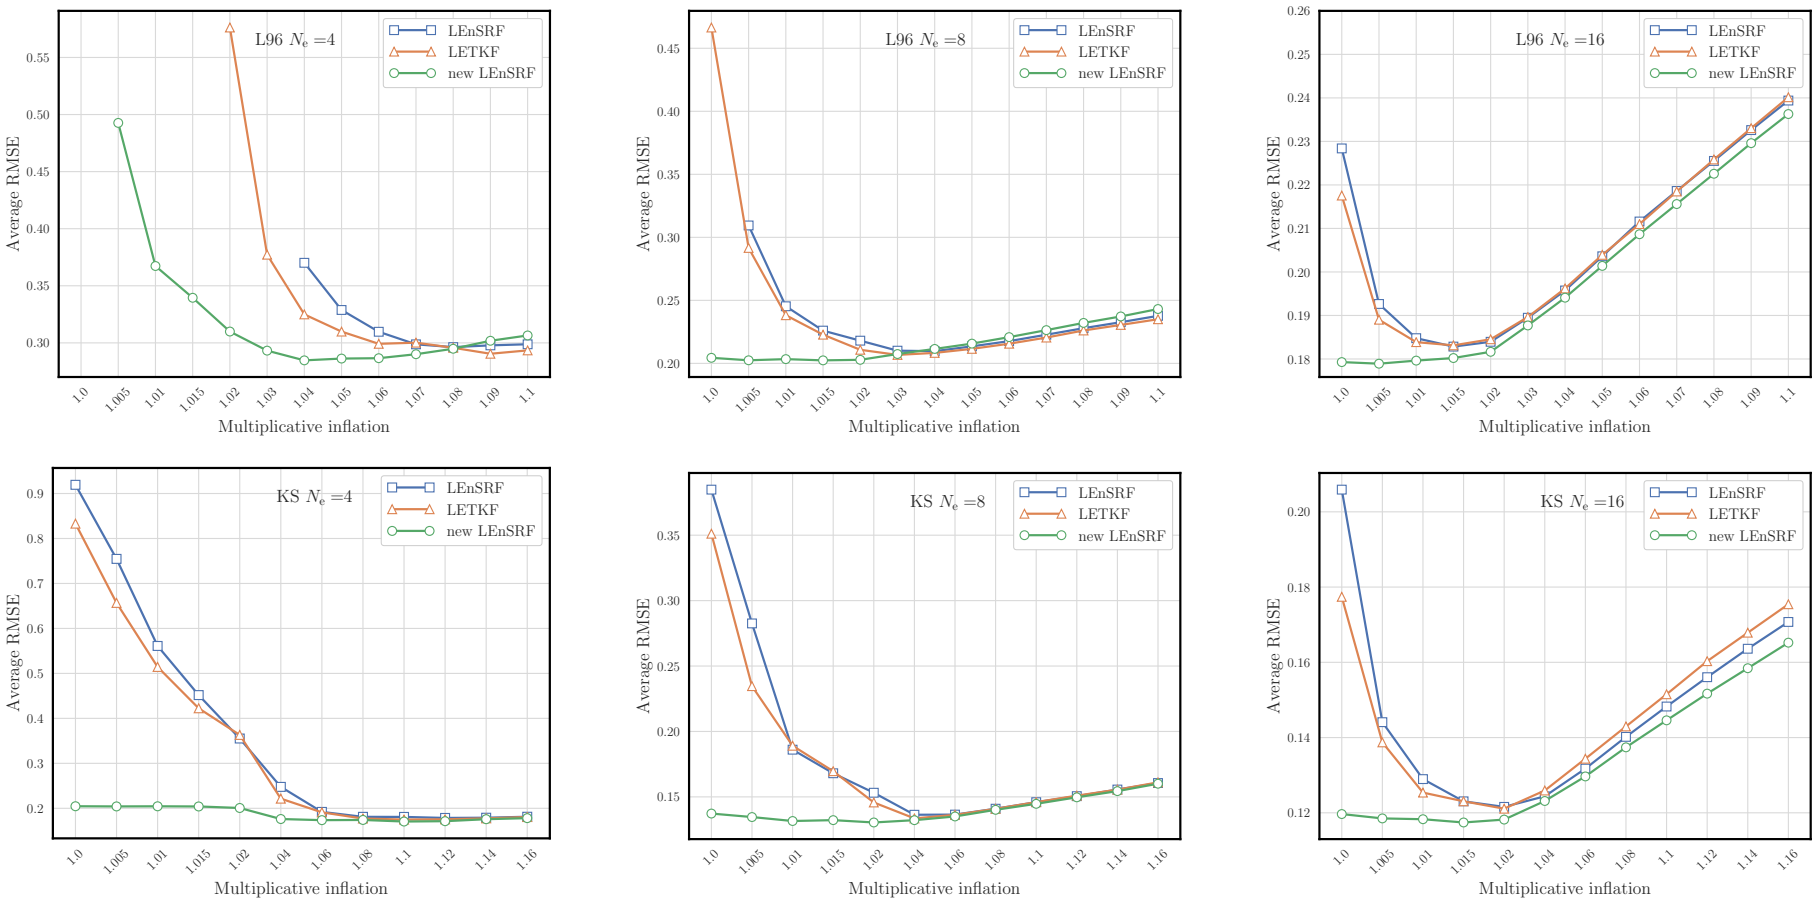

Time-averaged RMSE as a function of the multiplicative inflation, the localisation length being tuned so as to minimise the RMSE for L96 (top) and KS (bottom), for $N_e = 4, 8, 16$.

[Bocquet and Farchi, 2019]

Conclusions related to the new perturbation update scheme

- The updated perturbations in the local EnKFs based on covariance localisation (in particular the LEnSRF) *are not* the main modes of $\mathbf{P}_\mathrm{a}$ ($\neq$ LETKF).
- We have proposed a perturbation update scheme potentially *more consistent* such that the perturbations $\mathbf{X}$ are related to the error covariance matrix by $\mathbf{P} \approx \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^\mathrm{T}\right)$ throughout the EnKF:

$$\mathbf{X}_\mathrm{a} = \underset{\mathrm{rank}(\mathbf{X}) \leq N_\mathrm{e}-1}{\arg\min} \mathcal{L}\left(\mathbf{X}\right), \quad \text{with} \quad \mathcal{L}\left(\mathbf{X}\right) = \ln \left\| \boldsymbol{\rho} \circ \left(\mathbf{X}\mathbf{X}^\mathrm{T}\right) - \mathbf{P}_\mathrm{a} \right\|_\mathrm{F}, \tag{26}$$

- We have compared numerically the new LEnSRF to the LETKF and to the LEnSRF (with the exact left-transform update), using the L96 and the KS models.
- We have shown that for both models, *the requirement for residual multiplicative inflation is much weaker with the new LEnSRF* than with both the LETKF and the LEnSRF: much *weaker imbalance* of the new update scheme?
- Moreover, there is an *accuracy improvement of up to* 6% *in the analysis RMSE* in mildly nonlinear conditions, which is significant in these very well tuned configurations.

References

Bocquet, M. and A. Farchi (2019). "On the consistency of the local ensemble square root Kalman filter perturbation update". In: *Tellus A*. In press.

Farchi, A. and M. Bocquet (2019). "On the Efficiency of Covariance Localisation of the Ensemble Kalman Filter Using Augmented Ensembles". In: *Front. App. M. Stat.* 5, p. 3.