Assessing the Performance of the Ensemble Kalman Filter for Land Surface Data Assimilation

YUHUA ZHOU, DENNIS MCLAUGHLIN, AND DARA ENTEKHABI

Ralph Parsons Laboratory, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts

(Manuscript received 4 February 2005, in final form 8 July 2005)

ABSTRACT

The ensemble Kalman filter provides an easy-to-use, flexible, and efficient option for data assimilation problems. One of its attractive features in land surface applications is its ability to provide distributional information about variables, such as soil moisture, that can be highly skewed or even bimodal. The ensemble Kalman filter relies on normality approximations that improve its efficiency but can also compromise the accuracy of its distributional estimates. The effects of these approximations can be evaluated by comparing the conditional marginal distributions and moments estimated by the ensemble Kalman filter with those obtained from a sequential importance resampling (SIR) particle filter, which gives exact solutions for large ensemble sizes. Comparisons for two land surface examples indicate that the ensemble Kalman filter is generally able to reproduce nonnormal soil moisture behavior, including the skewness that occurs when the soil is either very wet or very dry. Its conditional mean estimates are very close to those generated by the SIR filter. Its higher-order conditional moments are somewhat less accurate than the means. Overall, the ensemble Kalman filter appears to provide a good approximation for nonlinear, nonnormal land surface problems, despite its dependence on normality assumptions.

1. Introduction

Ensemble-based data assimilation methods are becoming popular in many of the earth sciences, largely because they are easy to use, flexible, and make relatively few restrictive assumptions (see the review by Evensen 2003). In particular, the ensemble Kalman filter has recently been suggested as a practical option for real-time estimation of land–atmosphere fluxes from remote sensing data (Reichle et al. 2002; Margulis et al. 2002; Crow and Wood 2003; Reichle and Koster 2003; Entekhabi et al. 2004). Ensemble estimation procedures have the advantage of providing distributional information about uncertain variables, including approximate marginal distributions, quantiles, and higher-order moments. This information is particularly useful in land surface applications, where variables such as soil moisture can be highly skewed toward the wet or dry ends and can even be bimodal, depending on the time and space scale considered (Rodriguez-Iturbe et al. 1991). In such cases, means and covariances alone may not adequately characterize variability.

In sequential filtering the distributional properties of an uncertain state $\mathbf{x}_t$, given a set of measurements $\mathbf{y}_{0:t}$ taken through time t, are conveyed by the conditional probability density $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$. The random replicates generated by ensemble methods may be used to compute finite sample approximations to this density and its moments. When new measurements become available some version of Bayes's theorem is typically used to update the replicates (and the corresponding distributional approximations). The accuracy of this update depends on the assumptions made when applying Bayes's theorem as well as the number of replicates. The ensemble Kalman filter is particularly efficient because it relies on normality assumptions that greatly simplify the update process. But this simplification can also limit the filter's ability to provide accurate distributional information. Here we evaluate the accuracy of the ensemble Kalman filter by comparing its distributional estimates to those of a less efficient ensemble method that relies on an exact Bayesian update. This is done for two examples that provide useful insight about the en-

Corresponding author address: Dennis McLaughlin, Bldg. 48-329, 15 Vassar Street, Cambridge, MA 02139.
E-mail: dennism@mit.edu

semble Kalman filter's performance in land surface applications.

2. Ensemble filtering

Many of the nonlinear filtering problems encountered in the earth sciences rely on discrete time state and measurement equations of the following form:

$$\mathbf{x}_t = f(\mathbf{x}_{t-1}, \mathbf{u}_t) \quad \text{and} \tag{1}$$

$$\mathbf{y}_t = g(\mathbf{x}_t) + \mathbf{v}_t, \tag{2}$$

where $\mathbf{x}_t$ is the system state vector with an uncertain initial condition $\mathbf{x}_0$, $\mathbf{u}_t$ is a vector of uncertain model inputs (not necessarily additive), $\mathbf{y}_t$ is the measurement vector, and $\mathbf{v}_t$ is a vector of additive random measurement errors. In a land surface problem $\mathbf{x}_t$ could be a vector of soil moisture values in different pixels and layers, $\mathbf{u}_t$ a vector of precipitation rates, and $\mathbf{y}_t$ a vector of microwave radiometer measurements indirectly related to soil moisture. The uncertain variables $\mathbf{x}_0$, $\mathbf{u}_t$, and $\mathbf{v}_t$ are assumed to have known prior probability distributions and the measurement error vectors at different times are assumed to be independent. The functions $f()$ and $g()$ represent spatially and temporally discretized models of the system dynamics and measurement process.

The objective of filtering is to characterize the current state $\mathbf{x}_t$ from $\mathbf{y}_{0:t}$, the set of all measurements obtained at discrete times in the interval $[1, t]$. The ideal probabilistic characterization is the conditional probability density $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$, which conveys everything known about $\mathbf{x}_t$ given $\mathbf{y}_{0:t}$. Since this multivariate density is difficult to compute or interpret for large problems we typically focus on particular properties of $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$, such as its moments and univariate marginal densities of $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$.

In filtering applications it is convenient to distinguish two sequential estimation operations: 1) propagation of the state from one measurement time to the next (forecasting) and 2) updating of the propagated state with the new measurement (analysis). If the complete density $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$ is desired, forecasting is carried out by deriving $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$ from $p(\mathbf{x}_{t-1} \mid \mathbf{y}_{0:t-1})$ (e.g., using the Fokker–Planck equation) and analysis is carried out by deriving $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$ from $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$ (e.g., using Bayes's theorem; Jazwinsky 1970). The required calculations are generally feasible only for very small problems.

Ensemble methods are able to provide a practical alternative to exact Bayesian solutions because they rely on discrete approximations of the densities $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$ and $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$. The approximations can be expressed as

$$p(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}) \approx \sum_{i=1}^{N} w^i_{t|t-1} \delta(\mathbf{x}_t - \mathbf{x}^i_{t|t-1}) \quad \text{and} \tag{3}$$

$$p(\mathbf{x}_t \mid \mathbf{y}_{1:t}) \approx \sum_{i=1}^{N} w^i_{t|t} \delta(\mathbf{x}_t - \mathbf{x}^i_{t|t}). \tag{4}$$

These approximations replace each continuous density by a sum of N Dirac delta densities located at the randomly generated state vectors, or replicates, $\mathbf{x}^i_{t|t-1}$ or $\mathbf{x}^i_{t|t}$, for $i = 1, \ldots, N$. The Dirac delta terms (and the corresponding replicates) for each approximation are assigned discrete probabilities (or weights) $w^i_{t|t-1}$ or $w^i_{t|t}$, respectively. If the weights in each expansion sum to unity, the integrals of (3) and (4) yield stepwise approximations of the continuous cumulative distribution functions for $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$ and $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$, respectively. The random replicates and corresponding weights can be generated in a variety of ways. The sequential importance resampling (SIR) particle filter and ensemble Kalman filter discussed in this paper are two particular alternatives.

The SIR and ensemble Kalman filters share the same forecasting step. To examine the mechanics of this step, suppose that replicate i at $t - 1$ has the value $\mathbf{x}^i_{t-1|t-1}$ with weight $w^i_{t-1|t-1} = 1/N$. The nonlinear state Eq. (1) can be used to compute the value of this replicate at time t from the value at $t - 1$, giving

$$\mathbf{x}^i_{t|t-1} = f(\mathbf{x}^i_{t-1|t-1}, \mathbf{u}^i_t) \quad \text{and} \tag{5}$$

$$w^i_{t|t-1} = 1/N. \tag{6}$$

Note that this operation requires generation of a random input replicate $\mathbf{u}^i_t$, which is a random sample drawn from the specified prior input probability density $p(\mathbf{u}_t)$. Equations (5) and (6) yield the following approximation for the forecast probability density:

$$p(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}) \approx \frac{1}{N} \sum_{i=1}^{N} \delta(\mathbf{x}_t - \mathbf{x}^i_{t|t-1}) = \frac{1}{N} \sum_{i=1}^{N} \delta[\mathbf{x}_t - f(\mathbf{x}^i_{t-1|t-1}, \mathbf{u}^i_t)]. \tag{7}$$

This forecasting step is just a Monte Carlo-based procedure for deriving $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$ from $p(\mathbf{u}_t)$ and $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$. The differing analysis steps for the SIR and ensemble Kalman filters are discussed in the following paragraphs.

a. The SIR particle filter

Particle filters are a class of sequential Bayesian ensemble algorithms that can be derived from a discrete version of Bayes's theorem. Arulampalam et al. (2002)

provide a useful tutorial that shows how several different particle filtering algorithms may be developed from the perspective of sequential importance sampling (SIS). We use the SIR filter here because it is easy to implement and converges to the exact Bayesian solution as the number of replicates approaches infinity. It is also well suited for the land surface application, where uncertain time-dependent inputs are generally more important than initial condition errors. In other applications, other types of particle filters may give better performance for a given number of particles.

The SIR algorithm adopts the approximation of (5) and (6) during the forecasting step of filtering. The analysis step is based on the following form of Bayes's theorem:

$$p(\mathbf{x}_t \mid \mathbf{y}_{1:t}) = \frac{p(\mathbf{y}_t \mid \mathbf{x}_t)p(\mathbf{x}_t \mid \mathbf{y}_{1:t-1})}{p(\mathbf{y}_t \mid \mathbf{y}_{1:t-1})} = cp(\mathbf{y}_t \mid \mathbf{x}_t)p(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}), \tag{8}$$

where c is a normalizing constant that ensures that $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$ integrates to one. If we substitute the Dirac expansions for $p(\mathbf{x}_t \mid \mathbf{y}_{0:t-1})$ and $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$ into (8) we can relate the analysis density replicate values and weights of the unknown analysis density (left-hand side) to those of the known forecast density (right-hand side). In the SIR filter the analysis replicate values are initially kept the same as the forecast values and only the analysis weights are changed. This gives

$$\mathbf{x}^i_{t|t} = \mathbf{x}^i_{t|t-1} \quad \text{and} \tag{9}$$

$$w^i_{t|t} = cp(\mathbf{y}_t \mid \mathbf{x}^i_{t|t})w^i_{t|t-1} = cp(\mathbf{y}_t \mid \mathbf{x}^i_{t|t-1})w^i_{t|t-1} = \frac{c}{N}p(\mathbf{y}_t \mid \mathbf{x}^i_{t|t-1}), \tag{10}$$

where $p(\mathbf{y}_t \mid \mathbf{x}^i_{t|t-1})$ is the likelihood function for the propagated replicate $\mathbf{x}^i_{t|t-1}$ The likelihood function can be readily computed if the measurement error is additive (as assumed here) since

$$p(\mathbf{y}_t \mid \mathbf{x}^i_{t|t-1}) = p_{vt}(\mathbf{y}_t \mid \mathbf{x}^i_{t|t-1}) = p_{vt}[\mathbf{y}_t - g(\mathbf{x}^i_{t|t-1})], \tag{11}$$

where p_{vt} is the known (e.g., normal) probability density of the measurement error $\mathbf{v}_t$. The likelihood function can be viewed as a measure of the "closeness" of the replicate $\mathbf{x}^i_{t|t-1}$ to the measurement $\mathbf{y}_t$.

We could substitute (9) and (10) directly into (4) to obtain an approximation of the analysis probability density but the result may be unsatisfactory unless the number of replicates is very large. This is because (10) gives replicates "closer" to the measurements much more weight than those that are "farther away." This can result in the "collapse" of the ensemble to a very small number of replicates with high weights, giving a very coarse discrete representation of the analysis probability density. To prevent this, the SIR filter resamples the ensemble with replacement N times. The probability that replicate i is selected on sample k is equal to its weight:

$$p(\text{replicate } i \text{ selected on sample } k) = w^i_{t|t}. \tag{12}$$

By construction, this resampling operation creates a new analysis ensemble of N equally likely replicates with the following values $\mathbf{x}^k_{t|t}$ and weights $w^k_{t|t}$ (for $k = 1, \ldots, N$):

$$\mathbf{x}^k_{t|t} = \text{replicate value selected on sample } k \tag{13}$$

and

$$w^k_{t|t} = \frac{1}{N}. \tag{14}$$

The new analysis ensemble is a subset of the old analysis ensemble. Old replicates with high weight are more likely to be repeated in the new ensemble and old replicates with low weight are more likely to be omitted. Once the resampling operation is completed (13) and (14) can be substituted into (4) to give

$$p(\mathbf{x}_t \mid \mathbf{y}_{1:t}) \approx \frac{1}{N}\sum_{k=1}^{N} \delta(\mathbf{x}_t - \mathbf{x}^k_{t|t}). \tag{15}$$

The new equally weighted resampled replicates can then be propagated from t to $t + 1$, following the procedure given in (5) and (6) (with t replaced by $t + 1$ and k by i). Although many of the resampled analysis replicates at t have the same value, these values diverge in the subsequent propagation to $t + 1$ because of the influence of the random input noise $\mathbf{u}_{t+1}$. This keeps the ensemble from collapsing and is why the SIR approach works best for problems with random inputs.

The SIR filter's ensemble statistics (marginal densities, moments, etc.) can be shown to converge to their exact counterparts as the number of replicates approaches infinity. The version of the SIR filter described here assumes that the measurement errors are additive and independent over time but does not restrict the form of the probability densities for $\mathbf{x}_t$, $\mathbf{u}_t$, or $\mathbf{v}_t$ or the form of the functions $f()$ and $g()$. The primary disadvantage of the SIR filter is the large number of replicates required to accurately represent the multivariate conditional probability densities of $\mathbf{x}_t$. When the number of measurements exceeds a few hundred the SIR filter is not practical for land surface problems. However, it provides a very useful performance benchmark for small problems since it yields optimal condi-

tional densities (as well as conditional means and other moments) if the ensemble is sufficiently large.

b. The ensemble Kalman filter (EnKF)

Like the SIR filter, the ensemble Kalman filter uses the Dirac expansions of (3) and (4) to approximate the conditional probability densities of $\mathbf{x}_t$ and it adopts the approximation of (5) and (6) during the forecasting step. However, the ensemble Kalman filter makes more assumptions at the analysis step. The Kalman filter analysis step can be derived from various perspectives. Here we take a Bayesian or distribution-oriented perspective because we are interested in the filter's ability to estimate properties of the conditional density $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$.

It is generally very difficult to derive an exact closed form expression for $p(\mathbf{x}_t \mid \mathbf{y}_{0:t})$ from Bayes's theorem, especially for problems with nonlinear dynamics and measurement operators. However, it is possible to obtain an exact solution when the forecast states and measurements are jointly normal. This typically occurs only when the state and measurement equations are linear and all sources of uncertainty are normally distributed. In this special case the analysis density given by (8) is normal and completely defined by the following mean and covariance, which are the update expressions of the classical Kalman filter:

$$\bar{\mathbf{x}}_{t|t} = \bar{\mathbf{x}}_{t|t-1} + \mathbf{K}_t[\mathbf{y}_t - g(\bar{\mathbf{x}}_{t|t-1})] \quad \text{and} \tag{16}$$

$$\mathbf{C}_{xx,t|t} = \mathbf{C}_{xx,t|t-1} - \mathbf{K}_t\mathbf{C}_{yx,t|t-1}, \tag{17}$$

where the overbars indicate expected values, the $\mathbf{C}$s are covariances, and

$$\mathbf{K}_t = \mathbf{C}_{xy,t|t-1}(\mathbf{C}_{yy,t|t-1} + \mathbf{C}_{vv,t})^{-1}. \tag{18}$$

In practice, adopting a joint normality assumption is equivalent to assuming that the forecast and measurement densities are adequately characterized by their means and covariances (i.e., higher-order moments are ignored in the analysis step). It is possible to use the Kalman update expressions even when the joint normality assumption does not apply. In this case the conditional statistics produced by the Kalman filter may not match the true values but they may be close enough to be useful.

In an ensemble version of the Kalman filter we need to generate an ensemble of analysis replicates at t, for propagation from t to $t + 1$. The sample mean and covariance of this ensemble should converge to the mean and covariance of (16) and (17) in the limit as the number of replicates approaches infinity. There are many ways to generate analysis replicates that satisfy this requirement. In nonlinear applications it is best to use an ensemble generation method that preserves at least some of the nonnormal characteristics of the forecast ensemble when normality assumptions do not apply.

One way to accomplish this is to generate an analysis ensemble directly from the forecast ensemble, using the following algorithm (Evensen 1994, 2003):

$$\mathbf{x}^i_{t|t} = \mathbf{x}^i_{t|t-1} + \mathbf{K}_{s,t}[\mathbf{y}_t + \mathbf{v}^i_t - g(\mathbf{x}^i_{t|t-1})] \quad \text{and} \tag{19}$$

$$w^i_{t|t} = w^i_{t|t-1} = \frac{1}{N}, \tag{20}$$

where $\mathbf{v}^i_t$ is a sample drawn from the measurement error probability density $p(\mathbf{v}_t)$ and $\mathbf{K}_{s,t}$ is a sample estimate of the Kalman gain $\mathbf{K}_t$,

$$\mathbf{K}_{s,t} = \mathbf{X}_{t|t-1}\mathbf{Y}^{\mathrm{T}}_{t|t-1}(\mathbf{Y}_{t|t-1}\mathbf{Y}^{\mathrm{T}}_{t|t-1} + \mathbf{C}_{vv,t})^{-1}. \tag{21}$$

The columns of the sample matrices $\mathbf{X}_{t|t-1}$ and $\mathbf{Y}_{t|t-1}$ are constructed from the mean-removed replicates of $\mathbf{x}_{t|t-1}$, $g(\mathbf{x}_{t|t-1})$, and $\mathbf{v}_t$.

The ensemble Kalman filter algorithm of (19) through (21) produces analysis replicates that converge to the exact Bayesian solution for normal states and measurements. When there are deviations from normality the filter is suboptimal but the replicates are able to inherit nonnormal properties from the forecast.

There are a number of other versions of the ensemble Kalman filter that use different approaches for generating nonnormal ensembles that conform to (12) and (13) (Tippett et al. 2003). Although each of these has advantages in certain situations the basic concepts are similar to Evensen's (1994) classical ensemble Kalman filter. Moreover, we have found that the performance differences for land surface applications are not substantial. For that reason we focus on the classical version of the filter described above.

3. A simple nonlinear land surface data assimilation example

Soil moisture is one of the key states controlling the partitioning of water and energy fluxes at the land–atmosphere boundary. It is likely to be skewed to the wet end (after precipitation) or the dry end (after a prolonged drydown period). Here we use a simple scalar example motivated by soil moisture behavior to illustrate the two nonlinear filters described in the previous section. Suppose a scalar soil moisture value x at a particular measurement time has the following forecast probability density:

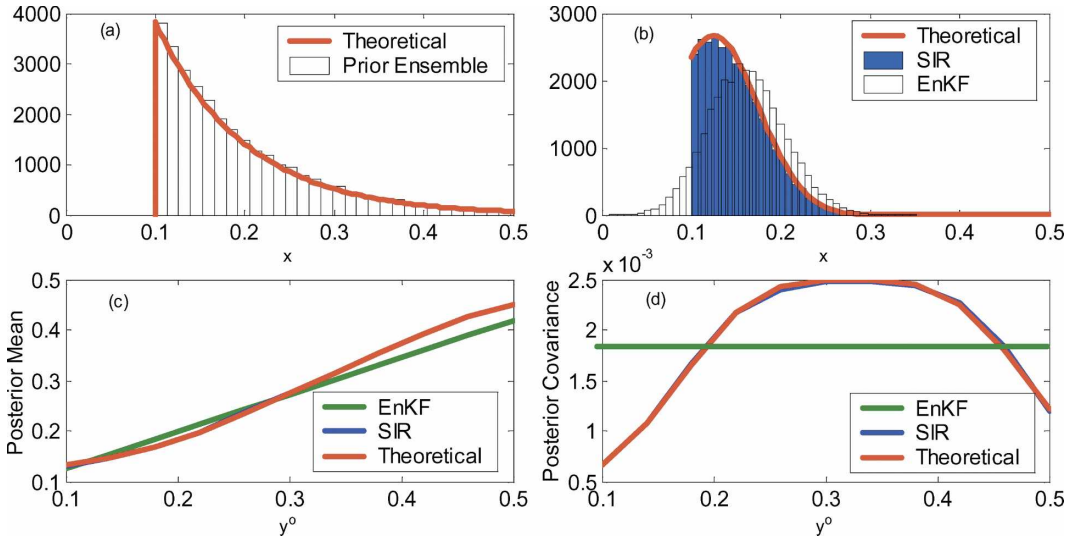

FIG. 1. Estimates of scalar soil moisture state statistics for a skewed prior probability density where y° is the actual observation. (a) Prior probability density function (pdf): $p(x) = 27.7 \exp(-x/0.1)$, $0.1 \le x \le 0.5$; (b) posterior pdf for $y^\circ = 0.15$, $R = 0.05$, estimated with SIR and ensemble Kalman filters; (c) Posterior mean vs y°; and (d) posterior covariance vs y°. The theoretical Bayesian solution is also plotted.

$$p(x) = 27.7 \exp\left(-\frac{x}{0.1}\right) \quad 0.1 \le x \le 0.5. \tag{22}$$

This truncated exponential density is shown in Fig. 1a. The associated mean x is 0.19 and the variance is $C_{xx} = (0.08)^2$. We suppose that a single measurement $y = x + w$ is taken, where w is a zero mean normally distributed additive error independent of x with standard deviation $C_{ww} = (0.05)^2$.

For this problem the analysis probability density $p(x\,|\,y)$ may be derived in closed form from Bayes's theorem,

$$p(x\,|\,y) = cp(y\,|\,x)p(x). \tag{23}$$

This exact analysis density is plotted in Fig. 1b for a measurement value $y_0 = 0.15$, together with the results obtained from an SIR filter and an ensemble Kalman filter, each using 30 000 replicates (this large sample size essentially eliminates sampling error problems). The SIR filter closely approximates the skewed exact analysis density. The ensemble Kalman filter analysis density is much more normal in shape, reflecting the influence of the normally distributed measurement perturbations $\mathbf{v}^i$ added in the update step. As the measurement error becomes larger the Kalman gain eventually becomes very small, the forecast replicates dominate, and the Kalman analysis density becomes more skewed.

The exact, SIR filter and ensemble Kalman filter analysis means are plotted versus the measurement value y_0 in Fig. 1c. The Kalman filter analysis mean deviates only slightly from the exact and SIR filter means. Figure 1d shows that the analysis standard deviations for the two filters behave quite differently. The SIR filter standard deviation depends on the measurement value while the ensemble Kalman filter standard deviation does not. So the Kalman filter underestimates uncertainty for midrange measurements and overestimates uncertainty for low or high measurement values.

Although this scalar example is very simple it suggests that differences between the SIR and ensemble Kalman filters for land surface problems may be more apparent in the higher-order moments than in the analysis means. We investigate this hypothesis further in the next section.

4. Formulation of an Observing System Simulation Experiment (OSSE)

In this section we describe a land surface simulation experiment that enables us to compare the performance of the suboptimal ensemble Kalman filter to an optimal SIR filter for a realistic land surface application. The problem is to characterize soil moisture and evapotranspiration from remotely sensed passive microwave (radiometer) measurements. Land surface dynamics are described by the Community Land Model (CLM, version 2.0; Bonan 1996; Bonan et al. 2002). Radiometer measurements are described by a nonlinear radiative transfer model (Njoku et al. 2002). Input uncertainties and measurement errors are described by

statistical models that are intended to realistically represent natural variability. These models determine how the replicates of the ensemble filters are generated.

a. The land surface and radiative transfer models

The CLM is a nonlinear spatially distributed model that describes energy, momentum, water, and CO_2 exchange between the land and the atmosphere. Dynamic inputs to the model include precipitation, wind speed, air temperature, pressure, humidity, and solar radiation. Time-invariant inputs include soil and vegetation classifications. The model is discretized into square pixels that are each divided into several soil layers. Moisture and heat can only move vertically within individual pixels. Further details are discussed in Bonan (1996) and Bonan et al. (2002). Although moisture does not flow between pixels the states in different pixels are correlated by virtue of their dependence on spatially correlated inputs such as precipitation and vegetation.

The study region for our computational experiment reflects conditions at the southern Great Plains (SGP97) site in eastern Oklahoma. This 18.75 km × 18.75 km region is shown in Fig. 2. It is discretized over a 6 × 6 grid of (approximately) 3.12 km × 3.12 km estimation pixels with eight soil layers in each pixel. The study region is small enough to be feasible for a SIR filter assimilation while large enough to reveal the impacts of horizontal correlation. The land use is assumed to be cropland with a loam soil and the soil layers have thicknesses (from top to bottom) of 2, 3, 5, 8, 12, 20, 57, and 88 cm, respectively. The CLM model states include soil moisture and soil temperature in the center of each soil layer as well as surface soil temperature, canopy temperature, and canopy-intercepted water, for a total of 684 states in our 36-pixel grid. The CLM derives evapotranspiration from these states. The study period corresponds to a 28-day field campaign conducted from 0000 UTC 19 June 1997 through 1500 UTC 16 July 1997 (Margulis et al. 2002). Input data are generated and the CLM is run for a 1-h time step.

Synthetic radiobrightness measurements can be related to soil moisture through soil reflectivity, as described by the Fresnel equation. For our experiment this process is described by the following expression for brightness temperature (Njoku et al. 2002):

$$T_b = T_s(1 - r_H)e^{-\tau} + T_c(1 - \omega)(1 - e^{-\tau})(1 + r_H e^{-\tau}), \tag{24}$$

where T_s and T_c are the surface and canopy temperature (K), respectively, and r_H is the horizontal polarization soil reflectivity. For L-band (1.4 GHz) microwave, the vegetation can be considered predominantly absorbing with a small single scattering albedo ω, and the vegetation opacity along the slant path is given by (Jackson and Schmugge 1991)

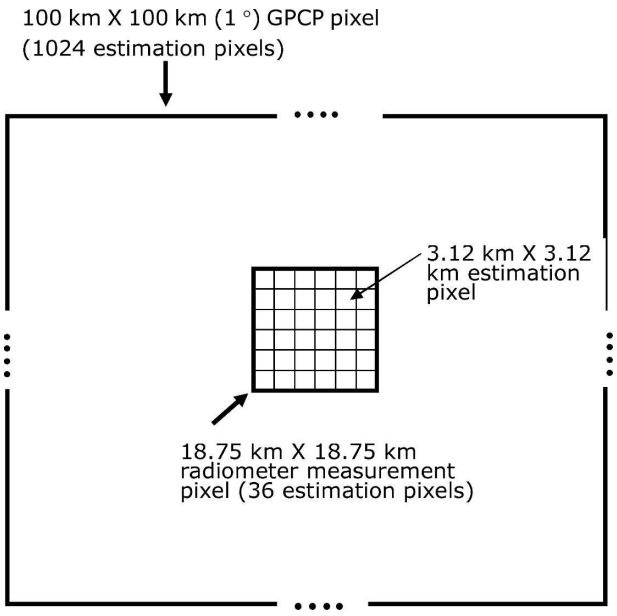

FIG. 2. Multiple scales used in the land surface OSSE. Precipitation data are available in a single 100 km × 100 km (GPCP) pixel, synthetic radiometer measurements are generated in a single 18.75 km × 18.75 km pixel, and estimates are computed in 36 pixels, each 3.12 km × 3.12 km, nested inside the radiometer pixel.

$$\tau = bw/\cos\theta, \tag{25}$$

where w is vegetation water content (kg m^{-2}), b is a vegetation-specific parameter, and θ is the incidence angle. The vegetation water content is derived from normalized difference vegetation index data (Jackson et al. 1999). Rough surface reflectivity is derived from the procedure described by (Choudhury et al. 1979)

$$r'_H = r_H \exp(-h \cos^2\theta), \tag{26}$$

where r_H is the smooth surface reflectivity and h is a vegetation-specific parameter. In our experiment, w, h, and b have the values 0.3 kg m^{-2}, 0.1, and 0.04, respectively. The view angle θ is set to zero and the scattering albedo is 0.03.

b. Uncertain model inputs and measurement errors

The primary sources of uncertainty in land surface applications are time-invariant soil properties and time-dependent meteorological inputs, including precipitation and initial conditions. In the ensemble approach, random replicates for each of the uncertain inputs are provided to the CLM, which generates random replicates of the land surface states. Corresponding radio-

TABLE 1. Summary of uncertain inputs and measurement errors for the land surface simulation experiment.

Variable	Specified nominal value	Uncertainties in replicates
Soil fractions (sand–silt–clay)	Loam over entire study region	Uniformly distributed points in loam section of soil triangle
Vegetation	Cropland with LAI = 1.6 and SAI = 0.4 (June); LAI = 1.3 and SAI = 0.8 (July)	Spatially uncorrelated multiplicative uniform noise $U(0.85, 1.15)$ for LAI
Humidity, solar radiation, wind speed	Oklahoma Mesonet time series at El Reno, assumed to apply over entire study region	Spatially and temporally uncorrelated multiplicative uniform noise: relative humidity: $U(0.9, 1.1)$; solar radiation: $U(0.9, 1.1)$; wind speed: $U(0.7, 1.3)$
Air temperature	Oklahoma Mesonet time series at El Reno, assumed to apply over entire study region	Spatially and temporally uncorrelated additive uniform noise $U(-4, +4$ K)
Precipitation	GPCP 1° daily data for SGP97 region	Nominal GPCP values downscaled in time from daily to hourly values with random pulses model. Temporally downscaled replicates downscaled in space from 100- to 3.1-km pixels with multiplicative cascade model
Initial soil moisture (at $t = -10$ days)	Specified soil moisture profiles	Spatially uncorrelated additive Gaussian noise $N(0.0, 0.3)$
Initial soil temperature (at $t = -10$ days)	Specified temperature profiles	Spatially uncorrelated additive Gaussian noise $N(0.0, 4$ K)
Radiobrightness measurement	Simulated true value at 18.3 km × 18.3 km scale	Temporally uncorrelated additive Gaussian noise $N(0.0, 3$ K)

brightness values at the estimation pixel scale are generated by the radiative transfer model of (24). The time-dependent random inputs can cause the ensemble to spread during the propagation step while assimilation of radiobrightness measurements can cause the ensemble to narrow at the analysis step. These effects are moderated by the physics of the problem, which constrains the states to lie in limited ranges (e.g., the volumetric soil moisture must lie between 0.0 and the porosity, which is less than 1.0).

The uncertain inputs are generated by transforming nominal input values to obtain sets of physically realistic replicates. This is done in various ways, depending on the variable. Table 1 lists the uncertain inputs and measurement errors considered in our simulation experiment. Note that different methods are used to introduce randomness for different inputs. The soil, vegetation, and precipitation inputs deserve some elaboration.

The nominal soil is assumed to be loam throughout the study region. Loam corresponds to a certain section of the classical silt–sand–clay soil triangle. The soil properties associated with different replicates and different pixels are obtained by selecting random points in the loam section and then reading off the corresponding silt, sand, and clay fractions that are used by CLM. Different independent random samples are taken in different pixels and soil layers so soil property fluctuations are not correlated over space.

The vegetation type is assumed to be cropland through the study region. The CLM characterizes land-use types in terms of the leaf area index (LAI) and the stem area index (SAI). It uses these indices to compute various model vegetation parameters that control net radiation, energy partitioning, and intercepted water capacity. In our experiment the nominal LAI for cropland is 1.6 for June and 1.3 for July. The nominal SAI is 0.4 for June and 0.8 for July. Individual LAI replicates are generated by multiplying the nominal value by a uniformly distributed random variable in the range (0.85, 1.15). SAI is treated as a deterministic input.

Precipitation displays significant correlation in time and space and has a patchy pattern that cannot be reproduced with simple multiplicative or additive perturbations to spatially uniform nominal values. A more realistic option is to generate small-scale replicates by downscaling (or disaggregating) larger-scale measurements of real precipitation over both time and space (S. A. Margulis 2004, personal communication). Downscaling relies on statistical models of small-scale variability.

Our simulation experiment uses nominal precipitation data from the Global Precipitation Climatology Project (GPCP). These daily data are available in the SGP97 region at a spatial resolution of 1° (100 km × 100 km). The GPCP time series for SGP97 during the 28-day time period of interest in our experiment is shown in Fig. 3d. We need to downscale this GPCP data from daily to hourly values in time and from 100 km × 100 km to 3.1 km × 3.1 km values in space, as indicated in Fig. 2.

Our temporal downscaling procedure is based on a

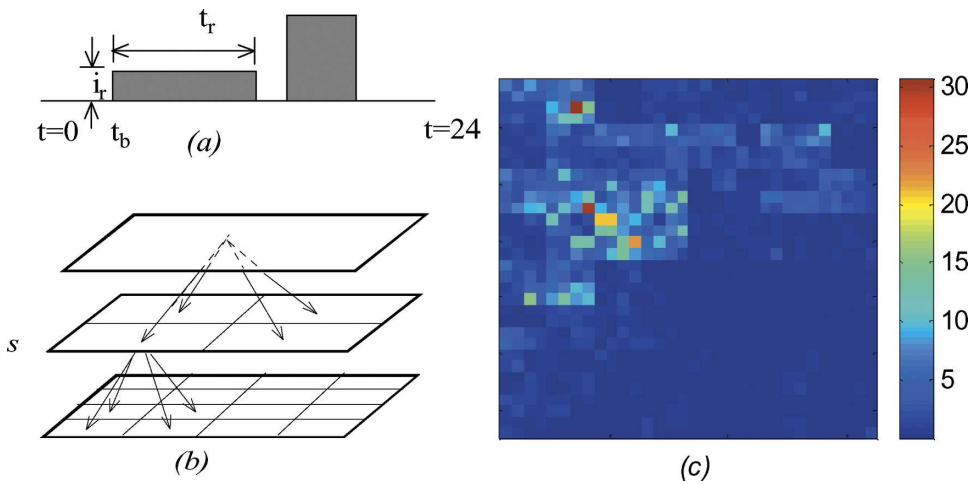

FIG. 3. Spatial and temporal rainfall disaggregation model. (a) RPM for temporal disaggregation of daily rainfall, (b) multiplicative cascade model for rainfall spatial disaggregation, and (c) one realization from the random cascade model over a 32 × 32 grid.

probabilistic rectangular pulses model (RPM; Margulis and Entekhabi 2001) that is constrained to reproduce observed daily totals (see Fig. 3a). The RPM treats rainfall events as random rectangular pulses with an exponentially distributed constant intensity i_r, and duration t_r and a uniformly distributed arrival time between 0 and 24 − t_r. Different RPM replicates have different hourly rainfall values. A given replicate may have no rainfall in any particular hour but its 24 hourly values must add up to the observed GPCP daily value. In our experiment the RPM mean intensity is 3.2 mm h^{-1} for June and 2.3 mm h^{-1} for July, and the mean time between storms is 5.0 h for June and 8.0 h for July. These were estimated from climatological data (Hawk and Eagleson 1992).

The temporal downscaling procedure provides 1-h precipitation replicates at the 100 km × 100 km GPCP measurement scale. These coarse-resolution replicates can be downscaled to the 3.1 km × 3.1 km estimation pixel scale if we suppose that rainfall follows a multiplicative cascade model that relates intensities at different scales (Gupta and Waymire 1993; Gorenburg et al. 2001). This model can be portrayed as a six-level tree composed of groups of pixels (nodes) covering regions of different areas (see Fig. 3b). The top (root) node defines the coarsest scale (one GPCP pixel) while the bottom nodes define the finest scale (the 1024 estimation pixels contained in the GPCP pixel). The rainfall value at a given node is obtained by multiplying the value at the next coarsest node (the parent) by a random lognormally distributed coefficient $W(s)$,

$$W(s) = \exp[w(s) - \sigma_w^2(s)/2], \tag{27}$$

where $w(s) \sim N[0, \sigma_w^2(s)]$, $\sigma_w(s) = 2^{-0.3(s-1)}$, and the scale index s increases from 1 at the root node of the cascade to 6 at the finest scale. A typical realization from this random multiplicative cascade is shown in Fig. 3c. The cascade model generates rainfall that has a patchy pattern and is correlated over space.

In our simulation experiment the cascade model generates spatially downscaled rainfall on a 32 × 32 grid with a finest scale resolution of about 3.1 km, enforcing the same spatial pattern for each replicate in each hour of a given rainy day but allowing the rainfall intensity to change every hour. Rainfall intensities at the finest scale are normalized for each replicate to ensure that the total rainfall at this scale is equal the total rainfall at the GPCP scale. A 6 × 6 portion of this grid provides the 3.1 km × 3.1 km rainfall data needed by the CLM model.

The CLM model is started at $t = -10$ days with random initial conditions generated by perturbing uniform soil moisture and temperature profiles. Each replicate is run forward with the model for 10 days to $t = 0$ to allow moisture in individual pixels to redistribute in accordance with local soil properties. The resulting soil moisture and temperature replicates initialize the SIR and Kalman filter ensemble simulations.

c. *Simulation experiment specifications*

For our synthetic experiment "truth" is defined by the state from a single CLM run obtained for a particular set of soil, vegetation, meteorological, and initial condition replicates, as described above. The CLM states and associated soil properties for this truth replicate are then used in (24) to generate a synthetic

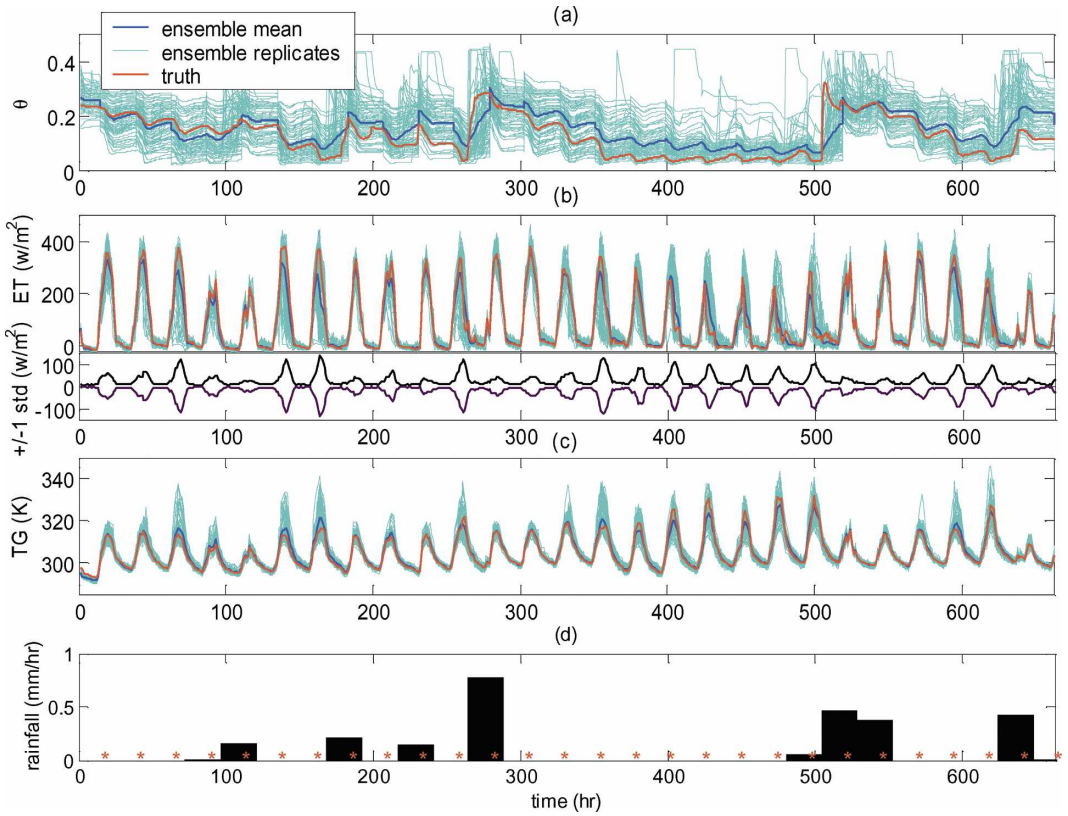

FIG. 4. Ensembles from SIR solution at pixel 9 and the associated GPCP rainfall data: (a) ensemble of the first-layer soil moisture θ; (b) ensemble of evapotranspiration and ± 1 std dev of the ensemble; (c) ensemble of surface temperature; and (d) GPCP 1DD daily rainfall data time series. The asterisks on time axis of (d) represent the measurement times.

brightness temperature measurement at 1500 UTC every day during the 28-day simulation period. This measurement is defined at a coarser scale than the model states, reflecting the lower resolution of anticipated satellite microwave radiometer measurements. In particular, we assume that the microwave measurement covers an 18.3 km $\times$ 18.3 km area (6 $\times$ 6 pixels) and is the average of the 36-pixel-scale brightness values computed from (24). At each measurement time a zero-mean normally distributed random perturbation is added to the averaged brightness temperature to account for the effect of measurement noise. This set of noisy measurements is provided to the two ensemble filters.

5. Results of the simulation experiment

It is useful to start our comparison of the SIR and ensemble Kalman filters by examining surface (top layer) soil moisture, surface soil temperature, and evapotranspiration replicates produced by the SIR filter at a typical pixel (pixel 9). These time series are shown in Fig. 4, together with the applicable 1° daily (1DD) GCPC precipitation record for the study period. The red line is the true replicate, the thick blue line is the mean of the SIR filter ensemble, and the thin cyan lines are the individual SIR filter replicates.

Before $t = 100$ uncertainty in soil moisture, indicated by the ensemble spread, is mainly due to uncertainties in initial conditions, soil properties, and LAI. After rainfall events occur, the uncertainties in surface layer soil moisture primarily reflect uncertainties in precipitation. Note that the ensemble spread is narrower during the dry periods, when the absence of rainfall makes it easy to infer that soil moisture values are low, even without the added information provided by radiobrightness measurements. By contrast, the ground temperature ensemble spread is wider during dry periods and narrower during wet periods. The spread of the evapotranspiration ensemble depends strongly on time of day, peaking just after noon. This is more apparent in the evapotranspiration ensemble standard deviation plot included just below the replicate plot. Replicates from the ensemble Kalman filter have a very similar structure.

It is important that our comparison of the SIR and ensemble Kalman filters is based on enough replicates to ensure that sampling error is not a significant factor. Figure 5 shows the spatial root-mean-squared error (rmse) for the top-layer soil moisture, computed over all analysis times, where error is defined as the difference between the analysis ensemble mean and the true value. Also plotted are error bars that show plus or minus one standard deviation of the rmse, computed over q filter runs started with different filter ensemble random seeds. The truth and measurements are kept the same for these runs. Clearly, the SIR filter needs more replicates to converge, although it eventually gives nearly the same rmse as the Kalman filter. This is not surprising, considering that the converged SIR filter needs to resolve higher-order distributional properties that are ignored by the Kalman filter.

Figure 6 shows marginal forecast (left) and analysis (right) probability densities for pixel 9 surface soil moisture for some typical analysis times. Open-loop (unconditional) densities are also shown for comparison. At the first analysis time, just prior to the first measurement ($t = 15$ h), both filters and the open loop share the same forecast density with a skewness of 0.2 and kurtosis of 3.5 (since there have not yet been any measurements). The difference in the SIR and Kalman filter analysis densities at this time is minimal. The beneficial effect of the measurement is best revealed by a comparison of the open-loop and analysis densities.

The densities plotted at times $t = 231$ and 279 show conditions during two rainy periods. The skewness to the left in both of the forecast densities reflects the effects of the preceding drydown. The measurements at both times move the density noticeably toward the wet end, producing significant differences between the open-loop and filtered densities. Here again, differences between the SIR and Kalman filter analysis densities are minor. Also, at $t = 351$, after a drydown period all of the forecast densities reveal bimodal behavior. This bimodality is likely because of the properties of the multiplicative cascade rainfall model, which tends to produce replicates with wet or dry patches. The analysis densities for the SIR and ensemble filters at $t = 351$ are noticeably different.

After a long period of drydown, at $t = 471$, the forecast and analysis densities are all skewed to the dry end and the radiobrightness measurement does not provide much additional information about the surface soil moisture.

The marginal densities shown in Fig. 6 illustrate the advantages of taking a distributional perspective in data assimilation. Ensemble means and even means plus

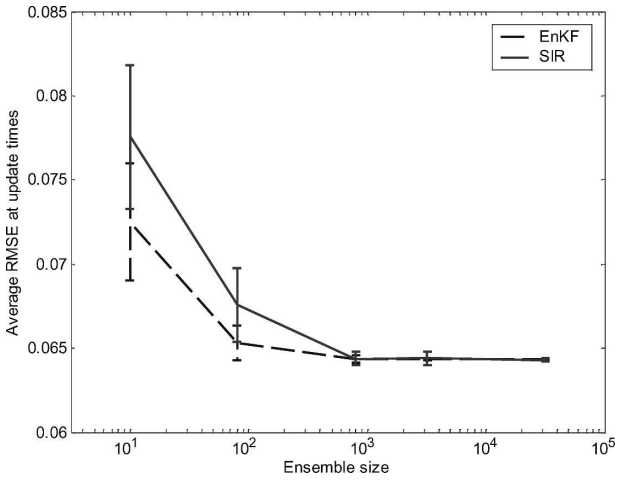

FIG. 5. Averaged spatial rmse of surface layer soil moisture at measurement times vs replicate numbers. Error bars show ±1 std dev of the rmse, computed over q filter runs started with different filter ensemble random seeds. For ensemble size $n = 10$, 80, 800, 3200, and 32 000, and the run times $q = 40$, 10, 8, 6, and 4, respectively.

variances do not always tell the whole story. Physical conditions such as prolonged wetting or drying can lead to skewed densities where the means are much different than the most probable values (modes). Although SIR filters provide accurate information on marginal distributions they are not practical for large problems. Fortunately, the ensemble Kalman filter seems able to convey much of this distributional information, despite its simplifying normality assumptions. This is a direct result of the ensemble Kalman filter's ability to update each replicate rather than just the ensemble mean. Individual replicate updating is able to preserve some skewness and multimodality, even when the analysis step is suboptimal.

To assess global performance, rather than performance at a single pixel, we examine in Fig. 7 the time series of the differences between the ensemble mean and the true replicates for surface soil moisture, evapotranspiration, and ground temperature, all averaged over the entire domain. The errors are shown for the SIR filter, ensemble Kalman filter, and open-loop estimates. The abrupt change in soil moisture error due to assimilation of brightness temperature can be observed at analysis times for both filters but the impact of measurements is less clear for the ground temperature. This reflects the fact that brightness temperature is more sensitive to soil moisture than to ground temperature.

Although evapotranspiration is a diagnostic variable rather than an updated state we can see that the SIR and ensemble Kalman estimates of evapotranspiration benefit from radiobrightness measurements. Both of these generally have lower errors than the open-loop

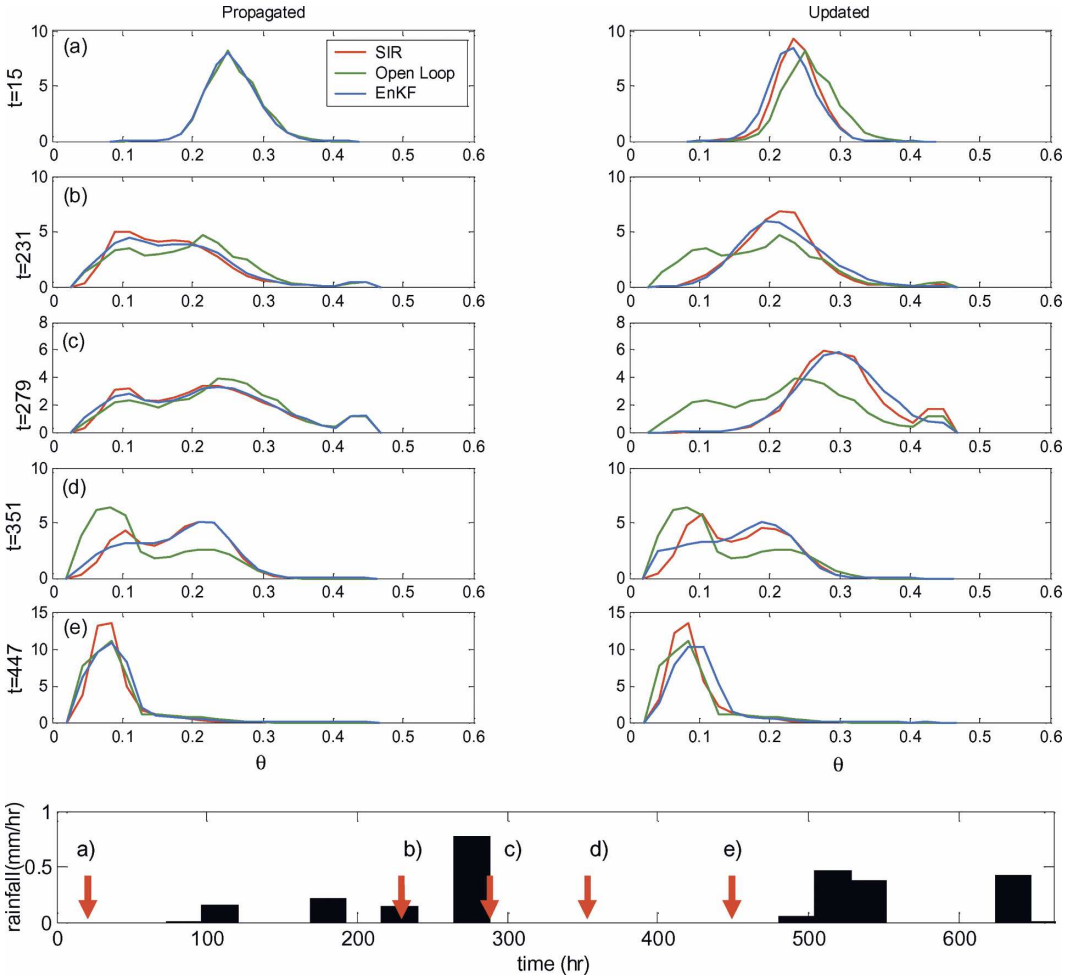

FIG. 6. (left) Marginal forecast, (right) analysis probability densities for pixel 9 surface soil moisture for some typical analysis times, and (bottom) daily rainfall series. Open-loop (unconditional) densities are also shown for comparison.

estimate. A closer look at the plots suggest that the study period can be roughly divided into two stages, before and after at $t = 250$. During the first stage, soil moisture is relatively high, so the evapotranspiration is controlled by available energy rather than soil moisture. Hence, the open-loop estimates of evapotranspiration and ground temperature are nearly as good as the filter estimates. During the second stage, there is a long drydown period, evapotranspiration is moisture limited, and open-loop errors are larger than filter errors. Assimilation of brightness temperature is clearly more beneficial during this stage.

Surface brightness temperatures can be used to estimate subsurface soil moisture profiles that are difficult to observe at large scales. Figure 8 shows the ensemble mean of the integrated soil water depth above 50 cm deep over the entire domain. The integrated soil water depth could be viewed as a rough measure of the water available to a plant with a root depth of 50 cm. Here again, the ensemble Kalman filter gives results that are nearly as good as the SIR filter.

Figure 9 provides some indication of the distributional differences between the two filters by comparing time series of the higher-order moments (standard deviation, skewness, and kurtosis) of the surface soil moisture for pixel 9. Differences between the higher-order moments produced by the two filters are greater than differences between the means. Both filters are able to capture the significant reduction in variance and increase in skewness experienced during the extended drydown period after $t = 300$. Heavy rainfall events seem to reduce differences between the two filters. It should be noted that the random pulse and multiplicative cascade models tend to generate very nonnormal surface soil moisture density functions, as shown by the skewness and kurtosis in Figs. 9b and 9c. The ensemble

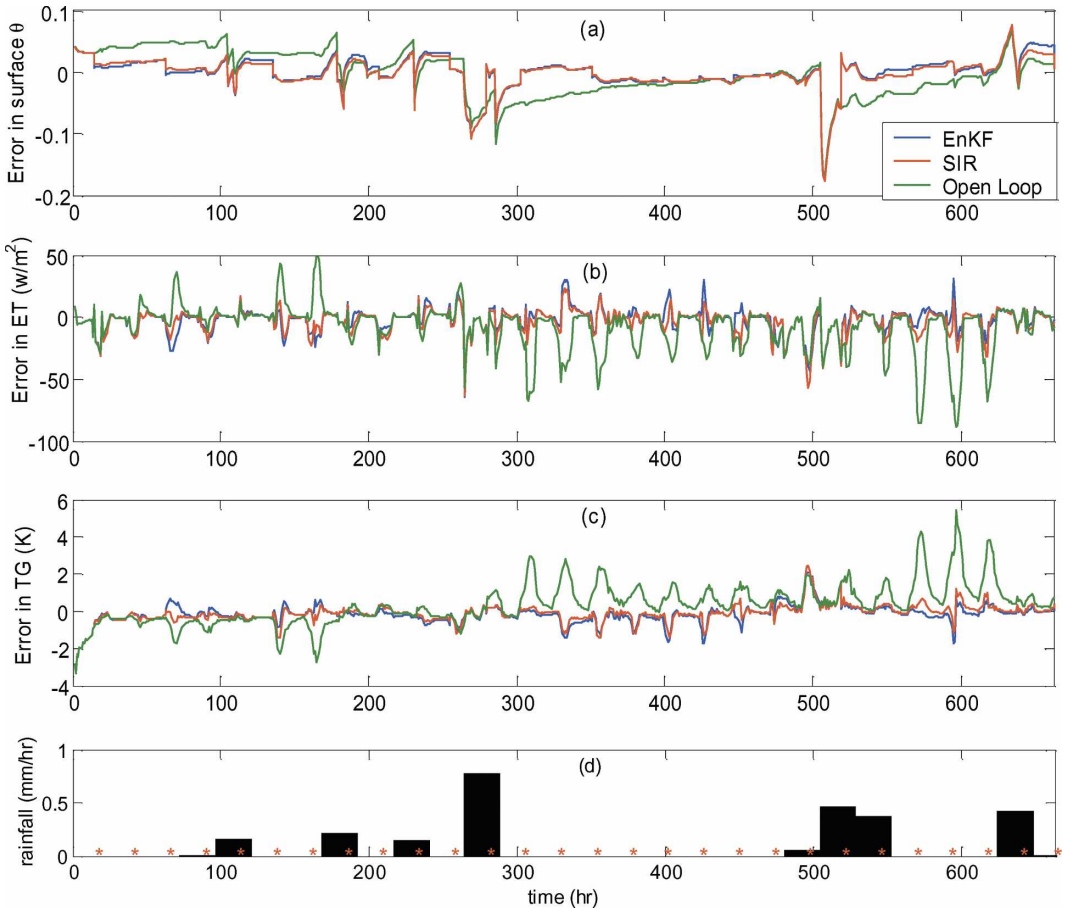

FIG. 7. Time series of the differences between the ensemble mean and the true replicates for (a) surface soil moisture, (b) evapotranspiration, (c) ground temperature, all averaged over the entire domain, and (d) GPCP rainfall time series.

Kalman filter captures much of this nonnormal behavior, at least for our application.

The results of our land surface data assimilation experiment are summarized in Table 2, which lists rmse values obtained from SIR filter, ensemble Kalman filter, and open-loop means for the four variables of most interest. It is obvious that the SIR and ensemble Kalman filter errors are comparable in all cases. Taken

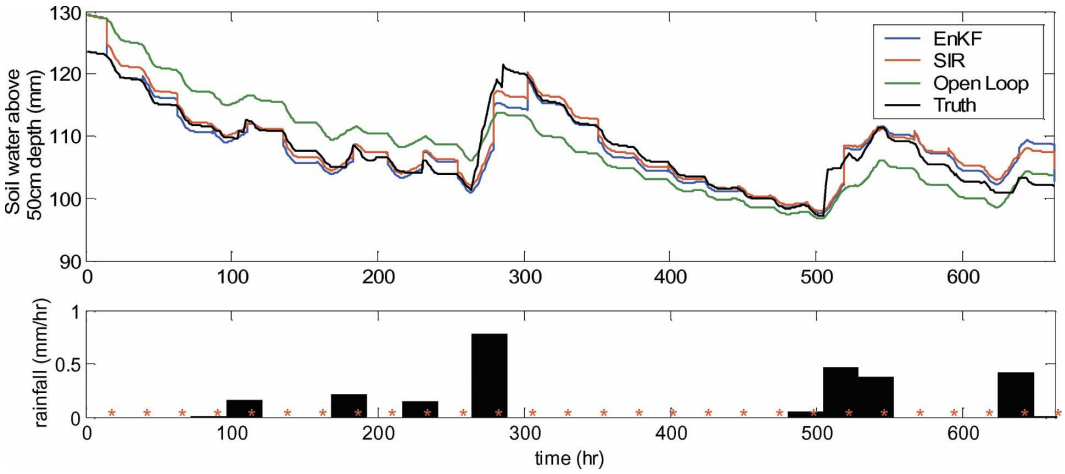

FIG. 8. The ensemble mean of the integrated soil water depth above 50-cm depth over the entire domain.

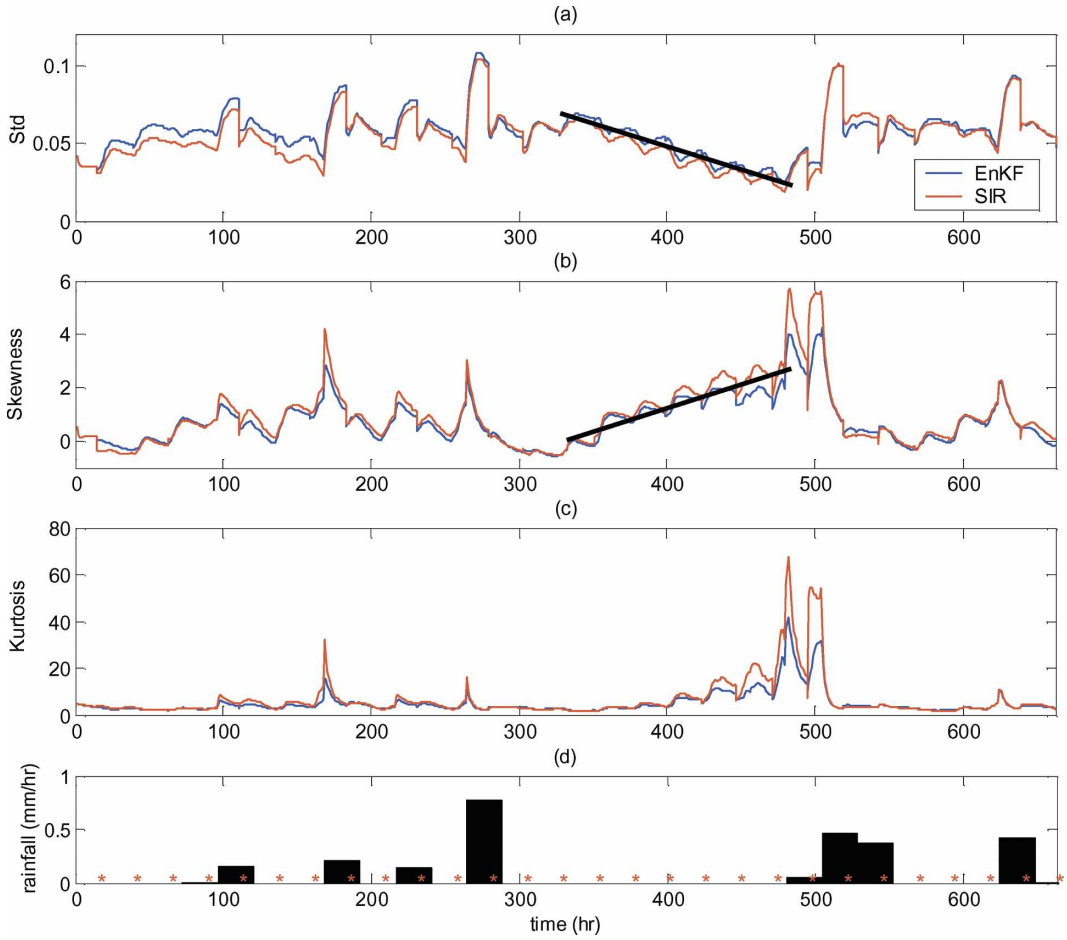

FIG. 9. Time series of (a) standard deviation, (b) skewness, (c) kurtosis of surface layer soil moisture at pixel 9, and (d) GPCP rainfall time series. The thick straight line in (a) and (b) shows the trend of standard deviation and skewness during drydown period.

together, our results strongly support the use of the Kalman approximation in land surface applications of ensemble data assimilation.

6. Discussion and conclusions

This paper considers the performance of the ensemble Kalman filter in a particular context: land surface data assimilation. Land surface problems have several distinctive characteristics. In particular, the state equation is nonlinear and dissipative and the states are confined to relatively small ranges, with probability distributions that change over time and are often nonnormal. Precipitation inputs are intermittent and highly variable over space and time, and other inputs, such as soil properties, are uncertain and difficult to observe over large regions. The measurement equation is also nonlinear. Ensemble Kalman filters have been applied to land surface data assimilation with reasonable success, despite their dependence on assumptions that may not apply. Our objective here has been to better understand the reasons for this success and to obtain a more complete picture of the strengths and weaknesses of the ensemble Kalman filtering approach.

The simple example described in section 3 shows that the SIR filter's conditional moment and marginal density estimates are very close to their exact counterparts if the replicate size is large enough. The ensemble Kal-

TABLE 2. Rmse (over time) of spatially averaged top-layer soil moisture, evapotranspiration, ground temperature, and water depth above 50 cm with respect to the truth.

	SIR	EnKF	Open loop
Top-layer soil moisture	0.026	0.027	0.036
Evapotranspiration (W m^{-2})	10.3	10.1	19.5
Ground temperature (K)	0.5	0.5	1.1
Water depth above 50 cm (mm)	2.2	2.4	4.3

man filter's conditional mean estimate is also quite close to the exact value but its marginal density and variance are noticeably different. This example suggests that a converged SIR filter provides a good basis for evaluating the ensemble Kalman filter when an exact solution is not available.

Sections 4 and 5 describe a more realistic land surface estimation example that relies on state and measurement models used in operational settings. In this OSSE we generate hypothetical true states and measurements so that filter estimation errors can be evaluated exactly. The example problem is kept small so that the SIR filter is computationally feasible. The number of replicates needed for this filter to converge becomes very large when the state and measurement dimensions increase much beyond the values used in our example. This is why the ensemble Kalman filter, which converges for much smaller ensemble sizes, is preferable to the SIR filter in practical applications.

The results of our land surface example reveal that the ensemble Kalman filter performs nearly as well as the SIR filter for most conditions simulated. The surface soil moisture forecast densities obtained from the Kalman filter can be quite skewed and even multimodal and are generally similar to those obtained from the SIR filter. The univariate densities of Fig. 6 make it clear that the normality assumptions that must be met in order for the ensemble Kalman filter to yield optimal point estimates do not prevent it from generating nonnormal ensembles. This is further emphasized in Fig. 9, which shows that the skewness and kurtosis of the ensemble Kalman filter soil moisture ensembles can differ significantly from those associated with normally distributed variables.

The ensemble Kalman filter is especially good at reproducing the correct soil moisture conditional mean. This appears to be a consistent result at all times and pixels in our experiment and it is observed both at the surface (Fig. 7) and integrated over the soil column (Fig. 8). Similar performance is obtained for evapotranspiration, which benefits most from radiobrightness measurements when it is limited by soil moisture.

It is worth noting that the structure and timing of precipitation exert a dominant influence on the land surface system. This influence tends to reduce differences between alternative data assimilation algorithms that make similar assumptions about rainfall. The RPM and multiplicative cascade disaggregation models used here tend to create very nonnormal soil moisture during rainy periods. In these periods soil moisture is skewed to the high end. As the surface moisture decreases through infiltration and evaporation, the skewness and kurtosis tend to decrease, making the ensemble filter's normality assumptions more appropriate. However, the skewness and kurtosis tend to increase again when soil dries and soil moisture is limited at the low end.

Soil properties also have a strong influence on the behavior of the land surface system and the performance of alternative filters. Open-loop (unconditional) predictions of soil moisture are usually better for rapidly infiltrating sandy soils than for less permeable loam or clay soils. Also, soil moisture updates have less impact on evapotranspiration for sandy soils. In such situations differences between optimal and suboptimal filtering algorithms are less likely to be dramatic.

Even taking these distinctive problem features into account, our overall conclusion is that the ensemble Kalman filter provides surprisingly good performance in the land surface application. This applies both to the filter's ability to characterize nonnormal distributional properties and its ability to provide accurate conditional means. We believe these results support previous studies that indicate the ensemble Kalman filter is a good estimation option for land surface applications. It would be useful to see the results of computational experiments similar to ours in other application areas. Such experiments could provide better understanding of when and why the ensemble Kalman filter can deal with nonlinearities and nonnormal uncertainties.

Acknowledgments. This research was sponsored by the U.S. National Science Foundation (EAR-0003361, ACI-0121182).

REFERENCES

Arulampalam, M. S., S. Maskell, N. Gordon, and T. Clapp, 2002: A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans. Signal Process.,* **50,** 174–188.

Bonan, G. B., 1996: A land surface model (LSM version 1.0) for ecological, hydrological, and atmospheric studies: Technical description and user's guide. NCAR Tech. Note NCAR/TN-4171STR, 150 pp. [Available online at http://www.cgd.ucar.edu/tss/lsm/availability/technote.tar.Z.]

——, K. W. Oleson, M. Vertenstein, S. Levis, X. B. Zeng, Y. J. Dai, R. E. Dickinson, and Z. L. Yang, 2002: The land surface climatology of the community land model coupled to the NCAR community climate model. *J. Climate,* **15,** 3123–3149.

Choudhury, B. J., T. J. Schmugge, A. Chang, and R. W. Newton, 1979: Effect of surface-roughness on the microwave emission from soils. *J. Geophys. Res.,* **84,** 5699–5706.

Crow, W. T., and E. F. Wood, 2003: The assimilation of remotely sensed soil brightness temperature imagery into a land surface model using ensemble Kalman filtering: A case study based on ESTAR measurements during SGP97. *Adv. Water Resour.,* **26,** 137–149.

Entekhabi, D., and Coauthors, 2004: The Hydrosphere State (Hydros) Satellite Mission: An Earth System Pathfinder for glob-

al mapping of soil moisture and land freeze/thaw. *IEEE Trans. Geosci. Remote Sens.,* **42,** 2184–2195.

Evensen, G., 1994: Sequential data assimilation with a nonlinear QG model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.,* **99,** 10 143–10 162.

——, 2003: The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.,* **53,** 343–367.

Gorenburg, I. P., D. McLaughlin, and D. Entekhabi, 2001: Scale-recursive estimation of precipitation at the TOGA COARE site. *Adv. Water Resour.,* **24,** 941–953.

Gupta, V. K., and E. C. Waymire, 1993: A statistical analysis of mesoscale rainfall as a random cascade. *J. Appl. Meteor.,* **32,** 251–267.

Hawk, K. L., and P. Eagleson, 1992: Climatology of station storm rainfall in the continental U.S.: Parameters of the Bartlett-Lewis and Poisson Rectangular Pulses models. Tech. Rep. 336, Department of Civil Engineering, Massachusetts Institute of Technology, 35 pp.

Jackson, T. J., and T. J. Schmugge, 1991: Vegetation effects on the microwave emission of soils. *Remote Sens. Environ.,* **36,** 203–212.

——, D. M. Le Vine, A. Y. Hsu, A. Oldak, P. J. Starks, C. T. Swift, J. D. Isham, and M. Haken, 1999: Soil moisture mapping at regional scales using microwave radiometry: The Southern Great Plains Hydrology Experiment. *IEEE Trans. Geosci. Remote Sens.,* **37,** 2136–2151.

Jazwinsky, A. H., 1970: *Stochastic Processes and Filtering Theory*. Academic Press, 376 pp.

Margulis, S. A., and D. Entekhabi, 2001: Temporal disaggregation of satellite-derived monthly precipitation estimates and the resulting propagation of error in partitioning of water at the land surface. *Hydrol. Earth Syst. Sci.,* **5,** 27–38.

——, D. McLaughlin, D. Entekhabi, and S. Dunne, 2002: Land data assimilation and estimation of soil moisture using measurements from the Southern Great Plains 1997 Field Experiment. *Water Resour. Res.,* **38,** 1299, doi:10.1029/2001WR001114.

Njoku, E. G., W. J. Wilson, S. H. Yueh, S. J. Dinardo, F. K. Li, T. J. Jackson, V. Lakshmi, and J. Bolten, 2002: Observations of soil moisture using a passive and active low-frequency microwave airborne sensor during SGP99. *IEEE Trans. Geosci. Remote Sens.,* **40,** 2659–2673.

Reichle, R. H., and R. D. Koster, 2003: Assessing the impact of horizontal error correlations in background fields on soil moisture estimation. *J. Hydrometeor.,* **4,** 1229–1242.

——, D. B. McLaughlin, and D. Entekhabi, 2002: Hydrologic data assimilation with the ensemble Kalman filter. *Mon. Wea. Rev.,* **130,** 103–114.

Rodriguez-Iturbe, I., D. Entekhabi, and R. L. Bras, 1991: Nonlinear dynamics of soil moisture at climate scales. 1: Stochastic analysis. *Water Resour. Res.,* **27,** 1899–1906.

Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, 2003: Ensemble square root filters. *Mon. Wea. Rev.,* **131,** 1485–1490.