EnKF and Catastrophic filter divergence

David Kelly Andrew Stuart Kody Law

Mathematics Institute
University of Warwick
Coventry UK CV4 7AL

dtbkelly@gmail.com

November 3, 2013

DAday, November 2013, UNC Chapel Hill.

Outline

- What is the **motivation** behind EnKF?
- What can we **prove** about EnKF?

The filtering problem

We have a **deterministic model**

$$\frac{dv}{dt} = F(v) \quad \text{with } v_0 \sim N(m_0, C_0)\ .$$

We will denote $v(t) = \Psi_t(v_0)$. Think of this as **very high dimensional** and **nonlinear**.

We want to **estimate** $v_j = v(jh)$ for some $h > 0$ and $j = 0, 1, \ldots, J$ given the **observations**

$$y_j = Hv_j + \xi_j \quad \text{for } \xi_j \text{ iid } N(0, \Gamma).$$

We can write down the condtional density using **Bayes' formula** ...

But it's horrible.

For linear models, one can draw samples, using the **Randomized Maximum Likelihood** method.

RML method

Let $u \sim N(\widehat{m}, \widehat{C})$ and $\eta \sim N(0, \Gamma)$. We make an observation

$$y = H(u) + \eta \, .$$

We want the conditional distribution of u given y. This is called an **inverse problem**.

For linear models, RML takes a sample

$$\{\widehat{u}^{(1)}, \dots, \widehat{u}^{(K)}\} \sim N(\widehat{m}, \widehat{C})$$

and turns them into a sample

$$\{u^{(1)}, \dots, u^{(K)}\} \sim u | y$$

RML method: How does it work?

Along with the prior sample $\{\widehat{u}^{(1)}, \dots, \widehat{u}^{(K)}\}$, we create **artificial observations** $\{y^{(1)}, \dots, y^{(K)}\}$ where

$$y^{(k)} = y + \eta^{(k)} \qquad \text{where } \eta^{(k)} \sim N(0, \Gamma) \text{ i.i.d}$$

Then define $u^{(k)}$ using the **Bayes formula** update, with $(\widehat{u}^{(k)}, y^{(k)})$

$$u^{(k)} = \widehat{u}^{(k)} + G(\widehat{u})(y^{(k)} - H\widehat{u}^{(k)}) \ .$$

Where the "Kalman Gain" $G(\widehat{u})$ is computing using the **real** covariance of the prior $\widehat{u}$.

EnKF uses the same method, but with an **approximation** of the covariance in the Kalman gain.

The set-up for EnKF

Suppose we are given the ensemble $\{u_j^{(1)}, \dots, u_j^{(K)}\}$ at time j. For each particle, we create an **artificial observation**

$$y_{j+1}^{(k)} = y_{j+1} + \xi_{j+1}^{(k)} \quad , \quad \xi_{j+1}^{(k)} \text{ iid } N(0, \Gamma).$$

We update each particle using the **Kalman update**

$$u_{j+1}^{(k)} = \Psi_h(u_j^{(k)}) + G(u_j) \left(y_{j+1}^{(k)} - H\Psi_h(u_j^{(k)}) \right) \; ,$$

where $G(u_j)$ is the **Kalman gain** computed using the **forecasted ensemble covariance**

$$\widehat{C}_{j+1} = \frac{1}{K} \sum_{k=1}^{K} (\Psi_h(u_j^{(k)}) - \overline{\Psi_h(u_j)})^T (\Psi_h(u_j^{(k)}) - \overline{\Psi_h(u_j)}) \; .$$

There aren't many **theorems** about EnKF.

Filter divergence

It has been observed ($\star$) that the ensemble can **blow-up** (ie. reach machine-infinity) in **finite time**, even when the model has nice bounded solutions.

This is known as **catastrophic filter divergence**.

It is suggested in ($\star$) that this is caused by numerically integrating a stiff-system. Our aim is to "prove" this.

$\star$ Harlim, Majda (2010), Gottwald (2011), Gottwald, Majda (2013).

Assumptions (†)

1 - We make a **dissipativity** assumption on F. Namely that

$$F(\cdot) = A \cdot + B(\cdot, \cdot)$$

with A linear elliptic and B bilinear, satisfying certain estimates and symmetries.

Eg. 2d-Navier-Stokes, Lorenz-63, Lorenz-96.

2 - The **observation operator** $H = Id$ and the **noise covariance** $\Gamma = \gamma Id$

Discrete time results

For a fixed time-step $h > 0$ we can prove

Theorem (AS,DK)

If (†) then there exists constant β such that

$$\mathbf{E}|u_j^{(k)}|^2 \leq e^{2\beta jh}\mathbf{E}|u_0^{(k)}|^2 + 2K\gamma^2 \left(\frac{e^{2\beta jh} - 1}{e^{2\beta h} - 1}\right)$$

Rmk. This becomes useless as $h \to 0$

For observations with $h \ll 1$, we need another approach.

The EnKF equations look like a discretization

Recall the ensemble update equation

$$
\begin{aligned}
u_{j+1}^{(k)} &= \Psi_h(u_j^{(k)}) + G(u_j)\left(y_{j+1}^{(k)} - H\Psi_h(u_j^{(k)})\right) \\
&= \Psi_h(u_j^{(k)}) + \widehat{C}_{j+1}H^T(H^T\widehat{C}_{j+1}H + \Gamma)^{-1}\left(y_{j+1}^{(k)} - H\Psi_h(u_j^{(k)})\right)
\end{aligned}
$$

Subtract $u_j^{(k)}$ from both sides and divide by h

$$
\begin{aligned}
\frac{u_{j+1}^{(k)} - u_j^{(k)}}{h} &= \frac{\Psi_h(u_j^{(k)}) - u_j^{(k)}}{h} \\
&\quad + \widehat{C}_{j+1}H^T(hH^T\widehat{C}_{j+1}H + h\Gamma)^{-1}\left(y_{j+1}^{(k)} - H\Psi_h(u_j^{(k)})\right)
\end{aligned}
$$

Clearly we need to rescale the noise (ie. Γ).

Continuous-time limit

If we set $\Gamma = h^{-1}\Gamma_0$ and substitute $y_{j+1}^{(k)}$, we obtain

$$\frac{u_{j+1}^{(k)} - u_j^{(k)}}{h} = \frac{\Psi_h(u_j^{(k)}) - u_j^{(k)}}{h} + \widehat{C}_{j+1} H^T (h H^T \widehat{C}_{j+1} H + \Gamma_0)^{-1}$$
$$\Big(Hv + h^{-1/2}\Gamma_0^{1/2}\xi_{j+1} + h^{-1/2}\Gamma_0^{1/2}\xi_{j+1}^{(k)} - H\Psi_h(u_j^{(k)}) \Big)$$

But we know that

$$\Psi_h(u_j^{(k)}) = u_j^{(k)} + O(h)$$

and

$$\widehat{C}_{j+1} = \frac{1}{K}\sum_{k=1}^{K} (\Psi_h(u_j^{(k)}) - \overline{\Psi_h(u_j)})^T (\Psi_h(u_j^{(k)}) - \overline{\Psi_h(u_j)})$$
$$= \frac{1}{K}\sum_{k=1}^{K} (u_j^{(k)} - \overline{u_j})^T (u_j^{(k)} - \overline{u_j}) + O(h) = C(u_j) + O(h)$$

Continuous-time limit

We end up with

$$\frac{u_{j+1}^{(k)} - u_j^{(k)}}{h} = \frac{\Psi_h(u_j^{(k)}) - u_j^{(k)}}{h} - C(u_j)H^T\Gamma_0^{-1}H(u_j^{(k)} - v_j)$$
$$+ C(u_j)H^T\Gamma_0^{-1}\left(h^{-1/2}\xi_{j+1} + h^{-1/2}\xi_{j+1}^{(k)}\right) + O(h)$$

This looks like a **numerical scheme** for

$$\frac{du^{(k)}}{dt} = F(u^{(k)}) - C(u)H^T\Gamma_0^{-1}H(u^{(k)} - v) \qquad (\bullet)$$
$$+ C(u)H^T\Gamma_0^{-1/2}\left(\frac{dW^{(k)}}{dt} + \frac{dB}{dt}\right) .$$

Properties of the limiting equation

$$\frac{du^{(k)}}{dt} = F(u^{(k)}) - C(u)H^T\Gamma_0^{-1}H(u^{(k)} - v) \qquad (\bullet)$$
$$+ C(u)H^T\Gamma_0^{-1/2}\left(\frac{dW^{(k)}}{dt} + \frac{dB}{dt}\right) .$$

1 - Extra dissipation term only sees **differences in observed space**

2 - Extra dissipation only occurs in the **space spanned by ensemble**

3 - If F were linear, the equation for the mean is the equation for the **classical Kalman filter**, with an added sampling error.

Continuous-time results

Theorem (AS,DK)

Suppose the framework satisfies (†) and $\{u^{(k)}\}_{k=1}^{K}$ satisfy (•). Let

$$e^{(k)} = u^{(k)} - v \ .$$

Then there exists constant β such that

$$\mathbf{E} \sum_{k=1}^{K} |e^{(k)}(t)|^2 \leq \left(\mathbf{E} \sum_{k=1}^{K} |e^{(k)}(0)|^2 \right) \exp\left(\beta t\right) \ .$$

Summary + Future Work

(1) Writing down an SDE/SPDE allows us to see the **important quantities** in the algorithm.

(2) Does not "prove" that catastrophic filter divergence is a numerical phenomenon, but is a decent starting point.

(1) Improve the condition on H.

(2) If we can **measure** the important quantities, then we can test the performance during the algorithm.

(3) Suggests new EnKF-like algorithms, for instance by discretising the stochastic PDE in a more **numerically stable** way.