EnKF and filter divergence

David Kelly Andrew Stuart Kody Law

Courant Institute
New York University
New York, NY

dtbkelly.com

December 12, 2014

Applied and computational mathematics seminar, **NIST**.

Talk outline

1. What is EnKF?

2. What is known about EnKF?

3. How can we use stochastic analysis to better understand 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 conditional density using **Bayes' formula** ...

But for high dimensional nonlinear systems it's horrible.

Bayes' formula filtering update

Let $Y_j = \{y_0, y_1, \ldots, y_j\}$. We want to compute the conditional density $\mathbf{P}(v_{j+1}|Y_{j+1})$, using $\mathbf{P}(v_j|Y_j)$ and y_{j+1}.

By Bayes' formula, we have

$$\mathbf{P}(v_{j+1}|Y_{j+1}) = \mathbf{P}(v_{j+1}|Y_j, y_{j+1}) \propto \mathbf{P}(y_{j+1}|v_{j+1})\mathbf{P}(v_{j+1}|Y_j)$$

But we need to compute the integral

$$\mathbf{P}(v_{j+1}|Y_j) = \int \mathbf{P}(v_{j+1}|Y_j, v_j)\mathbf{P}(v_j|Y_j)dv_j \ .$$

For high dimensional nonlinear systems, this is computationally infeasible.

The **Ensemble Kalman Filter** (EnKF) is a lower dimensional algorithm. (Evensen '94)

EnKF generates an ensemble of **approximate samples** from the posterior.

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

$$y = Hu + \eta \, .$$

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

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 covariance of the prior $\widehat{u}$.

The set $\{u^{(1)}, \dots, u^{(K)}\}$ are exact samples from $u|y$.

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 ensemble member, 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**

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

What do we know about EnKF?
Not much.

Theorem : For linear forecast models,
ENKF $\to$ KF as $N \to \infty$

(Le Gland et al / Mandel et al. 09').

Ideally, we would like a theorem about **long time behaviour** of the filter for a finite ensemble size.

Filter divergence

In certain situations, 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**.

We would like to investigate whether this has a **dynamical justification** or if it is simply a **numerical artefact**.

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

Assumptions on the dynamics

We make a **dissipativity** assumption on the model. Namely that

$$\frac{dv}{dt} + Av + B(v, v) = f$$

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

This guarantees **uniformly bounded** solutions.

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

Discrete time results

For a fixed observation frequency $h > 0$ we can prove

Theorem (AS,DK,KL)

If $H = \Gamma = Id$ then there exists constant $\beta > 0$ 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 $h \to 0$

Discrete time results with variance inflation

Suppose we replace

$$\widehat{C}_{j+1} \mapsto \alpha^2 I + \widehat{C}_{j+1}$$

at each update step. This is known as **additive variance inflation**.

Theorem (AS,DK,KL)

If $H = Id$ and $\Gamma = \gamma^2 Id$ then there exists constant $\beta > 0$ such that

$$\mathbf{E}|e_j^{(k)}|^2 \leq \theta^j \mathbf{E}|e_0^{(k)}|^2 + 2K\gamma^2 \left(\frac{1-\theta^j}{1-\theta}\right)$$

where $\theta = \frac{\gamma^2 e^{2\beta h}}{\alpha^2+\gamma^2}$. In particular, if we pick α large enough (so that $\theta < 1$) then

$$\lim_{j\to\infty} \mathbf{E}|e_j^{(k)}|^2 \leq \frac{2K\gamma^2}{1-\theta}$$

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 **Itô S(P)DE**

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

Nudging

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

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

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

Kalman-Bucy limit

If F were **linear** and we write $m(t) = \frac{1}{K}\sum_{k=1}^{K} u^{(k)}(t)$ then

$$\frac{dm}{dt} = F(m) - C(u)H^T\Gamma_0^{-1}H(m - v) + C(u)H^T\Gamma_0^{-1/2}\frac{dB}{dt} + O(K^{-1/2}) \,.$$

This is the equation for the **Kalman-Bucy** filter, with empirical covariance $C(u)$. The remainder $O(K^{-1/2})$ can be thought of as a **sampling error**.

Continuous-time results

Theorem (AS,DK)

Suppose that $\{u^{(k)}\}_{k=1}^{K}$ *satisfy* ($\bullet$) *with* $H = \Gamma = Id$. *Let*

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

Then there exists constant $\beta > 0$ *such that*

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

Why do we need $H = \Gamma = Id$?

In the equation

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

The **energy** pumped in by the noise must be balanced by **contraction** of $(u^{(k)} - v)$. So the operator

$$C(u)H^T\Gamma_0^{-1}H$$

must be **positive-definite**.

Both $C(u)$ and $H^T\Gamma_0^{-1}H$ are pos-def, but this doesn't guarantee the same for the **product**!

Testing stability on the fly

Suppose we can actually measure the spectrum of the operator

$$C(u)H^T\Gamma_0^{-1}H$$

whilst the algorithm is running. If we know that it is pos-def, then the filter must not be blowing up.

If we knew that

$$C(u)H^T\Gamma_0^{-1}H \geq \lambda(t) > 0 \, .$$

Then we can say even more (eg. stability).

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) Make use of **controllability** and **observability**.

Thank you!

Well-posedness and accuracy of the ensemble Kalman filter in discrete and continuous time.

D. Kelly, K.Law, A. Stuart.

Nonlinearity 2014.

www.dtbkelly.com