

A two-stage ensemble Kalman filter for smooth data assimilation

Craig J. Johns · Jan Mandel

Received: 1 March 2005 / Revised: 1 September 2005 / Published online: 14 September 2007
© Springer Science+Business Media, LLC 2007

Abstract The ensemble Kalman Filter (EnKF) applied to a simple fire propagation model by a nonlinear convection-diffusion-reaction partial differential equation breaks down because the EnKF creates nonphysical ensemble members with large gradients. A modification of the EnKF is proposed by adding a regularization term that penalizes large gradients. The method is implemented by applying the EnKF formulas twice, with the regularization term as another observation. The regularization step is also interpreted as a shrinkage of the prior distribution. Numerical results are given to illustrate success of the new method.

Keywords Data assimilation · Ensemble Kalman filter · State-space model · Penalty · Tikhonov regularization · Wildfire · Convection-reaction-diffusion · Shrinkage · Bayesian

1 Introduction

The discrete time state-space model in its most general form is an application of the Bayesian update problem: the modeled system is advanced in time until an *analysis* time, when the distribution of the system state before the update, called the *prior* or the *forecast* distribution, and the *data likelihood* are combined to give the new system state distribution, called the *posterior* or the *analysis* distribution. The system is then advanced until the next analysis time. Kalman (1960) and Kalman and Bucy (1961) provided simple recursive formulas for the system mean and covariance under the assumptions that the probability distributions are normal and the system is linear. The Kalman filter is popular in areas as diverse as medicine (Jones 1984), economics (Shumway and Stoffer 1982) and geosciences (Evensen 2003). Variants were devised for, e.g., nonlinear problems (Jazwinski 1970), missing observations

C. J. Johns
Milliman, Inc., Denver, CO, USA

J. Mandel (✉)
Department of Mathematical Sciences, University of Colorado at Denver and Health Sciences Center,
Denver, CO, USA
e-mail: Jan.Mandel@cudenver.edu

(Shumway and Stoffer 1982), censored observations (Johns and Shumway 2005), irregular observation times (Jones 1984), nonlinear updates (Harvey 1989), and non-Gaussian distributions (Carlin et al. 1992; Kitagawa 1987; Meinhold and Singpurwalla 1989).

Traditional Kalman filters implicitly manipulate the covariance matrix of the state, and they are thus unsuitable for systems with a large number of degrees of freedom, such as in computational models in geophysics. Ensemble Kalman Filters (EnKFs) were developed (Evensen 1994; Houtekamer and Mitchell 1998) that represent the distribution of the system state using a random sample, called an ensemble, and do not use the covariance matrix explicitly. The benefit of the EnKF comes in situations where the eigenvalues of the covariance matrix rapidly decay. In that case, even a few ensemble members can reproduce the large-scale behavior of the covariance behavior of the system. This situation is typical of models of governed by partial differential equations, such as in geophysical systems. For related filters relaxing or removing the Gaussian assumption, see Anderson and Anderson (1999), Bengtsson (2003), Doucet et al. (2001), van Leeuwen (2003). For comprehensive surveys, see Evensen (2003, 2004), Tippett (2003).

This work is part of an effort to build a Dynamic Data Driven Application System (DDDSS) for wildfires. The method proposed in this paper was motivated by the observation that straightforward application of EnKF to a simple wildfire model (Mandel et al. 2004a) always fails within a few analysis cycles. Due to statistical variability of the ensemble members, locations with large temperature gradients develop, resulting in bigger fires and some ensemble members move away from the truth. The EnKF update formulas, trying hopelessly to match the observations within the span of the ensemble, result in states that are nonphysical (too large, too small, or too rough) in some places, which causes a complete breakdown of the simulations in subsequent advancements.

We use a Bayesian approach similar to Tikhonov regularization (Hansen 1998) to force the analysis solution to be more spatially smooth. Tikhonov regularization techniques were used in Kalman filters in a different way than described here, to stabilize ill-conditioned parameter identification (Johansen 1997; Kim et al. 2002). The beauty and utility of the method comes from the implementation; regularization requires simply using the EnKF update formulas twice, thus avoiding the need for a new code.

The paper is organized as follows. In Sect. 2, we set the stage by briefly reviewing the Kalman filter. The Ensemble Kalman filter and its implementation using contemporary numerical software are considered in Sect. 3. In Sect. 4, we add regularization as an independent observation, leading to a two stage EnKF. Finally, results for a simple fire model problem are presented in Sect. 5.

2 The Kalman filter

We consider the state space model, in which the modeled quantity is the probability distribution of the state vector \mathbf{x} . The probability distribution is evolved in time by running the model until the end of an analysis cycle, when it is updated to account for new data. At the end of the cycle, the probability density $p(\mathbf{x})$ of the system state \mathbf{x} before the update (the prior) and the probability density $p(\mathbf{y}|\mathbf{x})$ of the data \mathbf{y} given an assumed value of the system state \mathbf{x} (the data likelihood) are combined to give the new probability density of the system state $p(\mathbf{x}|\mathbf{y})$ (the posterior) by the Bayes theorem,

$$p(\mathbf{x}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x}) p(\mathbf{x}), \quad (1)$$

where α means proportionality. Equation 1 determines the posterior density $p(\mathbf{x}|\mathbf{y})$ completely because $\int p(\mathbf{x}|\mathbf{y}) d\omega(\mathbf{x}) = 1$. Consider the case of linear observation operator H : given system state \mathbf{x} , the data value, \mathbf{y} , would be $H\mathbf{x}$ if the model and the data were perfect with no errors. Of course, in general, the given data $\mathbf{y} \neq H\mathbf{x}$, so discrepancies are modeled with the likelihood $p(\mathbf{y}|\mathbf{x})$. Assume that the prior has normal distribution with mean μ and covariance Q , and the data likelihood is normal with mean $H\mathbf{x}$ and covariance R ,

$$p(\mathbf{x}) \propto \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T Q^{-1}(\mathbf{x} - \mu)\right),$$

$$p(\mathbf{y}|\mathbf{x}) \propto \exp\left(-\frac{1}{2}(\mathbf{y} - H\mathbf{x})^T R^{-1}(\mathbf{y} - H\mathbf{x})\right).$$

Denote the posterior system state by $\hat{\mathbf{x}}$ instead of $\mathbf{x}|\mathbf{y}$. It can be shown by algebraic manipulations (Anderson and Moore 1979) that the posterior is also normal,

$$p(\hat{\mathbf{x}}) \propto \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mu})^T P^{-1}(\hat{\mathbf{x}} - \hat{\mu})\right),$$

where the posterior mean $\hat{\mu}$ and covariance P are given by the update formulas

$$\hat{\mu} = \mu + K(\mathbf{y} - H\mu), \quad P = (I - KH)Q, \tag{2}$$

$$K = QH^T(HQH^T + R)^{-1}. \tag{3}$$

The matrix K is called the Kalman gain matrix. The observation (Meinhold and Singpurwalla 1983; Paige and Saunders 1977) in the following Lemma interprets the Kalman filter as least squares: the posterior mean $\hat{\mu}$ is obtained by trying to match the observation, $H\hat{\mu} \approx \mathbf{y}$, as well as to preserve the mean, $\hat{\mu} \approx \mu$.

Lemma 1 *If $\hat{\mu}$ is defined by (2) and (3), then $\hat{\mu}$ is the solution \mathbf{x} of the least-squares problem*

$$S(\mathbf{x}) = (\mathbf{x} - \mu)^T Q^{-1}(\mathbf{x} - \mu) + (\mathbf{y} - H\mathbf{x})^T R^{-1}(\mathbf{y} - H\mathbf{x}) \rightarrow \min_{\mathbf{x}}. \tag{4}$$

Proof At the minimum,

$$\nabla S(\mathbf{x}) = 2Q^{-1}(\mathbf{x} - \mu) - 2H^T R^{-1}(\mathbf{y} - H\mathbf{x}) = 0,$$

which gives $\mathbf{x} = \tilde{P}(Q^{-1}\mu + H^T R^{-1}\mathbf{y})$, where

$$\tilde{P} = (Q^{-1} + H^T R^{-1}H)^{-1} = [Q - QH^T(HQH^T + R)^{-1}HQ] = (I - KH)Q = P.$$

Consequently, $\mathbf{x} = P(Q^{-1}\mu + H^T R^{-1}\mathbf{y}) = \mu + K(\mathbf{y} - H\mu)$. □

The next Lemma is an elementary consequence of the Bayes theorem (1).

Lemma 2 *Let \mathbf{y} and \mathbf{z} be observations such that, conditional on \mathbf{x} , the error distributions are independent. Then assimilating the observations \mathbf{y}, \mathbf{z} jointly gives the same result as first assimilating the observation \mathbf{y} to obtain the posterior $\propto p(\mathbf{y}|\mathbf{x}) p(\mathbf{x})$, then taking this posterior to be the new prior and assimilating \mathbf{z} .*

Proof From the Bayes theorem,

$$p(\mathbf{x}|\mathbf{y}, \mathbf{z}) \propto p(\mathbf{y}, \mathbf{z}|\mathbf{x}) p(\mathbf{x}) = p(\mathbf{z}|\mathbf{x}) p(\mathbf{y}|\mathbf{x}) p(\mathbf{x}) = p(\mathbf{z}|\mathbf{x}) [p(\mathbf{y}|\mathbf{x}) p(\mathbf{x})], \tag{5}$$

since \mathbf{y} and \mathbf{z} are conditionally independent random variables. □

3 Ensemble Kalman filter

The EnKF is a Monte Carlo implementation of the Kalman filter, which avoids evolving the covariance matrix of the distribution of the state vector \mathbf{x} . Instead, the distribution is represented by a sample, called an ensemble. Ensemble members are evolved in time until the end of the analysis cycle, when the ensemble is updated from the Bayes theorem.

Let \mathbf{X} be a p by n matrix whose columns are ensemble members before the Bayesian update, that is, a random sample from the prior distribution for the state vector \mathbf{x} . The prior distribution is assumed to be normal with covariance \mathcal{Q} . Further, replicate the observations \mathbf{y} into matrix \mathbf{Y} with d rows and n columns so that each column \mathbf{y}_k consists of the observed vector \mathbf{y} plus a random vector from $N(0, R)$. Then it follows from (2) and (3) that the columns of

$$\hat{\mathbf{X}} = \mathbf{X} + K(\mathbf{Y} - H\mathbf{X})$$

form a random sample from the posterior distribution.

The ensemble filter involves two approximations. First, in the Kalman gain matrix $K = \mathcal{Q}H^T(H\mathcal{Q}H^T + R)^{-1}$, the state covariance \mathcal{Q} is unknown, so it is replaced by the sample covariance computed from the ensemble members,

$$\hat{\mathcal{Q}} = \frac{\mathbf{E}\mathbf{E}^T}{n-1}, \quad \mathbf{E} = \mathbf{X} - \mathbf{X} \frac{\mathbf{e}_n \mathbf{e}_d^T}{n},$$

where \mathbf{e}_k is column vector of all ones of size k . This gives the matrix form of the analysis ensemble

$$\hat{\mathbf{X}} \approx \hat{\mathbf{X}}_{ens} = \mathbf{X} + \hat{\mathcal{Q}}H^T(H\hat{\mathcal{Q}}H^T + R)^{-1}(\mathbf{Y} - H\mathbf{X}). \tag{6}$$

See Burgers et al. (1998) for more details.

The second approximation results from the fact that if the state evolution is nonlinear, then the prior distribution is not necessarily normal. Nevertheless, in practice, it is usually hoped that the distribution from which analysis ensemble members in $\hat{\mathbf{X}}_{ens}$ are drawn from is a good approximation of the posterior and the ensemble filter formulas are used anyway. In general, the approximations are practical in the setting where a model is approximately linear in each model advancement step.

Remark 3 For EnKF, observations with independent error distributions can be again assimilated sequentially according to Lemma 2, because the analysis ensemble is a sample from a distribution that approximates the posterior in the limit for a large ensemble. However, the approximation is different when the observations are assimilated sequentially and when they are assimilated jointly because of the sample covariance approximation. In particular, the resulting analysis ensembles are in general different. For the example presented, we found that differences were quantitative rather than qualitative.

We now consider an efficient implementation of the EnKF. We write (6) as

$$\hat{\mathbf{X}}_{ens} = \mathbf{X} + \frac{\mathbf{E}(H\mathbf{E})^T}{n-1} \underbrace{\left(\frac{(H\mathbf{E})(H\mathbf{E})^T}{n-1} + R \right)}_B^{-1} (\mathbf{Y} - H\mathbf{X}),$$

where the matrix

$$HE = H\mathbf{X} - H\mathbf{X} \frac{\mathbf{e}_n \mathbf{e}_d^T}{n}$$

can be computed using the already known product $\mathbf{H}\mathbf{x}$, and note that efficient Choleski decomposition of the symmetric matrix B is possible because R is positive definite. The dominant operations are full matrix-matrix operations efficiently implemented in the Level 3 BLAS (Dongarra et al. 1990) and LAPACK (Anderson et al. 1999) routines, and also readily parallelizable by SCALAPACK (Blackford et al. 1997). In many cases, the computational cost can be further reduced as H is usually sparse and/or highly structured, since a component of \mathbf{y} represents a characteristic of the state vector at a single point or the integrated value over some small region. In addition, if the observation errors are independent, the computation may be split into assimilating a part of the observations at a time as noted above.

4 Two-stage Kalman filter

In order to stabilize the EnKF, we now combine EnKF with a technique related to Tikhonov regularization. Tikhonov regularization for the algebraic least squares problem $\|Ax - b\|^2 \rightarrow \min$ consists of solving instead

$$\|Ax - b\| + \lambda^2 \|Lx\|^2 \rightarrow \min_x. \tag{7}$$

The added term $\lambda^2 \|Lx\|^2$ incorporates a priori assumptions about the size and smoothness of the desired solution x , in the form of the quadratic $\|Lx\|^2$ [21, p. 100]. The parameter λ should be chosen so that both $\|Ax - b\|^2$ and $\|Lx\|^2$ are close to their minimal values as functions of λ ; see [21, p. 84] for details. In practice, λ is often determined by trial and error and can be considered a smoothing parameter.

From (6) and Lemma 1, it follows that each column $\hat{\mathbf{x}}_k$ of the posterior ensemble $\hat{\mathbf{X}}_{ens}$ is the solution of the least squares problem

$$(\hat{\mathbf{x}}_k - \mathbf{x}_k)^T \hat{Q}^{-1} (\hat{\mathbf{x}}_k - \mathbf{x}_k) + (\mathbf{y}_k - H\hat{\mathbf{x}}_k)^T R^{-1} (\mathbf{y}_k - H\hat{\mathbf{x}}_k) \rightarrow \min_{\hat{\mathbf{x}}_k}. \tag{8}$$

Note that in deriving (8), Lemma 1 is applied purely algebraically, with \mathbf{x}_k , $\hat{\mathbf{x}}_k$, and \mathbf{y}_k playing the roles of μ , $\hat{\mu}$, and \mathbf{y} , respectively.

We want to add to EnKF the assumption that $L\mathbf{x}$ does not vary much from $L\mu$, which, analogously to (7), leads to a modification of (8), where the columns of the posterior ensemble are found as the solutions $\hat{\mathbf{x}}_k$ of

$$\begin{aligned} & (\hat{\mathbf{x}}_k - \mathbf{x}_k)^T Q^{-1} (\hat{\mathbf{x}}_k - \mathbf{x}_k) + (\mathbf{y}_k - H\hat{\mathbf{x}}_k)^T R^{-1} (\mathbf{y}_k - H\hat{\mathbf{x}}_k) \\ & + (\mathbf{r}_k - L\hat{\mathbf{x}}_k)^T D^{-1} (\mathbf{r}_k - L\hat{\mathbf{x}}_k) \rightarrow \min_{\hat{\mathbf{x}}_k}. \end{aligned} \tag{9}$$

Here, D is a given symmetric positive definite covariance matrix, D^{-1} plays the generalized role of the parameter λ , and \mathbf{r}_k is sampled from $N(L\mu, D)$. According to Lemma 1, (9) is equivalent to assimilating two independent observations, $\mathbf{y} = H\mathbf{x}$ with the error covariance R , and $L\mathbf{x} = L\mu$ with error covariance D . From Lemma 2 and Remark 3, the Bayesian update corresponding to (9) can be implemented as a two stage EnKF simply by applying the EnKF formulas (6) twice for the two observations $H\mathbf{x} = \mathbf{y}$ and $L\mathbf{x} = L\mu$.

Remark 4 Assimilating the observation $L\mathbf{x} = L\mu$ first, it is easy to see that the proposed two-stage Kalman filter is equivalent to a hierarchical Bayesian update using the observation $L\mathbf{x} = L\mu$ to modify the prior. This first update step has the effect of “shrinkage,” reducing the spread of $L\mathbf{x}$ around $L\mu$. The observation $\mathbf{y} = H\mathbf{x}$ is then applied to the modified prior in the usual manner.

5 The fire model

A simple example useful for demonstrating the utility of the regularized algorithm proposed above is the simplified model of wildfire by a reaction-convection-diffusion equation (Mandel et al. 2004b),

$$\frac{\partial T}{\partial t} = -\nabla \cdot (k\nabla T) - c_1 \cdot \nabla T - c_2(T - T_a) + c_3 \frac{\partial S}{\partial t}$$

$$\frac{\partial S}{\partial t} = -c_4 \max \{0, T - T_i\}^\alpha S$$

on the spatial domain $[0, 1]$, with Dirichlet boundary conditions $T(0) = T(1) = T_a$. The first equation is the heat balance, where T is the temperature, $-\nabla \cdot (k\nabla T)$ is the diffusion of heat, $-c_1 \cdot \nabla T$ is the heat transport by wind, $-c_2(T - T_a)$ is the heat escaping to the environment with the ambient temperature T_a , and $c_3 \frac{\partial S}{\partial t}$ is the heat generated by burning. The second equation models the fuel supply; its right-hand side is the intensity of burning. This is a very simplified model and we do not use any physical data, yet it appears to capture some essential qualitative fire behavior. All coefficients $c_1, c_2, c_2, c_4, \alpha$ are positive. The variables and constants are dimensionless. The model is discretized by standard finite differences on the mesh $[x_0 = 0, x_1, \dots, x_N = 1]$ with uniform spacing $h = x_{k+1} - x_k = 1/N$ and mesh size $N = 100$. The MATLAB code is available from <http://www-math.cudenver.edu/~cjohns/fire1d>.

In this paper, we consider the differential equation in one spatial dimension. The model has ignition point $T_i = 300$. The initial values of S is 1 on the spatial domain between 0 and 1, except that $S = 0$ at a fuel break between 0.45 and 0.50. The reference solution profile of the initial conditions for temperature is shown in the top panel of Fig. 1. Panels (b) and (c) show how the temperature profile of the reference solution propagates to time periods 0.05 and further to 0.10.

Fig. 1 Panel (a) shows the temperature profiles of the reference solution (solid, dark line) as well as initial distribution of ensemble members (dotted, light lines) for the fire model. Panels (b) and (c) are similar and show the effect of forward propagation on the reference solution (solid dark line) as well as the ensemble members (dotted, light lines) to the first ($t=10$) and second ($t=20$) analysis times, respectively

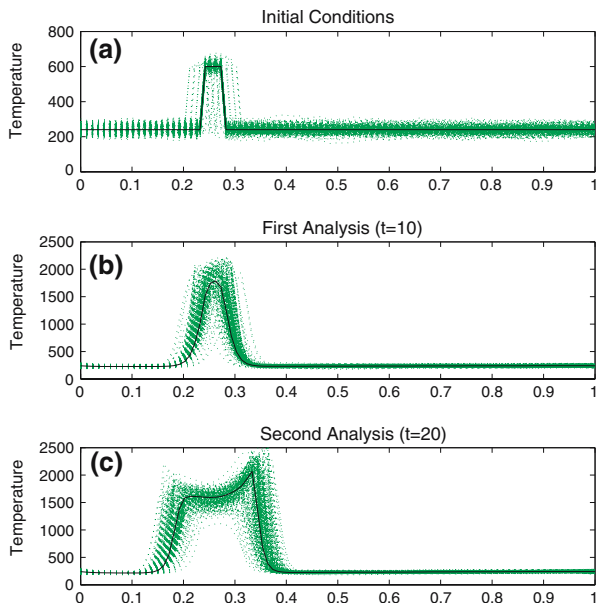


Fig. 2 The reference solution of temperature for the fire model over 30 time iterations starting from the first analysis time period ($t = 10$)


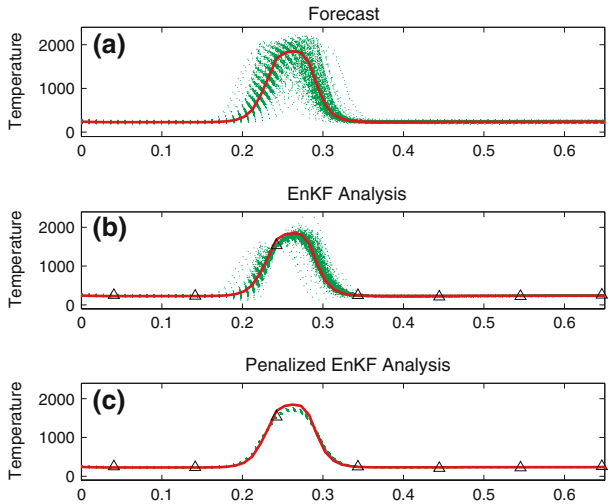


Fig. 3 Panel (a) shows the temperature profiles of 250 ensemble members (dotted lines) forecast to the first analysis cycle (10 time iterations from initial conditions) and the corresponding reference solution (solid line). Panel (b) shows the temperature profiles (dotted lines) in the analysis ensemble corresponding to the EnKF update and includes the information in the data (Δ). Panel (c) compares the regularized analysis ensemble with the reference solution where the second analysis smooths the temperature via an approximate spatial derivative constraint



An initial ensemble of size 250 was generated by perturbing the reference solution temperature and fuel supply values with spatially correlated normal deviates. Six of the ensemble members were generated by shifting the reference solution right and left spatially, and then adding perturbations. Each ensemble member is propagated forward ten time steps via the fire model to reach the time period shown in Fig. 2 to generate the forecast or prior ensemble. Panel (a) of Fig. 3 shows the 250 members of the forecast ensemble (dotted lines) and the associated true temperature profile (solid). Synthetic data (Δ) collected every 10 spatial units are shown in Panel (b) along with the ensemble (dotted lines) updated to the EnKF analysis stage. Note that the analysis step does a good job of cinching the temperature profiles to the few observations; however, the data information does not carry over to nearby locations (e.g. the analysis ensemble near 0.20 or 0.30 on the x -axis in panel (b).) This cinching phenomenon is a byproduct of the least-squares approach in the filter and can produce temperature profiles not compatible with common sense, theory, and numerical approximation schemes. Assuming a spatial correlation function amongst the observations can eliminate some of the cinching effect in the updates and produce smoother update ensemble members. However,

we found that the strength of correlation required to make significant improvements were impractical and counterintuitive. Furthermore, we are interested in the particular situation where observations are accurate but spatially sparse.


To force a measure of smoothness upon the analysis, we impose a penalty on the spread of the first spatial derivative on the temperature analysis field using a second application of the EnKF code as described in Sect. 4. We choose the observation function L by

$$L\mathbf{x} = L \begin{bmatrix} T \\ S \end{bmatrix} = \begin{bmatrix} (T_1 - T_0) / h \\ \vdots \\ (T_N - T_{N-1}) / h \end{bmatrix}$$

where $T_i \approx T(x_i)$ is the discrete temperature variable at node i . So, $L\mathbf{x}$ is the numerical derivative of the temperature field in the system state. In this example, D is diagonal with non-zero elements $d_{ii} = |z_i| / (2h^2)$, where $\mathbf{z} = L\boldsymbol{\mu}$. That is, at each node, the variance expected in the temperature gradient is proportional to the average temperature gradient at that node in the prior. The values of d_{ii} are large enough not to influence reasonable smooth simulations perceptibly but they will suppress simulations that start going unstable and exhibit large swings on the scale of the mesh step h .

The regularized analysis field is shown in panel (c) of Fig. 3. We now compare the forecasts from the EnKF analysis (A) and the regularized analysis (R). The ensembles shown in panels (b) and (c) of Figure 3 were propagated forward 30 time steps and the pointwise squared errors (MSE) from the reference solution were averaged over each spatial location to simulate a prediction. The log of the ratio of MSE's for the two sets of predictions, MSE_A / MSE_R , is shown in Fig. 4. In every case, the regularized predictions had a smaller MSE value than those based on the original EnKF analysis ensemble. Because the reference solution is the basis for this MSE analysis, the effects of bias are included in the calculations via the usual $bias^2 + variance$ formula. A graphical comparison of the bias showed that the bias for both methods were of roughly the same magnitudes and generally smaller for the regularized method.

Fig. 4 Log of the ratio of Mean-squared error (with respect to reference solution) of the analysis ensemble (MSE_A) and regularized analysis ensemble (MSE_R). Values greater than zero support the use of the two-stage shrinkage method



6 Conclusion

Including a regularization step with the EnKF can be considered changing the prior distribution and it was implemented by running EnKF update twice, once on the actual observation and once on the regularization term as another artificial observation. This additional update adds information just as the information added by the data likelihood, and it is justified by a belief about the properties of states that should result from the numerical simulations. In nonlinear problems, the region of valid simulation states (valid region) may be quite small and a linear ensemble Kalman filter update will take simulation states out of the valid region, particularly if the variance of the prior or the data likelihood are large, and a numerical breakdown results. The proposed regularization helps to keep the states within the valid region by statistically imposing restraints on the gradient. This two-step update is similar to a Bayesian hierarchical update and can be viewed as a shrinkage step.

We have applied the proposed improvement to EnKF for a simple fire model, which breaks down numerically because updated states under the usual Kalman filter results in large, non-physical values at a later time. Since a sign of the problem is the emergence of spikes in the fire temperature, and since detecting sharp gradients is more sensitive than detecting large values, we have chosen as the regularization term the numerical derivative of the temperature field. The covariance of the regularization term was chosen to penalize large swings on the spatial mesh scale.

In small systems with fast updates, the effect of choosing the covariance of the regularization term can be considered similar to choosing a smoothing parameter in a nonparametric regression setting.

Acknowledgements This research was supported by the National Science Foundation under grants 0325314, 0417909, and 0719641

References

- Anderson JL, Anderson SL (1999) A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon Weather Rev* 127:2741–2758
- Anderson BDO, Moore JB (1979) *Optimal filtering*. Prentice-Hall, Englewood Cliffs, NJ
- Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J, Du Croz J, Greenbaum A, Hammarling S, McKenney A, Sorensen D (1999) *LAPACK users' guide*, 3rd edn. Society for Industrial and Applied Mathematics, Philadelphia, PA
- Bengtsson T, Snyder C, Nychka D (2003) Toward a nonlinear ensemble filter for high dimensional systems. *J Geophys Res-Atmos* 108(D24):STS 2-1-10
- Blackford LS, Choi J, Cleary A, D'Azevedo E, Demmel J, Dhillon I, Dongarra J, Hammarling S, Henry G, Petitet A, Stanley K, Walker D, Whaley RC (1997) *ScaLAPACK users' guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA
- Burgers G, van Leeuwen PJ, Evensen G (1998) Analysis scheme in the ensemble Kalman filter. *Mon Weather Rev* 126:1719–1724
- Carlin JB, Polson N, Stoffer DS (1992) A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *J Am Stat Assoc* 418:494–500
- Dongarra JJ, Croz JD, Hammarling S, Duff I (1990) A set of level 3 basic linear algebra subprograms. *ACM Trans Math Softw* 16:1–17
- Doucet A, de Freitas N, Gordon N (2001) (eds), *Sequential Monte Carlo in practice*. Springer
- Evensen G (2004) Sampling strategies and square root analysis schemes for the EnKF. *Ocean Dyn* 54:539–560
- Evensen G (1994) Sequential data assimilation with nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J Geophys Res* 99(C5)(10):143–162
- Evensen G (2003) The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dyn* 53:343–367

- Hansen PC (1998) Rank-deficient and discrete ill-posed problems. SIAM Monographs on Mathematical Modeling and Computation. Society for Industrial and Applied Mathematics (SIAM), Philadelphia PA
- Harvey AC (1989) Forecasting, structural time series models and the Kalman filter. Cambridge
- Houtekamer P, Mitchell HL (1998) Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review* 126(3):796–811
- Jazwinski AH (1970) Stochastic processes and filtering theory. Academic Press, New York
- Johansen TA (1997) On Tikhonov regularization, bias and variance in nonlinear system identification. *Automatica* 33(3):441–446
- Johns CJ, Shumway RH (2005) A state-space model for censored data. *Environmetrics* 16:167–180
- Jones RH (1984) Fitting multivariate models to unequally spaced data. In: Parzen E (ed) Time series analysis of irregularly spaced data, no. 25 in Lecture Notes in Statistics. Springer-Verlag, pp 158–188
- Kalman RE (1960) A new approach to linear filtering and prediction problems. *Trans ASME - J Basic Eng Ser D* 82:35–45
- Kalman RE, Bucy RS (1961) New results in filtering and prediction theory. *Trans ASME - J Basic Eng* 83:95–108
- Kim KY, Kang SI, Kim MC, Kim S, Lee YJ, Vauhkonen M (2002) Dynamic image reconstruction in electrical impedance tomography with known internal structures. *IEEE Trans Magn* 38:1301–1304
- Kitagawa G (1987) Non-Gaussian state-space modeling of nonstationary time series. (with discussion). *J Am Stat Assoc* 400:1032–1044
- Mandel J, Chen M, Franca LP, Johns C, Puhalskii A, Coen JL, Douglas CC, Kremens R, Vodacek A, Zhao W (2004a) A note on dynamic data driven wildfire modeling. In: Computational Science - ICCS 2004 Bubak M, van Albada GD, Sloot PMA, Dongarra JJ (eds), vol 3038 of Lecture Notes in Computer Science. Springer pp. 725–731
- Mandel J, Chen M, Franca LP, Johns C, Puhalskii A, Coen JL, Douglas CC, Kremens R, Vodacek A, Zhao W (2004b) Dynamic data driven wildfire modeling. UCD CCM Report 208, March 2004, <http://www-math.cudenver.edu/ccm/reports/rep208.pdf>
- Meinhold R, Singpurwalla N (1983) Understanding the Kalman filter. *American Statistician* 37:123–127
- Meinhold R, Singpurwalla N (1989) Robustification of Kalman filter models. *J Am Stat Soc* 84:479–486
- Paige CC, Saunders MA (1977) Least squares estimation of discrete linear dynamic systems using orthogonal transformations. *SIAM J Numer Anal* 14(2):180–193
- Shumway RH, Stoffer DS (1982) An approach to time series smoothing and forecasting using the EM algorithm. *J Time Ser Anal* 3:253–264
- Tippett MK, Anderson JL, Bishop CH, Hamill TM, Whitaker JS (2003) Ensemble square root filters. *Monthly Weather Review* 131:1485–1490
- van Leeuwen P (2003) A variance-minimizing filter for large-scale applications. *Mont Weather Rev* 131(9):2071–2084

Author biographies

Craig Johns is a Statistician at Milliman, Inc. and previously worked at the University of Colorado Denver and Health Sciences Center and the National Center for Atmospheric Research. His interests include data mining, modeling rare events and problems related to filtering.

Jan Mandel is Professor of Mathematics and Computer Science at the University of Colorado at Denver and Health Sciences Center. He is currently the lead PI of an NSF sponsored project on dynamic data driven wildfire modeling. His interests include wildfire models, ensemble filters, parallel computing, multigrid, and domain decomposition.