Workshop: Stochastic Modelling in GFD, data assimilation, & non-equilibrium phenomena

2-6 November 2015, Imperial College London

Monday 2nd November

09:45 - 10:00	Registration
10:00 - 11:00	**Dr Alberto Carrassi, NERSC Norway**
11:00 - 11:30	*Break*
11:30 - 12:30	**Prof. Tim Palmer, University of Oxford**
12:30 - 14:00	*Lunch*
14:00 - 15:00	**Prof. Darryl Holm, Imperial College London**
15:00 - 15:30	*Break*
15:30 - 16:30	**Dr Chris Ferro, University of Exeter**

Tuesday 3rd November

10:00 - 11:00	**Prof. Valerio Lucarini, University of Hamburg**
11:00 - 11:30	*Break*
11:30 - 12:30	**Prof. Eric Vanden-Eijnden, Courant Institute**
12:30 - 14:00	*Lunch*
14:00 - 15:00	**Dr Piotr Smolarkiewicz, ECMWF**
15:00 - 15:30	*Break*
15:30 - 16:30	**Prof. Andrew Lorenc, Met Office**

Wednesday 4th November

09:30 - 10:30	**Prof. David Marshall, University of Oxford**
10:30 - 11:00	*Break*
11:00 - 12:00	**Prof. Tobias Kuna, University of Reading**
12:00 - 13:00	*Lunch*
13:00 - 14:00	**Prof. Chris Jones, University of North Carolina**
14:00 - 15:00	**Prof. Mike Cullen, Met Office**
15:00 - 15:30	*Break*
15:30 - 16:30	**Dr Marie Farge, LMD-IPSL-CNRS , Ecole Normale Supérieure**

Thursday 5th November

09:00 - 10:00	**Daniel Sanz-Alonso, University of Warwick**
10:00 - 10:30	*Break*
10:30 - 11:30	**Prof. Peter Jan van Leeuwen, University of Reading**
11:30 - 12:30	**Prof. Sebastian Reich, University of Potsdam**
12:30 - 14:00	*Lunch*
14:00 - 18:00	**Deutscher Wetterdienst (DWD)** Prof. Roland Potthast Dr Ana Fernadnez Dr Hendrik Reich Dr Andreas Rhodin
Workshop dinner	

Friday 6th November

09:30 - 10:30	**Dr Mike Fisher, ECMWF**
10:30 - 11:00	*Break*
11:00 - 12:00	**Prof. Johannes Zimmer, University of Bath**
12:00 - 12:15	Closing remarks

Speaker titles and abstracts

Data Assimilation to compute model evidence: The attribution problem
Dr Alberto Carrassi, NERSC Norway

A new approach allowing for near real time, systematic causal attribution of weather and climate-related events is described. The method is purposely designed to allow its operability at meteorological centers by synergizing causal attribution with Data Assimilation (DA). The concept of contextual model evidence is introduced and is explained its link to the causal attribution. It is then shown how causal attribution can be obtained as a side-product of the statistical inference performed for the assimilation of data. Three strategies are considered: DA-based ensemble forecasting, filtering and smoothing. The theoretical rationale of this approach is explained along with the most prominent features of a DA-based detection and attribution procedure. The proposal is illustrated numerically with low-order nonlinear models, and is compared with standard methods for detection and attribution showing promising performance. The method stresses on the concept of model evidence, and open questions on how to compute and interpret the response to forcing whose effects one wants to contrast with respect to model error and other source of uncertainties. Practical obstacles that need to be addressed to make the proposal readily operational within weather forecasting centers are finally laid out.

Prof. Tim Palmer, University of Oxford

Stochastic Variational Principles for GFD: Modelling the unknown unknowns
Prof. Darryl Holm, Imperial College London

This talk develops an approach to SPDEs in which noise terms have data driven coefficients, linked to a stochastic variational principle, constrained by stochastic advection equations. The approach combines geometric mechanics with probabilistic analysis, in a variational hierarchy of data driven GFD approximations which preserve the invariance properties of the original system. Specifically, it combines different rapidly developing methods (particle filters, stochastic calculus of variations, regularisation methods for singular SPDEs) with new ideas from geometric mechanics exploiting ``stochastic advection'' (Lie-derivative transport of advected quantities by a stochastic vector field, indexed by the Eulerian spatial position). It contributes to both applications and foundations of stochastic fluid dynamics and other complex nonlinear systems, to enhance numerical methods and statistical modelling. The approach is explained and examples are given for stochastic GFD equations in DD Holm, Variational principles for stochastic fluid dynamics, Proc Roy Soc A, 471: 20140963 (2015).

Evaluating forecasts and models
Dr Chris Ferro, University of Exeter

Fluctuations and Response in Geophysical Fluid Dynamics
Prof. Valerio Lucarini, University of Hamburg

Prof. Eric Vanden-Eijnden, Courant Institute

A hybrid all-scale finite-volume module for global NWP

Dr Piotr Smolarkiewicz, ECMWF

The talk introduces a global nonhydrostatic finite-volume module designed to enhance an established spectral-transform based numerical weather prediction (NWP) model. The module adheres to NWP standards, with formulation of the governing equations based on the classical meteorological latitude-longitude spherical framework. In the horizontal, a bespoke unstructured mesh with finite-volumes built about the reduced Gaussian grid of the existing NWP model circumvents the notorious stiffness in the polar regions of the spherical framework. All dependent variables are co-located, accommodating both spectral-transform and grid-point solutions at the same physical locations. In the vertical, a uniform finite-difference discretisation facilitates the solution of intricate elliptic problems in thin spherical shells, while the pliancy of the physical vertical coordinate is delegated to generalised continuous transformations between computational and physical space. The newly developed module assumes the compressible Euler equations as default, but includes reduced soundproof PDEs as an option. Furthermore, it employs semi-implicit forward-in-time integrators of the governing PDE systems, akin to but more general than those used in the NWP model. The module shares the equal regions parallelisation scheme with the NWP model, with multiple layers of parallelism hybridising MPI tasks and OpenMP threads. The efficacy of the developed nonhydrostatic module is illustrated with benchmarks of idealised global weather.

Modelling error distributions for NWP Data Assimilation

Prof Andrew Lorenc, Met Office

Nearly all operational Numerical Weather Prediction (NWP) systems base their data assimilation systems in principle on Bayes Theorem, however the prior comes from past data using a forecast model that is very large, complex, chaotic and erroneous. We can only afford one “best estimate” (or a small ensemble of forecasts) with such a model so the error distribution of the best estimate forecast is unknowable. We still need to perform data assimilation and issue an operational forecast, so we proceed with the Bayesian methodology anyway, making modelling assumptions about the error PDF. The Gaussian approximation is practically essential if we are to model the PDF for a large forecast with about a billion degrees of freedom; without it we have no hope of even sampling its billion-dimensional space.

Bayes Theorem and a Gaussian assumption for PDFs leads us to the extended Kalman filter equations. Even these are unaffordable for an NWP system since they require the propagation and inversion of non-sparse matrices of order a billion. Further assumptions about balance relationships, and stationarity in time space and direction, are often invoked. A recent focus has been the derivation of PDF information from an ensemble of forecasts. Although the basic Kalman filter can be applied sequentially, these additional approximations can introduce inaccuracies in the assimilation of observations which are near in space but not processed together - in this case the resulting approximate analsis need not fit accurate observations as it should. For this and more practical reasons operational NWP processes observations in batches each spanning a time-window (typically 6 hours for global NWP). Within each window, we get benefit from considering the observation times, i.e. we do an approximate Kalman smoother. Different approaches are presented for making this computationally feasible - hybrid-4DVar and hybrid-4DEnVar. The former uses a linear forecast model and its adjoint to construct the time-dimension of the PDF whereas the latter uses the time-evolution of the ensemble trajectories. Between batches we use the same methods as an ensemble Kalman filter to generate an ensemble of forecasts sampling the PDF. I show some recent results demonstrating the benefit of using improved ensembles in these methods.

In order to estimate covariances from a small (40~200) ensemble our hybrid-4DVar and hybrid-4DEnVar employ localisation similar to that in ensemble Kalman filter (EnKF) methods. The ensemble-Var methods allow a much more general interpretation of “localisation”, for example in spectral-space and time, as well as hybridisation with more stationary “climatological” estimates (these are much harder in the EnKF). Some examples of localised covariances from an NWP model are shown.

A geometric framework for parameterising ocean eddies and a simple model of eddy saturation.

Prof. David Marshall, University of Oxford

The ocean is populated by an intense geostrophic eddy field with a dominant energy-containing scale of order 100 km at midlatitudes. Ocean climate models are unlikely routinely to resolve geostrophic eddies for the foreseeable future and thus development and validation of improved parameterisations remains a vital task. Moreover, development and validation of improved eddy parameterizations is an excellent strategy for testing and advancing our understanding of how geostrophic ocean eddies impact the large-scale circulation. A geometric framework for parameterising ocean eddy fluxes is developed that is consistent with conservation of energy and momentum. The framework involves rewriting the residual-mean eddy force, or eddy potential vorticity flux, as the divergence of an eddy stress tensor. The magnitude of the tensor is bounded by the eddy energy, allowing its components to be rewritten in terms of the eddy energy and non-dimensional parameters describing the mean "shape" of the eddies, analogous to “eddy ellipses” used in observational oceanography. These non-dimensional geometric parameters have strong connections with classical stability theory, for example, the new framework preserves the functional form of the linear Eady growth rate and, with one additional ingredient, Arnold’s first stability theorem. In the second part of the talk, this framework will be used to developed a simple model of eddy saturation of the Antarctic Circumpolar Current (ACC) -the relative insensitivity of the ACC volume transport to the magnitude of the surface wind stress in ocean models with explicit eddies. The model predicts that the volume transport is independent of the surface wind stress but scales with the bottom drag, whereas the eddy energy scales with the wind stress, independent of the bottom drag. These theoretical predictions are confirmed in eddy-resolving numerical calculations for an idealised reentrant channel. The results suggest that the rate of eddy energy dissipation has a strong impact not only the volume transport of the ACC, but also on global ocean stratification and heat content through the thermal wind relation. These results offer the prospect of being able to capture eddy saturation in ocean models with parameterised eddies.

Typical behaviour of extremes of chaotic dynamical systems for general observables

Prof. Tobias Kuna, University of Reading

Hybrid EnKF and Particle Filter: Lagrangian DA and Parameter Estimation

Prof. Chris Jones, University of North Carolina

Dealing with high dimensional systems is one of the central problems of data assimilation. A strategy is proposed here for systems that enjoy a skew-product structure. This is joint work with Naratip Santitssadeekorn (Surrey).

Evaluation of model error using data assimilation

Prof. Mike Cullen, Met Office

Wavelet-based methods to analyse, compress and compute turbulent flows

Dr Marie Farge, LMD-IPSL-CNRS, Ecole Normale Superieure

We will discuss the pertinence of the wavelet representation to study turbulence. We will present a wavelet-based algorithm to extract coherent structures out of fully-developed turbulent flows. We will demonstrate it on two- and three-dimensional turbulent flows computed by direct numerical simulation. We will finally present the CVS (Coherent Vorticity Simulation) method which computes Navier-Stokes equation in an orthogonal wavelet basis, adapted at each time step to resolve all nonlinear interactions whatever their scale, and compute the time evolution of a mixing layer with it.

Importance Sampling: Computational Complexity and Intrinsic Dimension

Daniel Sanz-Alonso, University of Warwick

We study importance sampling and particle filters in high dimensions and link the collapse of importance sampling to results about absolute continuity of measures in Hilbert spaces and the notion of effective dimension.

Nonlinear Data Assimilation in very high dimensional systems

Prof. Peter Jan van Leeuwen, University of Reading

Data assimilation is the science of combining observations of a system with a numerical model of that system, including uncertainties in both. Using Bayes Theorem it becomes clear that it all boils down to finding an accurate representation of the posterior probability density function (pdf).
If the model or the function that connects the model variables to observation space is nonlinear the functional shape of the pdf is not known apriori and one typically uses Monte-Carlo samples for this representation. And if the model is high dimensional each sample is very expensive computationally. So we have to find a way to represent the posterior pdf with a minimal number of samples. We will explore so-called Particle Filters in this talk.

Particle filters have long be considered inappropriate for high-dimensional systems, but the freedom in so-called proposals densities, that allows one to modify the model to force the particles to the observations has been underexplored, and the hope has been that some smart proposal could solve the issue. Unfortunately, recent papers show that even the best proposal (which turns out to be a proposal density that has been called optimal without proper foundation) leads to particle filters that are hopelessly inefficient. We will first try to understand in detail what the issue is, and conclude that the class of particle filters typically considered is indeed too narrow for high-dimensional applications.

However, we can extend this class and generate particle filters that do not suffer from the curse of dimensionality. The new class opens up a whole new field of possibly very efficient particle filters, and a few examples on high-dimensional systems with millions degrees of freedom will be discussed. The new class also opens up the discussion of bias versus accuracy, namely that when the sample size is small, a small bias in the method might be allowed if that bias is smaller than the statistical noise in Monte-Carlo sampling. If time permits we will discuss the bias issue in more detail, exploring some recent comparisons between different nonlinear data-assimilation methods.

Non-Gaussian data assimilation via a localized hybrid ensemble transform filter

Prof. Sebastian Reich, University of Potsdam

Data assimilation is the task to combine evolution models and observational data in order to produce reliable predictions. In this talk, we focus on ensemble-based recursive data assimilation problems. Our main contribution is a hybrid filter that allows one to adaptively bridge between ensemble Kalman and particle filters. While ensemble Kalman filters are robust and applicable to high dimensional problems, particle filters are asymptotically consistent in the large ensemble size limit. We demonstrate numerically that our hybrid approach can improve the performance of both Kalman and particle filters at moderate ensemble sizes. We also show how to implement the concept of localisation into a hybrid filter, which is key to their applicability to high dimensional problems.

On the DWD Data Assimilation Environment and on Ensemble and Particle Filters for Large-Scale Data Assimilation

Prof. Roland Potthast, German Weather Service

Ensemble data assimilation techniques are of rapidly growing importance. Ensemble techniques allow to describe and forecast uncertainty of the analysis, but they also improve the assimilation result itself, by allowing estimates of the covariance or, more general, the prior and posterior probability distribution of atmospheric states.

In our talk, we will first give a survey about recent activities of the German Meteorological Service DWD. Then, we present recent work on the further development of the ensemble data assimilation towards a particle filter for large-scale atmospheric systems, which keeps the advantages of the LETKF, but overcomes some of its limitations.

Data Assimilation for the ICON global model: comparing 3D-VAR, LETKF and a hybrid setup.

Dr Ana Fernandez, German Weather Service

The ICOsahedral Nonhydrostatic general circulation (ICON) model has been jointly developed by the Max Planck Institute for Meteorology in Hamburg (MPI-M) and the German Weather Service (Deutscher Wetterdienst, DWD), and is the global model used at DWD in operations since January 2015. In this talk, we will briefly present three different data assimilation schemes that have been implemented and tested for ICON. Namely, the (currently operational) 3D-VAR, an LETKF and a hybrid system (soon to replace the current 3D-VAR in operations). We will then illustrate how they work differently using our system and compare the performance of forecasts initialized using the three different methods.

Multi-Scale Localisation in Ensemble Variational Data Assimilation

Dr Andreas Rhodin, German Weather Service

For estimating the uncertainty of the background, Ensemble Kalman Filter algorithms use the sample covariance matrix of a short range forecast ensemble. Due to the limited ensemble size (of order 100 in atmospheric data assimilation) these forecast error covariance matrix is rank deficient and compromised by sampling noise, especially in cases there the true correlations are small.

In order to suppress noise, localisation has to be applied: Instead of the raw sample covariance matrix the Schur product of this matrix with a localisation matrix is used. By this procedure matrix elements representing covariances at large distances, which are expected to be small, are forced to be zero. However it is difficult to choose an optimal covariance matrix in view of the large range of different scales involved.

In variational ensemble data assimilation it is possible to apply localisation not in physical space (i.e. based on the spatial distance) but in transformed representation. Here localisation in wavelet transformed representation is proposed leading to a scale selective localisation of covariance matrices.

Experience with Kilometer Scale Ensemble Data Assimilation for the COSMO Model (KENDA)

Dr Hendrik Reich, German Weather Service

An ensemble Kalman filter for kilometre-scale or convective scale data assimilation (KENDA) was developed for the COSMO model. The KENDA system includes a fully developed local ensemble transform Kalman filter (LETKF) including a deterministic analysis based on the Kalman gain for the analysis ensemble mean. The KENDA software suite includes tools for adaptive localization, multiplicative and additive covariance inflation, relaxation to prior perturbations, as well as adaptive observation errors.
We assimilate conventional data (radiosonde, aircraft, wind profiler, surface station data). Also, latent heat nudging of radar precipitation has been integrated into the KENDA system to be applied to the deterministic analysis or to all ensemble members. The performance of the different system components was investigated using a basic cycling environment (BACY) for a period of 30 days with 24h forecasts. The deterministic KENDA forecasts were compared with forecasts based on the current operational nudging data assimilation scheme at the German Weather Service (Deutscher Wetterdienst, DWD). For our experiments lateral boundary conditions for the regional model are given by the pre-operational global ensemble Kalman filter for the ICON model. The performance of the KENDA system proves to be overall superior to the forecast quality of the operational nudging scheme.

Dr Mike Fisher, ECMWF

Particle models, entropic flows and their stochastic correction

Prof. Johannes Zimmer, University of Bath

Large deviations provide a way to describe many particle dynamics by partial differential equations (PDEs); the PDE is a minimiser of the so-called rate function. However, the underlying microscopic process contains more information, notably fluctuations around the minimum state described by the deterministic PDE. Can stochastic terms be derived which model this additional information, in a way that is compatible to the limit passage via large deviations (and the geometric structure, such as the Wasserstein setting)? This question will be investigated for the simplest possible case of linear diffusion and some other cases. In a second part, the thermodynamic setting of a class of nonlinear diffusion problems will be discussed.