CopERnIcus climate change Service Evolution

D1.1 Preliminary assessment of ensemble perturbation methods for the land-surface assimilation systems

Due date of deliverable	31/12/2023
Submission date	14/12/2023
File Name	CERISE-D1-1-V1.0
Work Package /Task	WP1 Task 1.1
Organisation Responsible of Deliverable	Met Norway
Author name(s)	Jostein Blyverket, Åsmund Bakketun, Peter Weston, Ewan Pinnington, Jelena Bojarova, Patricia de Rosnay
Revision number	V1.0
Status	Issued
Dissemination Level	Public

Funded by the European Union

The CERISE project (grant agreement No 101082139) is funded by the European Union.
Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.

1 Executive Summary

The purpose of this work is to develop land-surface perturbation methods in the global IFS (Integrated Forecasting System) and regional HARMONIE-AROME systems used in C3S (Copernicus Climate Change Service). Currently, the land-surface model parameters are under-spread in the ensemble systems. This means that the use of ensemble information to inform background and analysis errors in the C3S land-surface data assimilation systems is sub-optimal.

In the ECMWF IFS system we assessed the impact of Stochastic Parameter Perturbations (SPP) of leaf area index (LAI) and vegetation fraction to generate realistic representations of the surface and near surface uncertainties.

In the HARMONIE-AROME land data assimilation system we have extended the forcing perturbations to include snow and tested a new remapping procedure to represent the spatial uncertainty of the offline forcing (not only accounting for the magnitude of the errors).
Tests are also ongoing for evaluating the spread in surface temperature and soil moisture using state perturbation of the aforementioned variables. In addition to this we have tested perturbation of surface parameters. This will allow us to do an analysis of e.g. LAI in the regional surface assimilation system.

Table of Contents

2 Introduction

2.1 Background

The scope of CERISE is to enhance the quality of the C3S reanalysis and seasonal forecast portfolio, with a focus on land-atmosphere coupling.

It will support the evolution of C3S, over the project's 4 year timescale and beyond, by improving the C3S climate reanalysis and the seasonal prediction systems and products towards enhanced integrity and coherence of the C3S Earth system Essential Climate Variables.

CERISE will develop new and innovative ensemble-based coupled land-atmosphere data assimilation approaches and land-surface initialisation techniques to pave the way for the next generations of the C3S reanalysis and seasonal prediction systems.

These developments will be combined with innovative work on observation operator developments integrating Artificial Intelligence (AI) to ensure optimal data fusion fully integrated in coupled assimilation systems. They will drastically enhance the exploitation of past, current, and future Earth system observations over land-surfaces, including from the Copernicus Sentinels and from the European Space Agency (ESA) Earth Explorer missions, moving towards an all-sky and all-surface approach. For example, land observations can simultaneously improve the representation and prediction of land and atmosphere and provide additional benefits through the coupling feedback mechanisms. Using an ensemble-based approach will improve uncertainty estimates over land and lowest atmospheric levels.

By improving coupled land-atmosphere assimilation methods, land-surface evolution, and satellite data exploitation, R&I inputs from CERISE will improve the representation of long-term trends and regional extremes in the C3S reanalysis and seasonal prediction systems.

In addition, CERISE will provide the proof of concept to demonstrate the feasibility of the integration of the developed approaches in the core C3S (operational Service), with the delivery of reanalysis prototype datasets (demonstrated in pre-operational environment), and seasonal prediction demonstrator datasets (demonstrated in relevant environment).

CERISE will improve the quality and consistency of the C3S reanalysis systems and of the components of the seasonal prediction multi-system, directly addressing the evolving user needs for improved and more consistent C3S Earth system products.

2.2 Scope of this deliverable

2.2.1 Objectives of this deliverables

This deliverable is a preliminary assessment of the ensemble perturbation methods for the global and regional land-surface assimilation systems.

2.2.2 Work performed in this deliverable

In this deliverable the work outlined in WP1 T1.1: *Investigate land-surface perturbation generation methods for Land DA systems* is described and evaluated. We describe methodology for representing uncertainties in land-surface models, and in particular how to obtain a more realistic ensemble spread in land and near surface variables. land-surface

models have several error sources; (i) error of representativeness, (ii) error in model parameters/parameterization of physical processes and (iii) external errors such as atmospheric forcing, land cover and soil classification. In this report we represent errors in the forcing (by perturbing the different forcing fields) and model parameter errors (by perturbing the model parameters, which partially encompass the sub-grid scale representativeness errors).

2.2.3 Deviations and counter measures

No deviations have been encountered.

2.2.4 Reference Documents

[1] Project 101082139- CERISE-HORIZON-CL4-2021-SPACE-01 Grant Agreement

2.2.5 CERISE Project Partners:

ECMWF	European Centre for Medium-Range Weather Forecasts
Met Norway	Norwegian Meteorological Institute
SMHI	Swedish Meteorological and Hydrological Institute
MF	Météo-France
DWD	Deutscher Wetterdienst
CMCC	Euro-Mediterranean Center on Climate Change
BSC	Barcelona Supercomputing Centre
DMI	Danish Meteorological Institute
Estellus	Estellus
IPMA	Portuguese Institute for Sea and Atmosphere
NILU	Norwegian Institute for Air Research
MetO	Met Office

3 Land-surface perturbations in the HARMONIE-AROME system

In the regional modeling system we develop our unified land data assimilation system (LDAS) using a sequential method; more specifically a flavor of the ensemble Kalman filter (EnKF). Ensemble based data assimilation relies on an ensemble of model trajectories for the computation of the background error covariance matrix. This error-covariance could be climatological or flow dependent. A climatological error-covariance could be derived by computing the difference between multiple cycles over a long time-period. This is a computationally cheap method, however we lose information about errors of the day. To compute a flow dependent error-covariance we can generate an ensemble of model trajectories that either start from different initial conditions, and/or, as is the case for land-surface models, receive perturbed forcing inputs. The ensemble of land-surface model runs could then be applied to describe the background error statistics in the data assimilation. In this report we present methodology for perturbing the atmospheric forcing, state variables and static parameters (such as leaf area index and vegetation fraction).

3.1 Perturbation of atmospheric forcing

The perturbations of the atmospheric forcing follows Blyverket et al. (2019). Time-series cross-correlated forcing fields are generated using an autoregressive lag-1 model (AR(1)). These perturbations in the atmospheric forcing allow for an ensemble of model runs, where the spread represents the model uncertainty. Precipitation and shortwave radiation have a lower bound of zero; the perturbations of these variables are therefore multiplicative. Perturbations in longwave radiation are additive. To avoid bias in the forcing we centered the perturbations around zero for the additive variables and around one for the multiplicative variables. We split the perturbation methodology into two: (i) spatial and temporal perturbation and (ii) random remapping of precipitation. Method (i) only affects the magnitude of the forcing, hence we are not able to represent e.g. misplaced precipitation fields, method (ii) addresses this issue and is able to represent the misplacement of precipitation (liquid and solid).

3.1.1 Temporal and spatial perturbation of forcing

The pseudo-random perturbation field, q_t, is modeled as:

$$q_t = \alpha q_{t-1} + \beta w_t$$

where $\alpha = \frac{\Delta t}{\tau}$ and $\beta = \sqrt{(1 - \alpha^2)}$ where $\tau = 24h$ and $\Delta t = 1h$. The w_t is a normally distributed random field. To ensure physical consistency in the perturbation parameters (e.g increase in longwave radiation gives a decrease in shortwave radiation), we impose cross-correlations on the pseudo-random fields using the correlations listed in Table 1. The ensemble is kept unbiased centered at the original variable, which implies that the original forcing variable is an ensemble member and at the same time is the ensemble mean. Note that, since the land-surface model is nonlinear, this needs not to be true for the prognostic model variables.

			Cross-correlation		
Variable	**Type**	**Std dev**	**PRECIP**	**SW**	**LW**
PRECIP	Multiplicative	0.5	1	-0.8	0.5
SW	Multiplicative	0.3	-0.8	1	-0.5
LW	Additive	30 W/m2	0.5	-0.5	1

Table 1: Fields that are perturbed using the spatial-temporal perturbation methodology. Precipitation (PRECIP), shortwave downward radiation (SW) and longwave downward radiation (LW).

In Fig. 1 we illustrate the resulting spread in the offline forcing variables. The solid line is the ensemble mean while the shading is the spread of a 16 member ensemble. This figure also shows an example of the limitation of this method, as there is no spread for shortwave radiation and precipitation when the original field is zero (multiplicative perturbations). As mentioned above we have addressed this problem by developing random remapping of precipitation (see Sec. 3.1.2).

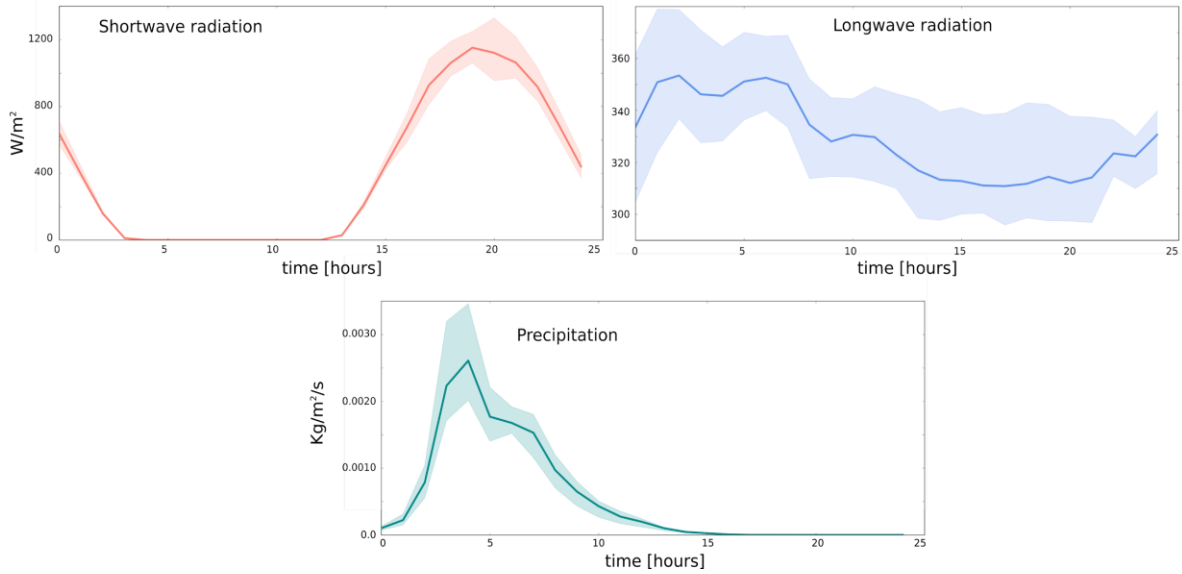

Figure 1: (Top left) Shortwave radiation, solid line is ensemble mean and ensemble spread is shading, (top right) longwave radiation and (bottom) precipitation for one day.

Figure 2 shows the time-series of the AR(1) cross-correlated perturbations for a variable (var 0-2). The blue line is without cycling of noise, which means that after a 3 hour cycle we apply a new random initial condition. To allow the temporal decorrelation to have an effect we need to make sure that we use the t-1 noise field as initial condition for our new 3 hour cycle. This results in larger values of the perturbations, as shown with the blue vs black line in the figure. Figure 2 also illustrates how the cross-correlation results in positive perturbations for var 1 and negative perturbations in var 2.

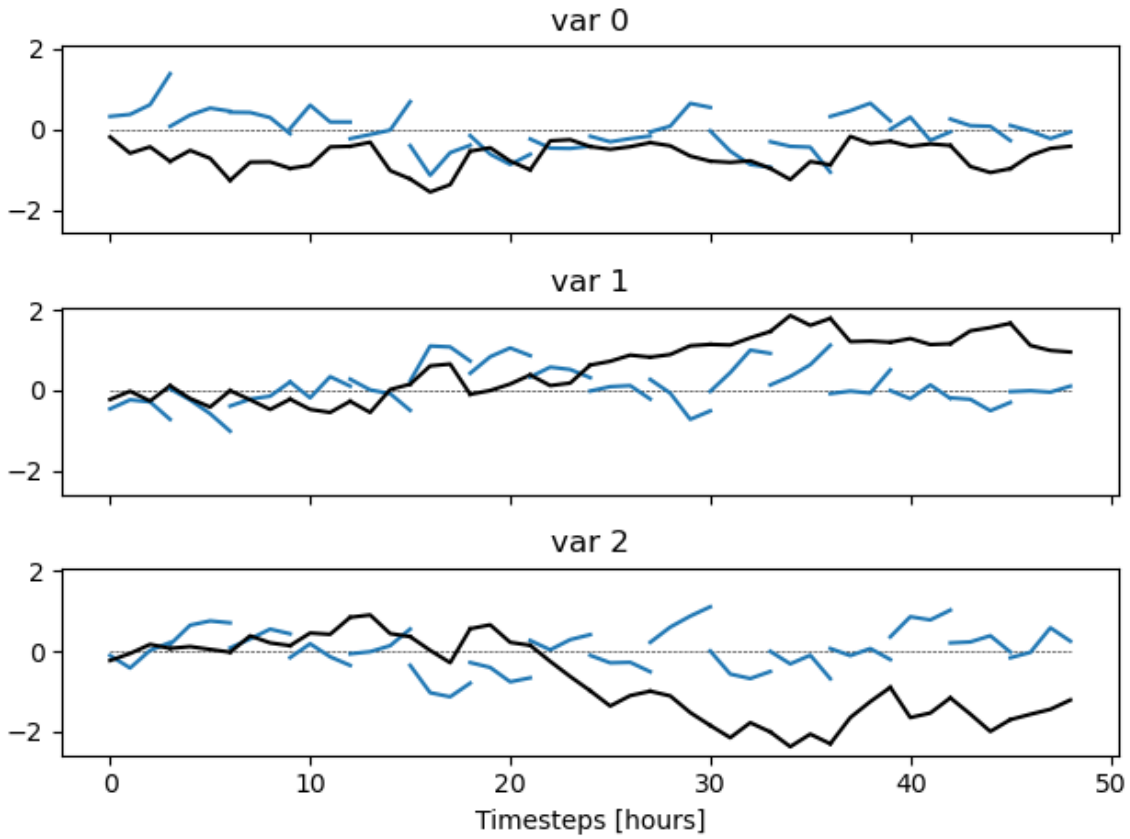

Figure 2: Time-series of AR(1) noise with correlation between variables (var 0-2). Blue line is noise without cycling of noise. Black line is with cycling.

Forcing perturbations are now correlated in time and between variables, the next step is to get spatially correlated noise. To impose spatial correlation on the noise we apply convolution with a Gaussian Kernel function, which has the form:

$$q_t{}' = \frac{1}{s} q_t \star g$$

where g is the Gaussian Kernel function and s is given by:

$$s = \int_k g \, dk.$$

Figure 3 shows the gridded noise before (left) and after (right) we apply the Gaussian Kernel convolution. In this particular case it is illustrated with a correlation length scale of 50 grid points.

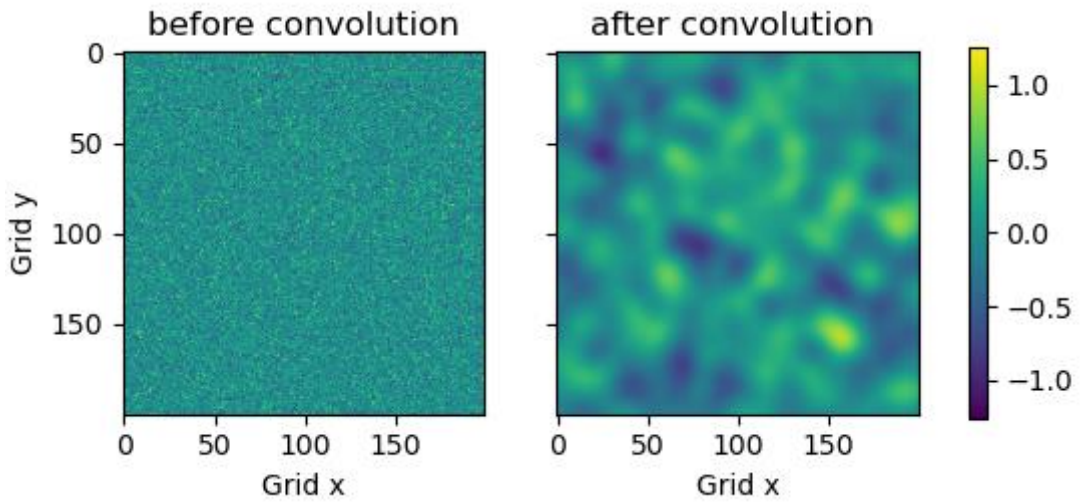

Figure 3: (Left) Gridded noise before Gaussian Kernel convolution, (right) gridded noise after convolution. Length scale is 50 grid points.

3.1.2 Random remapping of solid precipitation

As described in Sec. 3.1.1, we use a multiplicative method to perturb the precipitation forcing, which implies that no spread will be generated over areas with zero precipitation. To account for this limitation we implemented a random remapping of the forcing data. First, a random vector field is generated with some spatial correlation as in Sec. 3.1.1. Second, the vectors are used to advect the indices of the grid and the precipitation field is then interpolated according to the new indices. The remapping is performed on the total precipitation and a redistribution between rain and snow is done before the variables are outputted. The redistribution of phase is based on the ratio $r = snow/total$ or if no precipitation existed initially, a function based on 2m temperature is used to compute $snow = r \times total$, $rain = (1 - r) \times total$.

Figure 4 illustrates the remapping algorithm. Starting from a synthetic field in Fig. 4 a), we apply the remapping on this field and end up with the field in Fig. 4 b). The resulting difference field is depicted in Fig 4. c).

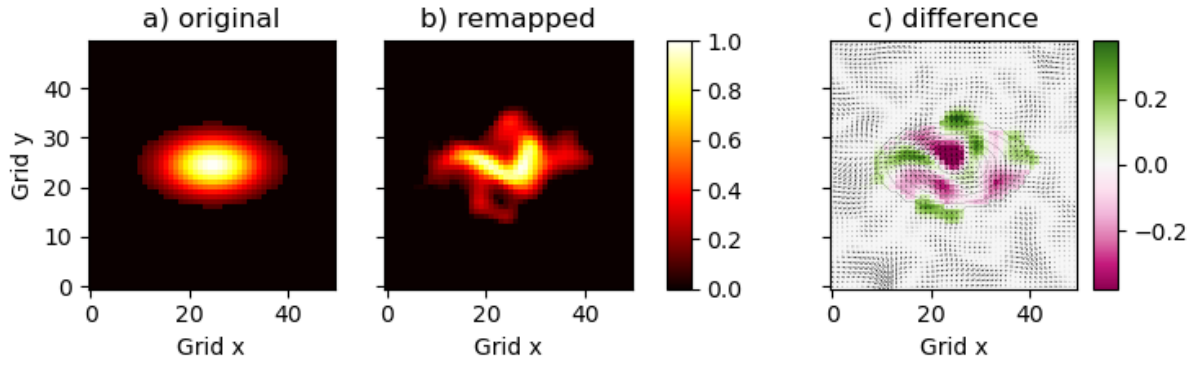

Figure 4: (a) Original synthetic field, (b) after applying the remapping and (c) the difference field (shading) and random vector field (arrows).

In Fig. 5 a) we show the forcing perturbation strategy on a precipitation field. First we apply spatially correlated multiplicative noise (b), before we remap this field (c). The figure clearly

illustrates how the remapping moves the precipitation field.

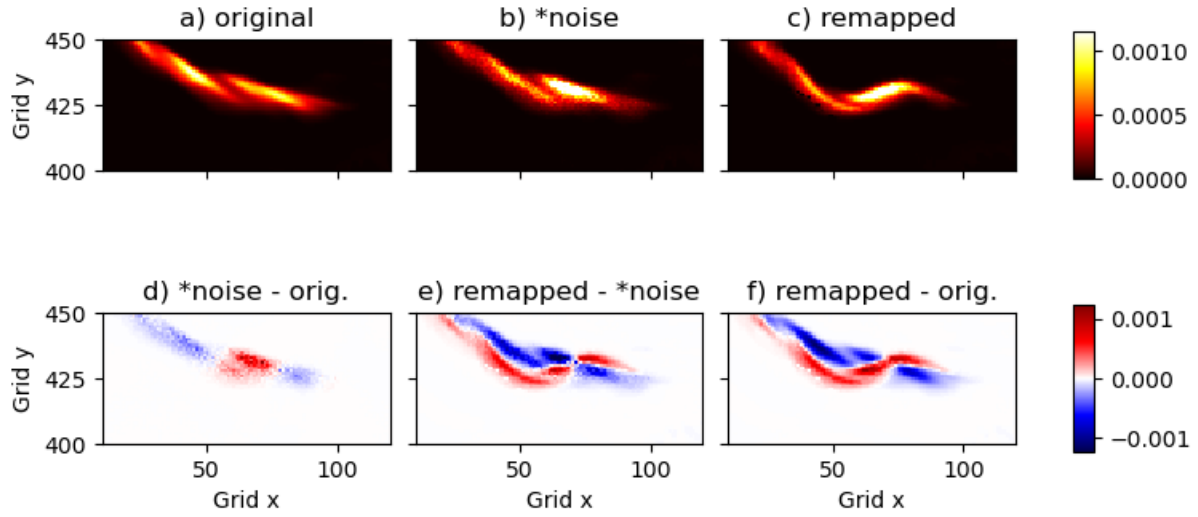

Figure 5: (a) Real case field, (b) with multiplicative perturbations, (c) multiplicative + remapped field. Difference fields are shown in (d) multiplicative perturbation minus original field, (e) remapped minus multiplicative noise and (f) final multiplicative + remapped minus original field.

3.2 Perturbation of model state and parameters

3.2.1 Soil moisture and soil temperature perturbations

For the regional setup, soil moisture and temperature perturbations are applied using the HARMONIE-AROME ensemble prediction system (EPS) code for surface perturbations. This code was originally developed for the 3 layer ISBA Force-Restore (ISBA-FR) surface scheme. Therefore we implemented an option to perturb more layers, as the ISBA-DF surface scheme has 14 soil layers. We added the option to perturb the top 3 layers (originally top 2 in ISBA-FR). The soil moisture perturbations are multiplicative while the soil temperature perturbations are additive. We apply a spatial correlation length of 150 km.

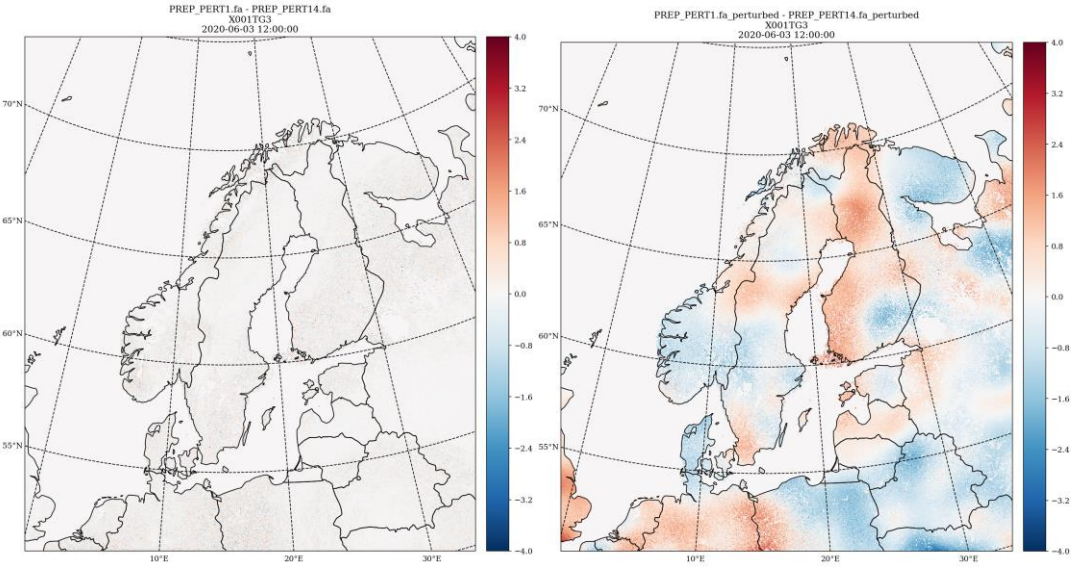

Figure 6: Illustration of soil temperature differences (K) in layer 3 without (left) and with (right) state perturbations of soil temperature.

Figure 6 shows the difference plot between two perturbed start files for the offline ensemble. We see how the state perturbation of soil temperature (right) increases the spread compared to forcing perturbations only (left). The noisy pattern in the figure to the left is because the experiment does not have any spatial correlation in the forcing perturbation.

3.2.2 Perturbation of surface parameters

LAI quantifies the amount of leaf material in a canopy and is measured as a ratio of a one-sided leaf area of the canopy to the ground surface area. In land-surface models LAI plays a crucial role in parameterizing impact from the biogeochemical and biophysical processes and modeling energy absorption and evapotranspiration. At the same time several sources of uncertainty are associated with the estimate of LAI. We have implemented a method for perturbing LAI in a physically-consistent way. The standard LAI perturbation technique applies a constant amplitude perturbation to the whole LAI field. As a result, an inconsistent LAI field might be obtained with such a perturbation technique. In certain occasions LAI might be unrealistically reduced over some forested areas, or LAI can be unrealistically increased over some bare soil areas. Such spurious LAI realizations lead to an unphysical response and degrade performance of the ensemble-based land-surface scheme.

We implement a robust LAI perturbation method where LAI is perturbed independently at each grid-box based on the statistical information specific for each grid-box. The rationale behind the method is the observation that variability of the LAI in a certain area depends on the temporal change rate in this area and that LAI changes rapidly only during some periods of the year (in spring and in fall) and is slowly changing during the rest of the year.

3.3 Resulting spread in surface and near surface variables

3.3.1 Soil temperature, moisture and snow

Here we present a preliminary assessment of the regional system ensemble perturbations in three set of experiments, (i) forcing perturbations to the offline runs; providing the ensemble spread in the HARMONIE-AROME EnKF surface analysis (Forcing pert.), (ii) an EPS run providing the surface ensemble to the surface analysis using IFSENS as the only perturbation method (EPS bdPert), and (iii) offline snow experiment using perturbed forcing and random remapping of precipitation. The latter being the experiments for developing the local ensemble transform Kalman filter (LETKF) for snow data assimilation in the regional offline system.

Figure 7 shows error estimates for layer 1 soil temperature for 00 and 12 UTC. We show results for the experiment perturbing forcing only (Forcing pert, blue) and the EPS experiment using IFSENS on the boundaries (EPS bdPert, orange). The target error values are taken from Holmes et al (2012). It is evident that only perturbing the forcing or applying bdPert in EPS mode underestimates the errors.

For soil moisture there is no truth with which to evaluate the ensemble spread (Draper 2021). An alternative error evaluation method is triple collocation, but that is outside the scope of this report. Error numbers reported in the literature are often around 0.05-0.06 m3/m3 in unbiased root-mean-square error (Reichle et al., 2017). When evaluating the same experiments as we did for soil temperature, we find a domain mean of 0.01 m3/m3 for the EPS bdPert experiment,

and 0.001 m3/m3 for the forcing perturbation experiment. The EPS bdPert experiment number is around the same value as found in Draper (2021). One thing to note is that this is only after 3 days and the forcing perturbation experiment uses a much larger domain than the EPS experiment. This would most likely result in large areas without any substantial precipitation, which is the main driver for soil moisture spread using this methodology.

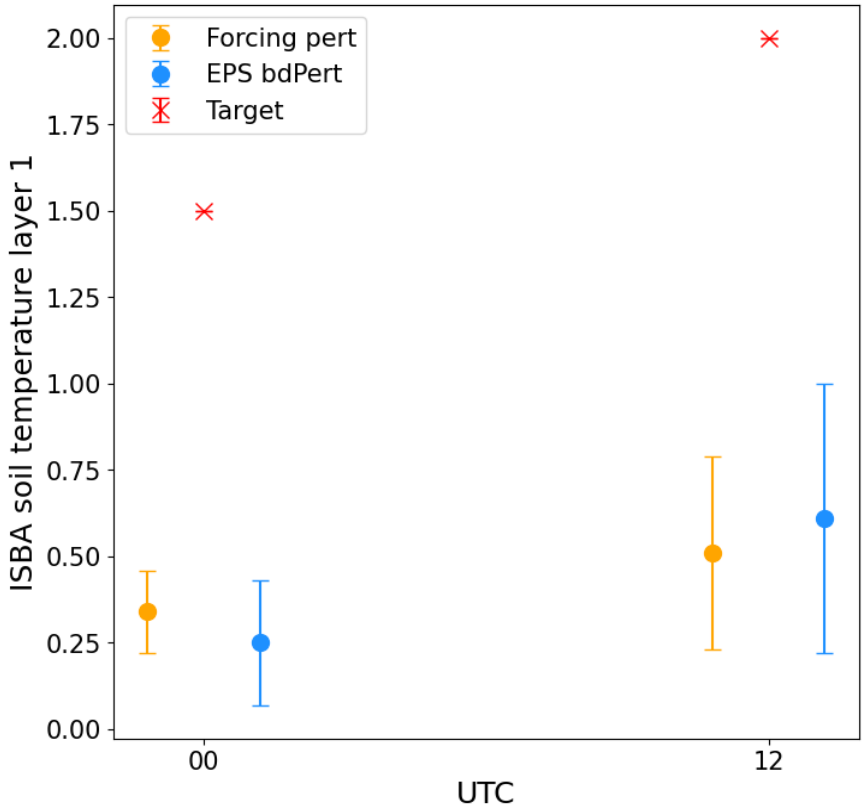

Figure 7: Domain mean ensemble spread for offline perturbation of forcing only (Forcing pert, blue) and ensemble mode (EPS bdPert, orange) using only different boundary conditions (IFSENS).

Exp	Forcing	Members	Assimilated obs.
ctrl	MEPS	1	-
openloop	MEPS	15	-
ref	Nordic analysis	1	-
daexp	MEPS (same as openloop)	15	ref (snow depth)

Table 2: Experimental setup for assessment of ensemble spread of snow properties. MEPS - MetCoOp ensemble prediction system with 2 km resolution, Nordic analysis is a post-processed 1 km product utilizing synop observations, citizen observations and radar to correct temperature, humidity and precipitation.

The impact of perturbed forcing on surface snow is evaluated through offline experiments using the 12 layer ISBA Explicit Snow (ISBA-ES) scheme. Four experiments covering one full year from 1st July 2021 to 1st July 2022 are set up as in Table 2. The domain covers northern

Scandinavia and includes 83 in situ snow depth stations. The snow depth ensemble is evaluated in these points.

A spread-skill time series plot (Fig. 8) indicates that the ensemble is underdispersive relative to the standard errors (orange dashed smaller than the orange solid line). We note that the increase in standard error of the openloop relative to ctrl is consistent with the ensemble spread. For the assimilation experiment (green lines) both the errors and ensemble spread are reduced and they are around the same size.

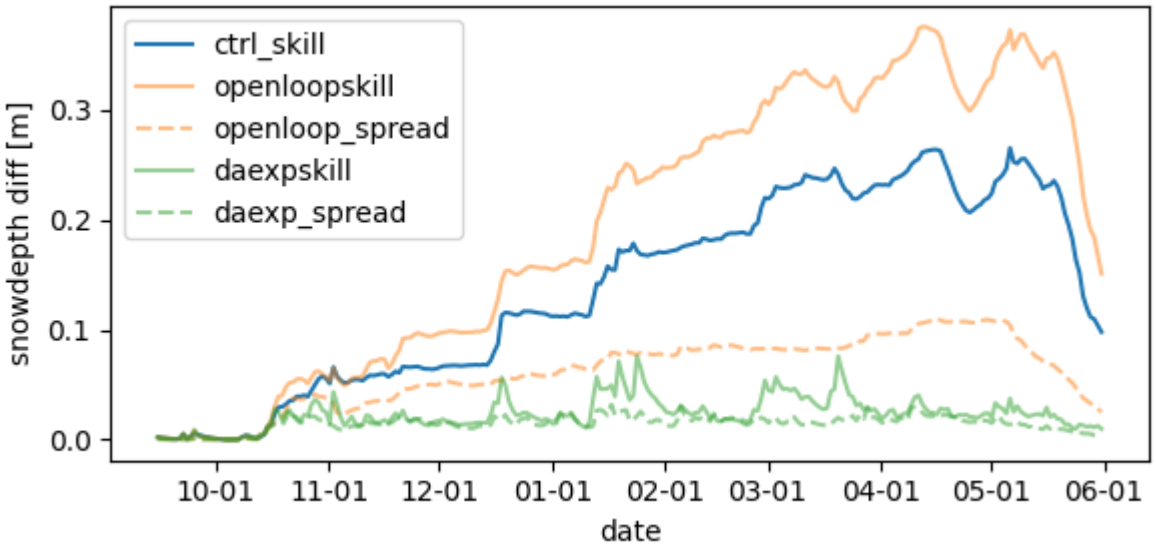

Figure 8: Ensemble spread (standard deviation of ensemble) and skill (error standard deviation) of ctrl (blue), openloop (orange), and daexp (green) relative to ref.

To evaluate the spatial properties of the surface snow ensemble, we show the ensemble covariance between observed snow depth and a control variable (snow water equivalent - SWE) in layer 6, with and without localization (see Fig. 9). Positive values indicate that positive innovations (model has too little snow) give positive increments (analysis adds SWE). We expect the ensemble covariances to be consistent (same sign) close to the observation. However, topography and variation of surface type can result in opposite signs. When no localization is applied, it is difficult to distinguish between real and spurious correlations; this might indicate that our ensemble is too small.

In Fig. 9 we have utilized the partial analysis increment (PAI), which can be derived from the LETKF (Diefenbach et al., 2023). The PAI is an approximation of the influence (potential increment) of a single station without the need for single observation experiments.

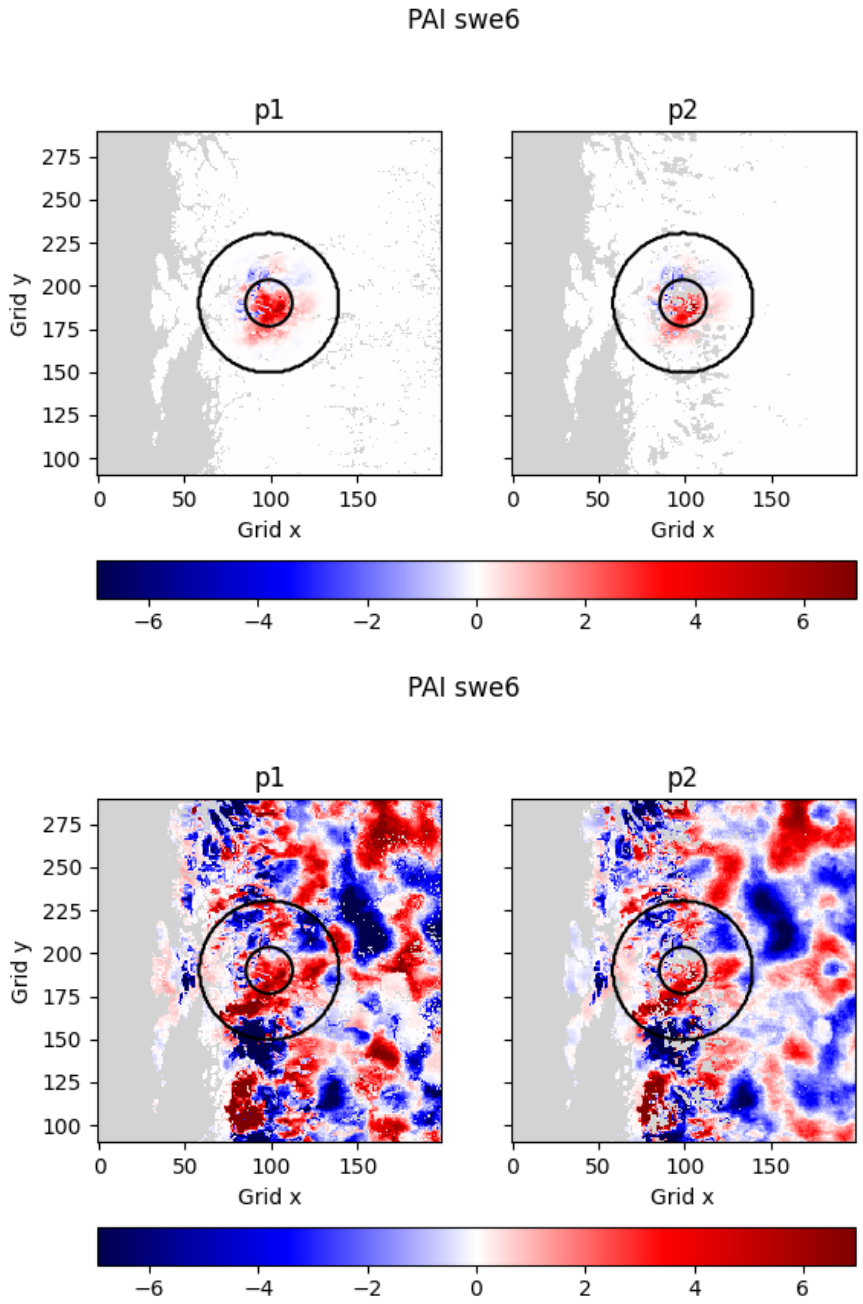

Figure 9: Structure function for one assimilation cycle for the snow water equivalent - SWE layer 6 control variable relative to snow depth observation, with (top) and without (bottom) localization. (Left) is for patch 1, open land, and (right) is for patch 2, vegetation. PAI - partial analysis increment is a bi-product of the local ensemble transform Kalman filter (LETKF).

3.3.2 Coupling between surface and near surface variables

In Fig. 10 we plot ensemble correlation between soil moisture layer 1 and 2m specific humidity for the EPS bdPert and Forcing pert. (offline) experiment and bin it by the soil wetness index (SWI) in the respective grid cell. The rationale behind these plots is to both evaluate our modeling system (that the connection between variables are what we expect), but also to see what kind of EnKF increments we could expect in both offline and coupled experiments.

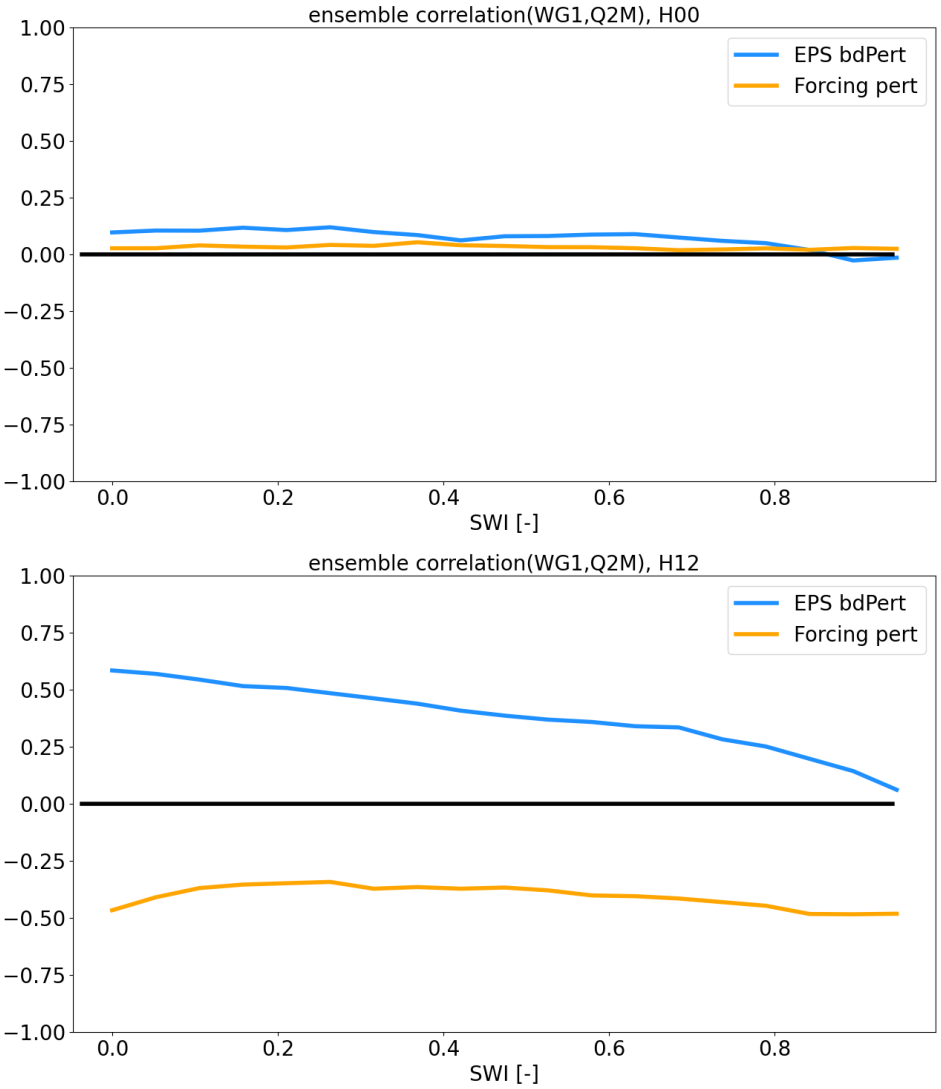

Figure 10: Ensemble correlation between surface soil moisture (WG1) for patch 1 and 2m specific humidity for patch 1 at 00 (top) and 12 (bottom). Blue is EPS mode where we use IFSENS boundaries while orange is for forcing perturbation of the offline runs.

Figure 10 (top) shows that the correlation at 00 UTC for the EPS bdPert is positive (but small) and in accordance with Draper (2021). This means that if we add soil moisture in the analysis we will get a positive 2m specific humidity increment and vice versa. For wet soil the correlation is close to zero, which most likely indicates that it is energy limited.

The 2m humidity and soil moisture is much more correlated at 12 UTC, Fig. 10 (bottom). Here we see that the EPS bdPert (blue) has a positive correlation, again decreasing to zero for wet soil (large SWI). For not too wet soils it means that if the analysis increment adds soil moisture it will increase the atmospheric humidity, which again seems reasonable.

In Forcing pert. we see a negative correlation, which means that a positive increment in specific humidity translates to a negative increment in soil moisture. If our first guess specific humidity is too dry (O-minus-F > 0) we will compensate by removing moisture from the soil. A negative soil moisture increment could then result in further drying in the subsequent coupled run, as there is less moisture available for evapotranspiration. Note that this is for layer 1 only

(top centimeter of the soil) and a further evaluation with deeper layers would need to be done in future.

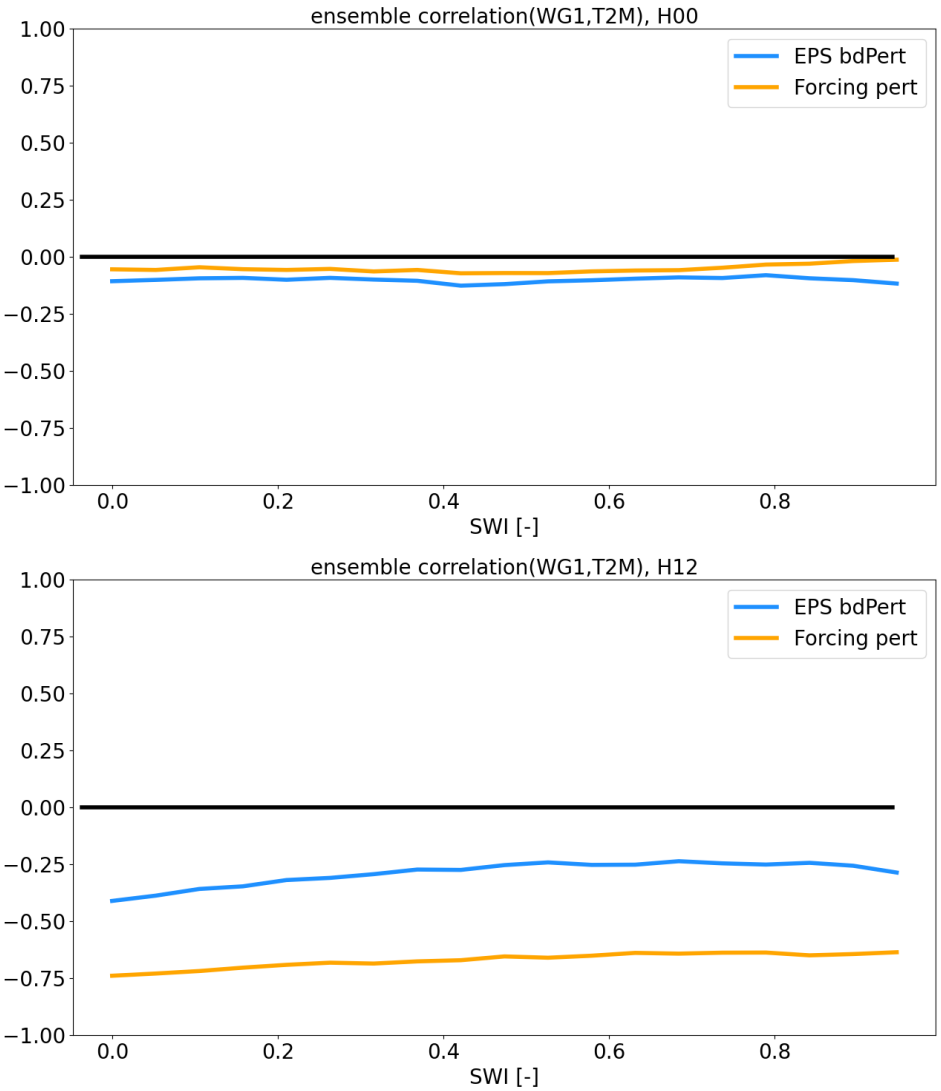

Figure 11: Ensemble correlation between surface soil moisture (WG1) for patch 1 and 2m temperature for patch 1 at 00 (top) and 12 (bottom). Blue is EPS mode where we use IFSENS boundaries while orange is for forcing perturbation of the offline runs.

For correlation between soil moisture and 2m temperature (Fig: 11), we see that for both the offline and online run the correlation is negative and largest at 12 UTC. This is reasonable as an increment in one of the variables results in an increment of the opposite sign in the second variable. The size of the correlation is around the same magnitude as found in Draper (2021).

4 Land-Surface Perturbations in the IFS System

4.1 Overview

The core land-surface data assimilation scheme at ECMWF is the Simplified Extended Kalman FIlter (SEKF) (de Rosnay et al. 2013). This scheme already uses ensemble information from the atmospheric Ensemble of Data Assimilations (EDA) (Bonavita et al. 2012) to diagnose the linearised observation operator, implemented in June 2019 as part of CY46R1, in place of the more traditional finite-difference method. The atmospheric EDA provides us with a great starting point to incorporate more ensemble information into the LDAS via background errors and spread in land-surface parameters. However, currently many land-surface parameters within the EDA are underdispersed. This can potentially lead to issues of overconfidence and ensemble collapse when moving towards increased utilisation of the EDA for the land-surface. The lack of spread in the surface fields also hampers our ability to apply the correct updates to fields such as 2m temperature and relative humidity. In order to remedy the issue of an under-spread ensemble at the surface we have explored perturbation schemes for land-surface parameters. It has been shown that land-surface variables of soil moisture and 2m temperature/relative humidity are sensitive to the parameters of leaf area index and vegetation fraction in work by Draper et al. (2021). They show that perturbing the parameters of vegetation fraction and leaf area index across an ensemble can help to significantly increase spread in a number of land-surface variables. We have set up an ensemble of ECLand (Boussetta et al. 2021) offline model runs using the different realisations of meteorological forcing from the EDA. These model runs give us a good benchmark for how much additional ensemble spread we can expect to achieve from land-surface parameter perturbations over the current spread with the different realisations of meteorological forcing.

4.2 Methodology

In initial results with the ECLand offline model (Boussetta et al. 2021) we found good increases in spread (10-40%) by varying the parameters of vegetation fraction and leaf area index with a simplistic additive Gaussian noise function, when compared to an ensemble with no such parameter perturbations. However, this simple function also led to some unwanted spatial artifacts due to the application of a blanket correction (completely spatially and temporally consistent) for any given ensemble member. It was therefore decided to move towards a methodology closer inline with the Stochastic Perturbed Parameter (SPP) scheme already implemented at ECMWF for certain atmospheric fields (Lang et al. 2021). This employs a Gaussian correlation function in time and space to generate fields of random noise with which to perturb the desired parameters. We chose a length scale of 3 months in time (consistent with leaf area index evolution) and have tested 3 differing spatial length scales as shown in Figure 12. Table 3 shows the different offline ECLand model ensemble experiment specifications (e.g. combinations of parameters and length scales).

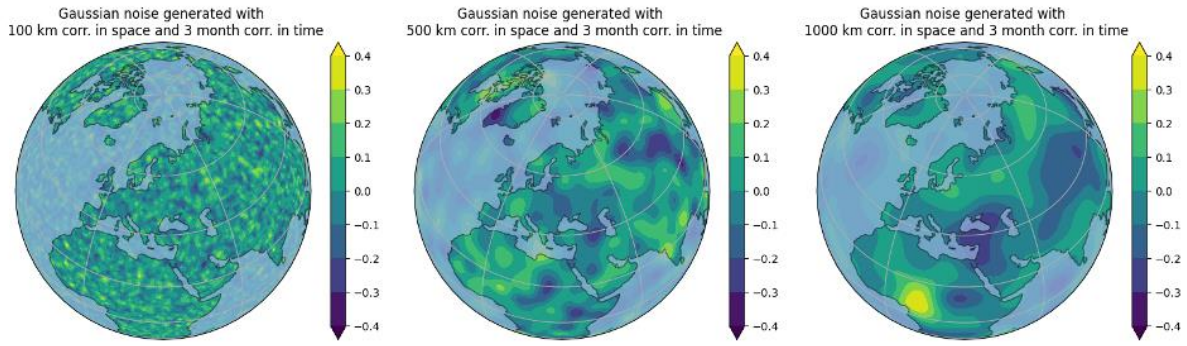

Figure 12: Spatial representation of different random fields generated with Gaussian noise of varying length scales (100, 500, 1000 km).

Experiment	Land parameters perturbed		Length scale (km)		
	LAI	Veg Frac	100	500	1000
EDA (control)	-	-	-	-	-
spp_lai_500km	x			x	
spp_cvt_500km		x		x	
spp_laicvt_100km	x	x	x		
spp_laicvt_500km	x	x		x	
spp_laicvt_1000km	x	x			x

Table 3: Table outlining the experiments conducted with the ensemble of ECLand offline simulations.

In Figure 13 we show what these perturbations and resulting model prognostic variables might look like at a single grid point; here we can see an increase in spread over the ensemble with just perturbed meteorological forcing from the EDA.

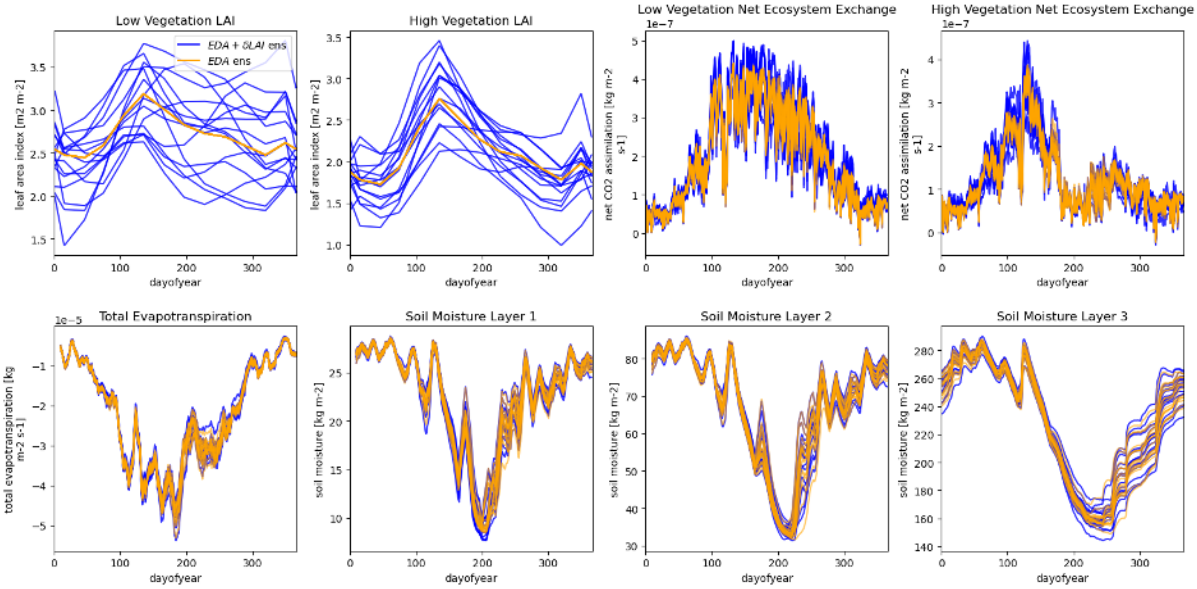

Figure 13: Representation of what the parameter perturbations and resulting model prognostic variables look like at a single grid point for a year-long model run. Orange shows the ECLand ensemble

run with only meteorological forcing perturbations from the EDA, Blue shows the ensemble when we add land-surface parameter perturbations (here for Leaf Area Index).

4.3 Results

In Figure 14 to 19 we show the results of the experiments outlined in Table 3, each Figure displays the spread from the control in the top left and then how the spread increases/decreases relative to this control in the subsequent 5 panels. For soil moisture, soil temperature, skin temperature and snow depth we find larger contributions to the increase in spread from perturbing vegetation fraction over leaf area index. This is due largely to changes in the bare soil fraction leading to an update in the model fluxes. In contrast, for net CO2 flux and evapotranspiration we find the largest contributions from perturbing leaf area index, which makes physical sense as leaf area index is one of the main controls over photosynthesis and transpiration. For the majority of variables we find slightly noisy results if we perturb a single parameter (LAI or vegetation fraction) but achieve a regularising effect by perturbing both LAI and vegetation fraction together.

Overall the difference in spread between the tested length scales is minor. For all variables we are finding positive increases in spread when perturbing both parameters over the control with just perturbed meteorology from the EDA. This is particularly encouraging for snow depth which is difficult to directly perturb in a physically consistent way. The largest increases in spread are found for net CO2 flux, this is expected as LAI is the main constraint on CO2 flux and the control has the same LAI profiles across all ensemble members and thus a very small spread in CO2 flux.

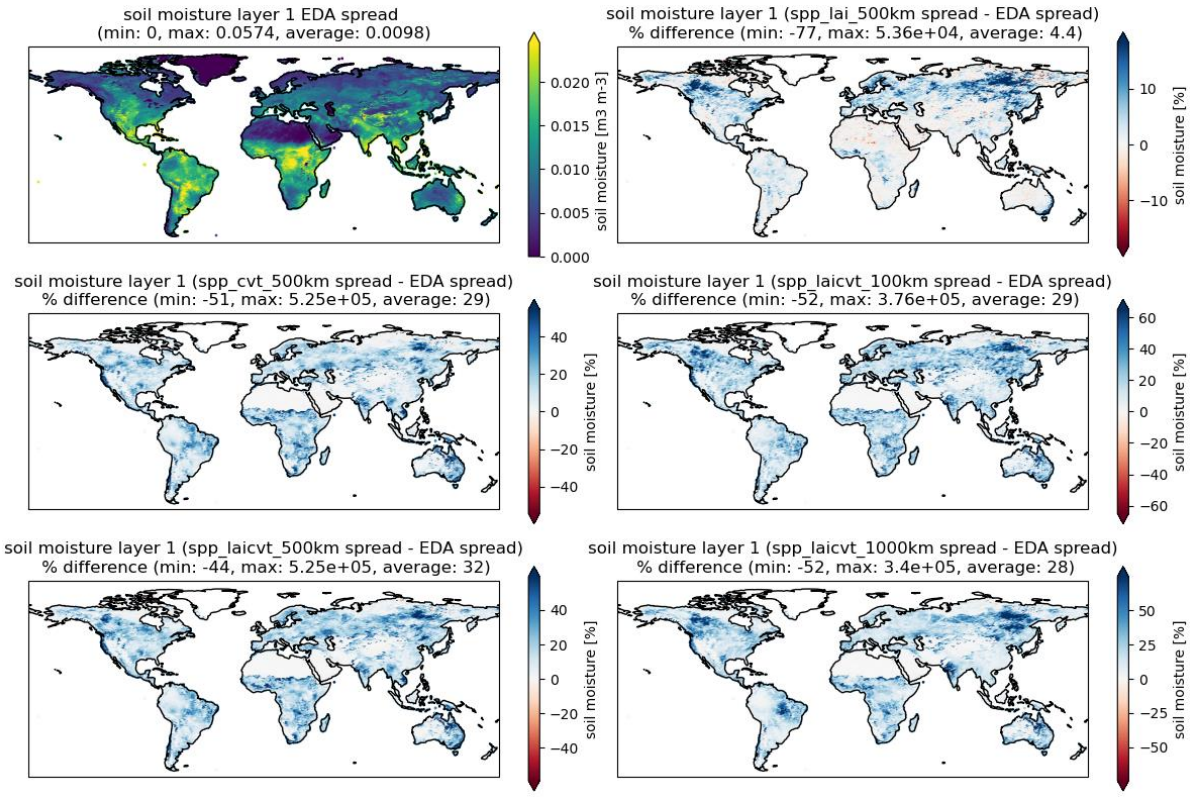

Figure 14: Figure showing the EDA spread for top layer soil moisture (top left) and the increase/decrease in spread for the different experiments outlined in Table 4.1 from the EDA benchmark.

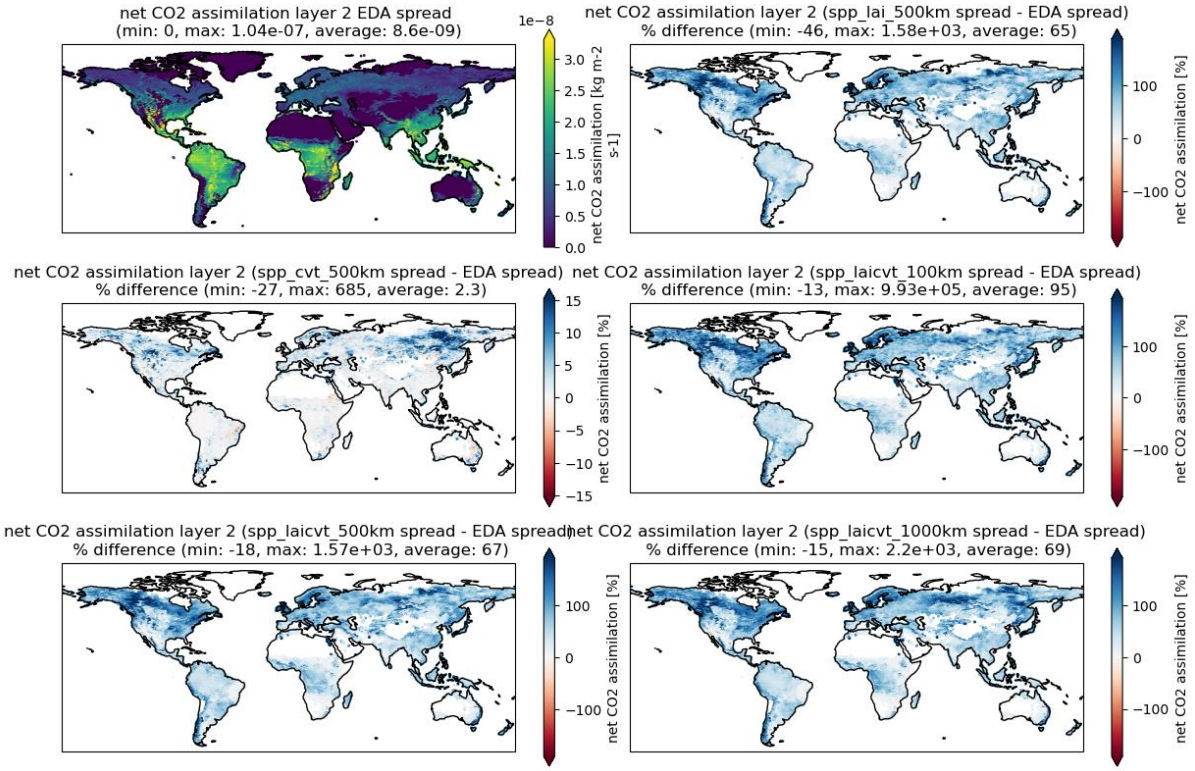

Figure 15: Figure showing the EDA spread for Net CO2 flux (top left) and the increase/decrease in spread for the different experiments outlined in Table 3 from the EDA benchmark.

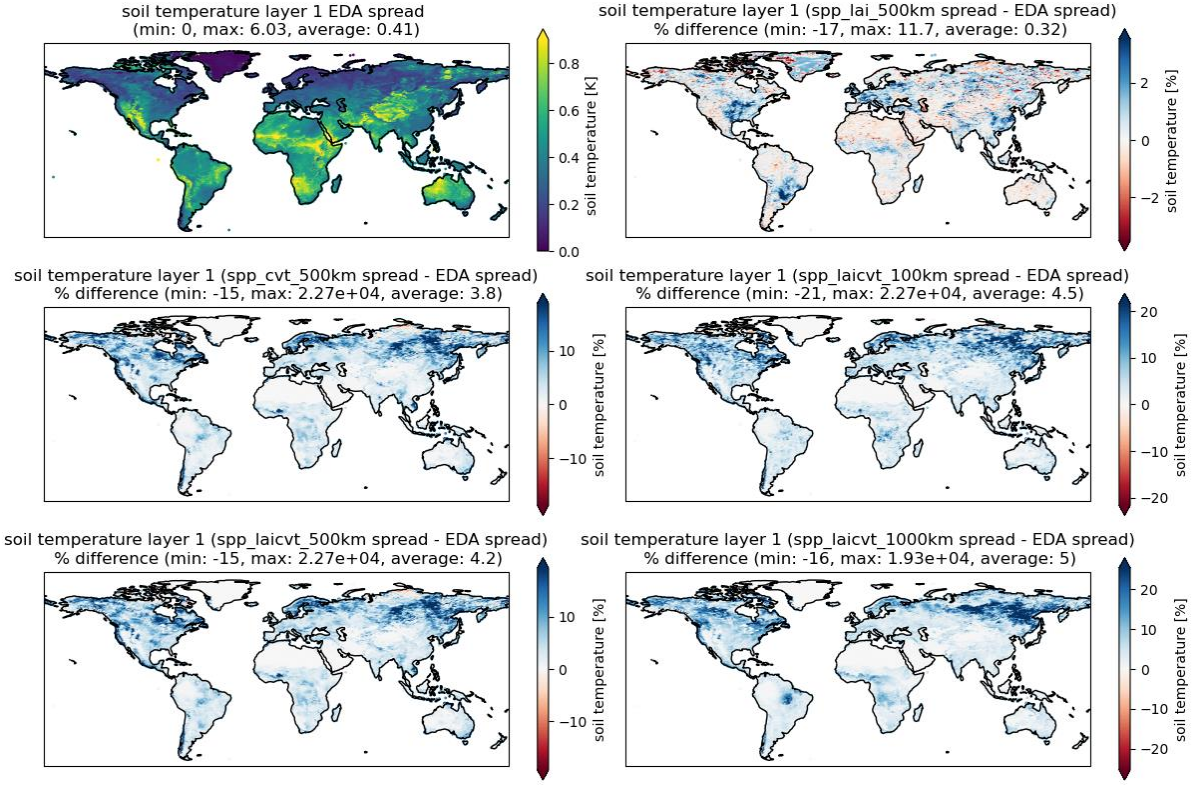

Figure 16: Figure showing the EDA spread for top layer soil temperature (top left) and the increase/decrease in spread for the different experiments outlined in Table 4.1 from the EDA benchmark.

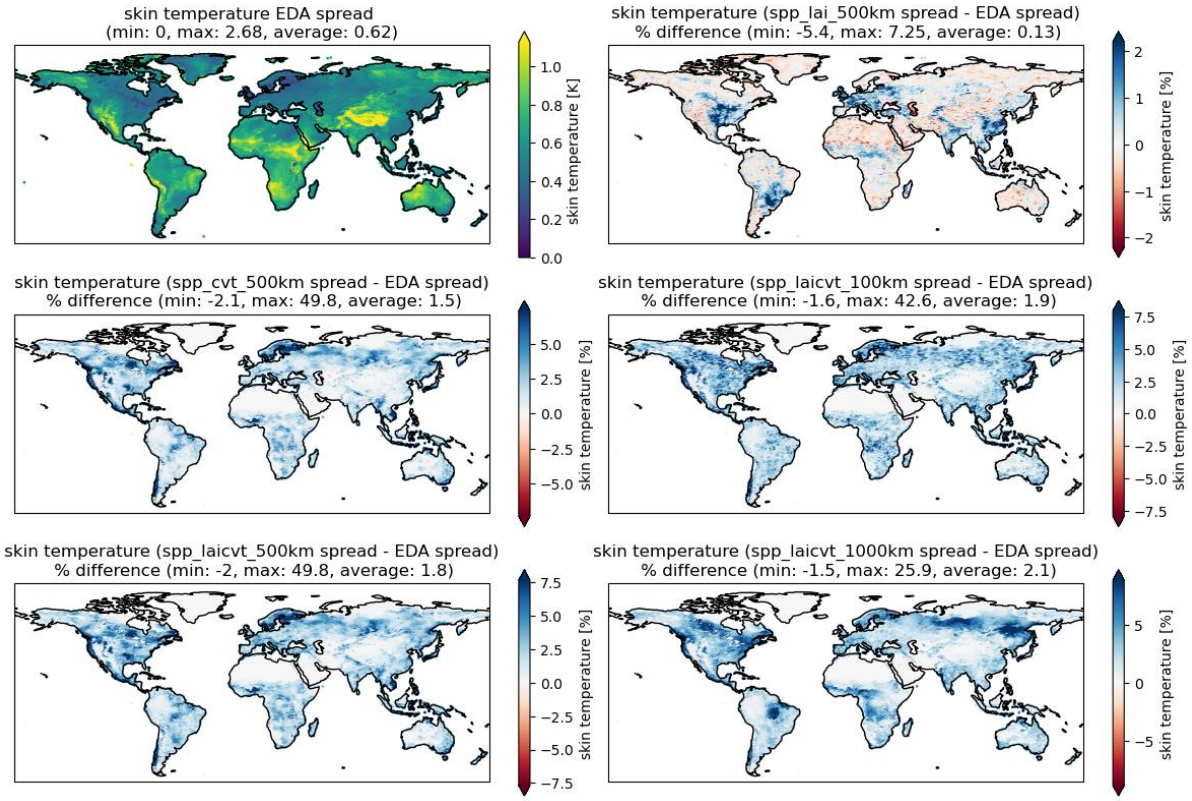

Figure 17: Figure showing the EDA spread for skin temperature (top left) and the increase/decrease in spread for the different experiments outlined in Table 3 from the EDA benchmark.

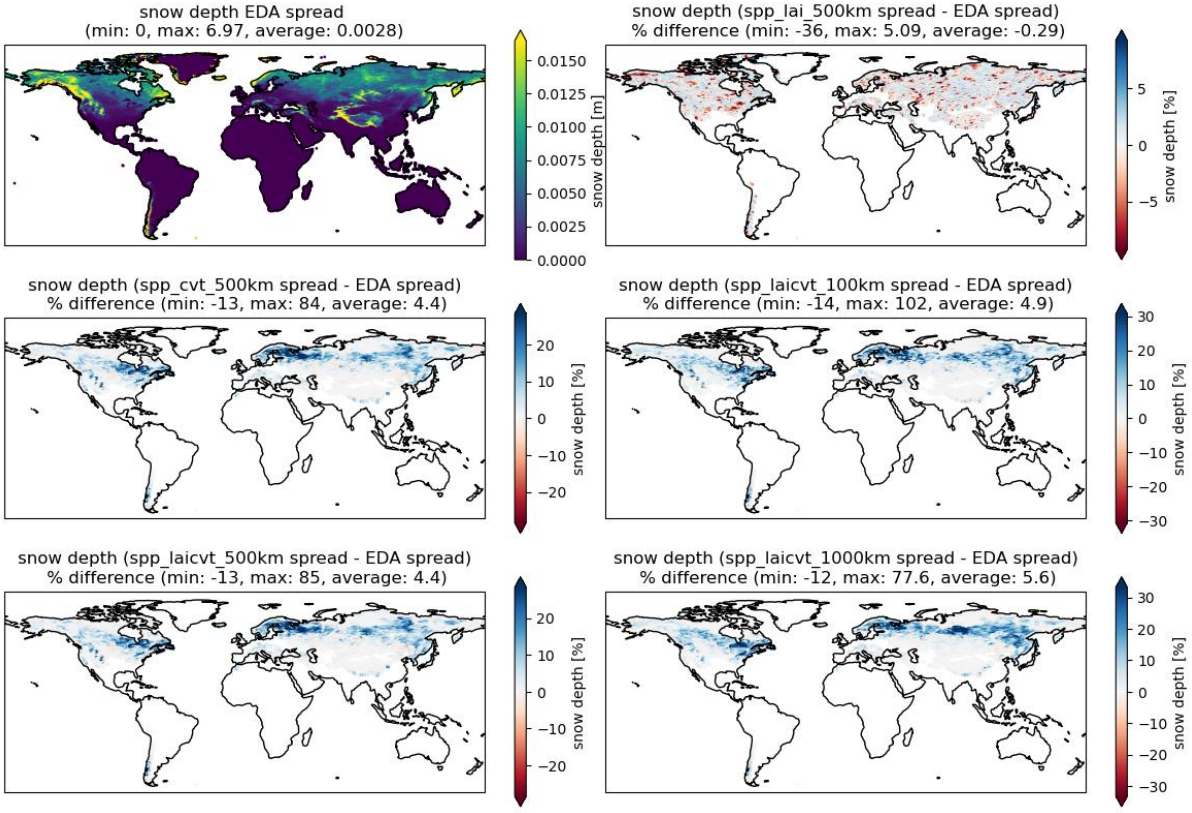

Figure 18: Figure showing the EDA spread for snow depth (top left) and the increase/decrease in spread for the different experiments outlined in Table 4.1 from the EDA benchmark.

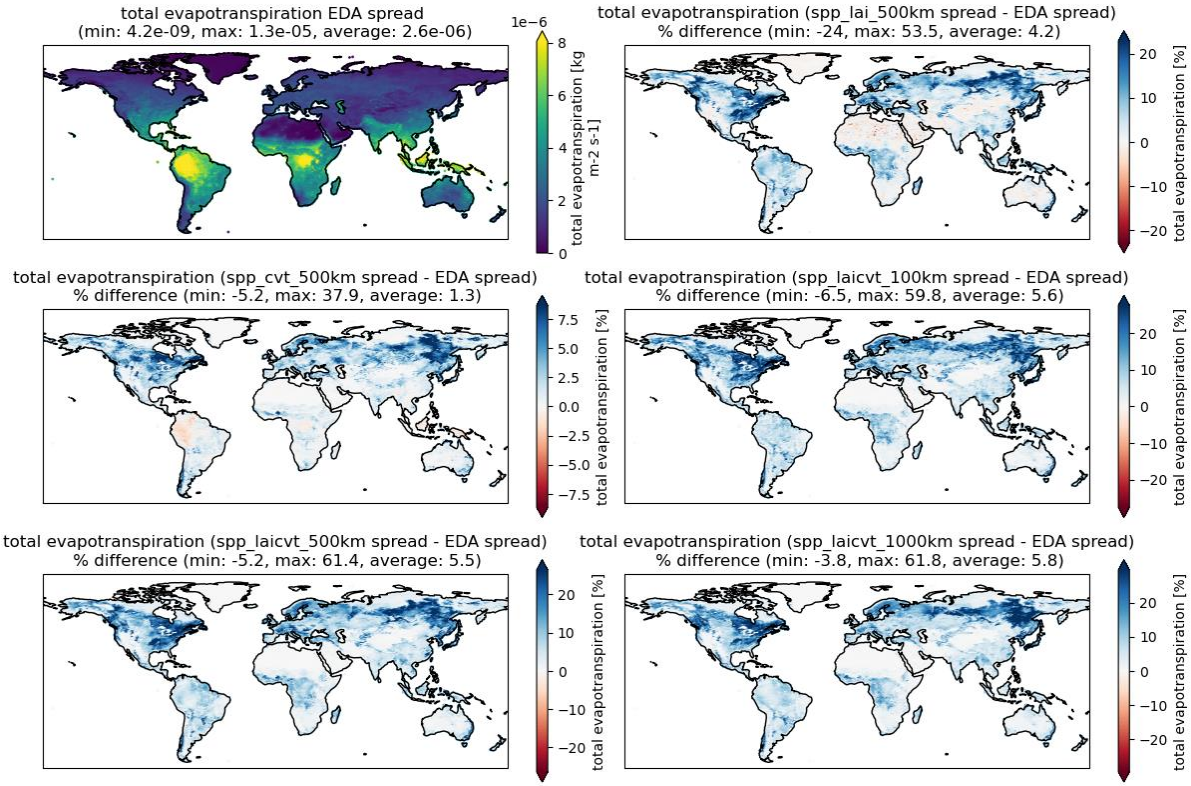

Figure 19: Figure showing the EDA spread for evapotranspiration (top left) and the increase/decrease in spread for the different experiments outlined in Table 3 from the EDA benchmark.

5 Summary and Conclusions

From the regional perspective we have extended the offline forcing perturbations to also have spatially correlated noise, and with cycling of perturbations (start from previous cycle noise). The forcing perturbations have been extended and tested for snow and we have implemented a new remapping procedure to represent the spatial uncertainty of the offline forcing (not only accounting for the magnitude of the errors). These developments are taken into CERISE WP1 Tasks 1.2 and 1.3, where we are working on a local ensemble transform Kalman filter (LETKF) for surface data assimilation. Tests are also ongoing for evaluating the spread in surface temperature and soil moisture using spatially correlated forcing and state perturbation of the same variables. In addition to this we have extended the perturbation of surface parameters to offline runs (a feature which was only available in coupled HARMONIE-AROME runs). This will allow us to e.g. add leaf area index (LAI) to the control variable in the surface assimilation system.

The multi-layer soil scheme we use here ISBA-DF (soil) has not earlier been evaluated in an ensemble prediction system (EPS). Here we have extended the surface perturbations to this scheme and started tests in a coupled system. We saw that the ensemble correlations between surface and near surface variables were reasonable. The EPS tests allow us to implement and use our land-surface EnKF to update the EPS surface ensemble (previously only based on an offline ensemble in the coupled runs).

From the global perspective we have outlined a number of experiments perturbing land-surface parameters across an ensemble of land-surface model simulations with the ECMWF ECLand system. We have judged the increase in land-surface variable spread that these perturbations offer versus a control ensemble forced with perturbed meteorological forcing from an atmospheric Ensemble of Data Assimilations (EDA). For all outputted model variables we find positive increases in spread when perturbing both Leaf Area Index and Vegetation Fraction. These positive increases have been made more robust by the Stochastically Perturbed Parameter (SPP) methodology, using a Gaussian correlation function in time and space to generate correlated fields of random noise. The largest increases in spread are found for net CO2 flux, this bodes well as ECMWF moves further towards the analysis and forecasting of carbon fluxes. The next steps for this work will be to test such perturbations within the coupled system to better understand the feedback with the atmosphere. This work is currently underway, where we have introduced such perturbations into the atmospheric IFS EDA.

6 References

Blyverket, J., Hamer, P. D., Bertino, L., Albergel, C., Fairbairn, D., & Lahoz, W. A. (2019). An Evaluation of the EnKF vs. EnOI and the Assimilation of SMAP, SMOS and ESA CCI Soil Moisture Data over the Contiguous US. *Remote Sensing*, *11*(5), Article 5. https://doi.org/10.3390/rs11050478

Bonavita, M., Isaksen, L. and Hólm, E. (2012), On the use of EDA background error variances in the ECMWF 4D-Var. Q.J.R. Meteorol. Soc., 138: 1540-1559. https://doi.org/10.1002/qj.1899

Boussetta S, Balsamo G, Arduini G, Dutra E, McNorton J, Choulga M, Agustí-Panareda A, Beljaars A, Wedi N, Munõz-Sabater J, et al. ECLand: The ECMWF land-surface Modelling System. *Atmosphere*. 2021; 12(6):723. https://doi.org/10.3390/atmos12060723

de Rosnay, P., Drusch, M., Vasiljevic, D., Balsamo, G., Albergel, C. and Isaksen, L. (2013), A simplified Extended Kalman Filter for the global operational soil moisture analysis at ECMWF. Q.J.R. Meteorol. Soc., 139: 1199-1213. https://doi.org/10.1002/qj.2023

Diefenbach, T., Craig, G., Keil, C., Scheck, L., & Weissmann, M. (2023). Partial analysis increments as diagnostic for LETKF data assimilation systems. *Quarterly Journal of the Royal Meteorological Society*, *149*(752), 740–756. https://doi.org/10.1002/qj.4419

Draper, C. S. (2021). Accounting for Land Model Uncertainty in Numerical Weather Prediction Ensemble Systems: Toward Ensemble-Based Coupled Land–Atmosphere Data Assimilation. *Journal of Hydrometeorology*, *22*(8), 2089–2104. https://doi.org/10.1175/JHM-D-21-0016.1

Frogner, I.-L., Andrae, U., Bojarova, J., Callado, A., Escribà, P., Feddersen, H., Hally, A., Kauhanen, J., Randriamampianina, R., Singleton, A., Smet, G., van der Veen, S., & Vignes, O. (2019). HarmonEPS—The HARMONIE Ensemble Prediction System.

Weather and Forecasting, *34*(6), 1909–1937. https://doi.org/10.1175/WAF-D-19-0030.1

Holmes, T. R. H., Jackson, T. J., Reichle, R. H., & Basara, J. B. (2012). An assessment of surface soil temperature products from numerical weather prediction models using ground-based measurements. *Water Resources Research*, *48*(2), 2011WR010538. https://doi.org/10.1029/2011WR010538

Lang, STK, Lock, S-J, Leutbecher, M, Bechtold, P, Forbes, RM. Revision of the Stochastically Perturbed Parametrisations model uncertainty scheme in the Integrated Forecasting System. Q J R Meteorol Soc. 2021; 147: 1364–1381. https://doi.org/10.1002/qj.3978

Reichle, R. H., De Lannoy, G. J. M., Liu, Q., Ardizzone, J. V., Colliander, A., Conaty, A., Crow, W., Jackson, T. J., Jones, L. A., Kimball, J. S., Koster, R. D., Mahanama, S. P., Smith, E. B., Berg, A., Bircher, S., Bosch, D., Caldwell, T. G., Cosh, M., González-Zamora, Á., … Zeng, Y. (2017). Assessment of the SMAP Level-4 Surface and Root-Zone Soil Moisture Product Using In Situ Measurements. *Journal of Hydrometeorology*, *18*(10), 2621–2645. https://doi.org/10.1175/JHM-D-17-0063.1

Document History

Version	Author(s)	Date	Changes
0.1	Jostein Blyverket, Åsmund Bakketun, Jelena Bojarova	01/12/2023	Initial version
0.2	Jostein Blyverket, Åsmund Bakketun, Jelena Bojarova, Ewan Pinnington, Peter Weston, Patricia de Rosnay	08/12/2023	First full draft
1.0	(As above)	14/12/2023	Issued version

Internal Review History

Internal Reviewers	Date	Comments
Patricia de Rosnay (ECMWF), Jelena Bojarova (SMHI)	Dec 2023	Initial version