Evaluating the utility of satellite soil moisture retrievals over irrigated areas and the ability of land data assimilation methods to correct for unmodeled processes

S. V. Kumar[1,2]**, C. D. Peters-Lidard**[2]**, J. A. Santanello**[2]**, R. H. Reichle**[3]**, C. S. Draper**[3,4]**, R. D. Koster**[3]**, G. Nearing**[1,2]**, and M. F. Jasinski**[2]

[1]Science Applications International Corporation, Beltsville, MD, USA
[2]Hydrological Sciences Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD, USA
[3]Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, MD, USA
[4]Universities Space Research Association, NASA Goddard Space Flight Center, Greenbelt, MD, USA

Correspondence to: S. V. Kumar (sujay.v.kumar@nasa.gov)

Abstract

The Earth's land surface is characterized by tremendous natural heterogeneity and human engineered modifications, both of which are challenging to represent in land surface models. Satellite remote sensing is often the most practical and effective method to observe the land surface over large geographical areas. Agricultural irrigation is an important human induced modification to natural land surface processes, as it is pervasive across the world and because of its significant influence on the regional and global water budgets. In this article, irrigation is used as an example of a human engineered, often unmodeled land surface process, and the utility of satellite soil moisture retrievals over irrigated areas in the continental US is examined. Such retrievals are based on passive or active microwave observations from the Advanced Microwave Scanning Radiometer for the Earth Observing System (AMSR-E), the Advanced Microwave Scanning Radiometer 2 (AMSR2), the Soil Moisture Ocean Salinity (SMOS) mission, WindSat and the Advanced Scatterometer (ASCAT). The analysis suggests that the skill of these retrievals for representing irrigation effects is mixed, with ASCAT based products somewhat more skillful than SMOS and AMSR2 products. The article then examines the suitability of typical bias correction strategies in current land data assimilation systems when unmodeled processes dominate the bias between the model and the observations. Using a suite of synthetic experiments that includes bias correction strategies such as quantile mapping and trained forward modeling, it is demonstrated that the bias correction practices lead to the exclusion of the signals from unmodeled processes, if these processes are the major source of the biases. It is further shown that new methods are needed to preserve the observational information about unmodeled processes during data assimilation.

1 Introduction

Examples of human induced land surface changes include urbanization, deforestation, and agriculture, all of which have significant impacts on local and regional water and energy

budgets and hydrologic and biogeochemical processes. The expansion of infrastructure and agriculture, necessitated by increasing societal demands has led to significant transformation of the natural features of the land surface, affecting more than 50 % of the land area (Hooke et al., 2012). Most current land surface models are not only severely deficient in representing the impacts of such engineered artifacts, but are also limited in representing features of many natural systems such as seasonal flood plains and wetlands. Remote sensing measurements offer a potential alternative for capturing the effects of such unmodeled processes. Moreover, data assimilation, which is a common approach to merge the information from observations with model estimates, may provide a possible mechanism for incorporating the effects of such unmodeled processes into model estimates.

Irrigation is an important land management practice that has had a significant impact on the global and regional water budgets. As noted in Gordon et al. (2005), the global increase in water vapor flows from irrigation is comparable to the decrease caused by deforestation. It has been estimated that as much as 87 % of the global fresh water withdrawals by humans have been used for agriculture (Douglas et al., 2009), which leads to significant alteration of the global and regional hydrological cycle. Though recent studies have reported the development of conceptual representations of irrigation in land surface models (de Rosnay et al., 2003; Ozdogan et al., 2010; Leng et al., 2013; Lawston et al., 2015), capturing and representing the nature of irrigation practices remains a hard problem. Therefore, in this article we focus on irrigation as an analog of a human engineered process that is typically not represented in land surface models.

There is a long legacy of retrieving estimates of surface soil moisture from satellite microwave radiometry using a variety of sensors (Jackson, 1993; Njoku and Entekhabi, 1995). In the past decade, near surface soil moisture retrievals have become available from a number of passive microwave and scatterometer-based platforms. They include Advanced Microwave Scanning Radiometer-Earth Observing System (AMSR-E) aboard the Aqua satellite, WindSat multifrequency polarimetric microwave radiometer aboard the Coriolis satellite, the Advanced Scatterometer (ASCAT), a C-band active microwave remote sensing instrument aboard the Meteorological Operational (METOP) satellites, the Advanced Microwave

3

Scanning Radiometer 2 (AMSR2) onboard the Global Change Observation Mission-Water (GCOM-W) satellite, and the Soil Moisture Ocean Salinity (SMOS) mission, and the Soil Moisture Active Passive (SMAP) mission. Except for AMSR-E, which stopped functioning in October 2011, all these instruments are currently providing measurements of surface soil
5 moisture. In this article, we first examine if the satellite soil moisture retrievals are effective in capturing the effects of irrigation. The comparison is performed through a quantitative comparison of the probability density functions (PDFs) of remote sensing datasets against those of land surface model simulations that do not include formulations of irrigation.

The second focus of the article is to examine whether the current data assimilation prac-
10 tices are adequate if unmodeled processes such as irrigation are present in observations. When irrigation practices are employed, they lead to wetter soil moisture relative to non-irrigated time periods. Assuming that the physical model used in data assimilation does not have irrigation formulations in it, the assimilated observations would have a systematic bias (relative to the model) during irrigation periods. In real data assimilation systems, biases
15 between model forecasts and observations are unavoidable, and they typically result from a combination of model deficiencies, instrument and retrieval errors. Proper treatment of these biases is important, as the assimilation methods are primarily designed to work with errors that are strictly random (Dee and da Silva, 1998). Here we evaluate the impact of using common bias correction practices when unmodeled processes are the primary source
20 of the biases between the model and observations.

Primarily, there are two approaches to handling biases in data assimilation systems (Dee, 2005): (1) "bias-aware" systems which are built to diagnose and correct the biases in the observations and/or the model forecasts during data assimilation integration and (2) "bias-blind" systems which assume the observations and model forecasts to be unbiased. Ideally,
25 biases must be estimated by comparing the observations and/or model states to the true mean states, derived for example from in-situ measurements. However, as noted in Draper et al. (2015), developing spatially distributed bias estimates is much harder for the land surface, compared to the atmosphere or ocean, since point scale in situ observations are generally not representative of the spatial scale of remotely sensed or modeled states, due

4

to the heterogeneity of land. Though there have been a number of studies that rely on online estimation of biases (De Lannoy et al. (2007); Reichle et al. (2010)), the common practice in land data assimilation studies is to remove the bias *between* the observations and the model, and use a bias-blind assimilation approach to correct only short-lived model errors.

This is typically achieved by rescaling the observations prior to assimilation, to have the same statistics as the model, using quantile mapping approaches so that the observational climatology matches that of the land model. This approach is easy to implement as a pre-processing step to the data assimilation system and has been used extensively in many land data assimilation studies (Reichle and Koster, 2004; Drusch et al., 2005; Crow et al., 2005; Reichle et al., 2007; Kumar et al., 2009; Liu et al., 2011; Draper et al., 2011, 2012; Hain et al., 2012; Kumar et al., 2012, 2014). A known disadvantage of the approach is that it assumes stationarity in model-observational biases and cannot easily adjust to dynamic changes in bias characteristics. Common quantile mapping approaches used for scaling observations into the model's climatology include the standard normal deviate based scaling (Crow et al., 2005) and the CDF-matching method (Reichle and Koster, 2004; hereafter referred to as RK04). The standard normal deviate based scaling matches the first and second moments of the observation and model distributions, whereas the CDF matching approach corrects all quantile dependent biases between the model and observations, regardless of the shape of the distributions.

When observations are rescaled prior to assimilation, standard normal deviates or percentiles (rather than the raw observations) are assimilated. This ensures that the model climatology is preserved in data assimilation and that assimilation only affects temporal patterns of the anomalies. In such cases, the influence of assimilation is likely to be greater at the shorter time scales (Kumar et al., 2014).

In this article, we argue that the approach of rescaling the observations could be problematic, particularly when the underlying distributions of the model estimates and the observations are different. Such differences in the distributions are possible when features from human induced activities such as irrigation are present in observations and are missing in modeled estimates. Through a suite of synthetic experiments, we demonstrate the limita-

tions of the rescaling approaches when the reference climatology is fundamentally limited in representing unmodeled processes whose effects nevertheless impact the observations. In such cases, stationarity assumptions about the climatologies also could lead to spurious, statistical features in the assimilation results. As a result, the use of rescaling would become

5 problematic for demonstrating short-term assimilation impacts.

The rescaling approach through CDF-matching for land data assimilation proposed by RK04 was motivated by the fact that the true climatology of soil moisture at the global scale remains unknown. The CDF-matching method is based on similar applications of the method for establishing rainfall-reflectivity relationships for the calibration of radar or

10 satellite observations of precipitation (Calheiros and Zawadzki, 1987; Atlas et al., 1990; Anagnostou et al., 1999). The quantile mapping methods are also widely used for correcting biases of regional climate model simulations relative to observational data (Salathe Jr. et al., 2007; Li et al., 2010). These studies assume that the probability density functions of radar reflectivity and in-situ rainfall are equivalent and therefore quantile mapping can

15 be used to translate one into the other. RK04 extended this approach to soil moisture data assimilation by transforming the satellite soil moisture retrievals into the model's climatology for removing the relative biases between the model and observations. The important difference between the precipitation/climate downscaling studies and RK04 is that in the former, the remotely sensed retrievals/climate model data were rescaled to observed data,

20 whereas in RK04, the satellite retrievals were rescaled to a modeled climatology. In this article, we demonstrate that rescaling to a model climatology that is not representative of the observations may distort the scale of the actual observational features and may lead to loss of valuable signals.

There are a number of alternative strategies for bias correction in bias-blind data assim-

25 ilation systems. Instead of employing a single CDF (at each grid cell) that encapsulates the soil moisture dynamics across all seasons (called lumped CDFs), temporally stratified (monthly or seasonally) CDFs can be used. The finer temporal stratification would help to reduce the impact of statistical artifacts of using lumped CDFs, but would also require sufficient sample sizes to accurately derive CDFs for each temporal window. As demonstrated

6

in Kumar et al. (2012), the land surface model could be calibrated against the retrieval products and the calibrated, unbiased model could then be used in assimilation. Though this strategy eliminates the biases in the variables being assimilated, the climatologies of other outputs from the model could be affected, unless additional constraints are included in calibration. De Lannoy et al. (2013) employed a similar strategy for assimilating SMOS L-band brightness temperatures, by calibrating the forward radiative transfer model parameters and by keeping the land surface model parameters unchanged. This strategy preserves the land surface model climatology, but only works when radiance measurements are being assimilated through a forward model. Forman et al. (2014) suggested a similar strategy for assimilating passive microwave brightness temperatures for snow data assimilation through the use of an artificial neural network (ANN). The ANN uses the inputs from the land surface model and is trained to the observed brightness temperatures to be assimilated. The results of Forman et al. (2014) indicated that the ANN could serve as a computationally efficient observation operator instead of more complex radiative transfer models. In the current study, we evaluate the effectiveness of a number of these strategies for land data assimilation when the observations include the effects of processes that are not included in the land surface model.

The article is organized as follows: First, we examine the effectiveness of satellite soil moisture retrieval products in their ability to capture the effects of irrigation (Sect. 2). The evaluation is conducted by quantitatively comparing the probability distribution functions from various remote sensing soil moisture datasets and land surface model simulations. The article then focuses on the impact of various a priori bias correction approaches in data assimilation when the distributions of the model and the observations are significantly different due to unmodeled irrigation processes (Sect. 3). Section 3 presents a synthetic data assimilation experiment that explores the limitations of a suite of a priori bias correction strategies in such scenarios. Finally, Sect. 4 presents a summary and discussion of major conclusions of the study.

7

2 Evaluation of satellite remote sensing data over irrigated areas

In this section, we examine the utility of modern soil moisture remote sensing datasets towards the detection of irrigation features. The irrigation practices over the world differ in the method of irrigation, the trigger used and the amount of water used in irrigation. A

[5] typical irrigation practice in the US is to apply irrigation throughout the growing season at a level where the plants are not under transpiration stress. The introduction of irrigation at the beginning of the growing season would lead to increased surface soil moisture and a significant dry down would only occur at the end of the growing season when irrigation controls are removed. Figure 1 shows the MODIS-based irrigation gridcell fraction map (%)

[10] derived by Ozdogan and Gutman (2008) and validated against USGS irrigation data. Some of the known hotspots of irrigation over the Continental US are highlighted in this map, which includes the plains of Nebraska, lower Mississippi basin and central California valley.

To demonstrate the impact of irrigation, land surface model simulations are conducted at a single grid point located in the plains of Nebraska (as shown in Fig. 1). Figure 2 presents

[15] the time series of surface soil moisture (using a 10 cm thick surface layer) for a representative year from two simulations of the Noah (version 3.3; Ek et al., 2003) land surface model. The simulations demonstrate the impact of irrigation at this location: (1) model forced with a given meteorological forcing data (SIM1) and (2) a seasonal irrigation scheme simulated on top of the SIM1 configuration (SIM2). The simulations use the modified 20 category MODIS

[20] landcover data (Friedl et al., 2002) and are forced with meteorological boundary conditions from the North American Land Data Assimilation System Phase 2 (NLDAS-2; Xia et al., 2012) data. The initial conditions for the model simulations are generated by spinning up the LSM from 1979 to 2000. The irrigation scheme employed here simulates a demand driven, sprinkler irrigation technique, based on Ozdogan et al. (2010). Irrigation is triggered

[25] when the root zone soil moisture falls below the transpiration stress threshold for a particular grid cell. The scheme computes the irrigation requirement as an equivalent height of water, which is then applied as an addition to the precipitation input to the model. The irrigation scheme is only applied to irrigated land types such as crops and grasslands and is

only enabled during growing seasons (when 40 % of the annual range of green vegetation fraction at a grid cell is exceeded). In addition, the irrigation requirement is enabled daily between 06:00 and 10:00 local time, similar to the approach used in Ozdogan et al. (2010).

As the seasonal irrigation picks up in April, surface soil moisture in SIM2 gets much wetter compared to the SIM1 integration. Towards the end of the summer, the imposed irrigation is removed, which then causes the soil moisture to dry down and approach the SIM1 based estimates. For the purpose of developing climatologies, the model integrations were conducted for several years (2000–2012). Similar features are repeated in other years, leading to a seasonal wet bias in SIM2 compared to SIM1, mainly in the summer months.

A comparison of the soil moisture distributions from the two integrations are shown in Fig. 3, which shows a plot of the quantiles of SIM1 data against the quantiles of SIM2 data. For comparison, the figure also shows a 45-degree reference line. If the two datasets come from similar distributions, the points in the q-q plot should fall approximately along this reference line and the departure from the reference line indicates differences in the distributions of the two datasets. Figure 3 indicates that there are significant differences in the SIM1 and SIM2 distributions with shifts in location, scale and symmetry. The data points are systematically above the 45-degree line indicating that the mean of the distributions are significantly different, with distribution for SIM2 higher valued than that of SIM1. The quantile-quantile plot also shows a bimodal nature due to the seasonal effect of irrigation. Finally, the slope of the points on the q-q plot is higher than 1, indicating that there are differences in the spread or variances of the two distributions as well.

A two sample Kolmogorov-Smirnov (K-S) test (Chakravarti et al., 1967) can be used to quantitatively compare the probability distributions of two datasets. The K-S statistic quantifies a distance between empirical distribution functions of two samples ($F(x)$ and $G(x)$ where x is the sampled variable), and is computed as follows:

$$D_{m,n} = \max_x |F(x) - G(x)|, \tag{1}$$

where m and n are the sample sizes of F and G.

The null distribution of the K-S statistic is calculated under the null hypothesis that samples are drawn from the same distribution. The null hypothesis is rejected at level α if

$$D_{m,n} > c(\alpha)\sqrt{\frac{m+n}{mn}}, \tag{2}$$

where $c(\alpha)$ is the inverse of the Kolmogorov distribution at α.

To examine the impact of irrigation on soil moisture distributions over a larger spatial domain, the SIM1 and the SIM2 experiments are extended to a larger domain, encapsulating the Continental US at $0.125°$ spatial resolution. The model integrations are conducted during the time period of 2000 to 2013. The K-S test is then applied to the probability distributions of the surface soil moisture estimates from the two integrations. The resulting values for the K-S statistic (D) are shown in Fig. 4. Only locations at which the null hypothesis of the K-S static is rejected are shown in Fig. 4. Values of D closer to zero indicate that the soil moisture distributions from the SIM1 and SIM2 integrations are similar. Conversely, larger D values indicate locations where the soil moisture distributions from the two integrations differ. As the difference between the two integrations in this example is only due to the simulation of seasonal irrigation, the locations with positive K-S metric values in Fig. 4 indicate areas where the irrigation artifacts are applied and are consistent with the input irrigation intensity data used in the simulations. The K-S metrics, therefore, can be used to detect instances where the distributions of soil moisture retrievals and the model estimates differ significantly, including differences due to the treatment of irrigation.

Figure 5 shows a quantitative comparison of the differences in soil moisture distributions using the K-S metrics from six remote sensing soil moisture retrievals and a land surface model simulation (SIM1 configuration), for the continental US. The remote sensing based products are: (1) the blended multi-sensor soil moisture product from the European Space Agency (ESA) known as the Essential Climate Variable (ECV) product (Liu et al., 2012b), (2) soil moisture retrievals from AMSR-E using the Land Parameter Retrieval Model (LPRM) algorithm (Owe et al., 2008), (3) soil moisture retrievals from WindSat, (4) soil moisture retrievals from the backscatter measurements acquired by ASCAT, (5) soil moisture retrievals from AMSR2 and (6) soil moisture retrievals from the SMOS mission. The WindSat

10

and ASCAT retrievals are obtained through the Soil Moisture Operational Products System (SMOPS; Liu et al., 2012a) of NOAA/NESDIS. The level 3 AMSR2 data from the Japan Aerospace Exploration Agency (JAXA; EORC, 2013) and the level 2 swath based SMOS products from ESA (Kerr et al., 2006) are used in these comparisons. The temporal extent of these datasets varies. The ECV data are available from January 1979–December 2013, AMSR-E from June 2002–October 2011, WindSat from January 2007–present, AS-CAT from January 2007–present, AMSR2 from July 2012–present and SMOS from April 2012 to present. The CDFs for each dataset are computed using all available data. The available quality control information in each remote sensing dataset is used to exclude data over regions with dense vegetation, RFI interference, precipitation and frozen ground. The model CDFs are computed using the simulated surface soil moisture estimates from 2000 to 2013. Differences in the dynamic range between observed and modeled soil moisture are normally removed prior to assimilation and here we remove the differences in the mean and variance prior to calculating the K-S metrics. The data values are normalized first with a standard score approach ($(x_i^t - \mu_i)/\sigma_i$ where x_i^t is the data value at time t and μ_i and σ_i are the mean and standard deviation of the data at grid point i) before computing the K-S metrics. Figure 5 shows the D estimates from the K-S test for each comparison at grid points where the the null hypothesis of the K-S static is not rejected. Not surprisingly, the ECV data comparison shows the lowest D values, possibly due to the fact that the ECV product was generated by CDF matching soil moisture estimates from different sensors to a simulation of Noah land surface model from the Global Land Data Assimilation System (GLDAS; Rodell et al., 2004). Comparatively, larger differences are seen in all other comparisons, which are likely caused by a mix of biases resulting from instrument error, retrieval algorithms errors, and unmodeled processes, and other representativeness differences.

The spatial patterns of these metrics shown in Fig. 5 can also be potentially used as a first measure of whether a sensor captures observational features such as irrigation. Specifically, a relatively small K-S metric at a location known to have irrigation suggests that the remotely sensed observations did not detect that irrigation. However, the converse is not necessarily true, in that a large K-S metric does not necessarily indicate successful detec-

tion of irrigation (since it could be caused by other model-remotely sensed discrepancies). For example, a strong signal of vegetation density in the eastern US can be noticed in the K-S metric map for AMSR2. Similarly, in the ASCAT K-S metric map, large differences can be observed around several major cities such as Dallas, Houston and Atlanta. We focus on three key hotspots of irrigation in the US shown in Fig. 1: the plains of Nebraska, lower Mississippi basin and central California valley. Of these three regions, only the lower Mississippi has relatively higher K-S metric values, and only for AMSR-E, AMSR2 and SMOS.

To examine if the spatial patterns of differences in K-S metrics from various remote sensing datasets are in fact representative of observational artifacts such as irrigation, we examine the time series of soil moisture over these regions. For an equivalent comparison given the possible differences in the respective dynamic ranges, each dataset is normalized first (using the standard score approach) before comparing them on the same graph. Figure 6 shows the normalized time series of soil moisture from the observations (from SMOS, AMSR2 and ASCAT), the Noah LSM driven with the NLDAS-2 forcing, with and without irrigation for the year 2013. A 5-day moving average is applied to the SMOS, AMSR2 and ASCAT retrievals to reduce the noise in the satellite retrieval time series. As the gridcell averages for the satellite retrievals reflect averages of the irrigated and non-irrigated pixels, a similar weighted estimate was produced for the LSM-based irrigated soil moisture values. The LSM time series representing irrigation in Fig. 6 is generated by weighting the model soil moisture estimate with and without irrigation for each gridcell by the irrigation fraction and (1 – irrigation fraction) of that gridcell, respectively. Finally, in each region, only gridcells with at least 30 % irrigation fraction indicated by the MODIS map are employed in computing the spatial averages.

Figure 6 indicates that the SMOS and AMSR2 retrievals agree more closely with the LSM estimate without irrigation, in all three regions. In particular, there are few indicators of systematic differences between the observations and the model without irrigation in the summer months, suggesting the limited skill of the SMOS and AMSR2 retrievals for detecting features of seasonal irrigation. In contrast, ASCAT retrievals show better agreement with the LSM estimate with irrigation in the summer and fall months, over the plains of Ne-

12

braska and lower Mississippi basin. In these regions, the ASCAT moisture signal shows a wetter trend in the late fall months, which are in agreement with the LSM estimate with irrigation. In the central California valley, however, no such distinct contrast due to irrigation is observed in all three satellite retrievals. Similar trends are seen in other years (not shown).

From these results, it appears that neither SMOS nor AMSR2 retrievals capture the effects of irrigation, whereas the ASCAT retrievals are somewhat effective in detecting irrigation features in the plains of Nebraska and the lower Mississippi basin.

It is important to note here that the apparent inability of SMOS and AMSR2 to capture the irrigation signal should not be assumed attributable to sensor deficiency; it may instead reflect their larger spatial footprints. The raw resolutions of SMOS and AMSR2 products used here are at least 40 km, much coarser than the 0.125 degree resolution employed in the LSM simulations. Thus, because Fig. 6 focuses on 0.125 grid cells with at least 30 % irrigation, the SMOS and AMSR2 data (interpolated to that resolution) will necessarily include some soil moisture information from areas outside those defined by the 30 % threshold – areas that are, almost by definition, drier. ASCAT, with a raw resolution of ∼25 km does not seem as affected by this, perhaps in part due to its finer base resolution. Another possible reason may be related to the influence of intercepted water, which has opposite effects on the active and passive sensors. More analysis is needed, however, to understand the different behaviors of the sensors.

Note that in the model formulations, irrigation is simulated consistently from the late spring months to early fall months, though these assumptions about the timing and duration of irrigation in these regions may be imperfect relative to the actual practices in the field. In the plains of Nebraska, the agreement between the ASCAT and model with irrigation are consistent throughout the summer and early fall months (from late June to early October). In lower Mississippi, on the other hand, the ASCAT time series indicates that the application of irrigation occurs in the later months (from late August onwards). The agreement between the model with irrigation and the ASCAT time series is lower in the early summer months. Though it is hard to ascertain the ability of ASCAT data for characterizing the timing of irrigation, it can be concluded that ASCAT retrievals perform better than the

13

SMOS and AMSR2 retrievals in terms of capturing the anomalous wet soil moisture signals from irrigation over these areas known to be irrigated.

3 Evaluation of bias correction strategies in the presence of unmodeled processes

This section presents an examination of the effectiveness of a number of a priori bias cor-
⁵ rection strategies in data assimilation when unmodeled processes (such as irrigation) are a major source of biases between the model and the observations. A synthetic experiment setup based on the SIM1 and SIM2 configurations presented in Section 2 is used to explore these issues.

If SIM2 represents the observations to be used in assimilation, the typical procedure in
¹⁰ data assimilation systems is to rescale SIM2 estimates to the model climatology (SIM1 in this example). Figure 7 illustrates the impact of rescaling SIM2 to SIM1 climatology with CDF matching (using both lumped and monthly CDFs), for the year 2000. When lumped CDFs are used, rescaling leads to a wetter soil moisture time series (compared to SIM1) during the summer months (but significantly lower than SIM2), whereas during the non-
¹⁵ irrigation months, rescaling leads to a much drier soil moisture time series, relative to SIM1. Lumped CDF-matching attempts to keep the climatology of the rescaled time series to be close to the overall SIM1 climatology. As a result, higher soil moisture values during irri-gation are compensated by lower soil moisture values during non-irrigated months to keep the overall climatology the same as that of SIM1. In this example, the lumped CDF-based
²⁰ rescaling approach introduces spurious statistical artifacts during non-irrigated periods. The statistical artifacts of rescaling during the non-irrigated months are greatly reduced if the CDF-matching is performed in a more temporally stratified manner. As indicated by Fig. 7, when rescaling uses monthly CDFs, the resulting time series remain close to SIM1 both during the irrigated and non-irrigated periods. Note that most data assimilation studies (Re-
²⁵ ichle et al., 2007; Kumar et al., 2009; Liu et al., 2011; Draper et al., 2012; Kumar et al., 2012, 2014) use the lumped CDF-scaling approach due to sampling density limitations of using temporally finer resolved CDFs.

14

3.1 Structure of synthetic data assimilation experiments

The suite of data assimilation experiments employs an identical twin experiment setup. The model simulations are conducted at the single grid point shown in Fig. 1. The Noah LSM simulation forced with the NLDAS-2 data is termed as the open loop (OL) integration. A scheme designed to mimic seasonal crop irrigation employed on top of the OL configuration is used as the "Control/Truth" simulation. All model integrations use the same forcing and parameter datasets as that of the experiment presented in Section 2. The time period from 2000 to 2012 is used here for various evaluations.

From the truth simulation, observations are generated after incorporating realistic errors and limitations of passive microwave remote sensing retrievals. To account for difficulties in retrieving soil moisture products from microwave sensors, the observations are masked out when the green vegetation fraction values exceed 0.7 and when snow or precipitation are present. Random gaussian noise with an error standard deviation of 0.02 $m^3 m^{-3}$ is added to the truth soil moisture values to mimic measurement uncertainties, which is an optimistic estimate of the error levels in the current space-borne L-band radiometers (SMOS and SMAP). Finally, a data assimilation (DA) integration that assimilates the simulated observations in the OL configuration is conducted. The DA and OL integrations are compared against the known truth to evaluate the impact of observations.

Most synthetic experiment studies (Reichle et al., 2002b; Crow et al., 2005; Crow and Reichle, 2008; Kumar et al., 2009, 2012) use different inputs and models in the control and OL configurations to simulate the systematic biases that are often present (between observations and the model) in real data assimilation scenarios. Here we intentionally use a setup where the only difference between the control run and OL is a process (irrigation) that is not modeled in the OL simulation, but included in the control run. One could envision similar issues in real data assimilation systems, where features from engineered systems will be present in observations, but not simulated in physical models. In this idealized scenario, biases between the model and the observations are purely from observational features that are not modeled.

15

Four different data assimilation integrations are conducted using the synthetic observations: (1) DA-NOBC, assimilating observations directly without any bias correction, (2) DA-CDFL, assimilating a priori scaled observations using CDF matching (using lumped CDFs representing all years and seasons), (3) DA-CDFM, assimilating a priori scaled observations using monthly CDF matching (the model and observation CDFs are generated separately for each calendar month) and (4) DA-ANN, assimilating the simulated observations directly and using a trained ANN as the observation operator in the data assimilation system (see Sect. 3.3 for details). In experiments DA-NOBC, DA-CDFL and DA-CDFM, the observation operator is the land surface model itself, whereas the observation operator is represented by the trained ANN in the DA-ANN experiment.

In the DA-CDFL experiment, the observation and model CDF are first computed independently for each grid cell using the 13 year (2000–2012) period. During data assimilation, the observations are rescaled (separately for each grid cell) using these lumped CDFs. As noted by Drusch et al. (2005), the climatologies between the model and observations may change with season, which is clearly the case in our synthetic experiment setup due to the influence of seasonal irrigation. In the DA-CDFM experiment, the observation and model CDFs are generated separately for each month and for each grid point. The 13 year record of data ensures that there is enough sampling density to accurately derive CDFs when the CDF calculation is stratified by calendar months.

3.2 Data assimilation method

The data assimilation integrations are conducted using a one-dimensional Ensemble Kalman Filter (EnKF; Reichle et al., 2002a) algorithm. An ensemble size of 12 is used in the simulations with perturbations applied to both meteorological fields and model prognostic fields to simulate uncertainty in the model estimates. The determination of 12 as the ensemble size was based on prior works (Reichle et al. (2007); Kumar et al. (2008, 2009, 2012)) and because the size of the model state vector is small (4 Noah soil moisture state variables). The EnKF alternates between an ensemble forecast step and a data assimilation step. An ensemble of model states is propagated forward in time using the land surface

model during the forecast step. In the update step at time k, the model forecast is adjusted toward the observation based on the relative uncertainties, with appropriate weights expressed in the "Kalman gain" $\mathbf{K}_k$:

$$x_k^{i+} = x_k^{i-} + \mathbf{K}_k[y_k^i - \mathbf{H}_k x_k^{i-}], \tag{3}$$

where x_k and y_k represent the model state and observation vectors, respectively. The observation operator $\mathbf{H}_k$ relates the model states to the observed variable. The superscripts $i-$ and $i+$ refer to the state estimates of the i-th ensemble member $(-)$ before and $(+)$ after the update, respectively. Equation (3) indicates that the analysis increments $(x_k^{i+} - x_k^{i-})$ are computed by multiplying the Kalman gain $\mathbf{K}_k$ with the innovations $(y_k^i - \mathbf{H}_k x_k^{i-})$. In "bias-blind" data assimilation systems, observations (y_k) and model forecasts $(\mathbf{H}_k x_k^{i-})$ are expected to be unbiased relative to each other, which presents two choices for bias correction: (1) rescale observations into the model climatology, so that the innovations are computed in the climatology of $\mathbf{H}_k x_k^{i-}$ or (2) compute the innovations in the observation space by having an operator $(\mathbf{H}_k)$ that translates the model states into the observation space. The quantile mapping approaches fall in the first category, whereas the use of trained forward models as observation operator represents the second category. We examine the impact of using both set of approaches when unmodeled processes dominate the sources of biases.

3.3 Use of a trained ANN as a forward observation operator

Artificial neural networks (ANNs) are data processing systems used for pattern matching applications and consist of a highly interconnected array of processing elements (called neurons), designed as a mathematical generalization of human cognition and learning. The basic architecture of an ANN consists of three layers: input, hidden and output layers. The inputs processed through the input layer are communicated to the hidden layers and the results are output through the output layer. The topology of the layers (defined by "activation functions") and the weights of the interconnections are used to develop accurate outputs.

During the training phase, the ANN is presented with a set of inputs and corresponding outputs. The trained ANN can then be used for generating new predictions when presented with a new set of inputs.

Figure 8 shows the structure of the ANN used in this study. The input layer consists of
5 six inputs, which are a combination of the meteorological inputs (rainfall and snowfall), land surface model parameters (green vegetation fraction) and land surface model estimates (surface soil temperature, snow water equivalent and surface soil moisture). Note that the surface soil moisture in the input layer is from the LSM integration without irrigation. Five neurons were employed in the hidden layer based on a similar approach used in Cao et al.
10 (2008) and Forman et al. (2014). For this study, a single output node that estimates surface soil moisture values is used. The ANN is trained to the simulated surface soil moisture observations at this grid point (generated from the truth integration) during the time period of 2000–2012. Since the entire observation record is used for estimating CDFs, we use the whole record for training the ANN as well, so that the experiments are comparable. Figure 9
15 shows a comparison of the simulated observations and the estimates from the trained ANN for the year 2000. It can be seen that the trained ANN model helps to capture the wetting of the soil due to seasonal irrigation during the spring and summer months. Similar patterns are observed in other years (not shown). The skill in turning on or off irrigation in the ANN is likely due to the incorporation of the information in the training inputs of soil temperature
20 and green vegetation fraction. The trained ANN is then used in the DA experiments.

3.4 DA experiment results

The evaluation of the four DA experiments is presented in Figures 10 to 12. Figure 10 shows daily averaged soil moisture estimates from various model integrations for the year 2000 as a representative time series. The DA-NOBC integration assimilates the raw observations
25 (shown in the figure) and as a result provides soil moisture estimates closer to the Control simulation during times when observations are available. DA-CDFL and DA-CDFM integrations ingest rescaled observations (not shown), which do not show a systematic increase in the soil moisture values during the spring and summer months. Similar behavior is seen

for the DA-ANN integration, which assimilates the raw observations, but does not represent the anomalously wet soil moisture of the Control simulation. The DA-CDFL, DA-CDFM and DA-ANN integrations do not deviate much from the open loop integrations as the size of the analysis increments ($\mathbf{K}_k[\boldsymbol{y}_k^i - \mathbf{H}_k\boldsymbol{x}_k^{i-}]$) in these integrations is small. In DA-CDFL and DA-CDFM, the rescaling causes the innovations to be computed in the climatology of the model states whereas in the DA-ANN experiment, the innovations are computed in the climatology of the observations. In either case, when these small increments generated by the assimilation system are applied back to the soil moisture forecast values (in the open loop climatology), the anomalous wet signals in the observations are removed as bias artifacts and are never included in the analysis.

Figure 11 shows the average seasonal cycle of RMSE (stratified monthly across the entire simulation period of 2000–2012) of surface soil moisture from various model integrations. Similar to the trends in Fig. 10, the DA-NOBC integration shows significant improvements from data assimilation except for August. The peak of vegetation (determined based on the green vegetation fraction) occurs in August leading to observations being excluded from the data assimilation system. As a result, the improvements through assimilation are small during this time period. The seasonal nature of the RMSE estimates from DA-CDFL, DA-CDFM, and DA-ANN is similar, and is close to the open loop RMSE estimates. The use of the scaled observations (in DA-CDFL and DA-CDFM) and the use of the trained forward model (in DA-ANN) causes the dampening and exclusion of the wet biases from irrigation in these DA integrations.

An important philosophical point, however, is warranted here. Implicit in the above discussion of Fig. 11 is the assumption that a higher RMSE reflects a poorer performance. Depending on application, this may not be true at all. It is a well-established fact that the soil moisture estimate from the model is essentially an index of wetness and a highly model-dependent quantity (Koster et al. (2009)). As a result, care must be exercised when comparing model soil moisture directly to in-situ or satellite measurements. The whole point of the scaling exercise is to convert a satellite-based soil moisture value, prior to its assimilation, to a value consistent with that of the LSM used. This allows the further use of the

assimilated soil moisture value in that LSM, e.g., to initialize a forecast. If, once the data assimilation process is finished, a soil moisture value is needed that reflects a more "correct" climatology (e.g. with an irrigation-influenced seasonal cycle, as in the Control simulation), the data assimilation product can easily be scaled back to that climatology using the reverse
5 of the original scaling approach. Viewed in this light, the data assimilation approach, with scaling, is essentially designed to capture the year-to-year or short-term variations in soil moisture anomalies rather than the structure of the seasonal cycle. Also note that though the seasonal cycle of RMSE is lowest in the DA-NOBC integration, this configuration is not really viable in real data assimilation systems where biases are unavoidable. The DA-NOBC
10 integration is included in the suite of experiments, as we have the knowledge of the exact sources and magnitudes that contribute to the biases in this synthetic configuration. For real data assimilation systems, metrics recommended by Entekhabi et al. (2010) that compute estimates of soil moisture accuracy while accounting for biases, may be more appropriate.

The EnKF algorithm assumes linear system dynamics. It further assumes model and
15 observation errors that are Gaussian and mutually and serially uncorrelated. If these assumptions hold, then the distribution of filter innovations (observation minus model forecast residuals) normalized by their expected covariance will follow a standard normal distribution $N[0,1]$ (Gelb, 1974). The deviations from the $N[0,1]$ of the normalized innovation distribution is typically used as a measure of the degree of suboptimality of the data assimilation
20 system (Reichle et al., 2002a; Crow and Van Loon, 2006; Reichle et al., 2007; Kumar et al., 2008). Figure 12 compares the distribution of the normalized innovations from DA-NOBC, DA-CDFL, DA-CDFM and DA-ANN to the $N[0,1]$. The mean and standard deviation of each distribution are also reported in the figure. Unsurprisingly, the DA-NOBC indicates the largest deviation in the mean among the experiments, indicating the presence of a bias.
25 The mean values of the distributions from DA-CDFL, DA-CDFM and DA-ANN are closer to zero, due to various a priori bias correction strategies employed.

These internal diagnostics are also often used for the estimation of input error parameters of the data assimilation system. For example, in adaptive filter implementations (Reichle et al., 2008), the model and observation error specifications are continually adjusted to

yield near optimal behavior of the internal diagnostics (i.e., close to $N[0, 1]$ response of normalized innovation distribution). The analysis presented above indicates that if unmodeled processes are present, these biases are reflected in the innovation diagnostics as deviations from expected optimal measures. In such cases, the reliance on these assimilation

5 diagnostics may be misleading if the end goal of the assimilation process is to correct the modeled seasonal cycle of soil moisture toward that of the observations. As noted above, though, the goal may instead be to capture, for various applications, year-to-year anomalies in soil moisture, and for this purpose, the innovation diagnostics reveal that the scaling approaches do provide superior behavior. Again, though, it should be pointed out that the

10 synthetic experiment was designed to isolate the impacts of unmodeled irrigation. In practice, the effects of unmodeled irrigation will be conflated with bias issues that result from differences in land surface parameters and differences in the very meaning (such as layer-depth) of the modeled and retrieved soil moisture values.

4 Summary

15 Due to the heterogeneity of the land surface and the large impact of human activities, quantifying the variability of water and energy budgets on the land surface presents unique challenges compared to the atmosphere and ocean components of the Earth system. Irrigation is one of the pervasive human induced land management practices that has a direct impact on the local and regional water budgets. In this article, we examine the utility of satellite

20 soil moisture retrievals to detect irrigated areas. In addition, the article also examines the limitations of current data assimilation practices when the observations are dominated by processes that are not included in the land model.

Application of seasonal irrigation is likely to introduce systematic differences in the soil moisture distributions. Therefore, if the remote sensing datasets are skillful in detecting

25 irrigation features, the resulting soil moisture distributions would be significantly different compared to a model simulation that does not simulate irrigation. We use this hypothesis to examine the effectiveness of modern remote sensing soil moisture products from ASCAT,

AMSR2 and SMOS in their ability to detect irrigation. A two sample Kolmogorov-Smirnov test is used to to quantify the systematic differences between distributions of model and remote sensing datasets, over a continental US domain. The analysis reveals systematic differences in spatial patterns of the distributions of model and remote sensing data. Additional analysis, however, suggests that these differences are not always related to the detection of irrigation artifacts. Generally ASCAT retrievals were found to be somewhat more skillful than the SMOS and AMSR2 retrievals in their ability to capture features of irrigation on the land surface.

Overall the analysis presented in the paper assumes a demand-driven irrigation scheme maintained throughout a growing season at a level where the plants are not under transpiration stress. In reality, however, the type and level of irrigation may not be seasonally persistent and therefore the nature of the expected biases in the soil moisture signal due to irrigation may not be systematic throughout a season. Further comparisons with in-situ soil moisture data at irrigated locations will be required to confirm and isolate the limitations of the remote sensing data over these areas. A major source of the biases between the satellite retrievals and the LSM estimates is the differences in the land surface parameters used in the respective models. The biases from these parameter differences are likely to dominate the more subtle effects of irrigation. In addition, the scale mismatches between the model and the observations are also likely to have an influence in the comparisons presented here. The spatial resolution of the model (0.125°) and the observations (∼25–40 km) can be considered relatively coarse for detecting uniformly and simultaneously irrigated areas.

The second focus of the article is on the limitations of various a priori bias correction strategies in land data assimilation towards representing unmodeled processes. This issue is explored through a suite of synthetic data assimilation experiments. A simulation of seasonal irrigation is used as analog for an engineered process that is typically not included in large-scale land surface model simulations. The data assimilation integrations merge the observations generated from the irrigation simulation into a model in which irrigation is absent that features a free running land surface model. The data assimilation integrations

22

include simulations that employ no bias correction, a lumped CDF-matching correction, a seasonally varying CDF-matching correction or ANN as a forward observation operator. As the a priori bias correction approaches make no distinction of the source of the biases (unmodeled or from other sources) they treat all systematic differences between the model and observations as biases. As a result, all a priori bias correction strategies considered above cause the signal from seasonal irrigation (or other unmodeled processes) to be excluded in the DA results, though the analysis of the DA internal diagnostics indicate near optimal performance for such configurations.

The challenge in data assimilation systems is to separate the biases that are due to instrument error and retrieval algorithm errors from the biases induced by unmodeled processes so that the true observational features are not excluded during data assimilation. Detecting spatial patterns of such differences could be useful in the utilization of soil moisture datasets in data assimilation systems. For example, even if a particular retrieval dataset is not skillful in detecting irrigation in a given region, it could still perform adequately over other regions. If knowledge of such limitations are known a priori, at the very least the assimilation system could simply exclude the use of the retrievals over known irrigated areas. Ancillary information from other remote sensing platforms could be useful in such scenarios. For example, estimates of land surface temperature and evapotranspiration measurements can be obtained from thermal remote sensing platforms, which can also be used as an analog for detecting the effects of irrigation. Similar to the wetting effect irrigation has on surface soil moisture, it also has a cooling effect on the surface temperature. Increased water availability from irrigation also leads to increased evaporation. Finally, irrigated time periods also correlate with an increase in vegetation indices such as Leaf Area Index (LAI) and Normalized Vegetation Difference Index (NDVI). Expected trends of anomalies in these datasets from irrigation could be used as added constraints in data assimilation to mask the known limitations of passive microwave soil moisture retrievals.

It is obviously difficult to attribute bias to unmodeled processes or other factors. Calibrating the land surface model parameters can be an effective a priori correction approach in this scenario. Land surface model calibration would incorporate the observational signals

23

by altering the default model behavior. When the calibrated model is subsequently used in data assimilation, the observational signal is preserved. In contrast, calibrating a forward radiative transfer model (or using a trained ANN) has a different impact in terms of representing unmodeled processes. Calibration of forward observation operators would attribute the bias to its parameters, but since the default land model behavior is unchanged, the unmodeled process is ultimately not represented. The land surface model parameter estimation in this context also has a number of disadvantages. The physical realism of the estimated parameters may be violated given that the calibration would attribute the error from all unmodeled processes to model parameters. If data from multiple sensors are being concurrently assimilated into such a "calibrated" model, the calibration approach would not be viable because a new set of calibrated model parameters would be needed for each sensor, leading to differing model climatologies and behaviors.

Another possible alternative may be to examine the characteristics of the differences between the model and observations over a larger domain and infer a general estimate of the relative biases. For example, instead of computing CDFs at each grid point, they can be computed by grouping and stratifying model and observation estimates based on the vegetation type across a larger domain. The grid points where the soil moisture PDFs differ significantly (based on the K-S metrics, for example) can be excluded in the computation of the CDFs. The model and observation CDFs computed for each vegetation type can then be used in the data assimilation system. The goal behind such an approach would be to develop an overall estimate of the true biases between the model and observations (i.e., an estimate that excludes biases due to unmodeled processes). The downside of such approaches is that they obviously disregard the importance of geographic specificity in bias correction strategies. The synthetic experiment that we have used above is not an ideal setup to examine this approach. Since the only difference between the observations and the open loop in this setup is the effect of irrigation, the use of the above mentioned approach would produce identical CDFs for the model and observations and therefore the assimilation approach would be equivalent to that of DA-NOBC. We therefore leave the evaluation of such alternate approaches to future work.

24

References

Anagnostou, E., Negri, A., and Adler, R.: Statistical adjustment of satellite microwave monthly rainfall estimates over Amazonia, J. Appl. Meteorol., 38, 1590–1598, 1999.

Atlas, D., Rosenfield, D., and Wolff, D.: Climatologically tuned reflectivity-rain rate relations and links to area-time integrals, J. Appl. Meteorol., 29, 1120–1135, 1990.

Calheiros, R. and Zawadzki, I.: Reflectivity-rain rate relationships for radar hydrology in Brazil, J. Clim. Appl. Meteorol., 118, 118–132, 1987.

Cao, Y., Yang, X., and Zhu, X.: Retrieval snow depth by artificial neural network methodology from integrated AMSR-E and in-situ data - A case study in QInghai-Tibet plateau, Chin. Geogra. Sci., 18, 356–360, 2008.

Chakravarti, I., Laha, R., and Roy, J.: Handbook methods of applied statistics, 392–394, John Wiley and Sons, New York, 1967.

Crow, W. and Reichle, R.: Comparison of adaptive filtering techniques for land surfacce data assimilation, Water Resour. Res., 44, W08423, doi:10.1029/2008WR006883, 2008.

Crow, W., Koster, R., Reichle, R., and Sharif, H.: Relevance of time-varying and time-invariant retrieval error sources on the utility of spaceborne soil moisture products, Geophys. Res. Lett., 32, L24405, doi:10.1029/2005GL024889, 2005.

Crow, W. T. and Van Loon, E.: Impact of Incorrect Model Error Assumptions on the Sequential Assimilation of Remotely Sensed Surface Soil Moisture, J. Hydrometeorol., 7, 421–432, 2006.

De Lannoy, G., Reichle, R., Houser, P., Pauwels, V., and Verhoest, N.: Correcting for forecast bias in soil moisture assimilation with the ensemble kalman filter, Water Resour. Res., 43, W09410, doi:10.1029/2006WR005449, 2007.

De Lannoy, G., Reichle, R., and Pauwels, V.: Global calibration of GEOS-5 L-band microwave radiative transfer model over nonfrozen land using SMOS observations, Journal of Hydrometeorology, 14, 765–785, doi:10.1175/JHM-D-12-092.1, 2013.

de Rosnay, P., Polcher, J., Laval, K., and Sabre, M.: Integrated parameterization of irrigation in the land surface model ORCHIDEE. Validation over Indian Peninsula, Geophysical Research Letters, 30(19), 1986, doi:10.1029/2003GL018024, 2003.

Dee, D.: Bias and data assimilation, Q. J. Roy. Meteor. Soc., 131, 3323–3343, 2005.

Dee, D. and da Silva, A.: Data assimilation in the presence of forecast bias, Q. J. Roy. Meteor. Soc., 124, 269–295, 1998.

Douglas, E., Beltran-Przekurat, A., Niyogi, D., Pielke Sr., R., and Vorosmarty, C.: The impact of agricultural intensification and irrigation on land-atmosphere interactions and Indian monsoon precipitation – A mesoscale modeling perspective, Global Planet. Change, 67, 117–128, doi:10.1016/j.gloplacha.2008.12.007, 2009.

Draper, C., Mahfouf, J., and Walker, J.: Root zone soil moisture from the assimilation of screen-level variables and remotely sensed soil moisture, J. Geophys. Res.-Atmos., 116, D02127, doi:10.1029/2010JD013829, 2011.

Draper, C., Reichle, R., De Lannoy, G., and Liu, Q.: Assimilation of passive and active microwave soil moisture retrievals, Geophysical Research Letters, 39, doi:10.1029/2011GL050655, 2012.

Draper, C., Reichle, R., De Lannoy, G., and Scarino, B.: A Dynamic Approach to Addressing Observation-Minus-Forecast Mean Differences in a Land Surface Skin Temperature Data Assimilation System, J. Hydrometeorol., 16, 449–464, doi:10.1175/JHM-D-14-0087.1, 2015.

Drusch, M., Wood, E., and Gao, H.: Observation operators for the direct assimilation of TRMM microwave imager retrieved soil moisture, Geophys. Res. Lett., 32, L15403, doi:10.1029/2005GL023623, 2005.

Entekhabi, D., Reichle, R., Koster, R., and Crow, W.: Performance metrics for soil moisture retrievals and applications requirements, J. Hydrometeorol., 11, 832–840, doi:10.1175/2010JHM1223.1, 2010.

Ek, M., Mitchell, K., Yin, L., Rogers, P., Grunmann, P., Koren, V., Gayno, G., and Tarpley, J.: Implementation of Noah land-surface model advances in the NCEP operational mesoscale Eta model, J. Geophys. Res., 108, 8851, doi:10.1029/2002JD003296, 2003.

EORC: Status of AMSR2 Level-3 products (Algorithms Version 1.00), Japan Aerospace Research Centre and GCOM-W1 Algorithm Development, 2013.

Forman, B., Reichle, R., and Derksen, C.: Estimating passive microwave brightness temperature over snow covered land in North America using a land surface model and an artificial neural network, IEEE Trans. Geosci. Remote Sens., 52, 235–248, doi:10.1109/TGRS.2013.2237913, 2014.

Friedl, M., McIver, D., Hodges, J., Zhang, X., Muchoney, D., Strahler, A., Woodcock, C., Gopal, S., Schneider, A., Cooper, A., Baccini, A., Gao, F., and Schaaf, C.: Global land cover mapping from MODIS: algorithms and early results, Remote Sens. Environ., 83, 287–302, doi:10.1016/S0034-4257(02)00078-0, 2002.

Gelb, A.: Applied Optimal Estimation, MIT Press, Cambridge, MA, 301–365, 1974.

Gordon, L., Steffen, W., Jonsson, B., Folke, C., Falkenmark, M., and Johannesen, A.: Human modification of global water vapor flows from the land surface, P. Natl. Acad. Sci., 102, 7612–7617, doi:10.1073/pnas.0500208102, 2005.

Hain, C., Crow, W., Anderson, M., and Mecikalski, J.: An ensemble Kalman filter dual assimilation of thermal infrared and microwave satellite observations of soil moisture into the Noah land surface model, Water Resour. Res., 48, W11517, doi:10.1029/2011WR011268, 2012.

Hooke, R., Martin-Duque, J., and Pedraza, J.: Land transformation by humans: a review, GSA Today, 22, 4–10, 2012.

Jackson, T. J.: Measuring surface soil moisture using passive microwave remote sensing, Hydrol. Proc., 7, 139–152, 1993.

Kerr, Y., Waldteufel, P., Richaume, P., Davenport, I., Ferrazzoli, P., and Wigneron, J.: SMOS level 2 processor soil moisture algorithm theoretical basis document (ATBD), Tech. Rep. SO-TN-ESL-SM-GS-0001, V5. a, 15/03, SM-ESL (CBSA), CESBIO, Toulouse, 2006.

Koster, R.D., Guo, Z., Yang, R., Dirmeyer, P.A., Mitchell, K., Puma, M.J.: On the Nature of Soil Moisture in Land Surface Models, J. Climate, 22, 4322–4335, doi:10.1175/2009JCLI2832.1, 2009.

Kumar, S., Reichle, R., Peters-Lidard, C., Koster, R., Zhan, X., Crow, W., Eylander, J., and Houser, P.: A land surface data assimilation framework using the Land Information System: Description and Applications, Adv. Water Res., 31, 1419–1432, doi:10.1016/j.advwatres.2008.01.013, 2008.

Kumar, S., Reichle, R., Koster, R., Crow, W., and Peters-Lidard, C.: Role of subsurface physics in the assimilation of surface soil moisture observations, J. Hydrometeorol., 10, 1534–1547, doi:10.1175/2009JHM1134.1, 2009.

Kumar, S., Reichle, R., Harrison, K., Peters-Lidard, C., Yatheendradas, S., and Santanello, J.: A comparison of methods for a priori bias correction in soil moisture data assimilation, Water Resour. Res., 48, doi:10.1029/2010WR010261, 2012.

Kumar, S., Peters-Lidard, C., Mocko, D., Reichle, R., Liu, Y., Arsenault, K., Xia, Y., Ek, M., Riggs, G., Livneh, B., and Cosh, M.: Assimilation of remotely sensed soil moisture and snow depth retrievals for drought estimation, J. Hydrometeor., 15, 2446–2469, doi:10.1175/JHM-D-13-0132.1, 2014.

Lawston, P. M., Santanello, J. A., Zaitchik, B. F., and Rodell, M.: Impact of Irrigation Methods on Land Surface Model Spinup and Initialization of WRF Forecasts, J. Hydrometeor., 16, 1135–1154, doi:10.1175/JHM-D-14-0203.1, 2015.

Leng, G., Huang, M., Tang, Q., Sacks, W., Lei, H., and Leung, R.: Modeling the effects of irrigation on land surface fluxes and states over the conterminous United States: Sensitivity to input data and model parameters, J. Geophys. Res. Atmos., 118, 9789–9803, doi:10.1002/jgrd.50792, 2013.

Li, H., Sheffield, J., and Wood, E.: Bias correction of monthly precipitation and temperature fields from Intergovernmental Panel of Climate Change AR4 models using equidistant quantile matching, J. Geophys. Res., 115, D10101, doi:10.1029/2009JD012882, 2010.

Liu, Q., Reichle, R., Bindlish, R., Cosh, M., Crow, W., de Jeu, R., De Lannoy, G. J. M., Huffman, G., and Jackson, T.: The contributions of precipitation and soil moisture observations to the skill of soil moisture estimates in a land data assimilation system, J. Hydrometeorol., 12, 750–765, doi:10.1175/JHM-D-10-05000, 2011.

Liu, J., Zhan, X., Zhao, L., Jensen, K., and Wang, X.: NOAA Soil Moisture Operational Product System (SMOPS) and its soil moisture retrieval from AMSR-E and WindSat, in: 92nd American Meteorological Society Annual Meeting, New Orleans, LA, 2012a.

Liu, Y., Dorigo, W., Parinussa, R., de Jeu, R., Wagner, W., McCabe, M., Evans, J., and van Dijk, A.: Trend-preserving blending of passive and active microwave soil moisture retrievals, Remote Sens. Environ., 123, 280–297, 2012b.

Njoku, E. G. and Entekhabi, D.: Passive microwave remote sensing of soil moisture, J. Hydrol., 184, 101–130, 1995.

Owe, M., de Jeu, R., and Holmes, T.: Multisensor historical climatology of satellite-derived global land surface soil moisture, J. Geophys. Res., 113, F01002, doi:10.1029/2007JF000769, 2008.

Ozdogan, M. and Gutman, G.: A new methodology to map irrigated areas using multi-temporal MODIS and ancillary data: An application example in the continental U.S., Remote Sens. Environ., 112, 3520–3537, 2008.

Ozdogan, M., Rodell, M., Beudoing, H., and Toll, D.: Simulating the effects of irrigation over the United States in a land surface model based on satellite-derived agricultural data, J. Hydrometeorol., 11, 171–184, doi:10.1175/2009JHM1116.1, 2010.

Reichle, R. and Koster, R.: Bias reduction in short records of satellite soil moisture, Geophys. Res. Lett., 31, L19501, doi:10.1029/2004GL020938, 2004.

Reichle, R., McLaughlin, D., and Entekhabi, D.: Hydrologic data assimilation with the ensemble Kalman Filter, Mon. Weather Rev., 130, 103–114, 2002a.

Reichle, R., Walker, J., Koster, R., and Houser, P.: Extended versus Ensemble Kalman Filtering for Land Data Assimilation, J. Hydrometeorol., 3, 728–740, 2002b.

Reichle, R., Koster, R., Liu, P. and. Mahanama, S., Njoku, E., and Owe, M.: Comparison and assimilation of global soil moisture retrievals from the Advanced Microwave Scanning Radiometer for the Earth Observing System (AMSR-E) and the Scanning Multichannel Microwave Radiometer (SMMR), J. Geophys. Res.-Atmos., 112, D09108, doi:10.1029/2006JD008033, 2007.

Reichle, R., Crow, W., and Keppenne, C.: An adaptive ensemble Kalman filter for soil moisture data assimilation, Water Resour. Res., 44, W03423, doi:10.1029/2007WR006357, 2008.

Reichle, R., Kumar, S. V., Mahanama, S., Koster, R. D. and Liu, Q.: Assimilation of satellite-derived skin temperature observations into land surface models, J. Hydrometeorol., 11, 1103–1122, doi:10.1175/2010JHM1262.1, 2010.

Rodell, M., Houser, P. R., Jambor, U., Gottschalck, J., Mitchell, K., Meng., C.-J., Arsenault, K., Cosgrove, B., Radakovich, J., Bosilovich, M., Entin, J. K., Walker, J. P., Lohmann, D., and Toll, D.: The Global Land Data Assimilation System, B. Am. Meteor. Soc., 85, 381–394, 2004.

Salathe Jr., E., Mote, P., and Wiley, M.: Review of scenario selection and downscaling methods for the assessment of climate change impacts on hydrology in the United States pacific northwest, Int. J. Climatol., 27, 1611–1621, 2007.

Xia, Y., Mitchell, K., Ek, M., Sheffield, J., Cosgrove, B., Wood, E., Luo, L., Alonge, C., Wei, H., Meng, J., Livneh, B., Lettenmaier, D.and Koren, V., Duan, Q., K., M., Fan, Y., and Mocko, D.: Continental-scale water and energy flux analysis and validation for the North American Land Data Assimilation System Project Phase 2 (NLDAS-2), Part 1: Comparison Analysis and Application of Model Products, J. Geophys. Res.-Atmos., 117, D03109, doi:10.1029/2011JD016048, 2012.
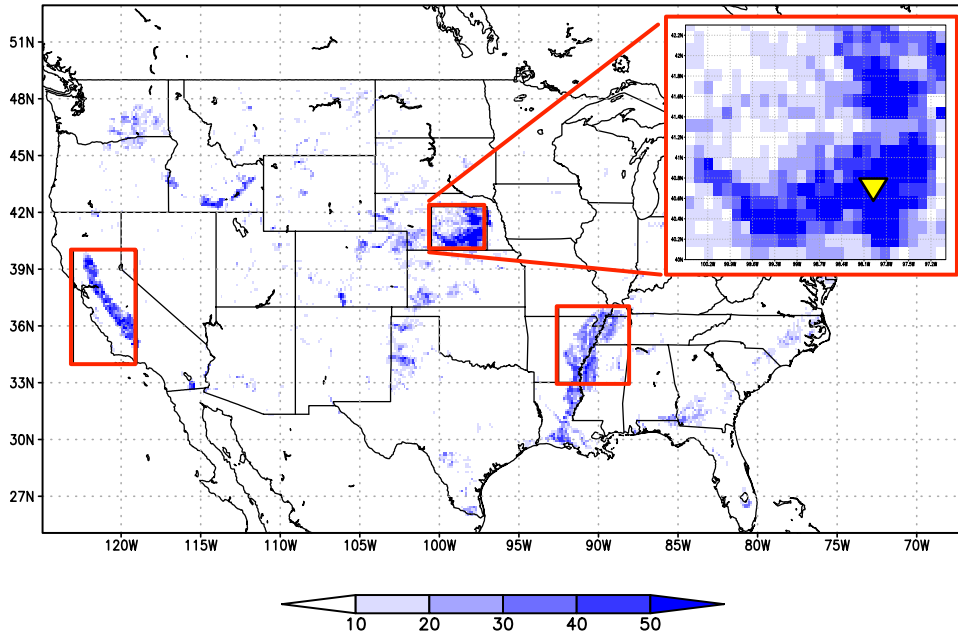
Figure 1. MODIS-based irrigated gridcell fraction (%) map of Ozdogan and Gutman (2008) over the continental U.S. The boxes (outlined in red color) highlight three known areas with large scale seasonal irrigation. The yellow triangle in the inset indicates the location of the grid cell used in point scale land surface model simulations.
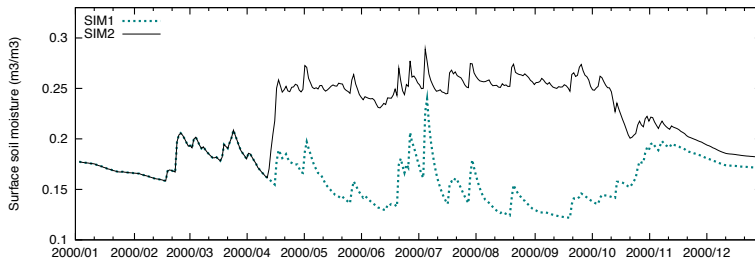
Figure 2. Surface soil moisture time series for the year 2000 from two simulations of a land surface model: (1) a free running model simulation with the NLDAS-2 forcing (SIM1) and (2) a seasonal irrigation scheme simulated on top of SIM1 (SIM2).
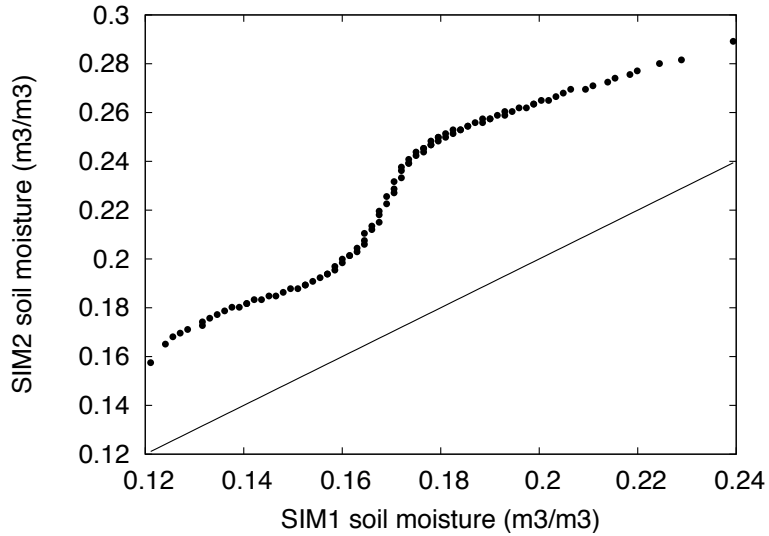
Figure 3. A comparison of the cumulative distribution functions (CDFs) from SIM1 and SIM2 integrations.
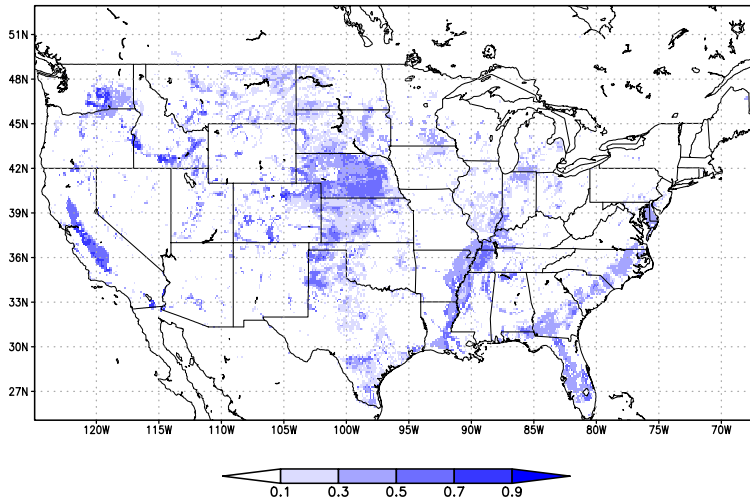
Figure 4. Kolmogorov-Smirnov distance (D) from comparison of soil moisture distributions from SIM1 and SIM2 integrations.
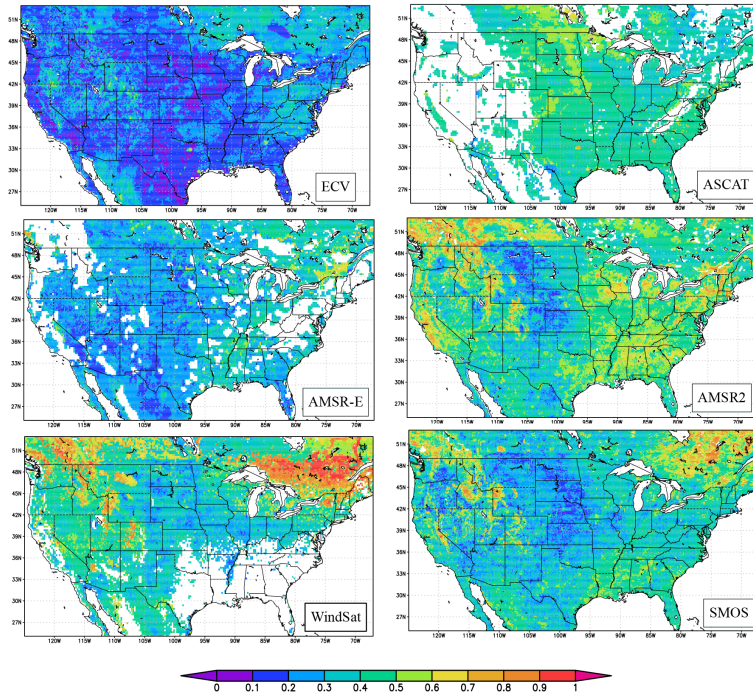
Figure 5. K-S distance (D) from the comparison of soil moisture distributions from land surface model and various satellite remote sensing datasets. Grid points with white color indicate locations that are omitted from the K-S comparisons due to insufficient, reliable data.
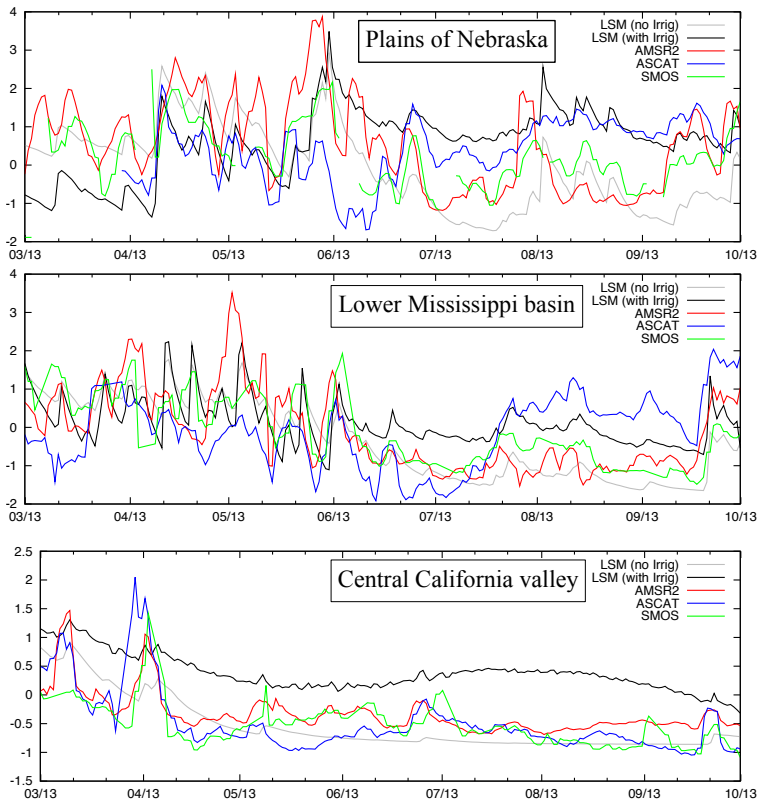
Figure 6. Comparison of normalized soil moisture values from Noah OL simulation, SMOS, AMSR2 and ASCAT retrievals, for three subregions shown in Fig. 1.
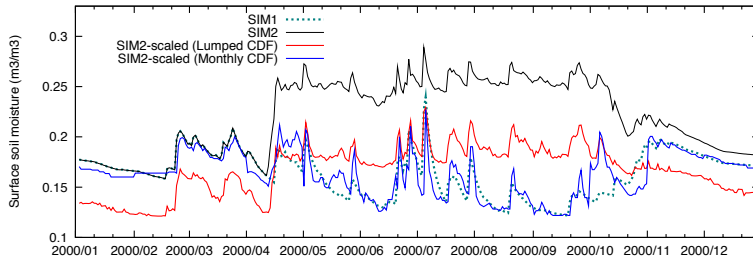
Figure 7. Same as Fig. 2 and including the surface soil moisture time series from rescaling SIM2 to the SIM1 climatology. The red and blue lines represent the SIM2 integration rescaled to the SIM1 climatology using lumped and monthly CDF-matching, respectively.
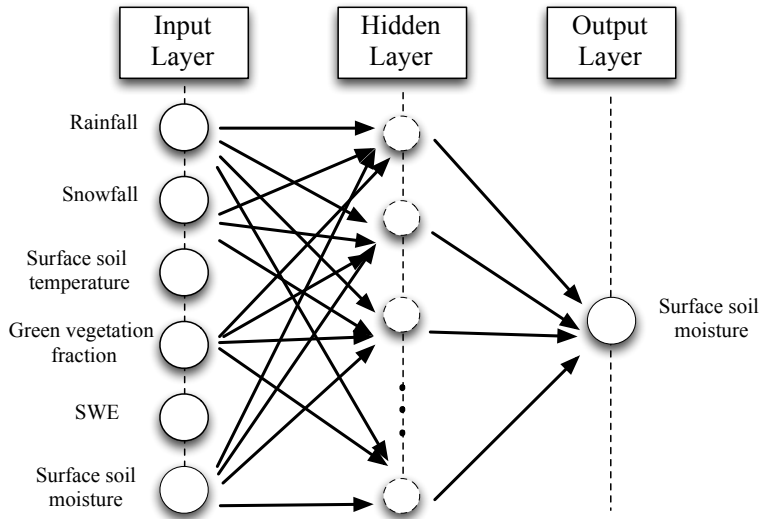
Figure 8. Structure of the artificial neural network employed in the synthetic DA integrations.
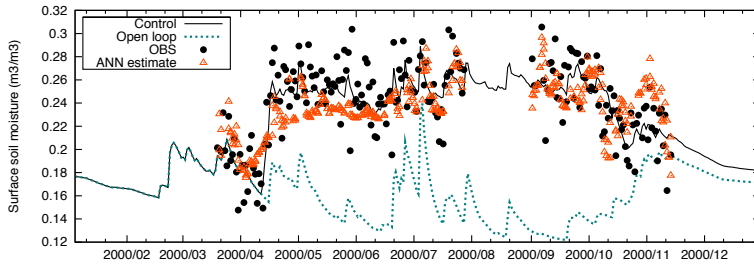
Figure 9. Time series of the Control simulation (black line), open loop (dashed line), simulated observations (filled circles) and estimates from the trained ANN (triangles) for the year 2000.
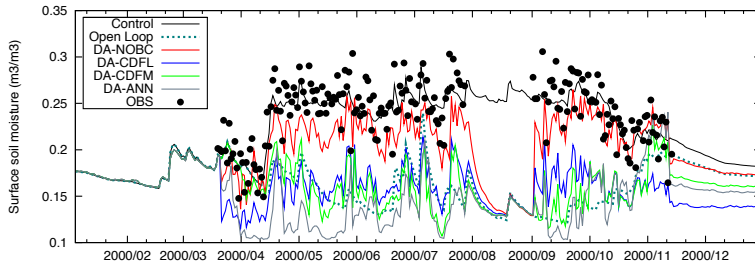
Figure 10. Time series of simulated observations, Control, open loop and various DA integrations for the year 2000. DA-NOBC assimilates observations directly without any bias correction, DA-CDFL assimilates a priori scaled observations using CDF matching (using lumped CDFs), DA-CDFM assimilates a priori scaled observations using monthly CDF matching and DA-ANN assimilates the simulated observations directly and using a trained ANN as the observation operator in the data assimilation system.
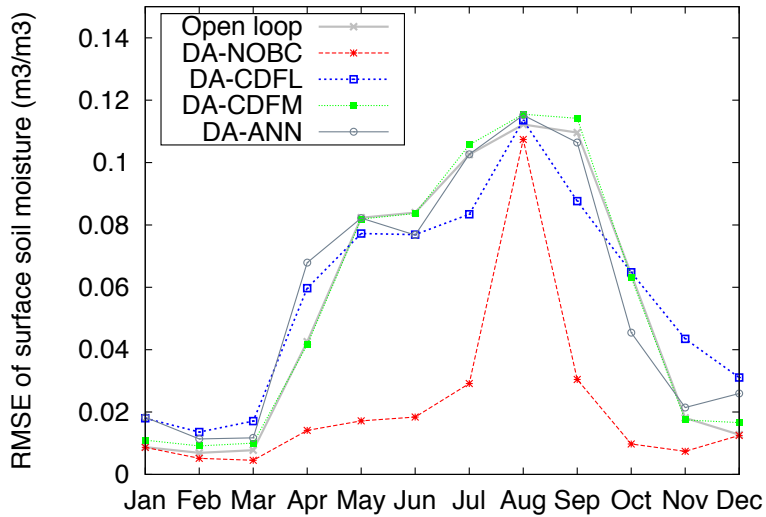
Figure 11. Average seasonal cycle of RMSE for surface soil moisture from the open loop, and various DA integrations relative to the Control simulation.
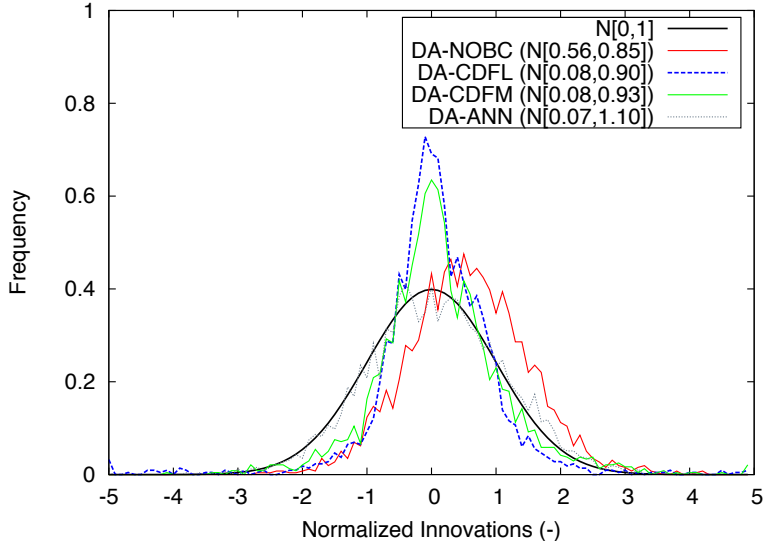
Figure 12. Distribution of normalized innovations from various DA integrations compared against the standard normal distribution.