Review
New and existing covariance estimation techniques for ensemble data assimilation
Shay Gilpin, Matthias Morzfeld, and Kevin K. Lin

The paper aims to compare different localisation methods for ensemble data assimilation in general and the ensemble Kalman filter (EnKF) in particular. It does so by first providing the equations behind the EnKF, followed by a short review of several spatially based localisation methods, several non-structured methods and several hybrid methods in which the localised ensemble covariance is combined with a climatological covariance. This is followed by a description of the experiments and the two Lorenz-96 models used in the experiments. In the first experiment, the localised covariances are compared to different realisations of a synthetically "true" covariance. In the second one, the two models are run forward in time with corrections produced by EnKFs using the different localisation methods, and the different methods are ranked based on analysis RMSE. It is found that, in the setups used, spatial localisation methods generally work better than unstructured ones and that, in line with previous work, the type of spatial localisation used is largely irrelevant to performance.

The topic and the more conceptual approach of the paper fit, in my view, the scope of NPG well, as the paper covers a data-driven approach for nonlinear systems. The work is not groundbreaking, but is a logical and very insightful extension of previous papers on this topic (e.g. Morzfeld & Hodyss, 2023). The structure of the text is clear and logical, and the text is easy to follow without being overly elaborate. Figures are legible and relevant. The results support the conclusions within the class of problems considered, and the work is a logical extension of previous work by the authors and others. The generality of some of the results for more realistic systems might be limited, as in some experiments covariances are constructed using the same model as the truth, and this is something that needs to be addressed in the revision. It is, in general, a good paper. There are some points that could/should be added (see below) and would, in my view, improve it even further. Based on this, my recommendation is to accept it with major (but limited) revisions.

Major points:

I) Equations 18, 21, 22, 25; "In contrast, the spatially-agnostic schemes we consider (e.g., NICE, thresholding, Ledoit-Wolf), are not competitive with localization or hybrid estimators in all experiments." (lines 544-546): The hyperparameters in the spatial localisation schemes in Section 5 are taken from the same function family as the "true" values. The experiments demonstrate the value of localisation when informative structural knowledge is available, but do not yet demonstrate that localisation remains superior when such information is only approximately known. In order to support the conclusion that the spatial localisation schemes outperform the spatially agnostic schemes, it needs to be shown (in an additional experiment) that they are robust if no information about the "true" covariance can be used in the calibration of the hyperparameters.

II) Section 5.2: Localisation acts as a convolution on the ensemble anomalies. By the convolution theorem, this is equivalent to element-wise multiplication/weighting of the Fourier spectrum and, by extension, the power spectrum of the covariance and hence, by construction of the Kalman gain, the DA corrections. In the paper, the EnKF performance is only evaluated based on RMSE. In geophysical models, the error spectrum is typically red, and hence the RMSE is dominated by the larger scales. On these large scales, the spatial localisation kernels will look quite similar, and hence I can understand that they all produce similar RMSEs. I expect that the differences between the different localisation kernels, and hence DA performance, become more pronounced towards small scales. Due to the lower signal-to-noise ratio at these scales, the distinction between the sampling error to be suppressed and the true covariance becomes even more challenging. To get to my point: I would recommend including error spectra alongside the already-present RMSEs in Section 5.2 to see whether they show more significant differences between the localisation methods at the smaller scales.

Minor points:

1) "by wavelet-based and spectral approaches" (line 158): Contrary to the other B-matrix construction techniques discussed in Section 1, there are no references (e.g. Deckmyn and Berre, 2005; Pannekoucke et al., 2007; Beezley et al., 2011; Bonavita et al., 2016; Pasmans et al., 2024), nor is there any mention that these methods can be used to radically simplify localisation (usually only block diagonals are retained).

2) "The AoD matrix is updated at the beginning of the next "day." (line 164): The update procedure is not completely clear to me. a) Is the AoD the average of the sample covariances over a single day, or does the number of sample covariances increase as the experiment moves forward in time, using all sample covariances up to that point? b) Are the covariances averaged always associated with times prior to the current DA correction time, or can the average also include covariances that are associated with the same day but a later DA correction time on that day? The latter would be problematic.

3) "Twenty observations y are collected in the center of the domain (every other grid point between coordinates 20 and 60) at each model integration time step ($\Delta t = 0.05$)." (lines 425-426): It is not clear to me how the "truth" from which these observations are generated is constructed. Is it just a free run of the model? The same applies to the setup in Section 5.2.2.

4) Reference section: The citation pattern is not consistent. Sometimes journal titles are abbreviated, and sometimes they are not. Information seems to be missing from Pourahmadi, M.: High-Dimensional Covariance Estimation, Wiley, 2013.

References:

Bonavita, M., Hólm, E., Isaksen, L., & Fisher, M. (2016). The evolution of the ECMWF hybrid data assimilation system. Quarterly Journal of the Royal Meteorological Society, 142(694), 287-303.

J. D. Beezley, J. Mandel and L. Cobb, "Wavelet ensemble Kalman filters," Proceedings of the 6th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems, Prague, Czech Republic, 2011, pp. 514-517, doi: 10.1109/IDAACS.2011.6072819.
Deckmyn, A., & Berre, L. (2005). A wavelet approach to representing background error covariances in a limited-area model. Monthly Weather Review, 133(5), 1279-1294.

Morzfeld, M., & Hodyss, D. (2023). A theory for why even simple covariance localization is so useful in ensemble data assimilation. Monthly Weather Review, 151(3), 717-736.

Pannekoucke, O., Berre, L., & Desroziers, G. (2007). Filtering properties of wavelets for local background-error correlations. Quarterly Journal of the Royal Meteorological Society: A Journal of the Atmospheric Sciences, Applied Meteorology and Physical Oceanography, 133(623), 363-379.

Pasmans, I., Chen, Y., Carrassi, A., & Jones, C. K. (2024). Tailoring data assimilation to discontinuous Galerkin models. Quarterly Journal of the Royal Meteorological Society, 150(762), 2820-2847.