ISDA-Online

March 12, 2021 / 15 – 17 UTC
(16 – 18 pm CET / 8 – 10 am MST / 12 pm – 2 am JST)

"Mathematics of Data Assimilation"

Session Chairs: Javier Amezcua, Steven Fletcher, Lars Nerger

Program:

15:00 – 15:10 **Welcome**

15:10 – 15:30 **Mitigating Sampling Error using Optimal Localisation in Ensemble Data Assimilation**
Rebecca Atkinson, Sue Hughes, Jonathan Flowerdew

15:30 – 15:50 **Ensemble-based data assimilation via nonlinear couplings**
Ricardo Baptista, Youssef Marzouk

15:50 – 16:10 **Continuum Covariance Propagation for Understanding Variance Loss in Advective Systems**
Shay Gilpin, Tomoko Matsuo, Stephen E. Cohn

16:10 – 16:30 **A subspace iterative ensemble smoother for solving DA and inverse problems**
Geir Evensen

16:30 – 16:55 **High-dimensional Data Assimilation using Regularization and Iterative Resampling with the Local Particle Filter**
Jonathan Poterjoy

16:55 – 17:00 **Closing; Information on upcoming sessions**

Please note:

- The times in UTC are approximate. In case of technical problems, we might have to change the order of the presentations.
- Login to the session is possible from 20 minutes before the event starts. When you login before 15:00 UTC, and everything is quiet, this is most likely because we muted the microphones.

Mitigating Sampling Error using Optimal Localisation in Ensemble Data Assimilation

Rebecca Atkinson[1], Ian Roustone[1], Sue Hughes[1], Jonathan Flowerdew[2]

[1] University of Surrey, UK
[2] Met Office, Exeter, UK

Ensemble data assimilation can be sensitive to sampling errors due to finite ensemble size. This is addressed via techniques such as localisation. Current localisation methods are typically ad hoc. They assume a fixed shape for the localisation function and may tune its parameters based on past data. We define an optimal localisation as one which produces the most accurate analysis state for an ensemble which perfectly samples the background uncertainty. An improved understanding of the form of and factors that determine the optimal localisation could inform better localisation techniques.

A series of theoretical optimal localisations were identified, which provide the optimal analysis state when assimilating a single observation. They are expected to remain applicable for sufficiently sparse observations. Optimal localisation methods were found for: a fixed single true covariance (OSTC), a variable true covariance (OVTC) (a covariance with a climatology) and a hybrid of the OVTC. Our hybrid localises the variation from the mean covariance. This is different from previous hybrid methods which combine the climatological covariance with a localised ensemble covariance using a fixed weight. Numerical experiments are performed using an Ensemble Kalman Filter in a framework where the true statistics were known. Details of the formulations will be presented along with some recent results.

This work is part of our agenda to apply these localisations to a balanced model and investigate the impacts on the accuracy and balance of the analysis state. The atmosphere is known to approximately obey balance equations and localisation can have a damaging effect on the balance of the analysis state. One hypothesis is that the hybrid OVTC localisation will impact the balance less negatively than the other localisations because it preserves the mean covariances.

Ensemble-based data assimilation via nonlinear couplings

Ricardo Baptista, Youssef Marzouk

Center for Computational Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

We consider the Bayesian inference step at the core of filtering and smoothing problems for high- dimensional and non-Gaussian state-space models. While the ensemble Kalman filter and smoother can yield robust estimates of the state in many settings, these algorithms are limited by linear transformations and are generally inconsistent with the Bayesian solution in the large-sample limit. To generalize these approaches, at each assimilation step we propose transforming the non-Gaussian prior ensemble into a col- lection of posterior samples via a deterministic coupling, represented as a composition of nonlinear transport maps. This approach avoids any form of importance sampling and is thus applicable to high-dimensional systems. In this presentation, we propose an adaptive algorithm for estimating transport maps, given prior ensembles with small sample sizes. Our algorithm gradually enriches the complexity of the map to carefully balance the bias due to a finite parameterization with the estimation variance. We also show how to construct nonlinear couplings for both filtering and smoothing problems in a scalable way, exploiting both conditional independence and decay of correlation/dependence. We demonstrate the performance of our algorithms on chaotic dynamical systems with non-Gaussian statistics.

Continuum Covariance Propagation for Understanding Variance Loss in Advective Systems

Shay Gilpin[1,2], Tomoko Matsuo[2,1], Stephen E. Cohn[3]

[1]Department of Applied Mathematics, University of Colorado Boulder, USA
[2]Smead Aerospace Engineering Sciences, University of Colorado Boulder, USA
[3]Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, Maryland, USA

At the heart of modern data assimilation schemes is covariance propagation. In this work, we show that for advective dynamics the covariance propagation by itself typically causes significant spurious loss of variance, even at full rank. Standard methods for covariance propagation, such as in the Kalman filter, difference across the diagonal of the covariance matrix. This results in spurious dissipation and loss of variance because its continuum counterpart, the diagonal of the kernel of the covariance operator, is a characteristic surface for advective dynamics. In essence, most methods are differencing across a shock-like surface.

To demonstrate this, we first study continuum covariance propagation by analyzing the covariance evolution equation for advective dynamics. The behavior of this evolution equation changes abruptly as the correlation length tends to zero since the diagonal is a characteristic surface. Our numerical experiments then show that the variance lost during numerical propagation greatly exceeds that traditionally expected from numerical dissipation. In certain cases, the numerical behavior for short correlation lengths approximates the limiting continuum behavior at zero correlation length well, while approximating the continuum behavior for short, nonzero correlation lengths poorly. Our results suggest that developing local covariance propagation methods may prove useful in ameliorating the variance loss observed in data assimilation schemes.

A subspace iterative ensemble smoother for solving DA and inverse problems

Geir Evensen

NORCE and NERSC, Norway

The use of iterative smoothers has led to significant improvements when solving nonlinear and high- dimensional inverse problems, e.g., within parameter estimation or history matching of oil-reservoir models. Thus, there is now extensive operational use of iterative smoothers in the petroleum industry. One can expect a similar level of improvement with the introduction of iterative smoothers in ensemble DA applications. This presentation will consider a new subspace formulation (Raanes et al., 2019; Evensen et al., 2019) of the EnRML method by Chen and Oliver (2013). The new formulation is a stochastic iterative ensemble smoother suitable for big models and data. Particularities of the method include the elimination of pseudo inversions of large matrices that were problematic in the formulation by Chen and Oliver (2013), fast convergence, the computation of the transform matrix is independent of the state vector, and it is suitable for use with local analysis. The new method includes an ensemble sub-space projection of measurement errors and measurements and the impact of this projection will be addressed.

Subspace EnRML algorithm

1:	Input: $\mathbf{X} \in \Re^{n\times N}$ (model ensemble)	
2:	Input: $\mathbf{D} \in \Re^{m\times N}$ (perturbed measurements)	
3:	$\mathbf{W}_0 = 0$	$\mathbf{W} \in \Re^{N\times N}$
4:	$\mathbf{\Pi} = \left(\mathbf{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^{\mathrm{T}}\right)\Big/\sqrt{N-1}$	$\mathbf{\Pi} \in \Re^{N\times N}$
5:	$\mathbf{E} = \mathbf{D}\mathbf{\Pi}$	$\mathbf{E} \in \Re^{m\times N}$
6:	i=0	
7:	**repeat**	
8:	$\mathbf{Y}_i = \mathbf{g}(\mathbf{X}_i)\mathbf{\Pi}$	$\mathbf{Y} \in \Re^{m\times N}$
9:	$\mathbf{\Omega}_i = \mathbf{I} + \mathbf{W}_i\mathbf{\Pi}$	$\mathbf{\Omega} \in \Re^{N\times N}$
10:	$\mathbf{S}_i = \mathbf{Y}_i\mathbf{\Omega}_i^{-1}$	$\mathbf{S} \in \Re^{m\times N}$
11:	$\mathbf{H}_i = \mathbf{S}_i\mathbf{W}_i + \mathbf{D} - \mathbf{g}(\mathbf{X}_i)$	$\mathbf{H} \in \Re^{m\times N}$
12:	$\mathbf{W}_{i+1} = \mathbf{W}_i - \gamma\Big(\mathbf{W}_i - \mathbf{S}_i^{\mathrm{T}}\big(\mathbf{S}_i\mathbf{S}_i^{\mathrm{T}} + \mathbf{E}\mathbf{E}^{\mathrm{T}}\big)^{-1}\mathbf{H}_i\Big)$	
13:	$\mathbf{T}_i = \Big(\mathbf{I} + \mathbf{W}_{i+1}\Big/\sqrt{N-1}\Big)$	$\mathbf{T} \in \Re^{N\times N}$
14:	$\mathbf{X}_{i+1} = \mathbf{X}\mathbf{T}_i$	
15:	i=i+1	
16:	**until** convergence	

The most expensive computation the evaluation of the model-ensemble prediction $\mathbf{Y}_i = \mathbf{g}(\mathbf{X}_i)$. All other computations in this algorithm are linear in the number of measurements m as well as the number of state variables n. The projection $\mathbf{\Pi}$ subtracts the ensemble mean. The evaluation of $\mathbf{S}_i$, is most easily computed by writing $\mathbf{\Omega}^{\mathrm{T}}\mathbf{S}_i^{\mathrm{T}} = \mathbf{Y}^{\mathrm{T}}$ and computing the LU factorization of $\mathbf{\Omega}_i^{\mathrm{T}}$ followed by m back substitutions to a cost $\mathcal{O}(mN^2)$. In line 10, one needs to compute the matrix multiplication $\mathbf{S}_i\mathbf{W}_i$ which is mN^2. It is possible to approximately compute the inversion in line 11 to a cost $\mathcal{O}(mN^2)$ using the ensemble subspace algorithm discussed in *Evensen* (2004); *Evensen et al.* (2019). Finally, in line 12, we obtain the updated model ensemble in nN^2 floating-point operations.

References

Chen, Y., and D. S. Oliver, Levenberg-Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification, Computat Geosci, 17, 689–703, 2013.

Evensen, G., Sampling strategies and square root analysis schemes for the EnKF, Ocean Dynamics, 54, 539–560, 2004.

Evensen, G., P. Raanes, A. Stordal, and J. Hove, Efficient implementation of an iterative ensemble smoother for data assimilation and reservoir history matching, Frontiers in Applied Mathematics and Statistics, 5, 47, 2019, https://www.frontiersin.org/article/10.3389/fams.2019.00047.

Raanes, P. N., A. S. Stordal, and G. Evensen, Revising the stochastic iterative ensemble smoother, Nonlin. Processes Geophys, 26, 325–338, 2019, https://doi.org/10.5194/npg-26-325-2019.

High-dimensional Data Assimilation using Regularization and Iterative Resampling with the Local Particle Filter

Jonathan Poterjoy

University of Maryland, Colleg Park, MD, USA

Particle filters (PFs) are sequential Monte Carlo methods that can solve data assimilation problems characterized by non-Gaussian error distributions for prior variables or measurements. Recent efforts to apply PFs for high-dimensional geophysical models have resulted in localized PFs, which significantly reduce the number of particles required for applications of large spatial dimension. Localization, however, is often insufficient for preventing particle weight collapse for real geophysical problems, like numerical weather prediction. For example, it does not prevent the local collapse of weights when provided with a dense network of accurate, independent measurements or when model error is not well characterized. Both situations can lead to filter divergence even for univariate problems. This presentation introduces several approaches for maintaining filter stability under the above circumstances, which can be characterized broadly as “large sampling error” regimes. The first set of approaches adopts regularization in a manner similar to past particle filtering studies; i.e., by including an extra term in the weight calculations to place a lower-bound on effective ensemble size or maximum particle weights. The second set of strategies extend regularization to factor the posterior density, thus allowing for a sequence of iterative resampling steps – each step using a larger effective ensemble size than if a single resampling were performed. In addition to preventing filter divergence in large sampling error regimes, iterative resampling helps alleviate some of the assumptions used to derive the Poterjoy (2016), Poterjoy et al. (2019) local PF algorithms. In the absence of localization or other particle mixing parameters, the iterative resampling converges to the Bayesian solution as sample size increases, thus maintaining this property of the original local PF. Results will be presented for low-dimensional applications, including one that mimics feature displacement errors in weather models, as well as a high-dimensional general circulation model in the NCAR Data Assimilation Research Testbed.

References:

Poterjoy, J., 2016: A localized particle filter for high-dimensional nonlinear systems. Mon. Wea. Rev., 144, 59 – 76.

Poterjoy, J., L. J. Wicker, and M. Buehner, 2019: Progress in the development of a localized particle filter for data assimilation in high-dimensional geophysical systems., Mon. Wea. Rev. 147, 1107 – 1126.