Introduction to Kalman filtering

Emma Horton
University of Melbourne

Turin-Bath PhD Workshop in Applied Probability and Statistics
19-21 June 2023

Part I: The Kalman filter

Motivating example: target tracking

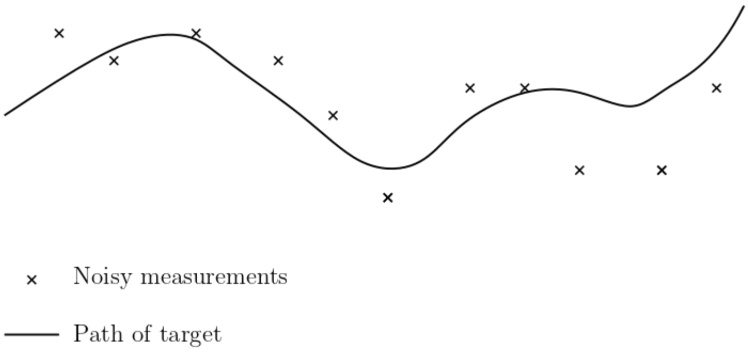

Motivating example: target tracking

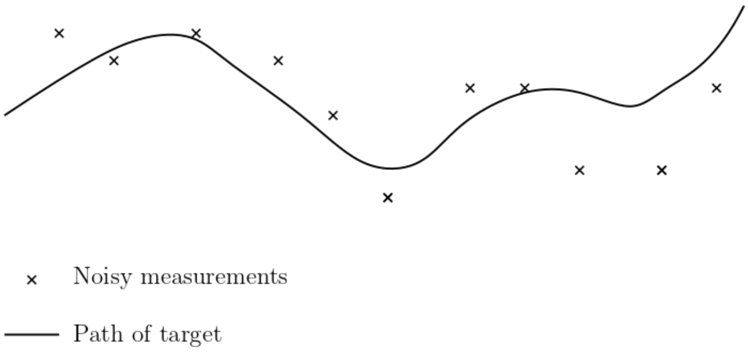

$\rightsquigarrow$ What is the best way to combine the model and measurements to estimate the path of the target?

Motivating example: radiotherapy

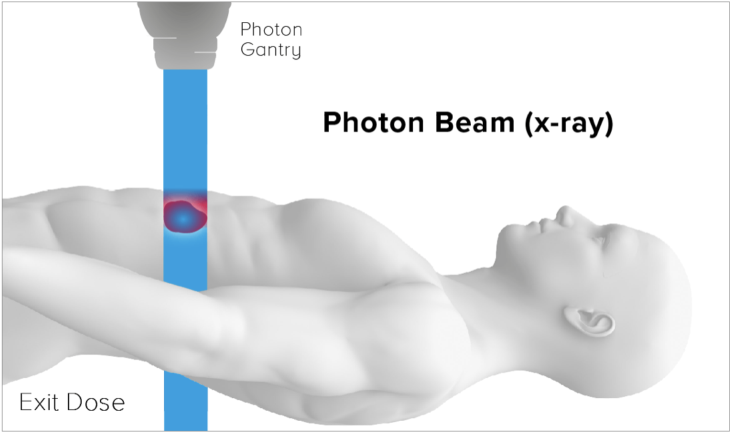

Motivating example: radiotherapy

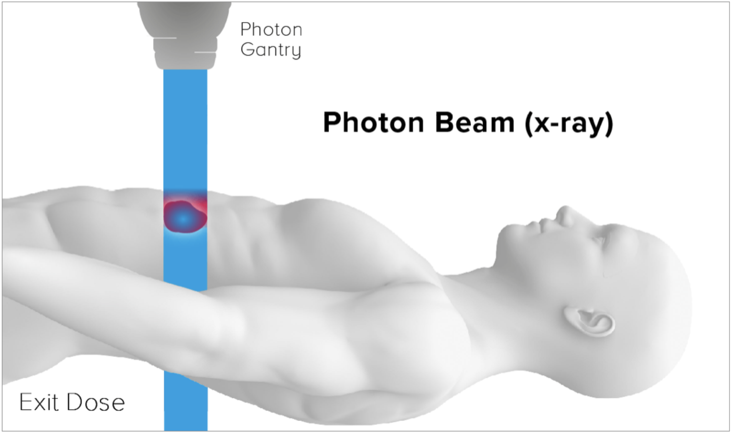

⇝ What is the best way to combine the model and measurements to estimate the dose?

The problem

Consider the following one-dimensional model,

$$X_{n+1} = AX_n + BW_{n+1}, \quad n \geq 1,$$

with noisy measurements

$$Y_n = CX_n + DV_n, \quad n \geq 0,$$

where

- $X_0 \sim \mathcal{N}(\widehat{X}_0^-, \widehat{P}_0^-)$,
- $V_n, W_{n+1} \sim \mathcal{N}(0, 1)$ are independent,
- $A, B, C, D \neq 0$.

The problem

Consider the following one-dimensional model,

$$X_{n+1} = AX_n + BW_{n+1}, \quad n \geq 1,$$

with noisy measurements

$$Y_n = CX_n + DV_n, \quad n \geq 0,$$

where

- $X_0 \sim \mathcal{N}(\widehat{X}_0^-, \widehat{P}_0^-)$,
- $V_n, W_{n+1} \sim \mathcal{N}(0, 1)$ are independent,
- $A, B, C, D \neq 0$.

$\rightsquigarrow$ Aim: to compute the distribution of X_n given the measurements, $Y_0, \ldots, Y_n$.

Some observations

- Write $x_{0:n}$ for the tuple $(x_0, \ldots, x_n)$ and similarly for y.
- From Bayes' rule, the Markov property and the fact that the errors are independent, we have

$$p(x_{0:n}|y_{0:n}) \propto p(y_{0:n}|x_{0:n})p(x_{0:n})$$

Some observations

- Write $x_{0:n}$ for the tuple $(x_0, \ldots, x_n)$ and similarly for y.
- From Bayes' rule, the Markov property and the fact that the errors are independent, we have

$$
\begin{aligned}
p(x_{0:n}|y_{0:n}) &\propto p(y_{0:n}|x_{0:n})p(x_{0:n}) \\
&= \prod_{k=0}^{n} p(y_k|x_k)p(x_0)\prod_{k=1}^{n} p(x_k|x_{k-1})
\end{aligned}
$$

Some observations

- Write $x_{0:n}$ for the tuple $(x_0, \ldots, x_n)$ and similarly for y.
- From Bayes' rule, the Markov property and the fact that the errors are independent, we have

$$
\begin{aligned}
p(x_{0:n}|y_{0:n}) &\propto p(y_{0:n}|x_{0:n})p(x_{0:n}) \\
&= \prod_{k=0}^{n} p(y_k|x_k)p(x_0)\prod_{k=1}^{n} p(x_k|x_{k-1}) \\
&= p(x_0)p(y_0|x_0)\prod_{k=1}^{n} p(y_k|x_k)p(x_k|x_{k-1})
\end{aligned}
$$

Some observations

- Write $x_{0:n}$ for the tuple $(x_0, \dots, x_n)$ and similarly for y.
- From Bayes' rule, the Markov property and the fact that the errors are independent, we have

$$
\begin{aligned}
p(x_{0:n}|y_{0:n}) &\propto p(y_{0:n}|x_{0:n})p(x_{0:n}) \\
&= \prod_{k=0}^{n} p(y_k|x_k)p(x_0)\prod_{k=1}^{n} p(x_k|x_{k-1}) \\
&= p(x_0)p(y_0|x_0)\prod_{k=1}^{n} p(y_k|x_k)p(x_k|x_{k-1}) \\
&= p(x_0)p(y_0 - Cx_0)\prod_{k=1}^{n} p(y_k - Cx_k)p(x_k|x_{k-1})
\end{aligned}
$$

Some observations

- Recall that $X_0 \sim \mathcal{N}(\widehat{X}_0^-, \widehat{P}_0^-)$ and the model

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1},$$

- Then $X_n | \mathcal{Y}_{n-1} \sim \mathcal{N}(\widehat{X}_n^-, \widehat{P}_n^-)$, where

$$\widehat{X}_n^- = \mathbb{E}[X_n | \mathcal{Y}_{n-1}], \quad \widehat{P}_n^- = \mathbb{E}[(X_n - \widehat{X}_n^-)^2].$$

- Similarly $X_n | \mathcal{Y}_n \sim \mathcal{N}(\widehat{X}_n, \widehat{P}_n)$, where

$$\widehat{X}_n = \mathbb{E}[X_n | \mathcal{Y}_n], \quad \widehat{P}_n = \mathbb{E}[(X_n - \widehat{X}_n)^2].$$

- $\rightsquigarrow$ Just need to know how to calculate the mean and covariance.

Some observations

- Recall that $X_0 \sim \mathcal{N}(\widehat{X}_0^-, \widehat{P}_0^-)$ and the model

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1},$$

- Then $X_n|\mathcal{Y}_{n-1} \sim \mathcal{N}(\widehat{X}_n^-, \widehat{P}_n^-)$, where

$$\widehat{X}_n^- = \mathbb{E}[X_n|\mathcal{Y}_{n-1}], \quad \widehat{P}_n^- = \mathbb{E}[(X_n - \widehat{X}_n^-)^2].$$

- Similarly $X_n|\mathcal{Y}_n \sim \mathcal{N}(\widehat{X}_n, \widehat{P}_n)$, where

$$\widehat{X}_n = \mathbb{E}[X_n|\mathcal{Y}_n], \quad \widehat{P}_n = \mathbb{E}[(X_n - \widehat{X}_n)^2].$$

- $\rightsquigarrow$ Just need to know how to calculate the mean and covariance.

Some observations

- Recall that $X_0 \sim \mathcal{N}(\widehat{X}_0^-, \widehat{P}_0^-)$ and the model

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1},$$

- Then $X_n | \mathcal{Y}_{n-1} \sim \mathcal{N}(\widehat{X}_n^-, \widehat{P}_n^-)$, where

$$\widehat{X}_n^- = \mathbb{E}[X_n | \mathcal{Y}_{n-1}], \quad \widehat{P}_n^- = \mathbb{E}[(X_n - \widehat{X}_n^-)^2].$$

- Similarly $X_n | \mathcal{Y}_n \sim \mathcal{N}(\widehat{X}_n, \widehat{P}_n)$, where

$$\widehat{X}_n = \mathbb{E}[X_n | \mathcal{Y}_n], \quad \widehat{P}_n = \mathbb{E}[(X_n - \widehat{X}_n)^2].$$

- $\rightsquigarrow$ Just need to know how to calculate the mean and covariance.

The Kalman filter

The Kalman filter consists of two steps (update and predict)

$$(\widehat{X}_n^-, \widehat{P}_n^-) \longrightarrow (\widehat{X}_n, \widehat{P}_n) \longrightarrow (\widehat{X}_{n+1}^-, \widehat{P}_{n+1}^-).$$

1. Update

$$\widehat{X}_n = (1 - G_n C)\widehat{X}_n^- + G_n Y_n = \widehat{X}_n^- + G_n(Y_n - C\widehat{X}_n^-)$$
$$\widehat{P}_n = (1 - G_n C)\widehat{P}_n^-,$$

where $G_n = C\widehat{P}_n^-(C^2\widehat{P}_n^- + D^2)^{-1}$.

2. Predict

$$\widehat{X}_{n+1}^- = A\widehat{X}_n$$
$$\widehat{P}_{n+1}^- = A^2\widehat{P}_n + B^2.$$

The Kalman filter

The Kalman filter consists of two steps (update and predict)

$$(\widehat{X}_n^-, \widehat{P}_n^-) \longrightarrow (\widehat{X}_n, \widehat{P}_n) \longrightarrow (\widehat{X}_{n+1}^-, \widehat{P}_{n+1}^-).$$

1. Update

$$\widehat{X}_n = (1 - G_n C)\widehat{X}_n^- + G_n Y_n = \widehat{X}_n^- + G_n(Y_n - C\widehat{X}_n^-)$$
$$\widehat{P}_n = (1 - G_n C)\widehat{P}_n^-,$$

where $G_n = C\widehat{P}_n^-(C^2\widehat{P}_n^- + D^2)^{-1}$.

2. Predict

$$\widehat{X}_{n+1}^- = A\widehat{X}_n$$
$$\widehat{P}_{n+1}^- = A^2\widehat{P}_n + B^2.$$

The Kalman filter

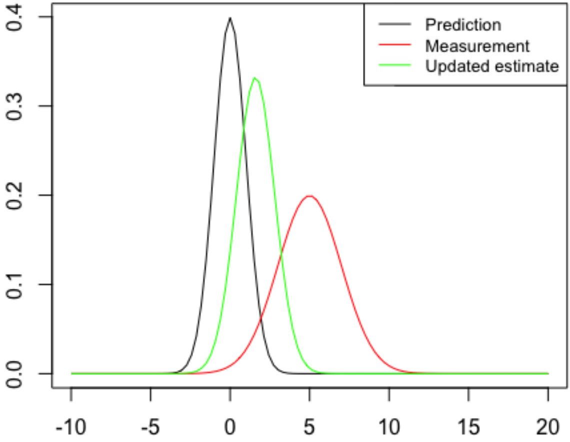

The Kalman gain

- The quantity
$$G_n = C\widehat{P}_n^-(C^2\widehat{P}_n^- + D^2)^{-1}$$
is called the Kalman gain.

- It represents the relative importance of the errors $Y_n - C\widehat{X}_n^-$ with respect to the prior estimate $\widehat{X}_n^-$.

- As $D \to 0$, $G_n \to C^{-1}$ and the update step becomes
$$\widehat{X}_n = \widehat{X}_n^- + C^{-1}(Y_n - C\widehat{X}_n^-) = C^{-1}Y_n.$$

- As $\widehat{P}_n^- \to 0$, $G_n \to 0$ and the update step converges to $\widehat{X}_n = \widehat{X}_n^-$.

The Kalman gain

- The quantity

$$G_n = C\widehat{P}_n^-(C^2\widehat{P}_n^- + D^2)^{-1}$$

is called the Kalman gain.

- It represents the relative importance of the errors $Y_n - C\widehat{X}_n^-$ with respect to the prior estimate $\widehat{X}_n^-$.

- As $D \to 0$, $G_n \to C^{-1}$ and the update step becomes

$$\widehat{X}_n = \widehat{X}_n^- + C^{-1}(Y_n - C\widehat{X}_n^-) = C^{-1}Y_n.$$

- As $\widehat{P}_n^- \to 0$, $G_n \to 0$ and the update step converges to $\widehat{X}_n = \widehat{X}_n^-$.

The Kalman gain

- The quantity
$$G_n = C\widehat{P}_n^-(C^2\widehat{P}_n^- + D^2)^{-1}$$
is called the Kalman gain.

- It represents the relative importance of the errors $Y_n - C\widehat{X}_n^-$ with respect to the prior estimate $\widehat{X}_n^-$.

- As $D \to 0$, $G_n \to C^{-1}$ and the update step becomes
$$\widehat{X}_n = \widehat{X}_n^- + C^{-1}(Y_n - C\widehat{X}_n^-) = C^{-1}Y_n.$$

- As $\widehat{P}_n^- \to 0$, $G_n \to 0$ and the update step converges to $\widehat{X}_n = \widehat{X}_n^-$.

Literature

- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. Journal of Basic Engineering. 82 (1): pp. 35–45
- Kalman, R.E., Bucy, R.S. (1961). New Results in Linear Filtering and Prediction Theory. Journal of Basic Engineering. 83: pp. 95–108
- Evensen, G., Vossepoel, F. C., & van Leeuwen, P. J. (2022). Data assimilation fundamentals: A unified formulation of the state and parameter estimation problem. Springer Nature.
- Simon, D. (2006). Optimal State Estimation: Kalman, $H\infty$, and Nonlinear Approaches. John Wiley & Sons.

The Kalman filter is BLUE

- The Kalman filter is the Best Linear Unbiased Estimator.
- Best = the estimator that minimises the MSE amongst all unbiased linear estimators.
- To prove this, consider the following estimator

$$\widehat{X}_n := H_n \widehat{X}_n^- + G_n Y_n.$$

The Kalman filter is BLUE

- The Kalman filter is the Best Linear Unbiased Estimator.
- Best = the estimator that minimises the MSE amongst all unbiased linear estimators.
- To prove this, consider the following estimator

$$\widehat{X}_n := H_n \widehat{X}_n^- + G_n Y_n.$$

The Kalman filter is BLUE

- The Kalman filter is the Best Linear Unbiased Estimator.
- Best = the estimator that minimises the MSE amongst all unbiased linear estimators.
- To prove this, consider the following estimator

$$\widehat{X}_n := H_n \widehat{X}_n^- + G_n Y_n.$$

The Kalman filter is BLUE: unbiasedness

Consider the bias

$$\mathbb{E}[\widehat{X}_n - X_n] = \mathbb{E}[H_n\widehat{X}_n^- + G_n Y_n - X_n]$$

The Kalman filter is BLUE: unbiasedness

Consider the bias

$$
\begin{aligned}
\mathbb{E}[\widehat{X}_n - X_n] &= \mathbb{E}[H_n \widehat{X}_n^- + G_n Y_n - X_n] \\
&= \mathbb{E}[H_n(X_n + \widehat{X}_n^- - X_n) + G_n(CX_n + DV_n) - X_n]
\end{aligned}
$$

The Kalman filter is BLUE: unbiasedness

Consider the bias

$$\begin{aligned}\mathbb{E}[\widehat{X}_n - X_n] &= \mathbb{E}[H_n\widehat{X}_n^- + G_n Y_n - X_n] \\ &= \mathbb{E}[H_n(X_n + \widehat{X}_n^- - X_n) + G_n(CX_n + DV_n) - X_n] \\ &= (G_n C + H_n - 1)\mathbb{E}[X_n] + G_n D\mathbb{E}[V_n] + H_n\mathbb{E}[\widehat{X}_n^- - X_n]\end{aligned}$$

The Kalman filter is BLUE: unbiasedness

Consider the bias

$$\begin{aligned}
\mathbb{E}[\widehat{X}_n - X_n] &= \mathbb{E}[H_n \widehat{X}_n^- + G_n Y_n - X_n] \\
&= \mathbb{E}[H_n(X_n + \widehat{X}_n^- - X_n) + G_n(CX_n + DV_n) - X_n] \\
&= (G_n C + H_n - 1)\mathbb{E}[X_n] + G_n D\mathbb{E}[V_n] + H_n \mathbb{E}[\widehat{X}_n^- - X_n] \\
&= (G_n C + H_n - 1)\mathbb{E}[X_n] + G_n D\mathbb{E}[V_n] + H_n A\mathbb{E}[\widehat{X}_{n-1} - X_{n-1}] \\
&\qquad - H_n B\mathbb{E}[W_n].
\end{aligned}$$

The Kalman filter is BLUE: unbiasedness

Consider the bias

$$\begin{aligned}\mathbb{E}[\widehat{X}_n - X_n] &= \mathbb{E}[H_n \widehat{X}_n^- + G_n Y_n - X_n] \\ &= \mathbb{E}[H_n(X_n + \widehat{X}_n^- - X_n) + G_n(CX_n + DV_n) - X_n] \\ &= (G_n C + H_n - 1)\mathbb{E}[X_n] + G_n D\mathbb{E}[V_n] + H_n \mathbb{E}[\widehat{X}_n^- - X_n] \\ &= (G_n C + H_n - 1)\mathbb{E}[X_n] + G_n D\mathbb{E}[V_n] + H_n A\mathbb{E}[\widehat{X}_{n-1} - X_{n-1}] \\ &\qquad - H_n B\mathbb{E}[W_n].\end{aligned}$$

In order for the estimator to be unbiased, need

$$H_n = (1 - G_n C),$$

which yields

$$\widehat{X}_n := \widehat{X}_n^- + G_n(Y_n - C\widehat{X}_n^-).$$

The Kalman filter is BLUE: optimality

First note that

$$(X_n - \widehat{X}_n)^2 = (X_n - \widehat{X}_n^- - G_n(Y_n - C\widehat{X}_n^-))^2$$
$$= [(1 - G_n C)(X_n - \widehat{X}_n^-) - G_n D V_n]^2.$$

The Kalman filter is BLUE: optimality

First note that

$$
\begin{aligned}
(X_n - \widehat{X}_n)^2 &= (X_n - \widehat{X}_n^- - G_n(Y_n - C\widehat{X}_n^-))^2 \\
&= [(1 - G_n C)(X_n - \widehat{X}_n^-) - G_n D V_n]^2.
\end{aligned}
$$

Expanding and taking expectations yields

$$
\widehat{P}_n = (1 - G_n C)^2 \widehat{P}_n^- + G_n^2 D^2.
$$

The Kalman filter is BLUE: optimality

First note that

$$(X_n - \widehat{X}_n)^2 = (X_n - \widehat{X}_n^- - G_n(Y_n - C\widehat{X}_n^-))^2$$
$$= [(1 - G_nC)(X_n - \widehat{X}_n^-) - G_nDV_n]^2.$$

Expanding and taking expectations yields

$$\widehat{P}_n = (1 - G_nC)^2\widehat{P}_n^- + G_n^2D^2.$$

Minimise w.r.t. G_n:

$$0 = \frac{\partial}{\partial G_n}\widehat{P}_n = -2C(1 - G_nC)\widehat{P}_n^- + 2D^2G_n,$$

The Kalman filter is BLUE: optimality

First note that

$$(X_n - \widehat{X}_n)^2 = (X_n - \widehat{X}_n^- - G_n(Y_n - C\widehat{X}_n^-))^2$$
$$= [(1 - G_nC)(X_n - \widehat{X}_n^-) - G_nDV_n]^2.$$

Expanding and taking expectations yields

$$\widehat{P}_n = (1 - G_nC)^2\widehat{P}_n^- + G_n^2D^2.$$

Minimise w.r.t. G_n:

$$0 = \frac{\partial}{\partial G_n}\widehat{P}_n = -2C(1 - G_nC)\widehat{P}_n^- + 2D^2G_n,$$

which implies

$$G_n = C\widehat{P}_n^-(C\widehat{P}_n^- + D^2)^{-1}.$$

Stability of the Kalman filter

- We may write

$$X_{n+1} - \widehat{X}_{n+1}^{-} = A(1 - G_n C)(X_n - \widehat{X}_n^{-}) + BW_{n+1} - AG_n DV_n.$$

- Can't hope to obtain a result of the form $\|X_n - \widehat{X}_n^{-}\| \to 0$ unless the measurement noise vanishes.

- However, we can study the homogeneous part of the above recursion:

$$Z_{n+1} = A(1 - G_n C)Z_n \qquad (1)$$

Stability of the Kalman filter

- We may write

$$X_{n+1} - \widehat{X}_{n+1}^{-} = A(1 - G_n C)(X_n - \widehat{X}_n^{-}) + BW_{n+1} - AG_n DV_n.$$

- Can't hope to obtain a result of the form $\|X_n - \widehat{X}_n^{-}\| \to 0$ unless the measurement noise vanishes.

- However, we can study the homogeneous part of the above recursion:

$$Z_{n+1} = A(1 - G_n C)Z_n \tag{1}$$

Stability of the Kalman filter

- We may write

$$X_{n+1} - \widehat{X}^-_{n+1} = A(1 - G_n C)(X_n - \widehat{X}^-_n) + BW_{n+1} - AG_n DV_n.$$

- Can't hope to obtain a result of the form $\|X_n - \widehat{X}^-_n\| \to 0$ unless the measurement noise vanishes.

- However, we can study the homogeneous part of the above recursion:

$$Z_{n+1} = A(1 - G_n C)Z_n \tag{1}$$

Stability of the Kalman filter

- We may write

$$X_{n+1} - \widehat{X}_{n+1}^{-} = A(1 - G_n C)(X_n - \widehat{X}_n^{-}) + BW_{n+1} - AG_n DV_n.$$

- Can't hope to obtain a result of the form $\|X_n - \widehat{X}_n^{-}\| \to 0$ unless the measurement noise vanishes.

- However, we can study the homogeneous part of the above recursion:

$$Z_{n+1} = A(1 - G_n C)Z_n \tag{1}$$

Stability of the Kalman filter

Theorem (Deyst & Price '68, Jazwinski '70)

If $A, B, C, D \neq 0$ and $\widehat{P}_0^- > 0$ then there exist constants $K > 0$, $\gamma \in (0,1)$, $n_0 \geq 0$ such that

$$\|Z_n\| \leq K\gamma^{n-n_0}\|Z_{n_0}\|.$$

Stability of the Kalman filter

Theorem (Deyst & Price '68, Jazwinski '70)

If $A, B, C, D \neq 0$ and $\widehat{P}_0^- > 0$ then there exist constants $K > 0$, $\gamma \in (0, 1)$, $n_0 \geq 0$ such that

$$\|Z_n\| \leq K\gamma^{n-n_0}\|Z_{n_0}\|.$$

Many criteria for proving exponential stability but we will focus on Lyapunov-type criteria.

Stability of the Kalman filter

Lyapunov Stability Theorem

The system (1) is exponentially stable if there exists a continuous scalar function such that

- $V(0) = 0$ and $V(x) > 0$ for $x \neq 0$,
- $V(x) \to \infty$ as $\|x\| \to \infty$, and
- $V(Z_{n+1}) - V(Z_n) < 0$.

- A Lyapunov function is a non-negative function of a system's state that decreases as the state changes.
- If a system is described by a set of differential equations and we can find a Lyapunov function for these equations, then local minima of the Lyapunov function are stable.
- A Lyapunov function V is an analogue of the energy of a physical system.

Stability of the Kalman filter

Lyapunov Stability Theorem

The system (1) is exponentially stable if there exists a continuous scalar function such that

- $V(0) = 0$ and $V(x) > 0$ for $x \neq 0$,
- $V(x) \to \infty$ as $\|x\| \to \infty$, and
- $V(Z_{n+1}) - V(Z_n) < 0$.

- A Lyapunov function is a non-negative function of a system's state that decreases as the state changes.
- If a system is described by a set of differential equations and we can find a Lyapunov function for these equations, then local minima of the Lyapunov function are stable.
- A Lyapunov function V is an analogue of the energy of a physical system.

Stability of the Kalman filter

Lyapunov Stability Theorem

The system (1) is exponentially stable if there exists a continuous scalar function such that

- $V(0) = 0$ and $V(x) > 0$ for $x \neq 0$,
- $V(x) \to \infty$ as $\|x\| \to \infty$, and
- $V(Z_{n+1}) - V(Z_n) < 0$.

- A Lyapunov function is a non-negative function of a system's state that decreases as the state changes.
- If a system is described by a set of differential equations and we can find a Lyapunov function for these equations, then local minima of the Lyapunov function are stable.
- A Lyapunov function V is an analogue of the energy of a physical system.

Stability of the Kalman filter

Lyapunov Stability Theorem

The system (1) is exponentially stable if there exists a continuous scalar function such that

- $V(0) = 0$ and $V(x) > 0$ for $x \neq 0$,
- $V(x) \to \infty$ as $\|x\| \to \infty$, and
- $V(Z_{n+1}) - V(Z_n) < 0$.

- A Lyapunov function is a non-negative function of a system's state that decreases as the state changes.
- If a system is described by a set of differential equations and we can find a Lyapunov function for these equations, then local minima of the Lyapunov function are stable.
- A Lyapunov function V is an analogue of the energy of a physical system.

Stability of the Kalman filter

- What about the covariance?
- We have the following recursion for $\widehat{P}_n^-$

$$\widehat{P}_{n+1}^- = A^2(1 - G_n C)\widehat{P}_n^- + B^2$$

- Again, we can't hope for a result of the form $\|\widehat{P}_n^-\| \to 0$.

Stability of the Kalman filter

- What about the covariance?
- We have the following recursion for $\widehat{P}_n^-$

$$\widehat{P}_{n+1}^- = A^2(1 - G_n C)\widehat{P}_n^- + B^2$$

- Again, we can't hope for a result of the form $\|\widehat{P}_n^-\| \to 0$.

Stability of the Kalman filter

- What about the covariance?
- We have the following recursion for $\widehat{P}_n^-$

$$\widehat{P}_{n+1}^- = A^2(1 - G_n C)\widehat{P}_n^- + B^2$$

- Again, we can't hope for a result of the form $\|\widehat{P}_n^-\| \to 0$.

Stability of the Kalman filter

- What about the covariance?
- We have the following recursion for $\widehat{P}_n^-$

$$\widehat{P}_{n+1}^- = A^2(1 - G_n C)\widehat{P}_n^- + B^2$$

- Again, we can't hope for a result of the form $\|\widehat{P}_n^-\| \to 0$.

Theorem (West & Harris '97)

For time-invariant systems, there exists $P_\infty^- > 0$ such that

$$\|\widehat{P}_n^- - P_\infty^-\| \to 0, \qquad \text{as } n \to \infty.$$

In addition, there exists G_∞ such that

$$\|\widehat{G}_n - G_\infty\| \to 0, \qquad \text{as } n \to \infty.$$

Riccati rational difference equation

- The evolution equation for $\widehat{P}_n^-$ can equivalently be written as

$$\widehat{P}_{n+1}^- = \phi(\widehat{P}_n^-) = \frac{a\widehat{P}_n^- + b}{c\widehat{P}_n^- + d},$$

where a, b, c, d are determined by A, B, C, D.

- The function ϕ is known as a Riccati map and the above equation as a Riccati rational difference equation.

- More on these maps later...

Riccati rational difference equation

- The evolution equation for $\widehat{P}_n^-$ can equivalently be written as

$$\widehat{P}_{n+1}^- = \phi(\widehat{P}_n^-) = \frac{a\widehat{P}_n^- + b}{c\widehat{P}_n^- + d},$$

  where a, b, c, d are determined by A, B, C, D.

- The function ϕ is known as a Riccati map and the above equation as a Riccati rational difference equation.

- More on these maps later...

Riccati rational difference equation

- The evolution equation for $\widehat{P}_n^-$ can equivalently be written as

$$\widehat{P}_{n+1}^- = \phi(\widehat{P}_n^-) = \frac{a\widehat{P}_n^- + b}{c\widehat{P}_n^- + d},$$

  where a, b, c, d are determined by A, B, C, D.

- The function ϕ is known as a Riccati map and the above equation as a Riccati rational difference equation.

- More on these maps later...

Assumptions

Recall the original system:

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1}.$$

Definition (Kalman '60)

The system is said to be observable if

$$[C, CA, CA^2, \ldots, CA^{d-1}]^T$$

has rown rank d, where d is the dimension of the state space.

Assumptions

Recall the original system:

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1}.$$

Definition (Kalman '60)

The system is said to be observable if

$$[C, CA, CA^2, \ldots, CA^{d-1}]^T$$

has rown rank d, where d is the dimension of the state space.

Assumptions

- Observability says that any state x can be inferred from a sufficient number of observations.

- Suppose now the state space has dimension d, so that $X_n \in \mathbb{R}^d$, for example and consider the average behaviour of the system:

$$Y_n = CX_n,$$
$$X_{n+1} = AX_n.$$

- If we knew X_0, then the second equation would give us complete knowledge of the state at any time $n \rightsquigarrow$ just need to determine X_0.

Assumptions

- Observability says that any state x can be inferred from a sufficient number of observations.
- Suppose now the state space has dimension d, so that $X_n \in \mathbb{R}^d$, for example and consider the average behaviour of the system:

$$Y_n = CX_n,$$
$$X_{n+1} = AX_n.$$

- If we knew X_0, then the second equation would give us complete knowledge of the state at any time $n \rightsquigarrow$ just need to determine X_0.

Assumptions

- Observability says that any state x can be inferred from a sufficient number of observations.
- Suppose now the state space has dimension d, so that $X_n \in \mathbb{R}^d$, for example and consider the average behaviour of the system:

$$\begin{aligned} Y_n &= CX_n, \\ X_{n+1} &= AX_n. \end{aligned}$$

- If we knew X_0, then the second equation would give us complete knowledge of the state at any time $n \rightsquigarrow$ just need to determine X_0.

Assumptions

- Since X_0 has d unknowns, we expect to need d observations in order to determine X_0:

$$\begin{aligned} Y_0 &= CX_0 \\ Y_1 &= CX_1 = CAX_0 \\ &\vdots \\ Y_{d-1} &= CX_{d-1} = \cdots = CA^{d-1}X_0. \end{aligned}$$

- Writing this in vector form,

$$[Y_0, \ldots, Y_{n-1}]^T = [C, CA, CA^2, \ldots, CA^{d-1}]^T X_0,$$

we see that this has a unique solution iff $[C, CA, CA^2, \ldots, CA^{d-1}]^T$ is invertible.

- In one dimension, this implies that $C \neq 0$.

Assumptions

- Since X_0 has d unknowns, we expect to need d observations in order to determine X_0:

$$
\begin{aligned}
Y_0 &= CX_0 \\
Y_1 &= CX_1 = CAX_0 \\
&\vdots \\
Y_{d-1} &= CX_{d-1} = \cdots = CA^{d-1}X_0.
\end{aligned}
$$

- Writing this in vector form,

$$[Y_0, \ldots, Y_{n-1}]^T = [C, CA, CA^2, \ldots, CA^{d-1}]^T X_0,$$

we see that this has a unique solution iff $[C, CA, CA^2, \ldots, CA^{d-1}]^T$ is invertible.

- In one dimension, this implies that $C \neq 0$.

Assumptions

- Since X_0 has d unknowns, we expect to need d observations in order to determine X_0:

$$\begin{aligned} Y_0 &= CX_0 \\ Y_1 &= CX_1 = CAX_0 \\ &\vdots \\ Y_{d-1} &= CX_{d-1} = \cdots = CA^{d-1}X_0. \end{aligned}$$

- Writing this in vector form,

$$[Y_0, \ldots, Y_{n-1}]^T = [C, CA, CA^2, \ldots, CA^{d-1}]^T X_0,$$

we see that this has a unique solution iff $[C, CA, CA^2, \ldots, CA^{d-1}]^T$ is invertible.

- In one dimension, this implies that $C \neq 0$.

Assumptions

Recall the original system:

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1}.$$

Definition (Kalman '60)

The system is said to be controllable if

$$[B, AB, A^2 B, \dots, A^{d-1} B]$$

has column rank d where d is the dimension of the state space.

- Controllability says that any state x can be reached from any initial condition x_0 in a finite number of steps.
- In one dimension, this implies that $B \neq 0$, i.e. it means that the state is affected by the noise W_n.

Assumptions

Recall the original system:

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1}.$$

Definition (Kalman '60)

The system is said to be controllable if

$$[B, AB, A^2B, \dots, A^{d-1}B]$$

has column rank d where d is the dimension of the state space.

- Controllability says that any state x can be reached from any initial condition x_0 in a finite number of steps.
- In one dimension, this implies that $B \neq 0$, i.e. it means that the state is affected by the noise W_n.

Assumptions

Recall the original system:

$$Y_n = CX_n + DV_n,$$
$$X_{n+1} = AX_n + BW_{n+1}.$$

Definition (Kalman '60)

The system is said to be controllable if

$$[B, AB, A^2B, \dots, A^{d-1}B]$$

has column rank d where d is the dimension of the state space.

- Controllability says that any state x can be reached from any initial condition x_0 in a finite number of steps.
- In one dimension, this implies that $B \neq 0$, i.e. it means that the state is affected by the noise W_n.

Assumptions

- The parameters A, B, C, D are non-zero.
- The parameters A, B, C, D are independent of the time step.
- The source is Gaussian.
- The errors are Gaussian
- The system is one-dimensional.
- The system is linear...

Example: Numerical weather prediction

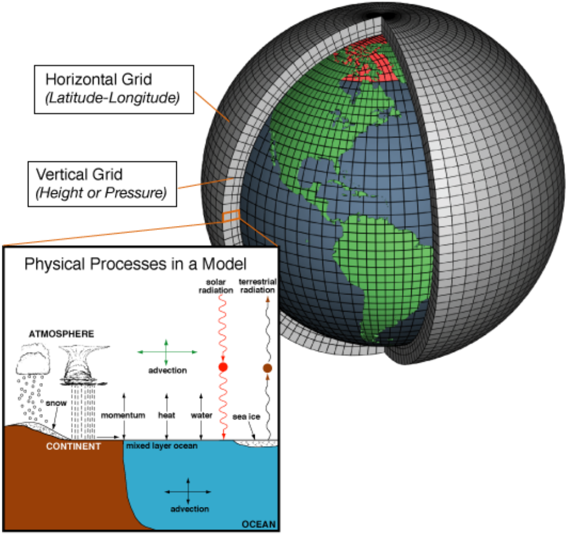

A nonlinear problem

Now suppose we are given the following dynamics

$$X_{n+1} = f(X_n) + BW_{n+1},$$
$$Y_n = h(X_n) + DV_n,$$

where $B, D \neq 0$, V_n, W_n are i.i.d. $\mathcal{N}(0,1)$ and the functions f, h are C^1.

A nonlinear problem

Now suppose we are given the following dynamics

$$X_{n+1} = f(X_n) + BW_{n+1},$$
$$Y_n = h(X_n) + DV_n,$$

where $B, D \neq 0$, V_n, W_n are i.i.d. $\mathcal{N}(0,1)$ and the functions f, h are C^1.

$\rightsquigarrow$ Can we still estimate the distribution of X_n given measurements $Y_1, \ldots, Y_n$?

Extended Kalman filter

The EKF is given by the following update-predict steps:

1. Update

$$\widehat{X}_n = \widehat{X}_n^- + G_n(Y_n - h(\widehat{X}_n^-))$$
$$\widehat{P}_n = (1 - G_n H_n)\widehat{P}_n^-,$$

where $H_n = \dfrac{\partial h}{\partial x}(\widehat{X}_n^-)$ and $G_n = H_n \widehat{P}_n^- (H_n^2 \widehat{P}_n^- + D^2)^{-1}$.

2. Predict

$$\widehat{X}_{n+1}^- = f(\widehat{X}_n)$$
$$\widehat{P}_{n+1}^- = F_n^2 \widehat{P}_n + B^2,$$

where $F_n = \dfrac{\partial f}{\partial x}(\widehat{X}_n)$.

Further reading

- Tine Lefebvre and Herman Bruyninckx. Kalman Filters for Nonlinear Systems: A Comparison of Performance.
- K.Reif, S.Gunther, E.Yaz, and R.Unbehauen. Stochastic Stability of the Discrete-Time Extended Kalman Filter. IEEE Trans.Automatic Control, 1999.
- John M. Lewis and S.Lakshmivarahan. Dynamic Data Assimilation, a Least Squares Approach. 2006.
- R. van der Merwe. Sigma-Point Kalman Filters for Probabilistic Inference in Dynamic State-Space Models. Technical report, 2003.

Remarks

- The EKF is not optimal.
- If the initial estimate is wrong, the filter quickly diverges. (However, this is still the "go to" filter for navigation systems and GPS... 😬)
- Further errors due to linearising the model.
- Computationally expensive in high dimensions...

Part II : The Ensemble Kalman filter[1]

[1] Based on joint work with Pierre Del Moral (Inria, Bordeaux)

EnKF

- The EnKF is a Monte Carlo implementation of the Kalman filter.
- Idea is to evolve the ensemble forward in time and estimate the mean and covariance from the evolved sample.
- Fix $N \geq 1$ and let $\xi^{-}_{0,i}$, $i = 1, \ldots, N$ be i.i.d. copies of $X_0 \sim \mathcal{N}(\widehat{X}^{-}_0, \widehat{P}^{-}_0)$.
- For $n \geq 1$, let W^i_n, $i = 1, \ldots, N$ be i.i.d. copies of $W_n \sim \mathcal{N}(0, 1)$. Similarly for V_n.

EnKF

- The EnKF is a Monte Carlo implementation of the Kalman filter.
- Idea is to evolve the ensemble forward in time and estimate the mean and covariance from the evolved sample.
- Fix $N \geq 1$ and let $\xi^{-}_{0,i}$, $i = 1, \ldots, N$ be i.i.d. copies of $X_0 \sim \mathcal{N}(\widehat{X}^{-}_0, \widehat{P}^{-}_0)$.
- For $n \geq 1$, let W^i_n, $i = 1, \ldots, N$ be i.i.d. copies of $W_n \sim \mathcal{N}(0, 1)$. Similarly for V_n.

EnKF

The EnKF algorithm is given by the following update-predict sequence:

$$\begin{aligned} \xi_{n,i} &= \xi^-_{n,i} + g_n(Y_n + DV^i_n - C\xi^-_{n,i}), &\quad \text{(update)} \\ \xi^-_{n+1,i} &= A\xi_{n,i} + BW^i_n, &\quad \text{(predict)} \end{aligned}$$

where

$$g_n = Cp^-_n(C^2p^-_n + D^2)^{-1},$$

is the sample Kalman gain, and

$$m^-_n = \frac{1}{N}\sum_{i=1}^N \xi^-_{n,i} \quad \text{and} \quad p^-_n = \frac{1}{N-1}\sum_{i=1}^N (\xi^-_{n,i} - m^-_n)^2$$

are the prior sample mean and covariance, respectively.

EnKF

Similarly, we can define the posterior sample mean and covariance

$$m_n = \frac{1}{N}\sum_{i=1}^{N}\xi_{n,i} \quad \text{and} \quad p_n = \frac{1}{N-1}\sum_{i=1}^{N}(\xi_{n,i} - m_n)^2.$$

Thus, we have the following update-predict steps of the EnKF

$$(m_n^-, p_n^-) \longrightarrow (m_n, p_n) \longrightarrow (m_{n+1}^-, p_{n+1}^-).$$

Literature

- G. Evensen. Sequential data assimilation with a non-linear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. J Geophys Res 99(C5): vol.10 pp. 143–162 (1994)
- G. Burgers, P. J. van Leeuwen, and G. Evensen. Analysis scheme in the ensemble Kalman filter. Monthly Weather Review, 126:1719–1724, (1998)
- F. Le Gland, V. Monbet, V.D. Tran. Large sample asymptotics for the ensemble Kalman filter The Oxford Handbook of Nonlinear Filtering, chapter 22, pp. 598–631 (2011).
- J. Mandel, L. Cobb, J. D. Beezley. On the convergence of the ensemble Kalman filter Applications of Mathematics, vol. 56, no. 6, pp. 533–541 (2011).
- A. J. Majda, X. T. Tong. Performance of ensemble Kalman filters in large dimensions. Communications on Pure and Applied Mathematics, vol. 71, no. 5, pp. 892–937 (2018).

Literature

- G. Evensen. Sequential data assimilation with a non-linear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. J Geophys Res 99(C5): vol.10 pp. 143–162 (1994)
- G. Burgers, P. J. van Leeuwen, and G. Evensen. Analysis scheme in the ensemble Kalman filter. Monthly Weather Review, 126:1719–1724, (1998)
- F. Le Gland, V. Monbet, V.D. Tran. Large sample asymptotics for the ensemble Kalman filter The Oxford Handbook of Nonlinear Filtering, chapter 22, pp. 598–631 (2011).
- J. Mandel, L. Cobb, J. D. Beezley. On the convergence of the ensemble Kalman filter Applications of Mathematics, vol. 56, no. 6, pp. 533–541 (2011).
- A. J. Majda, X. T. Tong. Performance of ensemble Kalman filters in large dimensions. Communications on Pure and Applied Mathematics, vol. 71, no. 5, pp. 892–937 (2018).

  "Why the EnKF works well with a small ensemble has remained a complete mystery."

EnKF: First results

Theorem (Le Gland et. al., '11, Mandel et. al. '11)

For each $n \geq 0$, as $N \to \infty$,

$$m_n \to \widehat{X}_n \quad \text{and} \quad p_n \to \widehat{P}_n,$$

in L^p at rate $1/\sqrt{N}$, and almost surely.

Theorem (Le Gland et. al. '11)

The EnKF does not converge to the optimal filter for non-linear or non-Gaussian filtering problems.

Local perturbations

Theorem (Del Moral, H., '22)

With initial conditions,

$$m_0^- = \widehat{X}_0^- + \frac{1}{\sqrt{N+1}}\upsilon_0 \quad \text{and} \quad p_0^- = \widehat{P}_0^- + \frac{1}{\sqrt{N}}\nu_0,$$

the ensemble Kalman filter update-predict transitions are as follows:

1. Update

$$m_n = m_n^- + g_n(Y_n - Cm_n^-) + \frac{1}{\sqrt{N+1}}\upsilon_n$$

$$p_n = (1 - g_n C)p_n^- + \frac{1}{\sqrt{N}}\nu_n,$$

2. Predict

$$m_{n+1}^- = Am_n + \frac{1}{\sqrt{N+1}}\upsilon_{n+1}^-$$

$$p_{n+1}^- = A^2 p_n + B^2 + \frac{1}{\sqrt{N}}\nu_{n+1}^-.$$

Local perturbations

- The previous result implies that the sample variance, p_n^-, of the EnKF satisfies the stochastic Riccati rational difference equation

$$p_{n+1}^- = \phi(p_n^-) + \frac{1}{\sqrt{N}}\delta_{n+1},$$

where $\delta_{n+1} = A^2\nu_n + \nu_{n+1}^-$.

- Let $\mathcal{P}(p, \mathrm{d}q)$ denote the Markov transitions associated with the Markov chain $(p_n^-)_{n\geq 0}$, i.e. $\mathcal{P}(p, \mathrm{d}q) = \mathbb{P}[p_{n+1}^- \in \mathrm{d}q | p_n = p]$.
For suitable test functions, we write $\mathcal{P}(f)(p) = \mathbb{E}[f(p_{n+1}^-)|p_n = p]$

- For a locally finite signed measure μ on $\mathbb{R}_+$ and functions $f : \mathbb{R}_+ \to \mathbb{R}$, $V : \mathbb{R}_+ \to \mathbb{R}_+$, define

$$\|f\|_V = \sup_{p\geq 0}\left|\frac{f(p)}{\frac{1}{2} + V(p)}\right| \quad \text{and} \quad \|\mu\| := \sup\{|\mu(f)| : \|f\|_V \leq 1\}.$$

Local perturbations

- The previous result implies that the sample variance, p_n^-, of the EnKF satisfies the stochastic Riccati rational difference equation

$$p_{n+1}^- = \phi(p_n^-) + \frac{1}{\sqrt{N}} \delta_{n+1},$$

where $\delta_{n+1} = A^2 \nu_n + \nu_{n+1}^-$.

- Let $\mathcal{P}(p, \mathrm{d}q)$ denote the Markov transitions associated with the Markov chain $(p_n^-)_{n \geq 0}$, i.e. $\mathcal{P}(p, \mathrm{d}q) = \mathbb{P}[p_{n+1}^- \in \mathrm{d}q | p_n = p]$.
For suitable test functions, we write $\mathcal{P}(f)(p) = \mathbb{E}[f(p_{n+1}^-)|p_n = p]$

- For a locally finite signed measure μ on $\mathbb{R}_+$ and functions $f : \mathbb{R}_+ \to \mathbb{R}$, $V : \mathbb{R}_+ \to \mathbb{R}_+$, define

$$\|f\|_V = \sup_{p \geq 0} \left| \frac{f(p)}{\frac{1}{2} + V(p)} \right| \quad \text{and} \quad \|\mu\| := \sup\{|\mu(f)| : \|f\|_V \leq 1\}.$$

Local perturbations

- The previous result implies that the sample variance, p_n^-, of the EnKF satisfies the stochastic Riccati rational difference equation

$$p_{n+1}^- = \phi(p_n^-) + \frac{1}{\sqrt{N}}\delta_{n+1},$$

where $\delta_{n+1} = A^2\nu_n + \nu_{n+1}^-$.

- Let $\mathcal{P}(p, \mathrm{d}q)$ denote the Markov transitions associated with the Markov chain $(p_n^-)_{n\geq 0}$, i.e. $\mathcal{P}(p, \mathrm{d}q) = \mathbb{P}[p_{n+1}^- \in \mathrm{d}q | p_n = p]$.
For suitable test functions, we write $\mathcal{P}(f)(p) = \mathbb{E}[f(p_{n+1}^-)|p_n = p]$

- For a locally finite signed measure μ on $\mathbb{R}_+$ and functions $f : \mathbb{R}_+ \to \mathbb{R}$, $V : \mathbb{R}_+ \to \mathbb{R}_+$, define

$$\|f\|_V = \sup_{p\geq 0}\left|\frac{f(p)}{\frac{1}{2} + V(p)}\right| \quad \text{and} \quad \|\mu\| := \sup\{|\mu(f)| : \|f\|_V \leq 1\}.$$

Stability results

Theorem (Del Moral, H., 2022)

There exists a unique invariant measure π such that $\pi\mathcal{P} = \pi$, a function $\mathcal{U}$ and a constant $\beta \in (0,1)$ such that for any function f satisfying $\|f\|_{\mathcal{U}} \leq 1$ and for any $p \in \mathbb{R}_+$, we have

$$|\mathcal{P}^n(f)(p) - \pi(f)| \leq \beta^n(1 + \mathcal{U}(p) + \pi(\mathcal{U})).$$

Stability results: idea of proof

For a function V, define

$$\beta_V(\mathcal{P}) := \sup_{p,q\geq 0} \frac{\|\mathcal{P}(p,\cdot) - \mathcal{P}(q,\cdot)\|_V}{1 + V(p) + V(q)}$$

Now suppose we can first prove the following result.

Proposition

There exists a function $\mathcal{U} : \mathbb{R}_+ \to \mathbb{R}_+$ such that $\beta_{\mathcal{U}}(\mathcal{P}) < 1$ and for any two probability measures μ_1, μ_2, we have

$$\|\mu_1\mathcal{P}^n - \mu_2\mathcal{P}^n\|_{\mathcal{U}} \leq \beta_{\mathcal{U}}(\mathcal{P})^n \|\mu_1 - \mu_2\|_{\mathcal{U}}.$$

Stability results: idea of proof

For a function V, define

$$\beta_V(\mathcal{P}) := \sup_{p,q\geq 0} \frac{\|\mathcal{P}(p,\cdot) - \mathcal{P}(q,\cdot)\|_V}{1 + V(p) + V(q)}$$

Now suppose we can first prove the following result.

Proposition

There exists a function $\mathcal{U} : \mathbb{R}_+ \to \mathbb{R}_+$ such that $\beta_{\mathcal{U}}(\mathcal{P}) < 1$ and for any two probability measures μ_1, μ_2, we have

$$\|\mu_1 \mathcal{P}^n - \mu_2 \mathcal{P}^n\|_{\mathcal{U}} \leq \beta_{\mathcal{U}}(\mathcal{P})^n \|\mu_1 - \mu_2\|_{\mathcal{U}}.$$

Stability results: idea of proof

- The existence of a unique invariant (probability) measure π follows from the fixed point theorem.

- From the definition of $\beta_{\mathcal{U}}(\mathcal{P})$, we have

$$|\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)| \leq \beta_{\mathcal{U}}(\mathcal{P}^n)(1 + \mathcal{U}(p) + \mathcal{U}(q)).$$

- Then

$$\begin{aligned}
|\mathcal{P}^n(f)(p) - \pi(f)| &= |\mathcal{P}^n(f)(p) - \pi\mathcal{P}(f)| \\
&= \left| \int_0^\infty \pi(\mathrm{d}q) \left(\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)\right) \right| \\
&\leq \int_0^\infty \pi(\mathrm{d}q) |\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)| \\
&\leq \beta_{\mathcal{U}}(\mathcal{P}^n)(1 + \mathcal{U}(p) + \pi(\mathcal{U})) \\
&\leq \beta_{\mathcal{U}}(\mathcal{P})^n(1 + \mathcal{U}(p) + \pi(\mathcal{U}))
\end{aligned}$$

Stability results: idea of proof

- The existence of a unique invariant (probability) measure π follows from the fixed point theorem.

- From the definition of $\beta_{\mathcal{U}}(\mathcal{P})$, we have
$$|\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)| \le \beta_{\mathcal{U}}(\mathcal{P}^n)(1 + \mathcal{U}(p) + \mathcal{U}(q)).$$

- Then
$$\begin{aligned}
|\mathcal{P}^n(f)(p) - \pi(f)| &= |\mathcal{P}^n(f)(p) - \pi\mathcal{P}(f)| \\
&= \left| \int_0^\infty \pi(\mathrm{d}q) \left(\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)\right) \right| \\
&\le \int_0^\infty \pi(\mathrm{d}q) |\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)| \\
&\le \beta_{\mathcal{U}}(\mathcal{P}^n)(1 + \mathcal{U}(p) + \pi(\mathcal{U})) \\
&\le \beta_{\mathcal{U}}(\mathcal{P})^n(1 + \mathcal{U}(p) + \pi(\mathcal{U}))
\end{aligned}$$

Stability results: idea of proof

- The existence of a unique invariant (probability) measure π follows from the fixed point theorem.
- From the definition of $\beta_{\mathcal{U}}(\mathcal{P})$, we have

$$|\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)| \leq \beta_{\mathcal{U}}(\mathcal{P}^n)(1 + \mathcal{U}(p) + \mathcal{U}(q)).$$

- Then

$$\begin{aligned}
|\mathcal{P}^n(f)(p) - \pi(f)| &= |\mathcal{P}^n(f)(p) - \pi\mathcal{P}(f)| \\
&= \left| \int_0^\infty \pi(\mathrm{d}q) \left(\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)\right) \right| \\
&\leq \int_0^\infty \pi(\mathrm{d}q) |\mathcal{P}^n(f)(p) - \mathcal{P}^n(f)(q)| \\
&\leq \beta_{\mathcal{U}}(\mathcal{P}^n)(1 + \mathcal{U}(p) + \pi(\mathcal{U})) \\
&\leq \beta_{\mathcal{U}}(\mathcal{P})^n(1 + \mathcal{U}(p) + \pi(\mathcal{U}))
\end{aligned}$$

Stability results: idea of proof

To prove the auxiliary result, it is sufficient to prove the following.

- For any compact set $K \subset \mathbb{R}_+$, there exists a constant $\varepsilon_K \in (0,1]$ and a probability measure ν_K on $\mathbb{R}_+$ such that for all $p \in K$,

$$\mathcal{P}(p, \mathrm{d}q) \geq \varepsilon_K \nu_K(\mathrm{d}q).$$

- There exists a non-negative function $\mathcal{U} : \mathbb{R}_+ \to [1, \infty)$ with compact level sets, such that

$$\mathcal{P}(\mathcal{U}) \leq \varepsilon \mathcal{U} + c,$$

for come $\varepsilon \in [0, 1)$ and $c < \infty$.

Uniform mean-error estimates

Theorem (Del Moral, H., 2022)

For any $k \geq 1$, there exists an integer $N_k \geq 1$ such that for any $N \geq N_k$ and $n \geq 0$, we have

$$\mathbb{E}\left[|p_n^- - \widehat{P}_n^-|^k\right]^{1/k} \vee \mathbb{E}\left[|p_n - \widehat{P}_n|^k\right]^{1/k} \vee \mathbb{E}\left[|g_n - \widehat{G}_n|^k\right]^{1/k} \leq \frac{C_k(1 \vee P_0)}{\sqrt{N}}.$$

Central Limit Theorem

Define the collection of stochastic processes $(\mathbb{Q}_{N,n}^{-}, \mathbb{Q}_{N,n+1})_{n\geq 0}$ defined via

$$\mathbb{Q}_{N,n}^{-} := \sqrt{N}(p_n^{-} - \widehat{P}_n^{-}) \quad \text{and} \quad \mathbb{Q}_{N,n} := \sqrt{N}(p_n - \widehat{P}_n).$$

Theorem (Del Moral, H., 2022)

The stochastic processes $(\mathbb{Q}_{N,n}, \mathbb{Q}_{N,n+1}^{-})$ converge in law in the sense of f.d.d., as the number of particles $N \to \infty$, to a sequence of centred stochastic processes $(\mathbb{Q}_n, \mathbb{Q}_{n+1}^{-})$ with initial condition $\mathbb{Q}_0^{-} = \mathbb{V}_0^{-}$ and update-predict transitions given by

$$\begin{aligned} \mathbb{Q}_n &= (1 - G_n C)\mathbb{Q}_n^{-} + \mathbb{V}_n \\ \mathbb{Q}_{n+1}^{-} &= A\mathbb{Q}_n + \mathbb{V}_{n+1}^{-}. \end{aligned}$$

Idea behind the proofs

- Let us consider $p_n^- - \widehat{P}_n^-$. The idea is to write this difference as a telescoping sum involving the increments of the Markov chain $(p_n^-)_{n\geq 1}$ and show that we can control these increments nicely.

- Recall from the evolution equation for p_n^-, the increments are related to the Riccati map

$$\phi(x) = \frac{ax+b}{cx+d}, \quad x \geq 0.$$

- Thus, we first need to look at the behaviour of these maps...

Idea behind the proofs

- Let us consider $p_n^- - \widehat{P}_n^-$. The idea is to write this difference as a telescoping sum involving the increments of the Markov chain $(p_n^-)_{n\geq 1}$ and show that we can control these increments nicely.

- Recall from the evolution equation for p_n^-, the increments are related to the Riccati map

$$\phi(x) = \frac{ax+b}{cx+d}, \quad x \geq 0.$$

- Thus, we first need to look at the behaviour of these maps...

Idea behind the proofs

- Let us consider $p_n^- - \widehat{P}_n^-$. The idea is to write this difference as a telescoping sum involving the increments of the Markov chain $(p_n^-)_{n\geq 1}$ and show that we can control these increments nicely.
- Recall from the evolution equation for p_n^-, the increments are related to the Riccati map
$$\phi(x) = \frac{ax+b}{cx+d}, \quad x \geq 0.$$
- Thus, we first need to look at the behaviour of these maps...

Idea behind proofs : Riccati maps

Lemma (Del Moral, H., 2022)

(i) For any $n \geq 1$, $b/d \leq \phi^n(x) \leq a/c$.

(ii) We have the Lipschitz estimates

$$|\phi^n(x) - \phi^n(y)| \leq C_1\lambda^n|x-y| \quad \text{and} \quad |\partial\phi^n(x) - \partial\phi^n(y)| \leq C_2\lambda^n|x-y|,$$

where $C_1, C_2 > 0$ and $\lambda \in (0,1)$.

(iii) Finally, we have the second order estimate

$$|\phi^n(x) - \phi^n(y) - \partial\phi^n(y)(x-y)| \leq C_3\lambda^n|x-y|^2,$$

where $C_3 > 0$.

Idea behind proofs : Riccati maps

- A.N. Bishop and P. Del Moral. On the stability of Kalman-Bucy diffusion processes. SIAM Journal on Control and Optimization. vol. 55, no. 6. pp 4015–4047 (2017); arxiv e-print arXiv:1610.04686.
- A.N. Bishop and P. Del Moral. An explicit Floquet-type representation of Riccati aperiodic exponential semigroups. International Journal of Control, pp. 1–9 (2019).
- P. Del Moral and E. Horton. A note on Riccati matrix difference equations. SIAM J. Control Optim., 60(3), pp. 1393-1409 (2022).

Idea behind proofs

- Consider the following decomposition

$$p_n^- - \widehat{P}_n^- = \phi^n(p_0^-) - \phi^n(\widehat{P}_0^-) + \sum_{k=1}^{n} \left(\phi^{n-k}(p_k^-) - \phi^{n-(k-1)}(p_{k-1}^-) \right).$$

- Use the Lipschitz estimates for ϕ^n to obtain bounds on the summands for the moment estimates.

- The CLT requires more delicate treatment: need to use the second order Taylor expansion type bounds and then the Lipschitz estimates for the first derivative.

Idea behind proofs

- Consider the following decomposition

$$p_n^- - \widehat{P}_n^- = \phi^n(p_0^-) - \phi^n(\widehat{P}_0^-) + \sum_{k=1}^{n}\left(\phi^{n-k}(p_k^-) - \phi^{n-(k-1)}(p_{k-1}^-)\right).$$

- Use the Lipschitz estimates for ϕ^n to obtain bounds on the summands for the moment estimates.

- The CLT requires more delicate treatment: need to use the second order Taylor expansion type bounds and then the Lipschitz estimates for the first derivative.

Idea behind proofs

- Consider the following decomposition

$$p_n^- - \widehat{P}_n^- = \phi^n(p_0^-) - \phi^n(\widehat{P}_0^-) + \sum_{k=1}^{n} \left(\phi^{n-k}(p_k^-) - \phi^{n-(k-1)}(p_{k-1}^-) \right).$$

- Use the Lipschitz estimates for ϕ^n to obtain bounds on the summands for the moment estimates.

- The CLT requires more delicate treatment: need to use the second order Taylor expansion type bounds and then the Lipschitz estimates for the first derivative.

Stability of the sample means

- Now define $M_n := m_n - X_n$.
- Observe that
$$M_{n+1} = \frac{A}{1+(C/D)^2 p_n} M_n + \Upsilon_{n+1},$$
where Υ_n is a conditionally centred Gaussian random variable.
- Understanding the stability of the sample means thus reduces to understanding the behaviour of the products
$$\mathcal{E}_{l,n} := \prod_{k=l}^{n} \frac{A}{1+(C/D)^2 p_k}.$$
- Similar theorems to those presented for the sample covariances and the corresponding Kalman gain hold for M_n.

Stability of the sample means

- Now define $M_n := m_n - X_n$.
- Observe that

$$M_{n+1} = \frac{A}{1 + (C/D)^2 p_n} M_n + \Upsilon_{n+1},$$

where Υ_n is a conditionally centred Gaussian random variable.

- Understanding the stability of the sample means thus reduces to understanding the behaviour of the products

$$\mathcal{E}_{l,n} := \prod_{k=l}^{n} \frac{A}{1 + (C/D)^2 p_k}.$$

- Similar theorems to those presented for the sample covariances and the corresponding Kalman gain hold for M_n.

Stability of the sample means

- Now define $M_n := m_n - X_n$.
- Observe that
$$M_{n+1} = \frac{A}{1 + (C/D)^2 p_n} M_n + \Upsilon_{n+1},$$
where Υ_n is a conditionally centred Gaussian random variable.
- Understanding the stability of the sample means thus reduces to understanding the behaviour of the products
$$\mathcal{E}_{l,n} := \prod_{k=l}^{n} \frac{A}{1 + (C/D)^2 p_k}.$$
- Similar theorems to those presented for the sample covariances and the corresponding Kalman gain hold for M_n.

Stability of the sample means

- Now define $M_n := m_n - X_n$.
- Observe that

$$M_{n+1} = \frac{A}{1 + (C/D)^2 p_n} M_n + \Upsilon_{n+1},$$

where Υ_n is a conditionally centred Gaussian random variable.

- Understanding the stability of the sample means thus reduces to understanding the behaviour of the products

$$\mathcal{E}_{l,n} := \prod_{k=l}^{n} \frac{A}{1 + (C/D)^2 p_k}.$$

- Similar theorems to those presented for the sample covariances and the corresponding Kalman gain hold for M_n.

Problem: filter divergence

- In geophysical models, the state dimension is $\mathcal{O}(10^8)$ but in the EnKF, the ensemble size is $\mathcal{O}(10^2)$.
- This means that the spread of the particles is not sufficient and can lead to underestimation of the covariance.
- This can causes filter divergence.

Problem: filter divergence

- In geophysical models, the state dimension is $\mathcal{O}(10^8)$ but in the EnKF, the ensemble size is $\mathcal{O}(10^2)$.
- This means that the spread of the particles is not sufficient and can lead to underestimation of the covariance.
- This can causes filter divergence.

Problem: filter divergence

- In geophysical models, the state dimension is $\mathcal{O}(10^8)$ but in the EnKF, the ensemble size is $\mathcal{O}(10^2)$.
- This means that the spread of the particles is not sufficient and can lead to underestimation of the covariance.
- This can causes filter divergence.

Problem: filter divergence

- In geophysical models, the state dimension is $\mathcal{O}(10^8)$ but in the EnKF, the ensemble size is $\mathcal{O}(10^2)$.
- This means that the spread of the particles is not sufficient and can lead to underestimation of the covariance.
- This can causes filter divergence.

Solution: Covariance inflation

- One solution to this problem is covariance inflation.
- The most common form of this is to make the following replacement

$$\xi_{N,i}^{-} - m_n^{-} \leftarrow r(\xi_{N,i}^{-} - m_n^{-}),$$

where r is the inflation factor. Similarly for $\xi_{N,i}$.

- This increases the sample covariances and thus stabilises the filter.
- Usually r is chosen to be slightly greater than 1 but the choice is somewhat heuristic.

Solution: Covariance inflation

- One solution to this problem is covariance inflation.
- The most common form of this is to make the following replacement

$$\xi^-_{N,i} - m^-_n \leftarrow r(\xi^-_{N,i} - m^-_n),$$

where r is the inflation factor. Similarly for $\xi_{N,i}$.

- This increases the sample covariances and thus stabilises the filter.
- Usually r is chosen to be slightly greater than 1 but the choice is somewhat heuristic.

Solution: Covariance inflation

- One solution to this problem is covariance inflation.
- The most common form of this is to make the following replacement

$$\xi^-_{N,i} - m^-_n \leftarrow r(\xi^-_{N,i} - m^-_n),$$

where r is the inflation factor. Similarly for $\xi_{N,i}$.

- This increases the sample covariances and thus stabilises the filter.
- Usually r is chosen to be slightly greater than 1 but the choice is somewhat heuristic.

Solution: Covariance inflation

- One solution to this problem is covariance inflation.
- The most common form of this is to make the following replacement

$$\xi_{N,i}^{-} - m_n^{-} \leftarrow r(\xi_{N,i}^{-} - m_n^{-}),$$

where r is the inflation factor. Similarly for $\xi_{N,i}$.

- This increases the sample covariances and thus stabilises the filter.
- Usually r is chosen to be slightly greater than 1 but the choice is somewhat heuristic.

Problem: spurious correlations

- With small sample sizes, the ranks of the covariance matrices are much smaller than the dimension of the state.
- This causes spurious correlations between unrelated state variables, e.g. a large correlation between the temperature at two distant locations on the globe.
- This then propagates through the update step $\rightsquigarrow$ state components that are uncorrelated to the observations, Y_n, are erroneously updated.

Problem: spurious correlations

- With small sample sizes, the ranks of the covariance matrices are much smaller than the dimension of the state.
- This causes spurious correlations between unrelated state variables, e.g. a large correlation between the temperature at two distant locations on the globe.
- This then propagates through the update step $\rightsquigarrow$ state components that are uncorrelated to the observations, Y_n, are erroneously updated.

Problem: spurious correlations

- With small sample sizes, the ranks of the covariance matrices are much smaller than the dimension of the state.
- This causes spurious correlations between unrelated state variables, e.g. a large correlation between the temperature at two distant locations on the globe.
- This then propagates through the update step $\rightsquigarrow$ state components that are uncorrelated to the observations, Y_n, are erroneously updated.

Solution: covariance localisation

- One solution is covariance localisation or covariance tapering.
- Replace the covariance matrices by $\rho \circ p_n$, where $\circ$ denotes the Schur product.
- Here ρ is a positive semi-definite matrix with ones on the diagonal and decays smoothly to zero for unwanted off-diagonal elements.
- Usually ρ is full-rank so that the resulting covariance matrix is full rank.

Solution: covariance localisation

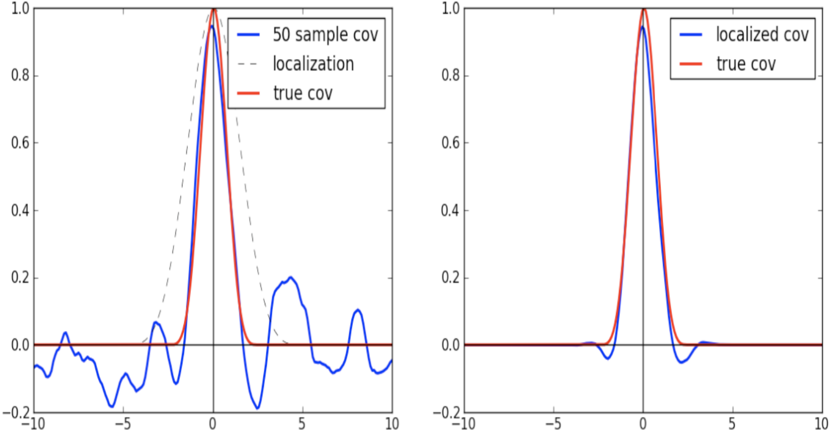

Solution: domain localisation

- Essentially divide and conquer.
- Idea is to split the state space into subdomains, update state estimates locally and then stick them back together.
- Applicable only if the long-range error correlations are negligible.

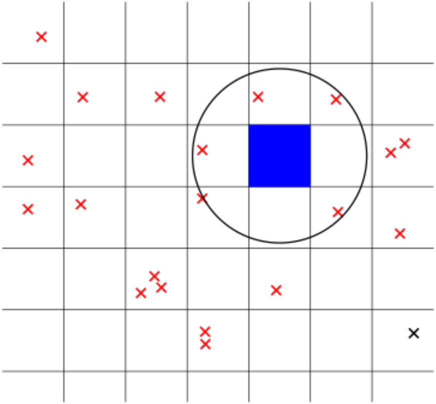

Future work/open problems

- Higher dimensions
- Stability analysis for unstable signals for other genetic-type particle filters
- Time-varying systems
- Genealogies of particle filters
- Plenty of other particle filters..!