Particle filters for Data Assimilation

Dan Crisan

Imperial College London

DataSig Seminar Series
1 December 2022
The Alan Turing Institute

- Motivation: Data Assimilation for GFD models
- The Stochastic Filtering Framework (Signal/Observation)
- Timeline for Particle Filter development
- Particle Filters for DA (model reduction, tempering, jittering, nudging)
- Numerical example: DA for a two-layer quasi-geostrophic model
- The nudging procedure
- Final Remarks

Joint work with Igor Shevchenko (Imperial)

https://www.imperial.ac.uk/ocean-dynamics-synergy/

◦ C Cotter, D Crisan, D Holm, W Pan, I Shevchenko, *Data assimilation for a quasi-geostrophic model with circulation-preserving stochastic transport noise*, Journal of Statistical Physics, 1-36, 2020.
◦ D Crisan, I Shevchenko, *Particle filters with nudging,* work in progress.

Data Assimilation for GFD models

Geophysical Fluid Dynamics (GFD) models are used extensively to describe the evolution of the atmosphere and the oceans and play a crucial role in numerical weather prediction (NWP).

What is Data Assimilation ?

- set of methodologies that combines past knowledge of a system in the form of a numerical model with new information about that system in the form of observations of that system.
- designed to improve forecasting, reduce model uncertainties and adjust model parameters.
- term used mainly in the computational geoscience community
- major component of Numerical Weather Prediction
  - Variational DA: combines the model and the data through the optimisation of a given criterion (minimisation of a so-called cost-function).
  - Sequential DA: uses a set of model trajectories/possible scenarios that are intermittently updated according to data and are used to infer the past, current or future position of a system.

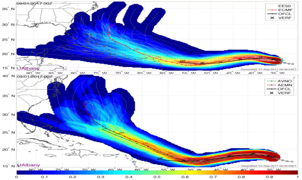

Hurricane Irma forecast: a. ECMWF, b. USA Global Forecast

DA recast as a stochastic filtering problem

Discrete framework: $\{X_t, Y_t\}_{t\geq 0}$ Markov process

- X the signal process - "hidden component", $X_{[0,t]} \triangleq (X_0, ..., X_t)$.
- Y the observation process - "the data" $Y_{[0,t]} \triangleq (Y_0, ..., Y_t)$.

The filtering problem: Find the conditional distribution of the *signal* X_t given the observation: $\pi_t(dx_t) = \mathbb{P}(X_t \in dx_t | Y_{[0,t]} = y_{[0,t]})$

The signal process $\{X_t\}_{t\geq 0}$ Markov chain, $X_0 \sim \pi_0(dx_0)$

- $\mathbb{P}(X_t \in dx_t | X_{t-1} = x_{t-1}) = K_t(x_{t-1}, dx_t)$
- Example : $X_t = b(X_{t-1}) + \sigma(X_{t-1}) B_t, \quad B_t \sim N(0,1)$ i.i.d.

The observation process

- $\mathbb{P}\left(Y_t \in dy_t | X_{[0,t]} = x_{[0,t]}, Y_{[0,t-1]} = y_{[0,t-1]}\right) = \mathbb{P}(Y_t \in dy_t | X_t = x_t) = g_t(y_t|x_t)dy$
- Example : $Y_t = h(X_t) + V_t, \quad V_t \sim N(0,1)$ i.i.d.

$$\pi_{t-1} \xrightarrow[\substack{\textit{model}\\ \textit{forecast}\\ \textit{prediction}}]{K_t} K_t\pi_{t-1} =: {\color{blue}p_t} \xrightarrow[\substack{\textit{assimilation}\\ \textit{analysis}\\ \textit{update}}]{\textbf{non-linear}:\ g_t *} g_t * p_t {\color{red}= \pi_t}$$

1. Initialisation $[t = 0]$.

- For $i = 1, ..., N$, sample $x_0^{(i)}$ from π_0,

2. Iteration $[t - 1 \text{ to } t]$.

Let $x_{t-1}^{(i)}$, $i = 1, \ldots, n$ be the positions of the particles at time $t - 1$.

- **Step 1.**

  For $i = 1, ..., n$, sample $\bar{x}_t^{(i)}$ from $K_t\left(x_{t-1}^{(i)}, dx_t\right)$.

  Compute the (normalized) weight $\bar{a}_t^{(i)} = g_t(\bar{x}_t^{(i)})/(\sum_{j=1}^{n} g_t(\bar{x}_t^{(j)}))$.
- **Step 2.**

  Sample n-times from $\bar{\pi}_t^N = \sum_{i=1}^{N} \bar{a}_t^{(i)} \delta_{\bar{x}_t^{(i)}}$

  Denote the positions of the particles by $x_t^{(i)}$, $i = 1, \ldots, n$.

$$\pi_t^N = \frac{1}{N} \sum_{i=1}^{N} \delta_{x_t^{(i)}} \simeq \pi_t$$

Asymptotic Consistency:

$$\lim_{N \to \infty} \pi_t^N = \pi_t.$$

Theoretical papers:

- D. Crisan, A. Lopez-Yela, J. Miguez, Stable approximation schemes for optimal filters. SIAM/ASA J. Uncertain. Quantification 2020.
- D. Crisan, J. Miguez, Nested particle filters for online parameter estimation in discrete-time state-space Markov models, 2018.
- A. Beskos, D. Crisan, A Jasra, K. Kamatani, Y. A. Zhou, A stable particle filter for a class of high-dimensional state-space models, 2017.
- A. Beskos, D. Crisan, A Jasra, On the stability of sequential Monte Carlo methods in high dimensions, 2014.

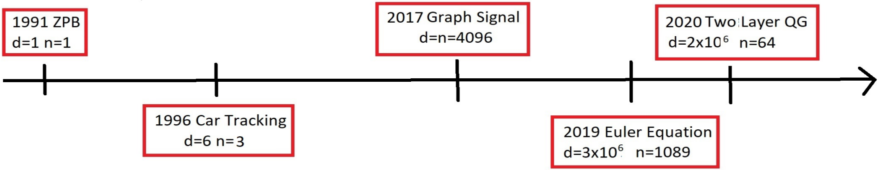

Car tracking using noisy video images

Signal $(\dot{v}_t, v_t, \theta_t, x_t, y_t, \tau_t)$

- (x, y) be the coordinates of the car (the center of the rear axle), v
- its tangential velocity ($\dot{v}$ is the rate of change of v),
- τ orientation (the angle between the direction of the car and the x-axis)
- θ the steering angle

Signal

$$\begin{cases} d\dot{v}_t &= \alpha dW_t^1 \\ dv_t &= v_t dt \\ d\theta_t &= -\beta\theta_t dt + \gamma dW_t^2 \\ dx_t &= v_t \cos(\tau_t) dt \\ dy_t &= v_t \sin(\tau_t) dt \\ d\tau_t &= a^{-1} v_t \theta_t dt \end{cases}$$

Observation

$$\begin{cases} d\bar{x}_t &= v_t \cos(\tau_t) dt & +\delta_3 dW_t^3 \\ d\bar{y}_t &= v_t \sin(\tau_t) dt & +\delta_4 dW_t^4 \\ d\bar{\tau}_t &= a^{-1} v_t \theta_t dt & +\delta_5 dW_t^5 \end{cases}$$

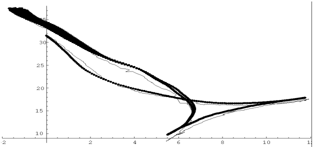

Data Assimilation - 24-28 September 2012

The third workshop on numerical methods for solving the filtering problem and high order methods for solving parabolic PDEs

Hosted by: Oxford-Man Institute, University of Oxford with support from the EPSRC

Held at: St. Anne's College, University of Oxford, Woodstock Road, Oxford, OX2 6HS

Organisers: D.Crisan, T.Lyons, A.Stuart

Monday, September 24, 2012	
08:30-09:20	Registration and coffee/tea, Mary Ogilvie Foyer, St. Anne's College
09:25-09:30	**Opening Remarks & Welcome, Dan Crisan** (Imperial) **Mary Ogilvie Lecture Theatre**
09:30	**Francois Le Gland** (IRISA) "Central limit theorem for the ensemble Kalman filter"
10:15-10:40	Refreshment Break, Mary Ogilvie Foyer
10:45	**Alex Beskos** (UCL) "A Sequential Monte Carlo sampler for target distributions in high dimensions"
11:30	**Jeff Anderson** (UCAR) "The Ensemble Kalman Filter: Data Assimilation and uncertainty quantification"
12:15-13:55	Lunch, Dining Hall
14:00	**Sumeetpal Singh** (Cambridge) "Estimating the Static Parameters in Linear Gaussian Multiple Target Tracking Models"
14:45	**Georg Gottwald** (USYD) "Controlling covariance overestimation in ensemble Kalman filters"
15:30-16.10	Refreshment Break, Mary Ogilvie Foyer
16:15	**Joaquín Miguez** (UC3M) "Iterative importance sampling with transformed weights and its application in stochastic filtering"
17:00	**Nikolaos Kantas** (Imperial) "A particle method for approximating principal eigen-functions and related quantities"
18:30-21:00	Drinks Reception and Buffet Supper in the Common Room at the Oxford-Man Institute, University of Oxford, Eagle House, Walton Well Road, Oxford OX2 6ED

Tuesday, September 25, 2012	
08:30-08:55	Coffee/Tea, Mary Ogilvie Foyer, St. Anne's College
09:00	**Istvan Gyongy** (Edinburgh) "On finite difference approximations for degenerate filtering equations"
09:45	**Christian Litterer** (Oxford) "A Hörmander theorem for stochastic differential equations driven by Gaussian noises"
10:30-10:55	Refreshment Break, Mary Ogilvie Foyer
11:00	**John Harlim** (NCSU) "Reduced stochastic models for filtering turbulent dynamical systems"
11:45	**Pierre Del Moral** (INRIA) "Particle methods for filtering and multiple object tracking"
12:30-13:55	Lunch, Dining Hall
14:00	**Hirromicki Nagao** (ISM) "Our attempts to broaden the application of data assimilation into many fields empowered by massively parallel computing"
14:45	**Sebastian Reich** (Uni-Potsdam) "Ensemble data assimilation: A coupling of measures perspective"
15:30-16:10	Refreshment Break, Mary Ogilvie Foyer
16:15	**Syoiti Ninomiya** (Tokyo Inst. Technology) "On the higher-order weak approximation of SDEs"
17:00	**Salvador Ortiz-Latorre** (Imperial) "A second order approximation of the continuous time filtering problem"
18:30-21:00	BBQ (weather permitting), The Lawn, St. Anne's College

DA is a hard problem

The likelihood function $x_k \mapsto g_t(x_k)$ can convey a lot of information about the signal, especially so in high dimensions. If this is the case, the problem becomes much harder. In particular, using the transition kernel as proposal, will be ineffective. Add-on techniques:

- Model Reduction (High $\mapsto$Low Res)
- Nudging
- Tempering
- Jittering

$$\pi_{t-1}^N \overbrace{\longrightarrow\longrightarrow}^{\text{nudging}} p_t^N \quad \overbrace{\longrightarrow\longrightarrow\longrightarrow\longrightarrow}^{\text{adaptive tempering+jiterring}} \quad \pi_t^N$$

- model reduction: coarsen the grid used for the numerical algorithm (the evolution of the particles) that approximates the dynamical system from $O(10^6)$ to $O(10^4)$.
- tempering: introduce a sequence of artificial densities, starting from a simple density and moving to the one of interest.
- jittering: reduce sample degeneracy through an MCMC procedure.
- nudging: correct the particle motion to bring them closer to the true state

Case study: two-layer quasi-geostrophic model for a β-plane channel flow with $O(10^6)$ degrees of freedom. The model is reduced by following the stochastic variational approach for geophysical fluid dynamics introduced in Holm[2015] as a framework for deriving stochastic parametrisations for unresolved scales.
Model reduction: We run the PF with $O(10^4)$ degrees of freedom.

$$\pi_{t-1}^N \overbrace{\longrightarrow\longrightarrow}^{\text{nudging}} p_t^N \quad \overbrace{\longrightarrow\longrightarrow\longrightarrow\longrightarrow}^{\text{adaptive tempering+jiterring}} \quad \pi_t^N$$

Example: 2-layer quasi-geostrophic model

The two-layer deterministic QG equations for the potential vorticity (PV) q:

$$\begin{aligned} \frac{\partial q_1}{\partial t} + \mathbf{u}_1 \cdot \nabla q_1 &= \nu \Delta^2 \psi_1 - \beta \frac{\partial \psi_1}{\partial x}, \\ \frac{\partial q_2}{\partial t} + \mathbf{u}_2 \cdot \nabla q_2 &= \nu \Delta^2 \psi_2 - \mu \Delta \psi_2 - \beta \frac{\partial \psi_2}{\partial x}, \end{aligned} \tag{1}$$

where ψ is the stream function, β is the planetary vorticity gradient, μ is the bottom friction parameter, ν is the lateral eddy viscosity, and $\mathbf{u} = (u, v)$ is the velocity vector. The computational domain $\Omega = [0, L_x] \times [0, L_y] \times [0, H]$ is a horizontally periodic flat-bottom channel of depth $H = H_1 + H_2$ given by two stacked isopycnal fluid layers of depth H_1 and H_2.

Forcing in (1) is introduced via a vertically sheared, baroclinically unstable background flow

$$\psi_i \to -U_i\, y + \psi_i, \quad i = 1, 2, \tag{2}$$

where the parameters U_i are background-flow zonal velocities. The PV anomaly and stream function are related through two elliptic equations:

$$q_1 = \Delta\psi_1 + s_1(\psi_2 - \psi_1), \tag{3a}$$

$$q_2 = \Delta\psi_2 + s_2(\psi_1 - \psi_2), \tag{3b}$$

with stratification parameters s_1, s_2. The system is augmented by the integral mass conservation constraint

$$\frac{\partial}{\partial t} \iint_{\Omega} (\psi_1 - \psi_2)\, dydx = 0, \tag{4}$$

by the periodic horizontal boundary conditions, $\boldsymbol{\psi}\Big|_{\Gamma_2} = \boldsymbol{\psi}\Big|_{\Gamma_4}, \quad \boldsymbol{\psi} = (\psi_1, \psi_2)$,

and no-slip boundary conditions $\boldsymbol{u}\Big|_{\Gamma_1} = \boldsymbol{u}\Big|_{\Gamma_3} = 0$ set at northern and southern boundaries of the domain.

Model reduction: The evolution of the particles satisfy a stochastic version of the QG equations (1) is given by:

$$
\begin{aligned}
dq_1 + \left(\mathbf{u}_1\, dt + \sum_{k=1}^{K} \xi_1^k \circ dW_t^k\right) \cdot \nabla q_1 &= \left(\nu \Delta^2 \psi_1 - \beta \frac{\partial \psi_1}{\partial x}\right) dt, \\
dq_2 + \left(\mathbf{u}_2\, dt + \sum_{k=1}^{K} \xi_2^k \circ dW_t^k\right) \cdot \nabla q_2 &= \left(\nu \Delta^2 \psi_2 - \mu \Delta \psi_2 - \beta \frac{\partial \psi_2}{\partial x}\right) dt.
\end{aligned}
\tag{5}
$$

The stochastic terms marked in red color is the only difference from the deterministic QG model (1), all other equations are the same as in the deterministic case.

- Stochastic solutions on signal grid $G_s = \{257 \times 129\}$.
- Observation Y velocity observed at on a grid $G_d = \{4 \times 8\}$.
- The size of the ensemble is taken to be $N = 100$ and the number of Brownian motion (independent sources of stochasticity) is taken to be $K = 32$. This is enough to reasonably quantify the uncertainty of the model: the spread of the ensemble will not increase substantially by taking more particles and/or sources of noise (BMs).

- The observations data Y_t is an 64-dimensional process that consists of noisy measurements of the velocity field **u** taken at a point belonging to the data grid G_d:
$$Y_t := \mathrm{P}_d^s(Z_t) + \eta,$$
where $\mathcal{P}_d^s : G_s \to G_d$ is a projection operator from the signal grid G_s to the data grid G_d, $\eta = \mathcal{N}(\mathbf{0}, I_\sigma)$ is a normally distributed random vector, with mean vector $\mathbf{0} = (0, \ldots, 0)$ and diagonal covariance matrix $I_\sigma = diag(\sigma_1^2, \ldots, \sigma_M^2)$.
- Rather than choosing an arbitrary $\sigma = (\sigma_1, \ldots, \sigma_M)$ for the standard deviation of the noise, we use the standard deviation of the velocity field computed over the coarse grid cell of the signal grid.

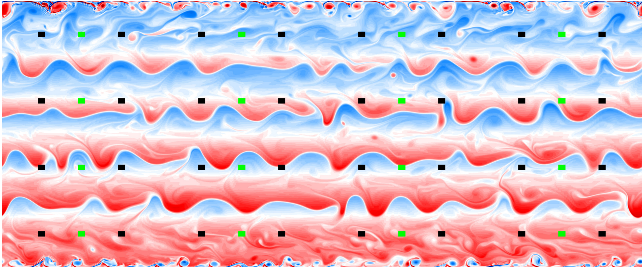

- We introduce the likelihood-weight function

$$\mathcal{W}(\mathbf{X},\mathbf{Y}) = \exp\left(-\frac{1}{2}\sum_{i=1}^{M}\left\|\frac{\mathrm{P}_d^s(X_i) - Y_i}{\sigma_i}\right\|_2^2\right), \qquad (6)$$

  with $M = 32$ being the number of grid points (weather stations).
- In order to measure the variability of the weights (6) of particles we use the effective sample size:

$$\mathrm{ESS}(\overline{\mathbf{w}}) = \left(\sum_{i=1}^{N}(\overline{w}_i)^2\right)^{-1}, \quad \overline{\mathbf{w}} := \mathbf{w}\left(\sum_{i=1}^{N} w_i\right)^{-1}, \qquad (7)$$

  which is close to the ensemble size N if the particles have weights that are close to each other, and decays to one, as the ensemble degenerates (i.e. there are fewer and fewer particles with large weights and the rest have small weights).
- One should resample for the weighted ensemble if the ESS drops below a given threshold, N^*,

$$\mathrm{ESS} < N^* = 80.$$

Nudging

- Nudging: a 'correction' applied to the state model evolution to bring the particles closer to the true state $\Rightarrow$ a nudging term $\mathcal{N}(\alpha)$ is added to the model
- $(t_{i-1}, t_i]$: Given the ensemble $\{X_n(t_{i-1})\}_{n=1}^N$ we want to assimilate observational data Y_{t_i} in order to obtain a new ensemble $\{X_n(t_i)\}_{n=1}^N$ that defines $\pi_{t_i}^N$:
  - Obtain observation Y_{t_i}
  - Evolve $X_n(t_{i-1}) \xrightarrow{modified \quad kernel} \tilde{X}_n(t_i)$

$$\tilde{X}_n(t_i) = X_n(t_{i-1}) + f(X_n(t_{i-1}), Y_{t_i}) + \sigma(X_n(t_{i-1})) B_{t_i}^n$$

  - Define new weights according to the ratio between the law of $\tilde{X}_n(t_i)$ and that of $X_n(t_i)$

We correct the solution of SPDE (5) so as to keep the particles closer to the true state. To do so, we add a 'nudging term' (marked in blue) to SPDE (5),

$$dq_i(\lambda) + \left(\mathbf{u}_i(\lambda)\,dt + \sum_{k=1}^{K} \xi_i^k \circ dW_t^k + \sum_{k=1}^{K} \xi_i^k \lambda_k\, dt\right) \cdot \nabla q_i(\lambda) = F_i\, dt, \quad i = 1, 2. \tag{8}$$

q depends on the parameter λ. The trajectories of the particles will be solutions of this perturbed SPDE (8). To account for the perturbation, the particles will have new weights given by

$$\exp\left(-\left(\left[\frac{1}{2}\sum_{i=1}^{M}\left\|\frac{P_d^s(\mathbf{q}_{t_{j+1}}(\lambda)) - \mathbf{Y}_{t_{j+1}}}{\sigma_i}\right\|_2^2 + \sum_k \int_{t_j}^{t_{j+1}} \left(\lambda_k^2 \frac{dt}{2} - \lambda_k dW_k\right)\right]\right)\right). \tag{9}$$

These weights measure the likelihood of the position of the particles given the observation, and the last term accounts for the change of probability distribution from q to $q(\lambda)$. We wish to choose λ so as to maximize these likelihoods.

In other words, we look to solve the equivalent minimization problem

$$\min_{\lambda_k,\ k\in[1..K]} \left[\frac{1}{2}\sum_{i=1}^{M} \left\| \frac{P_d^s(\mathbf{q}_{t_{j+1}}(\lambda)) - \mathbf{Y}_{t_{j+1}}}{\sigma_i} \right\|_2^2 + \int_{t_j}^{t_{j+1}} \left(\lambda_k^2 \frac{dt}{2} - \lambda_k dW_k \right) \right] \tag{10}$$

together with (8). In general this is a challenging nonlinear optimisation problem, especially if one allows the λ_k's to vary in time. For constant λ_k's the minimization problem (10) becomes

$$\min_{\lambda_k,\ k\in[1..K]} \left[\frac{1}{2}\sum_{i=1}^{M} \left\| \frac{P_d^s(\mathbf{q}_{t_{j+1}}(\lambda)) - \mathbf{Y}_{t_{j+1}}}{\sigma_i} \right\|_2^2 + \sum_{k=1}^{K} \left(\lambda_k^2 \frac{\delta t}{2} - \lambda_k \Delta W_k \right) \right], \tag{11}$$

where δt is the time step. Let us re-write

$$\mathbf{q}_{t_{j+1}}(\lambda) = A(q_{t_{j+1/2}}) + \sum_{k=1}^{K} B_k(\tilde{q}_{t_{j+1}})(\Delta W_k + \lambda_k \delta t),$$

where $q_{t_{j+1/2}}$ and $\tilde{q}_{t_{j+1}}$ are computed in the prediction and the extrapolation steps, respectively.

We can then re-write the minimisation problem (11) as

$$\min_{\lambda_k,\ k\in[1..K]} \mathcal{V}(\mathbf{q}(\lambda), \mathbf{Y}, \lambda), \tag{12}$$

where

$$\mathcal{V}(\mathbf{q}(\lambda), \mathbf{Y}, \lambda) = Q + Q_1(\lambda) + Q_2(\lambda, \Delta W_1, ..., \Delta W_K),$$

This is a quadratic minimization problem with the optimal value λ depending (linearly) on the increments $\Delta W_1, ..., \Delta W_K$. This optimal choice is not allowed as the parameter λ can only be a function of all the approximation $\tilde{q}_{t_{j+1}}$, $q_{t_{j+1/2}}$ and $Y_{t_{j+1}}$ (since it needs to be adapted to the forward filtration of the set of Brownian motions $\{W_k\}$). To ensure that this constraint is satisfied, we minimise *the conditional expectation* of $\mathcal{V}(\mathbf{q}(\lambda), \mathbf{Y}, \lambda)$ given the $\tilde{q}_{t_{j+1}}$, $q_{t_{j+1/2}}$ and $\mathbf{Y}_{t_{j+1}}$, that is

$$\min_{\lambda_k,\ k\in[1..K]} \mathbb{E}\left[\mathcal{V}(\mathbf{q}(\lambda), \mathbf{Y}, \lambda) | \tilde{q}_{t_{j+1}}, q_{t_{j+1/2}}, \mathbf{Y}_{t_{j+1}}\right].$$

This functional is quadratic in λ, and hence the optimization can be done by solving a linear system.

Important ! The resulting particle filter remains asymptotically consistent.

Twin experiment (signal is the solution of the SPDE run on the course grid)

Mean/Spread after combined N+T+J

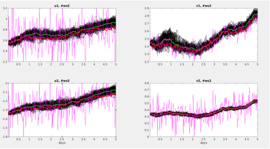

N+J compared with T+J

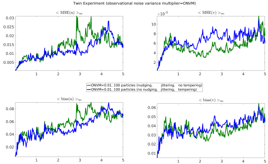

N+J (3 runs) compared with T+J

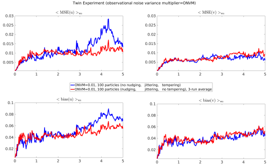

Outcome:

- Nudging reduces the MSE substantially.
- On their own, nudging and tempering have comparable results despite nudging being computationally cheaper.

True experiment (signal is the solution of the PDE run on the fine grid)

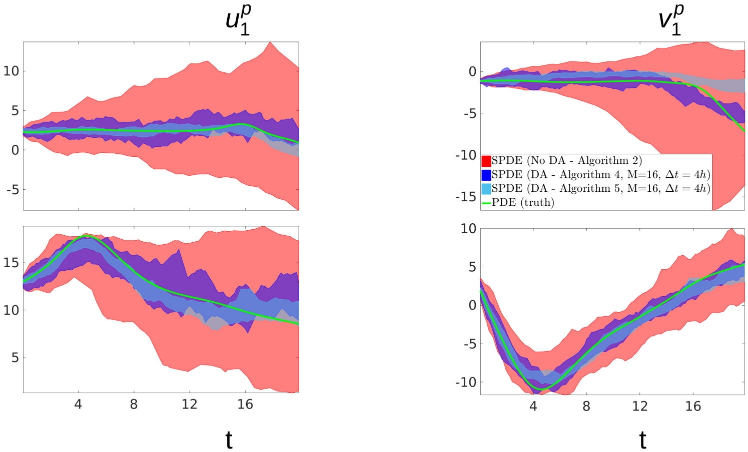

Figure: no DA vs DA (Te+Ji) vs Signal DA (Te+Ji+Nu)

Outcome: Nudging reduces the diameter of the cloud of particles.

Final remarks:

- We are developing particle filter based data assimilation methodology for high dimensional fluid dynamics models.
- The particle filter combines model reduction, tempering, jittering, and nudging and has been tested on the two layer quasi-geostrophic model with $O(10^6)$ degrees of freedom. Only a minute fraction of these are noisily observed.
- The nudging procedure brings improvements to the combinations of the tempering and jittering both in terms of the relative bias (RB) and ensemble mean relative l2-norm error.
- Further observation data sources are currently being explored:
  - Drifter Data (Lagrangian Data Assimilation)
  - Satelite observations https://www.youtube.com/watch?v=nTPTzHx3a74
  - Commercial aircraft data.

ERC Synergy grant 2020-2026 (https://www.imperial.ac.uk/ocean-dynamics-synergy/)

Stochastic Transport in Upper Ocean Dynamics

Bertrand Chapron
Dan Crisan
Darryl Holm
Etienne Mémin

Building new capabilities

- Decipher and forecast upper oceanic dynamics WP1+2+3+4
- Deliver a complete stochastic, data-driven theory for 3D-flows WP1+2+3+4

- Derive stochastic dynamics of the surface layer WP2+3
- Develop satellite and in situ data analysis techniques to calibrate our stochastic model WP1+2

STUOD Stochastic Transport in Upper Ocean Dynamics

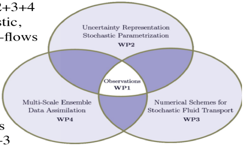

Ways to interact:

- Annual STUOD conferences
- Monthly sandboxes
- Annual Hackathons
- Postdocs/PhD students exchanges

Some of my papers on particle filters

Applications

• O. Lang, P. J. Van Leeuwen, D. Crisan, R. Potthast, Bayesian inference for fluid dynamics: A case study for the stochastic rotating shallow water model, Frontiers in Applied Mathematics and Statistics. 8, 2022.
• C Cotter, D Crisan, D Holm, W Pan, I Shevchenko, Data assimilation for a quasi-geostrophic model with circulation-preserving stochastic transport noise, Journal of Statistical Physics 179 (5), 1186-1221, 2020.
• C Cotter, D Crisan, DD Holm, W Pan, I Shevchenko, A Particle Filter for Stochastic Advection by Lie Transport: A Case Study for the Damped and Forced Incompressible Two-Dimensional Euler Equation, SIAM/ASA Journal on Uncertainty Quantification 8 (4), 1446-1492, 2020.
• C Cotter, D Crisan, D Holm, W Pan, I Shevchenko, Modelling uncertainty using stochastic transport noise in a 2-layer quasi-geostrophic model, Foundations of Data Science 2 (2), 173, 2020.
• C Cotter, D Crisan, DD Holm, W Pan, I Shevchenko, Numerically modeling stochastic Lie transport in fluid dynamics, Multiscale Modeling and Simulation 17 (1), 192-232, 2019.

General introduction and asymptotic consistency

• D Crisan, Particle filters - a theoretical perspective, Sequential Monte Carlo methods, 2001.
• D Crisan, A Doucet, A survey of convergence results on particle filtering methods for practitioners, IEEE Transactions on signal processing 50 (3), 736-746, 2002.
• A Bain, D Crisan, Fundamentals of stochastic filtering, Springer, 2008.

Stability in high-dimensions

• A Beskos, D Crisan, A Jasra, On the stability of sequential Monte Carlo methods in high dimensions, The Annals of Applied Probability 24 (4), 1396-1445, 2014.
•, A Beskos, D Crisan, A Jasra, K Kamatani, Y Zhou, A stable particle filter for a class of high-dimensional state-space models A Beskos, D Crisan, A Jasra, K Kamatani, Y Zhou, Advances in Applied Probability 49 (1), 24-48, 2017.
• H Ruzayqat, A Er-Raiy, A Beskos, D Crisan, A Jasra, N Kantas, A Lagged Particle Filter for Stable Filtering of certain High-Dimensional State-Space Models, arXiv:2110.00884, 2021.