Initial trials of convective-scale data assimilation with a cheaply tunable ensemble filter

Jonathan Flowerdew

Introduction

Ensemble methods are an attractive option for data assimilation (DA) in convection-permitting models, providing flow-dependent covariances which respect the complex balances that apply at these scales. This poster reports initial Met Office experiments in convective-scale ensemble DA (CsEnDA), exploring topics such as cycle length, filter type, localisation (Flowerdew, 2015) and inflation. All trials are one month long (June 2014).

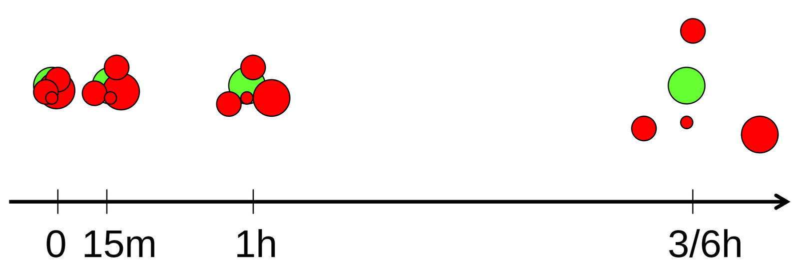

Figure 1: Short DA cycles are attractive to provide frequent updates to customers, particularly for very-short-range 'nowcasting' applications. They might also make better use of a finite ensemble. To assimilate the location of a feature such as a convective cell, we require multiple members with features that overlap the observed feature. Because ensemble spread tends to increase with lead time, this becomes more likely at shorter lead times. Short cycles may also improve linearity, reduce localisation problems associated with feature movement during the time window, and provide more nudges within the life cycle of the features being assimilated.

EnKF with analysis error diagnostic

We use the ensemble square root filter (EnSRF; Whitaker and Hamill, 2002), and also the serial perturbed observations filter. Gaspari-Cohn localisation is applied as a function of horizontal distance and vertical difference in the logarithm of pressure. The model equivalents of the observations are updated as additional elements of the state vector. This motivates a powerful diagnostic that is only available in serial filters. As usual, we can calculate the root mean square (RMS) innovation before any observations have been assimilated:

$$I_0^2 = \frac{1}{N_{obs}} \sum_{i=1}^{N_{obs}} (y_i - \bar{z}_{i,0})^2 / R_i,$$

where y_i is the ith observation, R_i is its estimated error variance, and $\bar{z}_{i,0}$ is the mean of the initial ensemble equivalents of the observed quantity. We can also calculate the RMS innovation just before each observation is assimilated:

$$I_{APO}^2 = \frac{1}{N_{obs}} \sum_{i=1}^{N_{obs}} (y_i - \bar{z}_{i,i-1})^2 / R_i,$$

where $\bar{z}_{i,i-1}$ is the model equivalent of the ith observation after applying the updates from all previous observations. I_{APO} is the RMS Innovation After assimilating Prior Observations (IAPO). Provided the observation errors are independent, this provides a direct verification of analysis quality without the need to involve the forecast model, and without the statistical problems associated with a 'fit' to observations which have already been assimilated.

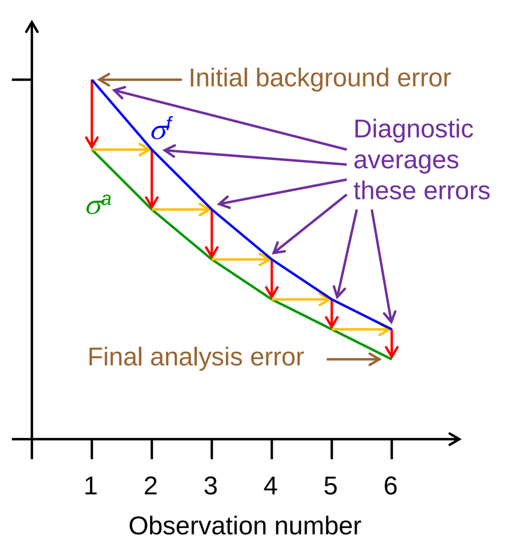

Figure 2: A schematic of the IAPO diagnostic. The initial background error is reduced by assimilating the first observation (first downward arrow); this updated state forms the background for the second observation (first rightward arrow), and so on. The diagnostic finds the RMS value of the innovations (normalised by the corresponding observation error estimates) just before each observation is assimilated, and therefore sees the impact of assimilating the preceding observations.

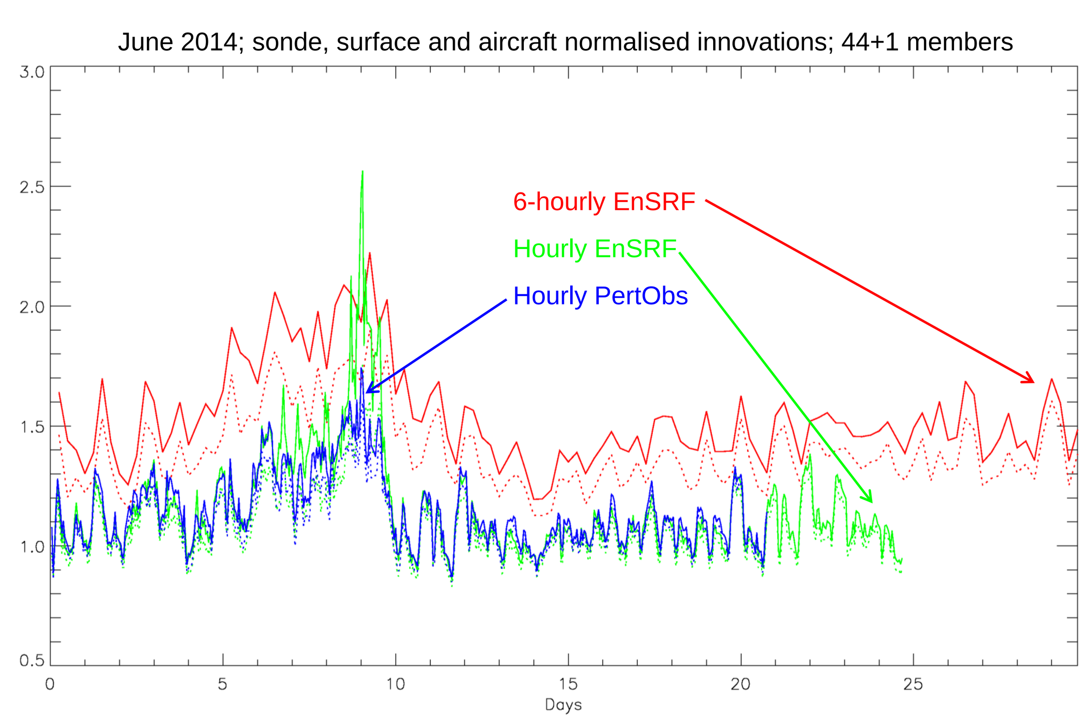

Figure 3: RMS normalised ensemble mean innovations before (solid) and after (dotted) assimilating prior observations, as a function of data time during the first three CsEnDA trials (colours), which assimilate sonde, surface and aircraft data. The IAPO provides immediate confirmation that the assimilation is 'working' in its own statistical terms, drawing towards observations which have yet to be assimilated. The hourly assimilations start from lower background errors as one would expect from the shorter lead time; the benefit of assimilation is smaller than in the 6h system, but recall there are six such increments for every 6h increment. The perturbed observations filter produces slightly larger errors that the EnSRF under normal circumstances, but is more robust in the period of larger errors around 5–10 days, which is consistent with theoretical expectations for 'deterministic' and 'stochastic' filters.

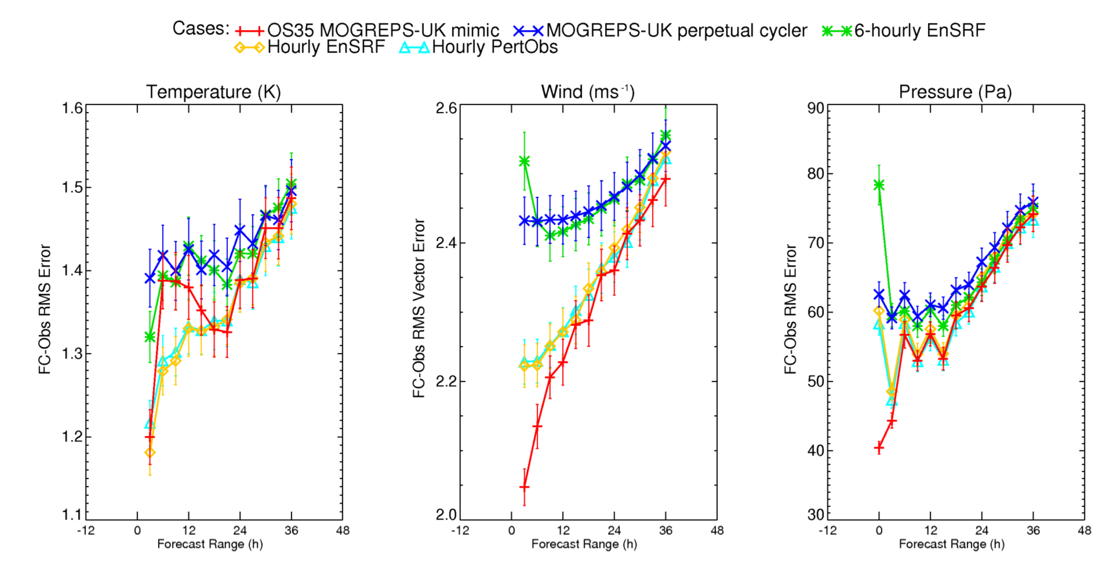

Figure 4: RMS difference between the unperturbed control forecast (updated using the same Kalman gain as the ensemble mean) and surface observations over the British Isles, as a function of lead time for the initial UK trials. Until recently, the Met Office UK ensemble (MOGREPS-UK) was a simple 'downscaler', interpolating each initial state from the corresponding member of the global ensemble (MOGREPS-G). The performance of this system is shown in red. The blue lines show a 'continuous cycler', where the initial states are taken from the previous cycle, with only the lateral boundaries being updated from the corresponding global member. This performs much worse than the downscaler, presumably due to the loss of the global DA in the interior of the domain. It is equivalent to CsEnDA with no observations, and thus provides the most direct baseline on which the CsEnDA should improve. Six-hourly assimilation (green) generally improves upon the continuous cycler, but not by much. By contrast, the hourly systems (yellow and cyan) achieve much superior performance from essentially the same total set of observations.

A cheap way to tune the EnKF

The IAPO diagnostic provides a cheap, independent measure of analysis error. This suggests it can be used to tune any parameter that directly affects analysis quality, such as localisation factors, observation error estimates, ensemble size, etc. This was tested by running all combinations of five horizontal and six vertical localisation radii against input data archived from one of the main trials every 1.25d. Whereas it would be completely impractical to run 30 month-long trials, these EnKF runs complete overnight and the core calculations could be performed in a matter of minutes. Based on results like Table 1, the horizontal and vertical localisation scales were changed from 333 km and 1.0 to 222 km and 0.406 respectively. The even tighter vertical localisation suggested by Table 1 was avoided for reasons such as concern over model balance, although later results suggest it does indeed bring further benefit.

V \ H (km)	333	222	111	56
1.000	1.1342	1.1339	1.1362	1.1398
0.406	1.1323	1.1323	1.1353	1.1393
0.214	1.1305	1.1310	1.1347	1.1391

Table 1: RMS IAPO from the EnKF tuner suite applied to input from the 19 June 21 UTC cycle of the Hourly PertObs trial. The horizontal/vertical localisation for each configuration uses a Gaspari-Cohn function where the c parameter is set to the value given in the corresponding column/row header. Although these results come from a single run, they show smooth variation with clean minima in each row/column. This suggests the signal-to-noise ratio is sufficiently good that some tuning could ultimately be built into the EnKF itself, adapting the assimilation parameters to the particular background fields and observations present on each individual cycle.

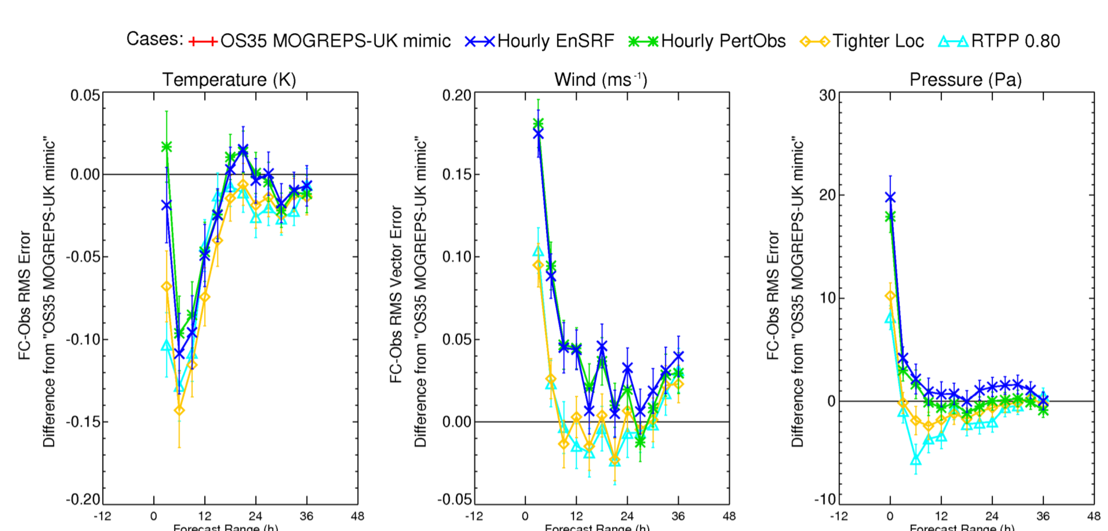

Figure 5: As Figure 4, but showing the difference between each RMS error and the corresponding result for the downscaler. Tighter localisation (yellow) improves all three variables, generally beating the downscaler. Changing from multiplicative inflation to Relaxation to Prior Perturbations (RTPP; cyan) improves pressure, with further benefits to the magnitude and distribution of spread (not shown).

Cloud-affected satellite radiances (with Pete Weston)

Cloud is a key forecast variable, linked to convective-scale features. The assimilation of cloud-affected satellite radiances also probes several important methodological issues:

- They are horizontally dense, so may require tighter horizontal localisation than conventional observations.
- The are vertically integrating, so may require broader vertical localisation than point observations.
- They measure nonlinear hydrometeor variables.

Description	Ctrl ID	Overall change	Cloud amount	Cloud base height
1. Conventional obs	N/A	N/A	N/A	N/A
2. Increment cloud water/ice	1	-0.44%	-0.056	-0.015
3. Only channel 5	2	+0.97%	+0.004	+0.029
4. Channels 5 & 9	2	-2.19%	-0.045	-0.045

Table 2: Trials assimilating data from the SEVIRI infrared geostationary imager. The third column shows the overall percentage change in a weighted sum of Equitable Threat Scores (ETS) for cloud amount, cloud base height, 6h-accumulated precipitation and visibility. Starting from the conventional observations RTPP trial, trial number 2 increments the cloud water and ice fields in addition to the standard dynamic variables, without adding any new observations. This is detrimental, suggesting the need for inter-variable localisation to reduce the impact of spurious inter-variable correlations. Adding SEVIRI channel 5 (which measures upper tropospheric humidity and thus has a weighting function that is fairly consistent between cases) brings a good positive benefit. Adding channel 9 (which measures the cloud top, or surface if there is no cloud) is detrimental. This probably indicates a need for more advanced vertical localisation where the peak moves up and down with the observation.

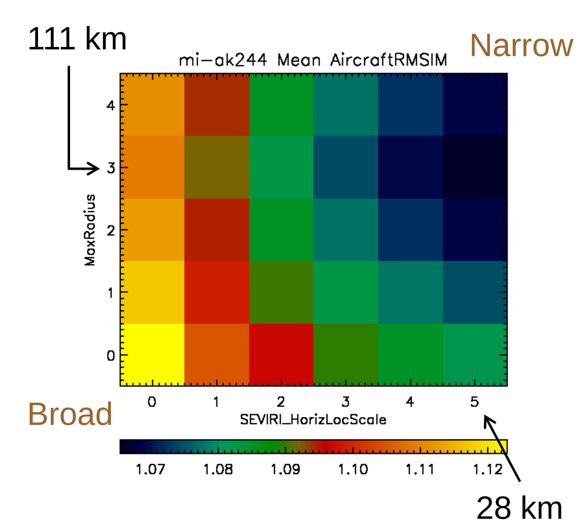

Figure 6: The mean over 15 cases spread throughout the month of I_{APO} for aircraft observations, as a function of the horizontal localisation scales applied to conventional and SEVIRI data (rows and columns respectively). The observation order has been randomised to ensure any subset of observations sees the impact of any other set of observations. The best configuration localises SEVIRI data about a factor of four more tightly than conventional data. Similar results support SEVIRI vertical localisation that is broadened with respect to point quantities to an extent consistent with the typical width of the observation operator.

Summary and future work

The serial ensemble filter is a promising technique for convective-scale DA. The best configurations assimilating sonde, surface and aircraft data are competitive with a downscaler, even though the latter benefits from global DA that has seen a far wider range of observation types.

The serial filter provides the unique ability to measure the error of the developing analysis against observations which have not yet been assimilated. This IAPO diagnostic can be used to quickly check an assimilation is drawing towards independent observations. It provides a cheap way to tune parameters such as localisation radii that directly affect analysis quality. It can also be used to make some inferences about the ensemble spread, for instance confirming that even the best configurations considered here are generally underspread (since the measured error reduction is greater than the initial spread).

Future work could extend these trials to cover a wider range of topics, including radar data, the handling of lateral boundaries, comparison to local variational assimilation, and further development of cloudy radiance assimilation. These results will then feed into the decision of what to implement operationally, beyond more immediate plans such as hourly UK 4DVar.

References

Flowerdew J. 2015. Towards a theory of optimal localisation. *Tellus A* **67**: 25257.

Whitaker JS, Hamill TM. 2002. Ensemble data assimilation without perturbed observations. *Mon. Weather Rev.* **130**: 1913–1924.

Met Office FitzRoy Road, Exeter, Devon, EX1 3PB United Kingdom
Tel: +44 1392 885680 Fax: +44 1392 885681
Email: jonathan.flowerdew@metoffice.gov.uk