



The Kalman filter or its ensemble version, the ensemble Kalman filter, is optimal for a linear model and -measurement operator. This chapter will comprehensively discuss the EnKF analysis scheme and its properties, focusing on an ensemble-subspace computation of the inverse. We demonstrate the importance of taking into account correlations in the model errors. Furthermore, we study the efficient ensemble-subspace inversion method that allows computing the analysis update at a linear cost in both the number of measurements and the state dimension. We also show how to reduce sampling errors by increasing the number of measurement-perturbations realizations used to represent the measurement-error-covariance matrix.

13.1 EnKF Update Example

We will use an example from Evensen (2021) in the following discussion. The purpose of this example is to illustrate the properties of the update scheme of the ensemble Kalman filter (EnKF) as described in Chap. 8, when using the measurement perturbations to represent the measurement error covariance matrix. In addition, this example verifies the robustness of the projection of the measurement error covariance matrix onto the ensemble of predicted measurements. Finally, the update is identical to the EnRML algorithm's solution in the linear case, and as such, the results are representative of the EnRML smoother update.

The test example uses a one-dimensional periodic domain with 1024 gridpoints and $\Delta x = 1$. In this domain, we simulate a smooth pseudo-random function with mean $\mu = 4$, variance $\sigma^2 = 1$, and decorrelation length $r_d = 40$, representing the unknown truth,

$$\mathbf{z}_{\text{true}} \sim \mathcal{N}(\mu = 4, \sigma^2 = 1, r_d = 40). \quad (13.1)$$

The constant $\mu = 4$, is just added for plotting purposes.

The first guess solution is generated by simulating another realization $\mathbf{z} \sim \mathcal{N}(0, 1, 40)$ and adding it to the truth, i.e.,

$$\mathbf{z}_{\text{fg}} = \frac{\mathbf{z} + \mathbf{z}_{\text{true}} - 4}{\sqrt{2}} + 4. \quad (13.2)$$

The factor $\sqrt{2}$ ensures that the variance of \mathbf{z}_{fg} is equal to one.

The initial ensemble is created by adding random realizations $\mathbf{z}_j \sim \mathcal{N}(0, 1, 40)$ to the first guess \mathbf{z}_{fg} ,

$$\mathbf{z}_j^f = \mathbf{z}_{\text{fg}} + \mathbf{z}_j. \quad (13.3)$$

The measurements are distributed uniformly over the domain and sampled from a perturbed true solution according to

$$\mathbf{d}_j = \mathbf{H}(\mathbf{z}_{\text{true}} + \mathbf{z}_j), \quad (13.4)$$

with either uncorrelated, $\mathbf{z}_j \sim \mathcal{N}(0, 0.25, 0)$, or Gaussian, $\mathbf{z}_j \sim \mathcal{N}(0, 0.25, 40)$, perturbations. Here, \mathbf{H} is the linear measurement operator that extracts the measurements from the functions $\mathbf{z}_{\text{true}} + \mathbf{z}_j$.

The following experiments use the same reference truth, measurements, initial ensemble, and random seed. However, when it is essential to eliminate sampling errors, we use extended ensemble sizes. Thus, although there are seed dependencies in the obtained solutions, the different methods should produce the same answer. Furthermore, in most experiments, we can attribute differences in the results to the methods used. Thus, our approach is different from running multiple data-assimilation experiments with varying seeds.

13.2 Solution Methods

In the following, we study two prominent cases, one with a diagonal measurement error-covariance matrix and another with correlated errors, \mathbf{e}_i , simulated from Eq. (13.4) and with $r_d = 40$. For the two cases of uncorrelated and correlated measurement errors, the EnKF computes the analysis using either an exactly specified measurement error-covariance matrix \mathbf{C}_{dd} or representing the measurement error covariance by the perturbations in \mathbf{E} . Since \mathbf{H} is a linear operator, we follow the EnKF update Eq. (7.3), rather than Eq. (7.12) because for this linear example, we do not need to consider the modification in Eq. (7.11).

The case with a full-rank measurement error-covariance matrix solves

$$\mathbf{X}^a = \mathbf{X}^f + \mathbf{A}\mathbf{S}^T(\mathbf{S}\mathbf{S}^T + \mathbf{C}_{dd})^{-1}(\mathbf{D} - \mathbf{H}\mathbf{X}), \quad (13.5)$$

where the ensemble perturbation matrix is $\mathbf{A} = \mathbf{X}\mathbf{\Pi}$. The matrix $\mathbf{C} = \mathbf{S}\mathbf{S}^T + \mathbf{C}_{dd}$ is formed and then inverted by computing an eigenvalue decomposition $\mathbf{C} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$. The inverse is just $\mathbf{C} = \mathbf{Q}\mathbf{\Lambda}^+\mathbf{Q}^T$ where the use of a pseudo inverse is needed in case the matrix \mathbf{C} is poorly conditioned.

When using an ensemble representation for the measurement error-covariance matrix $\mathbf{C}_{dd} = \mathbf{E}\mathbf{E}^T$, we can solve for the update from

$$\mathbf{X}^a = \mathbf{X}^f + \mathbf{A}\mathbf{S}^T(\mathbf{S}\mathbf{S}^T + \mathbf{E}\mathbf{E}^T)^{-1}(\mathbf{D} - \mathbf{H}\mathbf{X}). \quad (13.6)$$

In the examples below, the line labels used in the figures indicate the scheme used to compute the matrix inversion. The line label Cdd denotes using the standard EnKF analysis equation with a full rank measurement error-covariance matrix \mathbf{C}_{dd} , as explained above. The curves with line label EE correspond to the EnKF update when the samples in \mathbf{E} replace the “exact” analytic measurement error covariance matrix \mathbf{C}_{dd} , and we use the ensemble subspace scheme.

Using \mathbf{E} to represent the measurement error-covariance matrix introduces additional sampling errors. However, we will see below how to reduce these sampling errors to a negligible level with a simple algorithm modification. I.e., one uses a larger number of realizations in \mathbf{E} to represent \mathbf{C}_{dd} better. The code used is the test case from https://github.com/geirev/EnKF_analysis.git.

13.3 Example 1 (Large Ensemble Size)

The first example uses 50 measurements and a large ensemble size of 2000 to reduce sampling errors. Figure 13.1 shows the results for the two cases with either diagonal or correlated measurement errors.

The upper-left plot shows the EnKF estimates for the case with uncorrelated measurement errors for each grid cell numbered with indexes 1–1024. The two schemes, represented by the lines labeled Cdd and EE, give similar results in this case. The upper-right plot shows the prior and posterior error variances for the two updates, and again they are nearly identical. Finally, the lower-left plot presents the EnKF estimates for the case with correlated measurement errors. In this case, we also see that the results using the exact and approximate schemes (Cdd and EE) are nearly identical.

An apparent difference between the two cases is that, with uncorrelated errors, the measurements are scattered randomly about the correct solution. In contrast, with correlated measurement errors, successive measurements will have similar error values, and they follow a smooth curve. The nonzero measurement correlations’ role is to reduce the strength of the update, and the result is an update with a more substantial variance. By taking the measurement error correlations into account, we inform EnKF that neighboring measurements make the same error and reduce their accumulated impact.




Fig. 13.1 Simple update example: The upper plots present the results for a case with uncorrelated measurement errors, while the lower plots give the results when using measurements with correlated errors and decorrelation length $r_d = 40$. The left plots show the results for the posterior ensemble means, while the right plots provide the associated error estimates. The line labels Cdd, EE, denote different numerical implementations of the inversion scheme used, as is explained in the text. The ensemble size is 2000

13.4 Example 2 (Ensemble Size of 100)

We now repeat the previous experiment from Example 1 using a more common ensemble size of 100. The purpose is to illustrate the impact of sampling errors when using the measurement error perturbations in \mathbf{E} to represent \mathbf{C}_{dd} . Figure 13.2 shows that with 100 realizations, the additional sampling errors introduced by scheme EE lead to a slight deviation between the two estimates. More problematic is the underestimation of the ensemble variance. In a sequential data-assimilation context, this underestimation would have to be compensated for, e.g., by using inflation, to avoid possible filter divergence. In the following example, we will learn how to reduce these sampling errors to a negligible level.

13.5 Example 3 (Augmenting the Measurement Perturbations)

The benefit of using Eq. (13.6) over Eq. (13.5) is the reduced computational cost, but also the fact that it is easier to sample perturbations with accurate statistics than constructing a full-rank measurement error covariance matrix. An approach for re-




Fig. 13.2 Simple update example: Same as Fig. 13.1 but using an ensemble size of 100 realizations

ducing the sampling errors in scheme EE is to augment columns of new realizations of measurement perturbations to \mathbf{E} . This modification only slightly increases the computational cost of the algorithm when computing $\mathbf{\Sigma}^+ \mathbf{U}^T \mathbf{E}$ in Eq. (8.48) and is a simple modification of the code. Figure 13.3 shows the results using 100 realizations and correlated measurement errors, and when using 1000 samples in \mathbf{E} . The augmentation of additional columns to \mathbf{E} in Exp. 3 significantly reduces the errors in the estimated means and variances for the two cases with correlated and uncorrelated measurement errors compared with the results from Exp. 2. It is clear that the two schemes Cdd and EE, solving Eqs. (13.5) and (13.6) respectively, give almost identical results. In this case, the measurement error perturbations' projection onto the ensemble subspace does not significantly impact the results. Thus, the sampling errors introduced by using \mathbf{E} to represent \mathbf{C}_{dd} can be made negligible by increasing the sample size in \mathbf{E} to only a minor additional cost.

Evensen (2021) found that the algorithm works well with different measurement error decorrelation lengths. When the measurement perturbations include small scales not represented by the predicted measurements' ensemble, the projection onto the ensemble anomalies in \mathbf{S} introduces an approximation. The truncation of small scales in the measurement errors leads to a slight underestimation of the measurement error variance.




Fig. 13.3 Simple update example: Same as Fig. 13.2 but using an ensemble size of 1000 realizations to represent E

13.6 Example 4 (Large Number of Measurements)

Figure 13.4 shows results from the final example where the number of measurements increased from 50 to 200, i.e., twice the ensemble size. In this case, we apply a truncation at 99% of the variance when computing the inversion, retaining 29 singular values when computing the singular value decomposition of S . Again, the results obtained are very similar using the two algorithms. It is also interesting to see how




Fig. 13.4 Simple update example: Same as Fig. 13.3 but for a case with 200 measurements, which is twice the ensemble size, and with measurement-error correlations $r_d = 40$

the measurements' impact reduces at the grid cells with indices 400–500. Note that there is no indication of the so-called “ensemble degeneracy,” and the analysis ensemble retains a significant variance. The posterior variance using 200 measurements is similar to the one obtained using only 50 observations. This result indicates that including additional dependent measurements does not introduce much new information in this example.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

