

Scale-dependent background-error covariance localisation

By MARK BUEHNER* and ANNA SHLYAEVA, *Data Assimilation and Satellite Meteorology Research Section, Environment Canada, Dorval, Quebec H9P 1J3, Canada*

(Manuscript received 30 March 2015; in final form 5 November 2015)

ABSTRACT

A new approach is presented and evaluated for efficiently applying scale-dependent spatial localisation to ensemble background-error covariances within an ensemble-variational data assimilation system. The approach is primarily motivated by the requirements of future data assimilation systems for global numerical weather prediction that will be capable of resolving the convective scale. Such systems must estimate the global and synoptic scales at least as well as current global systems while also effectively making use of information from frequent and spatially dense observation networks to constrain convective-scale features. Scale-dependent covariance localisation allows a wider range of scales to be efficiently estimated while simultaneously assimilating all available observations. In the context of an idealised numerical experiment, it is shown that using scale-dependent localisation produces an improved ensemble-based estimate of spatially varying covariances as compared with standard spatial localisation. When applied to an ensemble of Arctic sea-ice concentration, it is demonstrated that strong spatial gradients in the relative contribution of different spatial scales in the ensemble covariances result in strong spatial variations in the overall amount of spatial localisation. This feature is qualitatively similar to what might be expected when applying an adaptive localisation approach that estimates a spatially varying localisation function from the ensemble itself. When compared with standard spatial localisation, scale-dependent localisation also results in a lower analysis error for sea-ice concentration over all spatial scales.

Keywords: *EnVar, EnKF, ensemble data assimilation, covariance localisation, sea-ice concentration*

1. Introduction

Continual improvements in high-performance computing will allow future global models and data assimilation systems for numerical weather prediction (NWP) to have spatial and temporal resolutions sufficient to resolve convective-scale phenomena. This is in contrast with the current situation in which most global models used for generating ensembles have relatively coarse resolution, while convective-scale ensemble systems often have spatial domains too small to fully represent the synoptic and global scales and instead partially rely on the lateral boundary conditions supplied by a larger domain (often global) system. Forecast quality for convective-scale phenomena strongly depends on having high accuracy at the larger scales, but also requires frequent assimilation of high-resolution, dense observation networks to continually correct the location and

intensity of small-scale features (e.g. fronts and individual thunderstorms). Consequently, future data assimilation systems will be required to estimate the synoptic and global scales at least as well as current global systems while also effectively making use of information from frequent and spatially dense observation networks to constrain convective-scale features.

Applications of data assimilation for NWP increasingly rely on ensemble approaches in which relatively small ensembles of short-range forecasts are used to estimate the background-error covariances for assimilating observations. This includes various types of the ensemble Kalman filter (EnKF; e.g. Houtekamer et al., 2014) and the more recent variational-ensemble (EnVar; e.g. Buehner et al., 2013) approaches. The computational cost of producing an ensemble of forecasts large enough to estimate sufficiently accurate and high-rank covariances is prohibitive. To overcome this obstacle, spatial covariance localisation is applied to the ensemble covariances, usually by assuming that distant correlations are nearly zero (Hamill et al., 2001;

*Corresponding author.
email: mark.buehner@canada.ca

Houtekamer and Mitchell, 2001). Spatial covariance localisation is an important factor that has enabled both EnKF and EnVar methods to produce more accurate analyses as compared with approaches such as three-dimensional variational data assimilation (3DVar) that use climatological error estimates usually with the assumptions of horizontally homogeneous and isotropic correlations (e.g. Buehner et al., 2010).

For a given system, the optimal amount of spatial localisation has been shown to depend on the ensemble size such that a smaller ensemble requires more localisation (Houtekamer and Mitchell, 2001; Lorenc, 2003). The amount of localisation is usually also dependent on the spatial resolution of the forecast model and the assimilated observations. Systems designed for initialising lower resolution global models typically use localisation that starts to decrease the ensemble covariances at much larger distances than systems designed for convective-scale models. For example, in a low-resolution global EnKF that assimilates simulated radiosonde and satellite thickness observations, Houtekamer and Mitchell (2001) determined the optimal distance for forcing covariances to zero to be between 3000 and 6000 km depending on the ensemble size. In contrast, in a high-resolution EnKF that assimilates simulated weather radar data, Caya et al. (2005) found that the optimal localisation forces covariances to zero at between 5 and 11 km. Such vastly different localisation strategies are currently effective since convective-scale ensemble-based data assimilation systems have relatively small spatial domains and therefore partially rely on a larger domain system to provide the large-scale component of the atmospheric state through the lateral boundary conditions.

Several approaches have been explored for directly estimating a spatially varying localisation function and thus avoid the time-intensive procedure of manually tuning the length scale of the chosen localisation function by performing a series of complete data assimilation experiments. In the ‘adaptive’ localisation approach of Bishop and Hodyss (2009), the spatially varying localisation function is derived from the ensemble covariances. Consequently, areas where the ensemble correlations have a shorter length scale will be more severely localised than areas with broader ensemble correlations. Alternatively, Anderson and Lei (2013) and Lei and Anderson (2014) compute ‘empirical localisation functions’ from minimising the error of the resulting analysis ensemble mean from an observing system simulation experiment. Flowerdew (2015) presents a similar approach for computing the optimal localisation function and also suggests an alternative approach that reverts the distant correlations toward a climatological mean rather than zero. Finally, Ménétrier et al. (2015) developed a practical approach for estimating the optimal localisation function

by merging the theories of centred moments estimation and optimal linear filtering.

It may be argued that data assimilation methods have not yet been designed to effectively assimilate simultaneously all available observations to correct the full range of scales that will be resolved by the future global convective-scale resolving models. The approach known as successive covariance localisation (Zhang et al., 2009) is applied to an EnKF to assimilate a subset of the observations with relatively weak localisation to estimate the large scales and another subset with more severe localisation to estimate the small scales. However, since each observation may contain useful information on the error at all spatial scales, this approach ignores potentially useful information that would be available if all observations could be assimilated to estimate all scales simultaneously. In another approach, multiple independent analyses are performed, each designed to estimate the correction to the background state considering a different range of scales with an appropriate amount of covariance localisation, which are later combined to produce the complete analysis (Miyoshi and Kondo, 2013). A somewhat related approach was developed in the context of 3DVar to treat different ranges of scales separately (e.g. Li et al., 2015). Similarly, the spatial/spectral localisation approach of Buehner (2012) uses, within an EnVar system, background-error covariances that have been separated according to overlapping spectral wavebands with an appropriate amount of spatial localisation applied to each waveband. Though all observations are assimilated simultaneously with this approach, the assumption is made that the covariances between the scales are zero. Approaches that treat different ranges of scales independently, either by performing independent analyses for each or by explicitly setting the between-scale covariances to zero in a single analysis, may lose useful information from the heterogeneity of covariances (e.g. related to the location of a storm or front), since the elimination of between-scale covariances is equivalent to local spatial averaging of the covariances (Buehner and Charron, 2007). Other approaches, such as those discussed by Berre and Desroziers (2010), have the explicit goal of applying local spatial averaging to covariance estimates to reduce sampling error, which may only be beneficial with a very small ensemble size [i.e. $O(10)$ members]. Therefore, it is desirable to develop and evaluate efficient approaches that allow the atmospheric state at all scales to be accurately estimated simultaneously using the heterogeneous covariance information present in relatively small ensembles [i.e. $O(100)$ members] to assimilate all available observations.

The goal of the present study is to evaluate a new approach for efficiently applying scale-dependent spatial

localisation to ensemble covariances. The next section presents the formulation of the approach, including how it can be implemented within an EnVar data assimilation system that employs the square-root of the background-error covariance matrix. Section 3 presents results from using scale-dependent covariance localisation to estimate known spatial covariances on a one-dimensional periodic domain from a small ensemble. In Section 4, an ensemble of the sea-ice concentration on a high-resolution grid over the Arctic region near North America and Greenland is used to evaluate the approach by performing idealised data assimilation experiments. In both applications, scale-dependent localisation is compared with using standard spatial localisation with either a relatively broad or severe localisation function. Finally, some conclusions are given in Section 5.

2. Scale-dependent covariance localisation

2.1. Basic approach

The approach presented here is closely related to the spatial/spectral localisation approach of Buehner (2012). The main difference is that in the current study we wish to retain, as much as possible, the original between-scale covariances present in the ensemble covariances while applying different amounts of spatial localisation for different ranges of spatial scales. As in Buehner (2012), the normalised ensemble member perturbations (that is, differences from the ensemble mean divided by $\sqrt{N_{ens} - 1}$, where N_{ens} is the ensemble size) are first each decomposed into a set of overlapping spectral wavebands, with the scale denoted by the index $j = 1, \dots, J$, as

$$\mathbf{e}_{j,k} = \mathbf{\Psi}_j \mathbf{e}_k, \quad (1)$$

where \mathbf{e}_k is the k th normalised ensemble perturbation, $\mathbf{\Psi}_j$ is the spectral filter that isolates the j th spectral waveband, and $\mathbf{e}_{j,k}$ is the k th ensemble perturbation containing only the j th waveband.

To retain the between-scale covariances, the filtered ensemble perturbations are not treated as independent, as they are in Buehner (2012). Instead, the contribution from each of J scales for each ensemble perturbation is concatenated to create extended vectors, and the resulting spatial/spectral ensemble covariance matrix can be written as

$$\mathbf{B}_{ss} = \sum_{k=1}^{N_{ens}} \begin{bmatrix} \mathbf{e}_{1,k} \\ \mathbf{e}_{2,k} \\ \vdots \\ \mathbf{e}_{J,k} \end{bmatrix} \begin{bmatrix} \mathbf{e}_{1,k}^T & \mathbf{e}_{2,k}^T & \cdots & \mathbf{e}_{J,k}^T \end{bmatrix}. \quad (2)$$

The original ensemble covariance matrix can be obtained by simply summing the contributions from each scale, including the between-scale covariances, giving

$$\begin{aligned} \mathbf{B}_{ens} &= [\mathbf{I} \ \mathbf{I} \ \cdots \ \mathbf{I}] \mathbf{B}_{ss} \begin{bmatrix} \mathbf{I} \\ \mathbf{I} \\ \vdots \\ \mathbf{I} \end{bmatrix} = \sum_{k=1}^{N_{ens}} \left(\sum_{j1=1}^J \mathbf{e}_{j1,k} \right) \left(\sum_{j2=1}^J \mathbf{e}_{j2,k}^T \right) \\ &= \sum_{k=1}^{N_{ens}} \mathbf{e}_k \mathbf{e}_k^T, \end{aligned} \quad (3)$$

where $j1$ and $j2$ denote the indices for each possible pair of scales and \mathbf{I} is the identity matrix of the same size as the original ensemble covariance matrix. For this to hold the scale-separated ensemble, perturbations must sum to equal the original member, that is

$$\mathbf{e}_k = \sum_{j=1}^J \mathbf{e}_{j,k}, \quad (4)$$

and therefore for each wavenumber the filter functions must sum to one [instead of their squared values summing to one as in Buehner (2012)]. This is demonstrated by combining eqs. (1) and (4):

$$\mathbf{e}_k = \sum_{j=1}^J \mathbf{e}_{j,k} = \sum_{j=1}^J \mathbf{\Psi}_j \mathbf{e}_k = \left(\sum_{j=1}^J \mathbf{\Psi}_j \right) \mathbf{e}_k. \quad (5)$$

The spatial/spectral covariance matrix can be rewritten in terms of submatrices representing the spatial covariances for each possible pair of scales ($j1, j2$), as in

$$\mathbf{B}_{ss,j1,j2} = \sum_{k=1}^{N_{ens}} \mathbf{e}_{j1,k} \mathbf{e}_{j2,k}^T. \quad (6)$$

Then, scale-dependent spatial covariance localisation can be applied to each submatrix by performing a Schur product with a spatial localisation matrix that varies as a function of the waveband indices $j1$ and $j2$:

$$\mathbf{B}_{ss,loc,j1,j2} = \mathbf{B}_{ss,j1,j2} \circ \mathbf{L}_{j1,j2} \quad (7)$$

The covariance matrix with scale-dependent spatial localisation is transformed back into the spatial domain by simply summing all of the localised submatrices:

$$\mathbf{B}_{sd,loc} = \sum_{j1=1}^J \sum_{j2=1}^J \mathbf{B}_{ss,loc,j1,j2} \circ \mathbf{L}_{j1,j2}. \quad (8)$$

It can be easily shown that if the same localisation matrix is applied to all within-scale and between-scale covariances, then the resulting matrix from eq. (8) will be identical to applying simple spatial localisation to the original ensemble covariance matrix.

2.2. Application to EnVar

To ensure that the complete spatial/spectral localisation matrix \mathbf{L} is positive semi-definite, and to facilitate implementation in an EnVar system that uses a square-root of \mathbf{B} and its transpose, the spatial localisation matrix for a given pair of scales ($j1, j2$) is defined as

$$\mathbf{L}_{j1,j2} = \left(\mathbf{L}_{j1,j1} \right)^{1/2} \left(\mathbf{L}_{j2,j2} \right)^{T/2}. \quad (9)$$

Consequently, the complete spatial/spectral localisation matrix can then be written as

$$\begin{aligned} \mathbf{L} &= \begin{bmatrix} \mathbf{L}_{1,1} & \mathbf{L}_{1,2} & \cdots & \mathbf{L}_{1,J} \\ \mathbf{L}_{2,1} & \mathbf{L}_{2,2} & \cdots & \mathbf{L}_{2,J} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{L}_{J,1} & \mathbf{L}_{J,2} & \cdots & \mathbf{L}_{J,J} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{L}_{1,1}^{1/2} \\ \mathbf{L}_{2,2}^{1/2} \\ \vdots \\ \mathbf{L}_{J,J}^{1/2} \end{bmatrix} \begin{bmatrix} \mathbf{L}_{1,1}^{T/2} & \mathbf{L}_{2,2}^{T/2} & \cdots & \mathbf{L}_{J,J}^{T/2} \end{bmatrix} \end{aligned} \quad (10)$$

The following variable transformation can then be easily implemented in an existing EnVar system for computing the analysis increment, $\Delta \mathbf{x}$, from the control vector (denoted by ξ_k for the portion corresponding to the k th ensemble perturbation):

$$\Delta \mathbf{x} = [\mathbf{I} \quad \mathbf{I} \quad \cdots \quad \mathbf{I}] \sum_{k=1}^{N_{ens}} \begin{bmatrix} \mathbf{e}_{1,k} \\ \mathbf{e}_{2,k} \\ \vdots \\ \mathbf{e}_{J,k} \end{bmatrix} \circ \left(\begin{bmatrix} \mathbf{L}_{1,1}^{1/2} \\ \mathbf{L}_{2,2}^{1/2} \\ \vdots \\ \mathbf{L}_{J,J}^{1/2} \end{bmatrix} \xi_k \right). \quad (11)$$

Alternatively, this can be written in the more compact form

$$\Delta \mathbf{x} = \sum_{j=1}^J \sum_{k=1}^{N_{ens}} \mathbf{e}_{j,k} \circ \left(\mathbf{L}_{j,j}^{1/2} \xi_k \right). \quad (12)$$

For comparison, the variable transformation in the standard EnVar approach with spatial localisation that does not depend on scale (Lorenc, 2003; Buehner, 2005) is expressed as

$$\Delta \mathbf{x} = \sum_{k=1}^{N_{ens}} \mathbf{e}_k \circ \left(\mathbf{L}^{1/2} \xi_k \right) \quad (13)$$

Note that the only differences between the current approach and the approach used by Buehner (2012) is that:

- (1) the same control vector is used for all scales here, whereas independent control vectors are used for each scale in Buehner (2012):

$$\Delta \mathbf{x} = \sum_{j=1}^J \sum_{k=1}^{N_{ens}} \mathbf{e}_{j,k} \circ \left(\mathbf{L}_{j,j}^{1/2} \xi_{j,k} \right) \quad (14)$$

and, related to this,

- (2) the spectral filter functions sum to one here, whereas the squares of the filter functions sum to one in Buehner (2012).

The use of independent control vectors in Buehner (2012) results from imposing spectral localisation by setting all of the between-scale covariances to zero, that is

$$\mathbf{L}_{j1,j2} = \mathbf{0}, \text{ for } j1 \neq j2. \quad (15)$$

3. Results from estimating covariances on a one-dimensional domain

Scale-dependent covariance localisation is first applied to the simple problem of estimating spatial covariances on a one-dimensional periodic domain of 60 grid points from a 50-member ensemble of random perturbations produced from the known true covariance matrix. The true heterogeneous covariance matrix (shown in Fig. 1) is constructed by taking a weighted average of two homogeneous Gaussian auto-covariance functions with different length scales. The large- and small-scale components have length scales of 7 and 0.7 grid points, respectively. The relative contribution of each component varies smoothly over the




Fig. 1. True covariance matrix used for the estimation of ensemble covariances on a one-dimensional periodic domain of 60 grid points.




Fig. 2. Spectral filter coefficients used to separate the background-error covariances into three wavebands: scales 1, 2 and 3 correspond to the large, medium and small scales, respectively.

domain such that the large-scale component constitutes 80 % of the variance at the two ends of the domain and only 20 % at the middle of the domain. The true variance equals one at all grid points. The random perturbations are generated by multiplying a vector of independent Gaussian

pseudorandom numbers with variance of one by the square-root of the true covariance matrix.

Figure 2 shows the spectral filter coefficients used to decompose the covariances into three ranges of scales (denoted as scales 1, 2 and 3 that correspond to the large, medium and small scales, respectively) used in eq. (1). More specifically, the individual ensemble perturbations are decomposed into the three scales by first transforming each into spectral (i.e. Fourier) space, then multiplying the resulting spectral coefficients by the filter coefficients shown in Fig. 2, and finally transforming the results back into grid-point space. To satisfy eq. (5), the spectral filter coefficients were chosen such that the three coefficients sum to one for each wavenumber (constructed by combining scaled and vertically shifted sine functions). Note that, at most, only two filter functions have non-zero values for each wavenumber, and therefore one of the functions can be constructed as one minus the other. Also, note that the width of the filter functions increases with increasing wavenumber, similar to most wavelet transforms (i.e. the transition between scales 1 and 2 is approximately half as wide as the transition between scales 2 and 3). Since there is a trade-off between spectral and spatial resolutions, this choice of filter functions provides a higher degree of spatial resolution for the covariances at the small scales than for

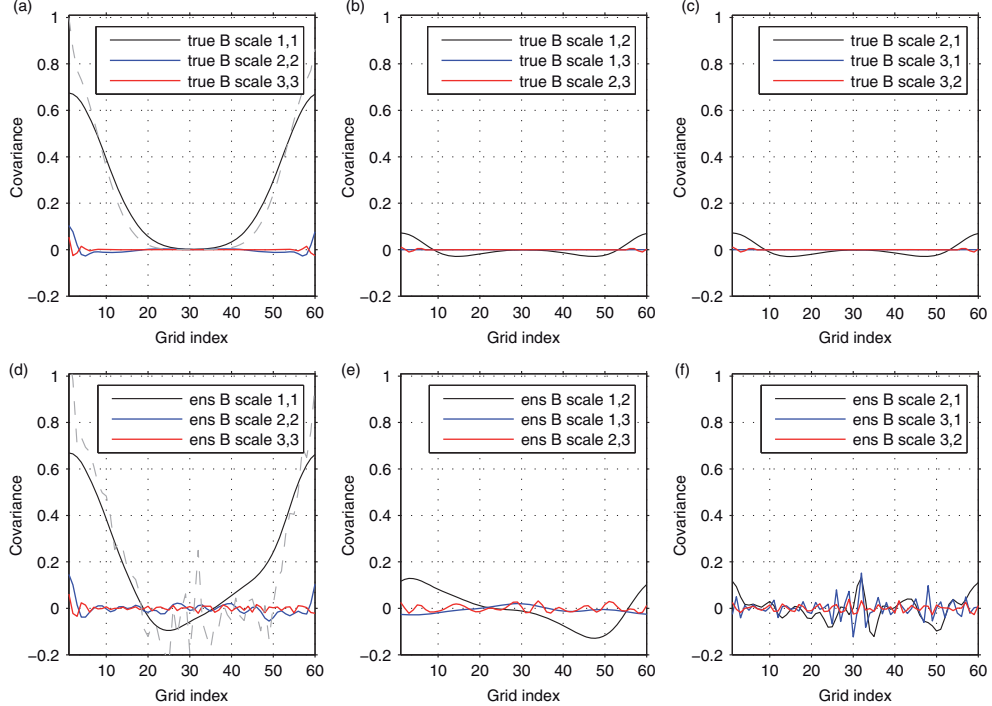


Fig. 3. The true (upper panels) and ensemble (lower panels) covariance functions after decomposition with respect to three ranges of scale for both the within-scale (left panels) and between-scale (middle and right panels) spatial covariances as described in the legends of each panel relative to the first grid point that is dominated by the large-scale covariances. For comparison, the complete true and ensemble covariance functions are also shown in dashed grey in the left panels only.

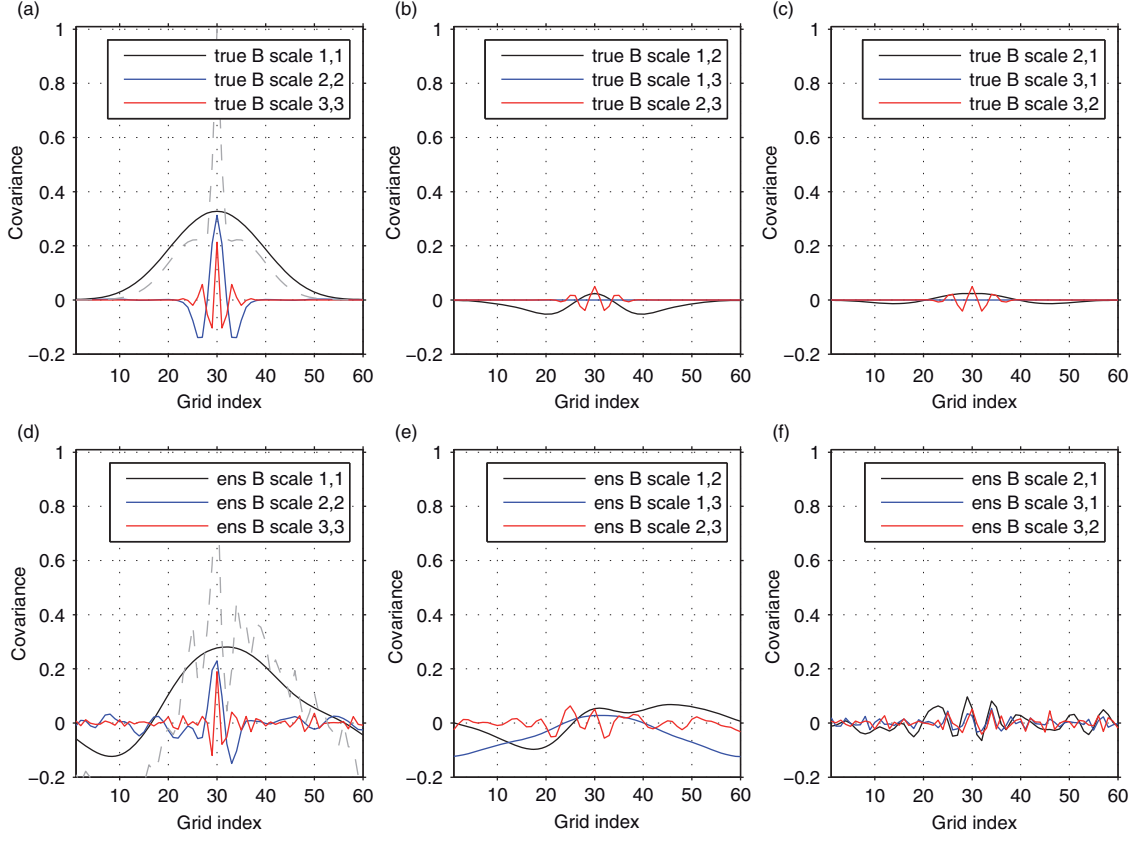


Fig. 4. Same as Fig. 3, except for covariances relative to the middle grid point that is dominated by the small-scale covariances.

the large scales. Figures 3 and 4 show the true covariances (upper panels) and a single realisation of the 50-member ensemble covariances (lower panels) with respect to the first grid point (Fig. 3) and the middle grid point (Fig. 4) after they have been decomposed into three ranges of scales

using these spectral filters. The within-scale covariances for each scale are shown (left panels) in addition to the between-scale covariances (middle and right panels). The between-scale covariances indicated in the legend as, for example, ‘scale 1, 3’ denote the covariances between all grid

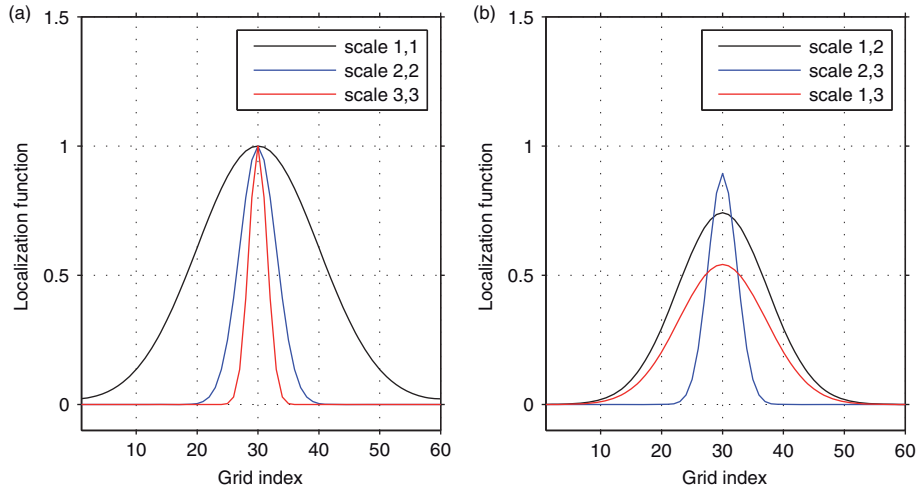


Fig. 5. Scale-dependent localisation functions (a) for each scale with itself and (b) for between every pair of scales for a single location at the middle of the spatial domain.




Fig. 6. Scale-dependent spatial localisation matrix for three scales and a one-dimensional spatial domain of 60 grid points.

points for scale 1 and a single grid point (either the first or middle grid point) for scale 3. The covariances of the small- and medium-scale components are clearly more local than those of the large-scale component. Also, the between-scale covariances have smaller amplitude than the within-scale covariances. Due to sampling error, the covariances estimated from 50 ensemble members are much less local for the small- and medium-scale components than those of the true covariances, which supports the idea of applying more severe localisation to these covariances than to those of the large scales. It is interesting to note that, from a comparison of the middle and right panels, the between-scale ensemble covariances are much less symmetric than the true between-scale covariances (though the true between-scale covariances with respect to the middle grid

point are also not perfectly symmetric). Note that, unlike a wavelet transform, we represent the scale-separated signal on the same 60 grid-point spatial domain for all scales, even though the spatial resolution could clearly be degraded for the large and medium scales without loss of information.

Figure 5a shows the chosen localisation functions for each scale with respect to the middle grid point. These are Gaussian functions modified to be periodic with spatially constant length scales of 10, 3 and 1.5 grid points. The between-scale localisation functions derived from these using eq. (9) are shown in Fig. 5b. Figure 6 shows the complete spatial/spectral localisation matrix for the three scales and 60 grid points. Note how the maximum of each between-scale localisation function is less than one, even though the specified localisation functions for the three scales all peak at one. This results from the requirement that the complete localisation matrix be positive-semidefinite to ensure that the resulting localised matrix is a valid covariance matrix (Gaspari and Cohn, 1999). The amount by which the between-scale covariances at zero separation distance (i.e., the diagonals of the off-diagonal blocks in Fig. 6) are reduced is directly related to the difference in the amount of localisation applied at the two scales. To illustrate, Fig. 7 shows the within-scale and between-scale localisation functions for a pair of scales when the localisation length scales are quite similar (10 and 8 grid points, Fig. 7a), somewhat different (10 and 3 grid points, Fig. 7b) and very different (10 and 0.2 grid points, Fig. 7c). The peak of the between-scale localisation function (shown in red) is almost equal to one when the localisation length scales are similar, but are reduced by nearly 75 % when they are very different. This unavoidable reduction of the between-scale covariances corresponds with a local spatial averaging of the covariances and therefore reduces to some extent the amount of heterogeneity in the ensemble covariances.

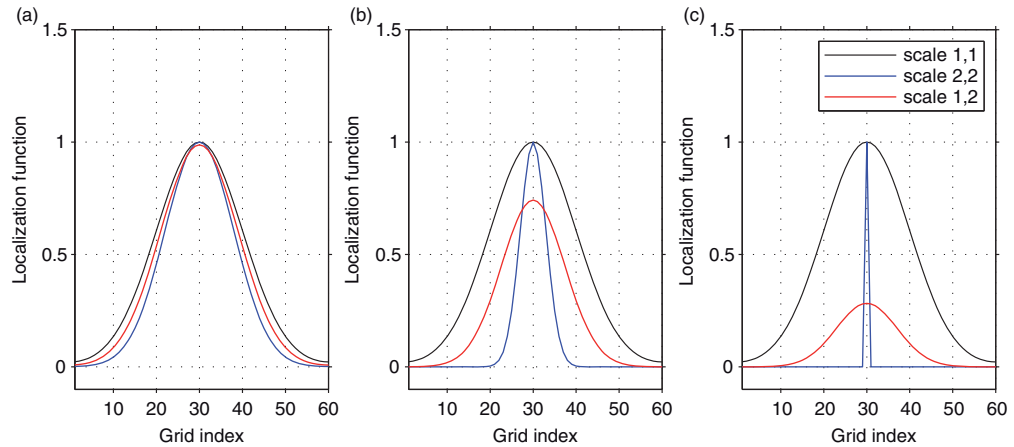


Fig. 7. The spatial localisation functions for two ranges of scales (black and blue) and the between-scale localisation function (red) when the two localisation length scales are (a) 10 and 8, (b) 10 and 3, or (c) 10 and 0.2 grid points.

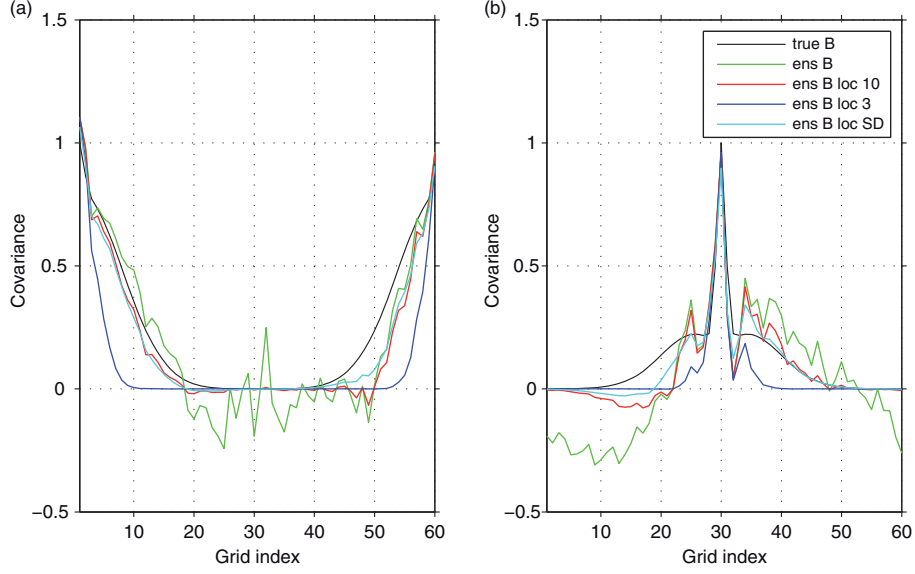


Fig. 8. Covariance functions relative to (a) the first grid point that is dominated by the large-scale covariances and (b) the middle grid point that is dominated by the small-scale covariances. The true covariances (black) are shown together with covariances estimated from a single ensemble of 50 random perturbations. The original ensemble covariances (green; same as the grey dashed curves in Fig. 3d and 4d) and the ensemble covariances with either broad localisation (red; 10 grid-point length scale), severe localisation (blue; 3 grid-point length scale) or scale-dependent localisation (cyan; 10, 3 and 1.5 grid-point length scales) are shown.

With an ensemble of only 50 members, the raw ensemble covariances differ significantly from the true covariances at all separation distances, as shown in Fig. 8. The application of a broad spatial localisation function (Gaussian function with a length scale of 10 grid points) results in a relatively close match with the true covariances for both short separation distances, where they closely resemble the raw ensemble covariances, and for large separation distances, where they are strongly attenuated. The application of more severe localisation (with a length scale of three grid points) results in a large underestimation of the true covariance, especially for the covariances with respect to the first grid point where they are dominated by the large-scale component (Fig. 8a). The use of scale-dependent localisation (with length scales of 10, 3 and 1.5 grid points, and denoted ‘SD’ in the legend) results in covariances that are quite similar to those with broad localisation for the location dominated by the large-scale component (Fig. 8a). This result is consistent with the projection of the covariances at the first grid point almost exclusively onto the largest scale waveband (as shown in Fig. 3a and d) for which the same 10 grid-point localisation function is applied as for the covariances with broad localisation. In contrast, the effect of scale-dependent localisation is noticeably stronger for the location dominated by the small-scale component (Fig. 8b). This is consistent with the similar projection onto all three wavebands of the covariances at this location (as shown in Fig. 4a and d), since the

covariances in the medium- and small-scale wavebands are more severely localised (with length scales of 3 and 1.5 grid points, respectively).

To quantitatively evaluate the accuracy of the ensemble covariances resulting from the application of each type of localisation used for Fig. 8, a set of 5000 independent 50-member ensembles were generated and the error was computed. Figure 9 shows the mean, stddev and root-mean-square of the error of each type of covariance estimate with respect to the first and middle grid points. The mean error (Fig. 9a and b) is essentially zero for the ensemble covariances with no localisation and mostly negative for each estimate with localisation. In other words, the more severe the localisation, the greater the amplitude of the mean error. Scale-dependent localisation results in a similar mean error as when using the broad localisation function for the covariances with respect to the first grid point, whereas the mean error is slightly smaller for the covariances with respect to the middle grid point. An exception to the mostly negative mean error is seen with scale-dependent localisation at the middle grid point (Fig. 9b) that results in positive mean errors at small separation distances. This positive mean error results from the more severe spatial localisation applied to the small and medium scales, which reduces the negative covariances that are present at short separation distances (see Fig. 4a and d). The stddev of the error (Fig. 9c and d) is also strongly affected by localisation. The ensemble covariance with no localisation

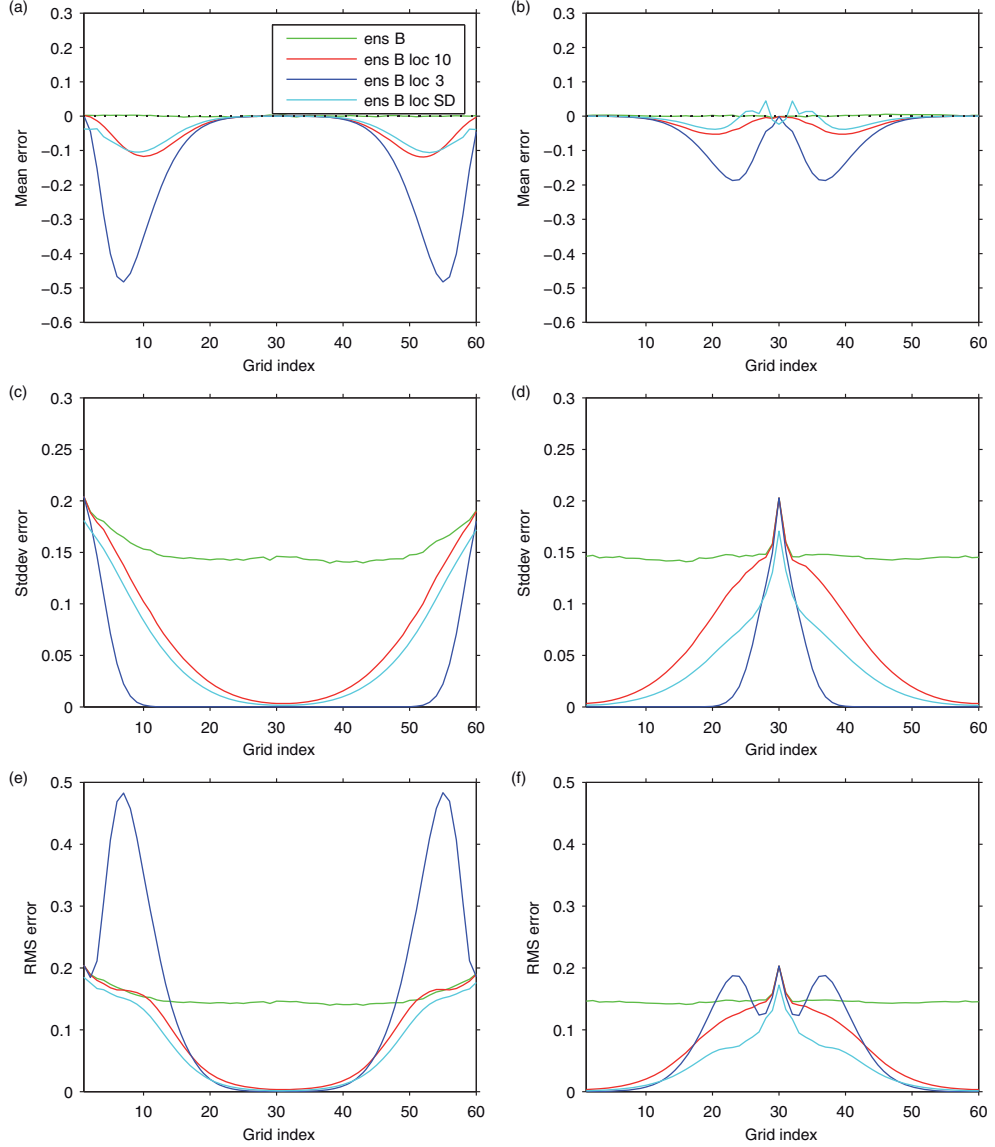


Fig. 9. Mean (upper panels), stddev (middle panels) and root-mean-square (lower panels) of the error in the estimated covariances with respect to the first (left panels) and middle (right panels) grid points computed from 5000 independent 50-member ensembles with no localisation (green), localisation with a large-length scale (red), localisation with a short-length scale (blue), and scale-dependent localisation using three different length scales (cyan).

generally has the highest stddev that is maximum at zero separation distance and decreases only slightly at large separation distances. The use of either a broad or sharp localisation function decreases the stddev of the error to a level that is proportional to the amplitude of the localisation function. Consequently, localisation results in zero stddev of the error for all grid points where the localisation function is zero. However, for separation distances where the true covariance is not close to zero, the use of severe localisation can result in a mean error larger than the stddev of the ensemble covariance with no localisation

(compare blue curves in upper and middle panels of Fig. 9). The use of scale-dependent localisation results in a small improvement relative to using the broad localisation function for the covariances with respect to the first grid point (Fig. 9c). For the middle grid point (Fig. 9d), where the covariances are dominated by the small-scale component, the use of scale-dependent localisation results in a more significant reduction to the error stddev of between 20 and 50 %. The lower panels of Fig. 9 show the root-mean-square of the error that combines the effects on the mean and stddev of the error shown previously.

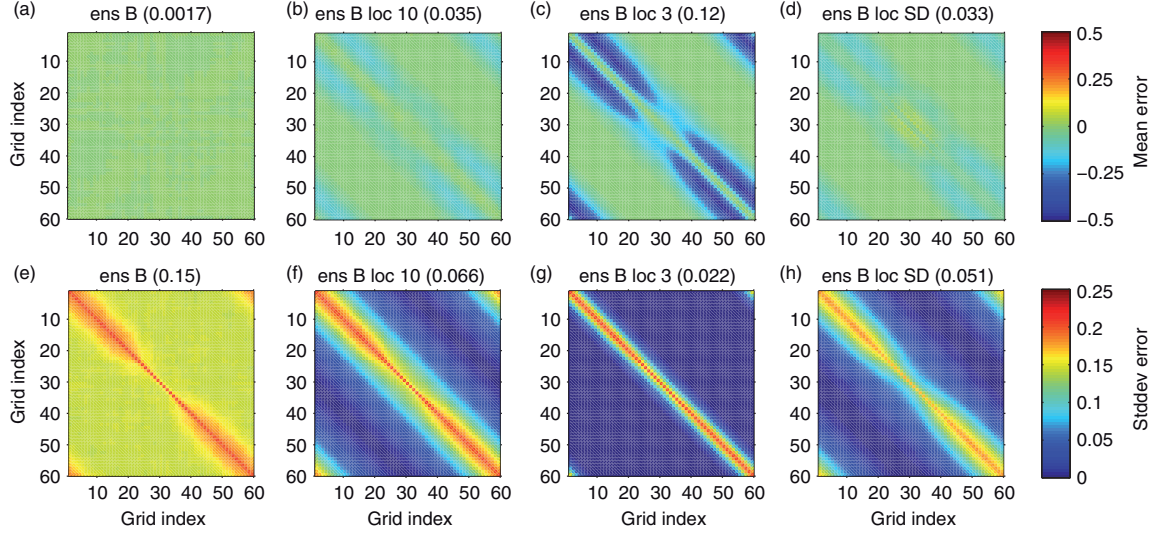


Fig. 10. The mean error (top panels) and stddev of the error (bottom panels) for (a, e) the raw ensemble covariances, the ensemble covariances with (b, f) the broad localisation function or (c, g) sharp localisation function applied, and (d, h) the ensemble covariances with scale-dependent localisation applied. The average magnitude of mean error and error stddev is shown in parentheses for each.

This demonstrates that scale-dependent localisation generally results in the lowest error, especially where the covariances are dominated by the small-scale component (Fig. 9f). Figure 10 presents essentially the same information as Fig. 9 except that the mean (upper panels) and stddev (lower panels) of the error are shown for the entire covariance matrix for each type of localisation estimate.




Fig. 11. Localisation functions used to isolate the impact of spectral localisation on the scale-dependent localisation results. The function labelled ‘scale j,j ’ refers to the localisation used for each scale with itself, while the other functions are for the between-scale localisation for a single location at the middle of the spatial domain. All functions have the same length scale of 10 grid points.

This more clearly shows how near the middle of the domain the stddev of the error is more effectively reduced with scale-dependent localisation (Fig. 10h) than with the use of broad localisation (Fig. 10f), without incurring the large mean error seen with the use of severe localisation (comparing Fig. 10c and d).

An additional localisation matrix was constructed to isolate only the influence of the unavoidable spectral localisation (i.e., the reduction in between-scale covariances) on the results from using scale-dependent localisation. This matrix has the same localisation function (with a length scale of 10 grid points) for all within-scale and between-scale covariances, except that the between-scale covariances are proportionately reduced, so that the diagonal values are equal to those of the scale-dependent localisation matrix. These localisation functions are shown in Fig. 11 and can be compared with the scale-dependent localisation functions in Fig. 5. The resulting mean and stddev of the error are shown in Fig. 12 comparing the result of using broad localisation and scale-dependent localisation with the result of using the same broad spatial localisation after it is combined with spectral localisation. This generally shows that less than half of the reduction in the error stddev for scale-dependent localisation as compared with simple spatial localisation is due to the unavoidable effect of spectral localisation. Though the reduction of the between-scale covariances reduces the sampling error, it also corresponds with a loss of some heterogeneity of the covariances. Because the true covariances used here have only a modest amount of heterogeneity, through the smoothly varying weighting of the small- and large-scale components,

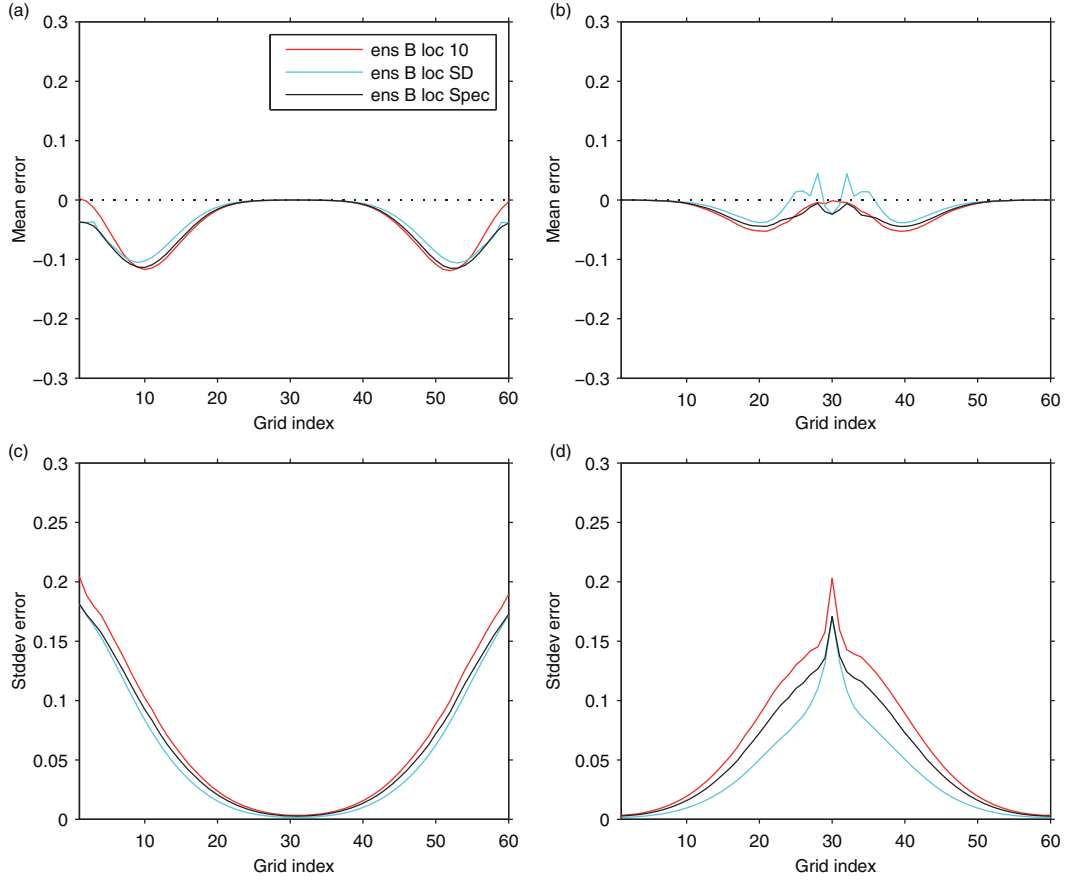


Fig. 12. Like Fig. 9, except for the estimated covariances computed using localisation with a large-length scale (red), scale-dependent localisation using three different length scales (cyan), and localisation with a large-length scale combined with the same amount of spectral localisation as for scale-dependent localisation (black).

spectral localisation leads to a reduction in the estimation error relative to using only spatial localisation. Spectral localisation may, however, result in increased error in cases with lower sampling error in the raw ensemble covariances (from using a much larger ensemble) or a higher degree of spatial heterogeneity in the true covariances.

4. Results from applying scale-dependent localisation to an EnVar sea-ice data assimilation system

Scale-dependent localisation is next applied to an ensemble of sea-ice concentration obtained from an ensemble of

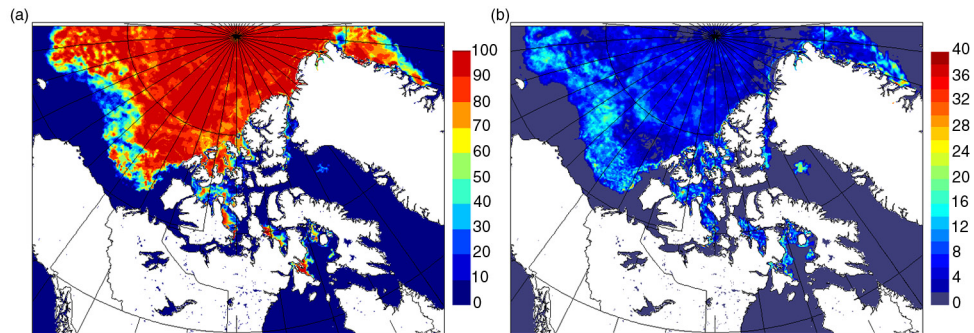


Fig. 13. The (a) background sea-ice concentration and (b) ensemble spread stddev of sea-ice concentration computed from the 60-member ensemble. Note that the units are in percent.




Fig. 14. Equivalent spectral filter coefficients computed from the application of the diffusion operator to perform the separation into four ranges of scales.

3DVar data assimilation experiments, described by Shlyueva et al. (2015), with the Environment Canada Regional Ice Prediction System. The ensemble was constructed using six-member ensembles of 6-hour forecasts from 10 consecutive assimilation times, separated by 6 hours (from 15 August 2011 at 1200 UTC to 17 August 2011 at 1800 UTC). All forecasts were taken from the configuration that employs a

sea-ice model with both dynamics and thermodynamics active. More details on the ensemble, including how the uncertainties due to model and external forcing error were simulated are given by Shlyueva et al. (2015). Figure 13 shows the background state and ensemble spread stddev of ice concentration. Note how the ensemble spread is mostly zero where the background ice concentration is zero and also generally small in the pack ice area where the background ice concentration is close to 100 %. In contrast, areas near the ice edge, where the ice concentration is between 0 and 100 %, have larger values of ensemble spread.

A diffusion operator is used for representing the spatial background-error correlations within this sea-ice 3DVar data assimilation system and not a spectral transform (Caya et al., 2010). Shlyueva et al. (2015) also used the diffusion operator to represent the spatial localisation function in preliminary experiments with EnVar. Therefore, it was decided to design a procedure that employs this diffusion operator for decomposing the ensemble perturbations with respect to scale. First, the k th ensemble perturbation is successively smoothed using the diffusion operator with a length scale of 10 km, followed by 30 and 100 km, resulting in:

$$\begin{aligned} \mathbf{f}_k^1 &= D_{10 \text{ km}} \mathbf{e}_k \\ \mathbf{f}_k^2 &= D_{30 \text{ km}} \mathbf{f}_k^1, \\ \mathbf{f}_k^3 &= D_{100 \text{ km}} \mathbf{f}_k^2 \end{aligned} \quad (16)$$

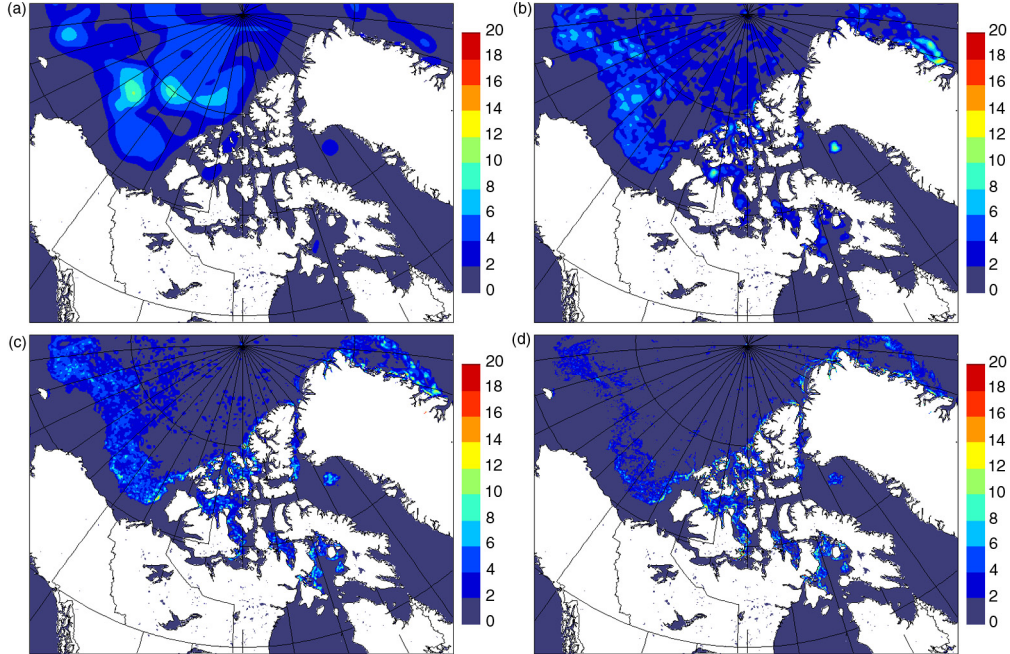


Fig. 15. The ensemble spread stddev for the 60-member sea-ice concentration ensemble for four different ranges of spatial scales, ordered from (a) largest to (d) smallest. Note that the units are in percent.

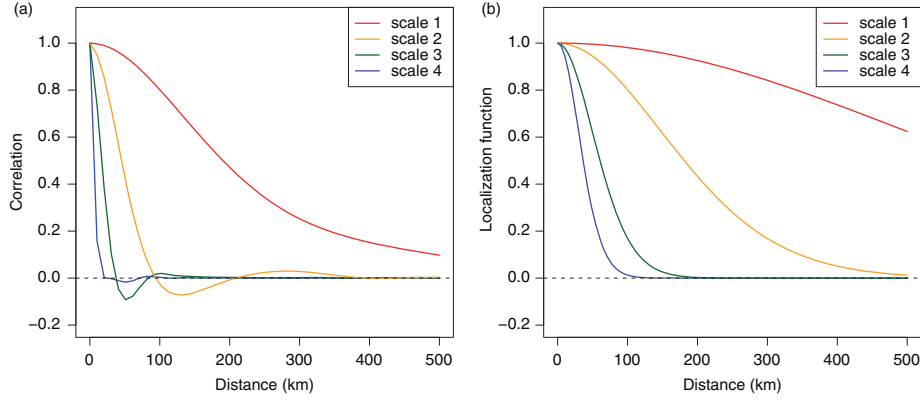


Fig. 16. (a) The homogeneous correlation functions computed from the ensemble covariances for each range of scales. (b) The specified spatial localisation functions chosen for each scale: Gaussian-like functions with length scales equal to 500, 150, 50, and 30 km for scales 1, 2, 3 and 4, respectively.

where, for example, $D_{10 \text{ km}}$ denotes the diffusion operator with a 10 km length scale. These smoothed versions of the original ensemble perturbation \mathbf{e}_k are then combined to obtain the decomposition in four ranges of horizontal scale as follows:

$$\begin{aligned} \mathbf{e}_{1,k} &= \mathbf{f}_k^3 \\ \mathbf{e}_{2,k} &= \mathbf{f}_k^2 - \mathbf{f}_k^3 \\ \mathbf{e}_{3,k} &= \mathbf{f}_k^1 - \mathbf{f}_k^2 \\ \mathbf{e}_{4,k} &= \mathbf{e}_k - \mathbf{f}_k^1 \end{aligned} \quad (17)$$

where the first subscript index for $\mathbf{e}_{j,k}$ indicates the range of scales, with 1 being the largest scale and 4 being the smallest. It is easy to show that the sum of the four scale-decomposed fields in eq. (17) results in exactly the original ensemble perturbation \mathbf{e}_k , thus satisfying eq. (5). It should be noted that this strategy for scale decomposition would not be applicable to the spectral localisation approach of Buehner (2012) since that approach requires that the sum of the squares of the filter functions sum to one. Consequently, no attempt is made to compare the current approach with that of Buehner (2012) in the context of the sea-ice data assimilation system.

Figure 14 shows the equivalent spectral filter coefficients for the resulting four ranges of scales. These were computed by applying the procedure for scale separation described above to 100 fields of spatially uncorrelated random values. The amplitudes of the two-dimensional spectral transforms were then computed for each of the four scale-decomposed fields and averaged over the 100 fields and as a function of total wavenumber. The resulting spectral filter functions generally resemble overlapping Gaussian functions with peaks at distinct wavenumbers. Unlike the functions chosen for the previous experiment (Fig. 2), however, most do not peak at a value near one due to a larger degree of overlap.

Figure 15 shows the ensemble spread computed from the 60 ensemble perturbations after they were decomposed with respect to the four ranges of scales. Note how most of the spread for scale 1 (Fig. 15a) occurs over the pack ice, whereas for the smaller scales the spread is increasingly concentrated near the ice edge.

The spatially averaged homogeneous and isotropic correlation function was computed separately from the ensemble perturbations for each range of scales (Fig. 16a). As expected, the correlation function is broadest for scale 1 and becomes increasingly local for scales 2, 3 and 4. Using these correlation functions as a guide, a length scale was chosen for the localisation function to be applied to each range of scales. Figure 16b shows the resulting localisation functions with chosen length scales of 500, 150, 50 and 30 km for scales 1, 2, 3 and 4, respectively.

To help in the interpretation of the results, two distinct regions are defined where there is sea ice in the background state. One region is referred to as the ‘pack ice’ area, defined as the area with ice concentration greater than 70 % (Fig. 17a). The other region is loosely referred to as the ‘marginal ice zone’ (MIZ), defined as the area with ice concentration between 1 and 70 %¹ (Fig. 17b). Figure 18 shows the ensemble spread averaged over these two regions as a function of scale. This demonstrates that the ensemble covariances in the pack ice area are dominated by the large scales, whereas in the MIZ area they are dominated by the smaller scales, consistent with Fig. 15.

A series of data assimilation experiments are performed using different types of covariance localisation applied to

¹Note that this definition of the MIZ is chosen for convenience, as it is more typically defined as the area where open ocean processes, including specifically ocean waves, alter significantly the dynamical properties of the sea-ice cover.

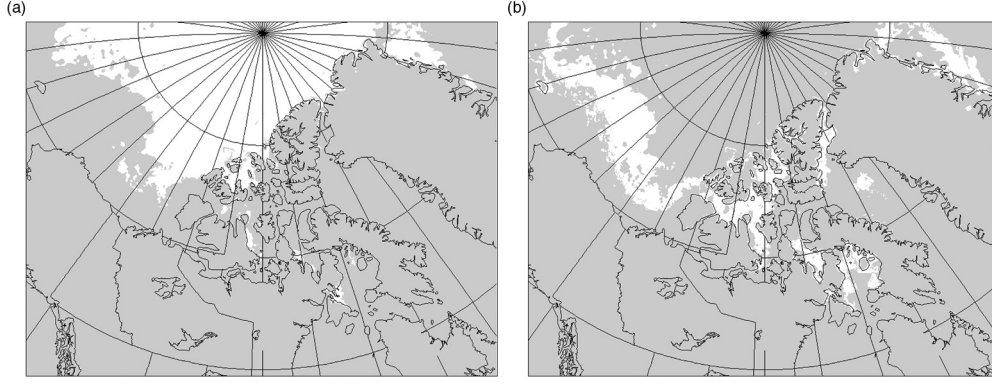


Fig. 17. (a) The pack ice area and (b) the marginal ice zone (shown in white).

the 60-member ensemble to assimilate two isolated observations located in the Beaufort Sea. The background ice concentration field and two observation locations are shown in Fig. 19a. The observations at the western and eastern locations are specified to have ice concentration values of 100 and 30 %, respectively. The analysis increment obtained when using a very broad spatial localisation with a length scale of 500 km is shown in Fig. 19b. From this result, it can be seen that the spatial covariances near the western observation location are more dominated by the large scales than at the eastern observation location. When using a more severe spatial localisation with a length scale of only 30 km, much of the increment resulting from the large-scale covariances near the western observation location are eliminated such that the spatial structure of the increment strongly resembles the localisation function (Fig. 19c). The use of scale-dependent localisation (Fig. 19d)




Fig. 18. The sea-ice concentration ensemble spread stddev (in percent) for each of four ranges of scales, averaged over the pack ice area (red) and the marginal ice zone (blue).

results in an analysis increment around the western observation location that is similar to using the broad localisation function, and near the eastern observation location the increment is much more similar to that obtained with severe localisation. We speculate that this feature of scale-dependent localisation is qualitatively similar to an adaptive localisation approach that attempts to estimate a spatially varying localisation function that is locally appropriate for application to the ensemble covariances.

A series of idealised experiments are conducted to quantitatively evaluate the accuracy of ice concentration analyses resulting from assimilating a realistic network of observations while using each type of covariance localisation. For each type of localisation, the data assimilation experiment is repeated 60 times, each time using a different member from the ensemble as the true state. The observations are simulated by extracting values from the true state at the chosen observation locations (every fourth grid point and with random gaps to simulate the effect of clouds on visible/infrared satellite observations) and perturbing them with a random Gaussian variable with stddev equal to 10 %. The stddev used for the perturbations are also used to specify the observation error covariance matrix for assimilation. The background state is taken as the true state plus the ensemble perturbation from another randomly chosen member, ensuring that the error in the background state is completely consistent with the ensemble covariances. The background-error covariance matrix used for assimilation is computed only from the remaining ensemble perturbations that are considered to be independent of the perturbation used to generate the background state. This independence is obtained by using the 45 perturbations from different analysis times (eliminating 6 of the 60 members) and from different members of the six-member ensemble (eliminating another nine members) as the member used to generate the background state.

The resulting root-mean-square error of the background state and analysis averaged both spatially and over all

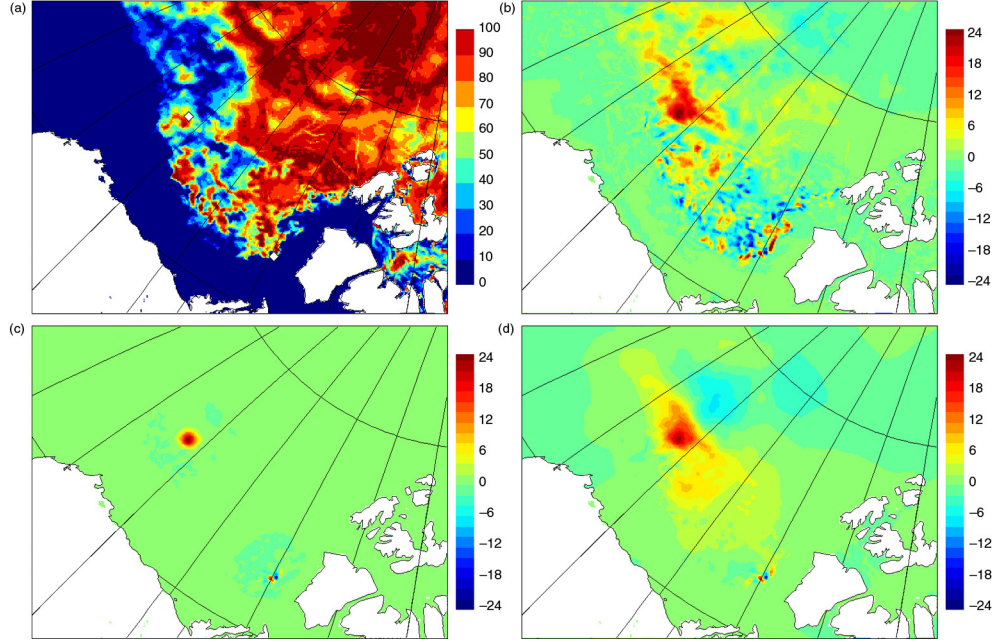


Fig. 19. (a) Background sea-ice concentration and locations of two observations (denoted by the white diamonds) in the Beaufort Sea that were assimilated. Resulting ice concentration analysis increment from assimilating these observations are shown when using (b) localisation with 500 km length scale, (c) localisation with 30 km length scale and (d) scale-dependent localisation with 500, 150, 50 and 30 km length scales. Note that the units are in percent.

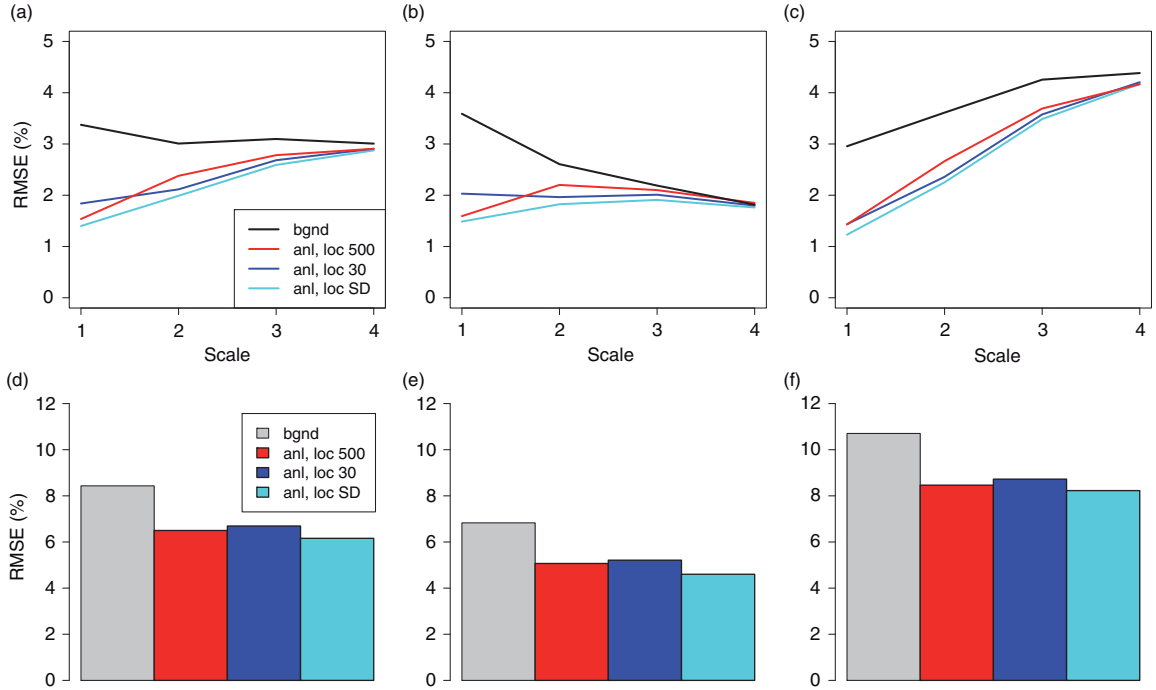


Fig. 20. Background (black/grey) and analysis root-mean-square error (in percent) when using localisation with 500 km (red) or 30 km (blue) length scale or scale-dependent localisation (cyan). The error is shown both as a function of horizontal scale (upper panels) and as the total error for all scales (lower panels). Left panels show the error averaged over the entire domain where the ice concentration is greater than 1 %, middle panels show the error averaged over only the pack ice area, and the right panels show the error averaged over only the marginal ice zone.

60 experiments for each type of covariance localisation is shown in Fig. 20. The top panels show the error decomposed with respect to the four scales and the lower panels show the total error. The error averaged over the entire domain where the ice concentration is greater than 1 % (Fig. 20a and d) shows an improvement, both overall and for all four scales, from using scale-dependent localisation as compared with using either broad (500 km length scale) or severe (30 km length scale) localisation. For the largest scales, the analysis error with scale-dependent localisation is similar to, but slightly lower than, when using broad localisation. For the smaller scales, scale-dependent localisation results in error that is similar to, but slightly lower than, when using severe localisation. These improvements from using scale-dependent localisation are also seen when evaluated for just the pack ice area (Fig. 20b and e) and the MIZ (Fig. 20c and f).

5. Conclusions

A new approach for applying scale-dependent spatial localisation to ensemble covariances within an EnVar data assimilation approach was presented. The mathematical formulation was compared with the formulations of standard spatial localisation and the spatial/spectral localisation of Buehner (2012). The application of the new approach was shown to result in more accurate covariance estimates than when using different amounts of standard spatial localisation in a one-dimensional idealised numerical experiment. The unavoidable effect of reducing the between-scale covariances (that is, partial spectral localisation) contributes a small amount to this improvement. Using a two-dimensional ensemble of sea-ice concentration, the approach was shown to produce lower analysis error as compared with standard spatial localisation. The assimilation of isolated individual observations also demonstrated that strong spatial gradients in the relative contribution of different spatial scales in the ensemble covariances results in strong spatial variations in the overall amount of spatial localisation, making this qualitatively similar to adaptive localisation approaches (e.g. Bishop and Hodyss, 2007, 2009).

Though the focus of this study was on spatial localisation in only one or two horizontal dimensions, the treatment of three- or four-dimensional covariances would be important for applications to NWP. Specifically, for three-dimensional covariances, the formulation given in eq. (12) would allow for a different amount of vertical localisation to be applied for each range of horizontal scales, if appropriate. In 4DEnVar (e.g. Buehner et al., 2013), a temporal localisation function could be introduced that was also dependent on the horizontal scale such that small-

scale errors could be forced to decorrelate more rapidly through time than large-scale errors.

6. Acknowledgements

The authors thank Jean-François Caron and two anonymous reviewers whose comments helped to improve an earlier version of the paper. M. B. also thanks Chris Snyder for helpful initial discussions.

References

- Anderson, J. and Lei, L. 2013. Empirical localization of observation impact in ensemble Kalman filters. *Mon. Weather Rev.* **141**, 4140–4153.
- Berre, L. and Desroziers, G. 2010. Filtering of background error variances and correlations by local spatial averaging: a review. *Mon. Weather Rev.* **138**, 3693–3720.
- Bishop, C. H. and Hodyss, D. 2007. Flow adaptive moderation of spurious ensemble correlations and its use in ensemble-based data assimilation. *Q. J. Roy. Meteorol. Soc.* **133**, 2029–2044.
- Bishop, C. H. and Hodyss, D. 2009. Ensemble covariances adaptively localized with ECO-RAP. Part 2: a strategy for the atmosphere. *Tellus A*. **61**, 97–111.
- Buehner, M. 2005. Ensemble-derived stationary and flow-dependent background-error covariances: evaluation in a quasi-operational NWP setting. *Q. J. Roy. Meteorol. Soc.* **131**, 1013–1043.
- Buehner, M. 2012. Evaluation of a spatial/spectral covariance localization approach for atmospheric data assimilation. *Mon. Weather Rev.* **140**, 617–636.
- Buehner, M. and Charron, M. 2007. Spectral and spatial localization of background-error correlations for data assimilation. *Q. J. Roy. Meteorol. Soc.* **133**, 615–630.
- Buehner, M., Houtekamer, P. L., Charette, C., Mitchell, H. L. and He, B. 2010. Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part II: one-month experiments with real observations. *Mon. Weather Rev.* **138**, 1567–1586. DOI: <http://dx.doi.org/10.1175/2009MWR3158.1>
- Buehner, M., Morneau, J. and Charette, C. 2013. Four-dimensional ensemble-variational data assimilation for global deterministic weather prediction. *Nonlin. Processes Geophys.* **20**, 669–682. DOI: <http://dx.doi.org/10.5194/npg-20-669-2013>
- Caya, A., Buehner, M. and Carrieres, T. 2010. Analysis and forecasting of sea ice conditions with three-dimensional variational data assimilation and a coupled ice–ocean model. *J. Atmos. Oceanic Technol.* **27**, 353–369. DOI: <http://dx.doi.org/10.1175/2009JTECHO701.1>
- Caya, A., Sun, J. and Snyder, C. 2005. A comparison between the 4DVAR and the ensemble Kalman filter techniques for radar data assimilation. *Mon. Weather Rev.* **133**, 3081–3094. DOI: <http://dx.doi.org/10.1175/MWR3021.1>
- Flowerdew, J. 2015. Towards a theory of optimal localization. *Tellus A*. **67**, 25257. DOI: <http://dx.doi.org/10.3402/tellusa.v67.25257>

- Gaspari, G. and Cohn, S. E. 1999. Construction of correlation functions in two and three dimensions. *Q. J. Roy. Meteorol. Soc.* **125**, 723–757. DOI: <http://dx.doi.org/10.1002/qj.49712555417>
- Hamill, T. M., Whitaker, J. S. and Snyder, C. 2001. Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Weather Rev.* **129**, 2776–2790.
- Houtekamer, P. L., Deng, X., Mitchell, H. L., Baek, S.-J. and Gagnon, N. 2014. Higher resolution in an operational ensemble Kalman filter. *Mon. Weather Rev.* **142**, 1143–1162. DOI: <http://dx.doi.org/10.1175/MWR-D-13-00138.1>
- Houtekamer, P. L. and Mitchell, H. L. 2001. A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Weather Rev.* **129**, 123–137.
- Lei, L. and Anderson, J. L. 2014. Empirical localization of observations for serial ensemble Kalman filter data assimilation in an atmospheric general circulation model. *Mon. Weather Rev.* **142**, 1835–1851.
- Li, Z., McWilliams, J., Ide, K. and Farrara, J. 2015. A multiscale variational data assimilation scheme: formulation and illustration. *Mon. Weather Rev.* **143**(9), 3804. DOI: <http://dx.doi.org/10.1175/MWR-D-14-00384.1>
- Lorenc, A. 2003. The potential of the ensemble Kalman filter for NWP – a comparison with 4D-Var. *Q. J. Roy. Meteorol. Soc.* **129**, 3183–3203. DOI: <http://dx.doi.org/10.1256/qj.02.132>
- Ménétrier, B., Montmerle, T., Michel, Y. and Berre, L. 2015. Linear filtering of sample covariances for ensemble-based data assimilation. Part I: optimality criteria and application to variance filtering and covariance localization. *Mon. Weather Rev.* **143**, 1622–1643. DOI: <http://dx.doi.org/10.1175/MWR-D-14-00157.1>
- Miyoshi, T. and Kondo, K. 2013. A multi-scale localization approach to an ensemble Kalman filter. *SOLA*. **9**, 170–173. DOI: <http://dx.doi.org/10.2151/sola.2013-038>
- Shlyueva, A., Buehner, M., Caya, A., Lemieux, J.-F., Smith, G. and co-authors. 2015. Ensemble data assimilation for the Environment Canada Regional Ice Prediction System. *Q. J. Roy. Meteorol. Soc.* DOI: <http://dx.doi.org/10.1002/qj.2712>
- Zhang, F., Weng, Y., Sippel, J. A., Meng, Z. and Bishop, C. H. 2009. Cloud-resolving hurricane initialization and prediction through assimilation of Doppler Radar observations with an ensemble Kalman filter. *Mon. Weather Rev.* **137**, 2105–2125. DOI: <http://dx.doi.org/10.1175/2009MWR2645.1>