14 Continuous-time Markov processes with partial observations

14.1 The Kalman–Bucy filter

Recall the LQR problem in continuous time: we have a diffusion process on $\mathbb{R}^n$ with controlled drift $b(x,u) = Ax + Bu$ and constant diffusion matrix $FF^\top$, with some given matrices $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$ and $F \in \mathbb{R}^{n\times n}$. Thus, the generator corresponding to a given action $u \in \mathbb{R}^m$ is

$$\mathcal{A}^u f(x) = (Ax + Bu)^\top \nabla f(x) + \frac{1}{2}\operatorname{tr}(FF^\top \nabla^2 f(x)). \tag{1}$$

In the preceding lecture, we have covered the finite-horizon control problem with running cost $c(x,u) = x^\top Qx + u^\top Ru$ with terminal cost $c_T(x) = x^\top Q_T x$, where the control at each time t was allowed to depend on the state trajectory $\{X_s : 0 \le s \le t\}$. We now consider the case of *partial observations*, where the state process $\{X_t\}$ is hidden from view and instead the control at time t must be based on the observation trajectory $\{Y_s : 0 \le s \le t\}$. More precisely, we consider a controlled continuous-time pair process $\{(X_t, Y_t)\}_{t\ge 0}$ described by the Itô SDEs

$$\mathrm{d}X_t = (AX_t + BU_t)\,\mathrm{d}t + F\,\mathrm{d}W_t \tag{2a}$$
$$\mathrm{d}Y_t = CX_t\,\mathrm{d}t + G\,\mathrm{d}V_t, \tag{2b}$$

where the $\mathbb{R}^p$-valued observation process is specified by the matrices $C \in \mathbb{R}^{p\times n}$ and $G \in \mathbb{R}^{p\times p}$, and the $\mathbb{R}^m$-valued control process $U = \{U_t\}_{t\ge 0}$ is assumed to be adapted to the filtration $\{\mathcal{F}^Y_t\}$ associated with the observation process, i.e., $\mathcal{F}^Y_t := \sigma(Y_s : 0 \le s \le t)$ is the σ-algebra generated by the observations up to time t. The state noise $W = \{W_t\}_{t\ge 0}$ and the observation noise $V = \{V_t\}_{t\ge 0}$ are two independent standard Wiener processes with values in $\mathbb{R}^n$ and $\mathbb{R}^p$, respectively. Moreover, we make the following assumptions:

1. The initial state X_0 is Gaussian with mean μ_0 and covariance matrix K_0.
2. The matrix F is invertible.

The objective is to compute the conditional mean $\widehat{X}_t := \mathbf{E}[X_t|\mathcal{F}^Y_t]$ and the error covariance $K_t := \mathbf{E}[(X_t - \widehat{X}_t)(X_t - \widehat{X}_t)^\top|\mathcal{F}^Y_t]$. We will prove the following result:

Theorem 14.1 (Kalman–Bucy) *For all sufficiently regular $\{\mathcal{F}^Y_t\}$-adapted control processes $U = \{U_t\}_{t\ge 0}$, the linear filter $\{(\widehat{X}_t, K_t)\}_{t\ge 0}$ is given by*

$$\mathrm{d}\widehat{X}_t = A\widehat{X}_t\,\mathrm{d}t + BU_t\,\mathrm{d}t + K_t(G^{-1}C)^\top\,\mathrm{d}\bar{V}_t, \qquad \widehat{X}_0 = \mu_0 \tag{3a}$$
$$\frac{\mathrm{d}K_t}{\mathrm{d}t} = AK_t + K_tA^\top - K_tC^\top(GG^\top)^{-1}CK_t + FF^\top \tag{3b}$$

where the innovations process $\mathrm{d}\bar{V}_t = G^{-1}(\mathrm{d}Y_t - C\widehat{X}_t\,\mathrm{d}t)$ is an n-dimensional Wiener process.

The overall idea is, more or less, the same as in the discrete-time case: We first consider the case of no controls (i.e., $U_t \equiv 0$ for all t) and show that the problem of computing the linear filter can be reduced to a *deterministic* continuous-time LQR problem. Then we will show that, under suitable regularity conditions on U, the same solution still works. The main concept here is that the conditional mean $\widehat{X}_t = \mathbf{E}[X_t|\mathcal{F}^Y_t]$ is the best estimate of X_t given $\{Y_s : 0 \le s \le t\}$ in the mean-square sense, and the computation of $\widehat{X}_t$ can be thus turned into a deterministic optimal control problem. Another useful fact that we will exploit is the following:

Lemma 14.1 *Given any admissible (i.e., $\{\mathcal{F}^Y_t\}$-adapted) control U, the state process X can be written as*

$$X_t = e^{At}X_0 + \int_0^t e^{A(t-s)}BU_t \, \mathrm{d}t + \int_0^t e^{A(t-s)}F \, \mathrm{d}W_t. \tag{4}$$

Proof: Consider the process $Z_t := e^{-At}X_t$ and apply Itô's lemma. ■

14.1.1 The case of no controls

When $U_t \equiv 0$ for all t, Lemma 14.1 tells us that the state process is given by

$$X_t = e^{At}X_0 + \int_0^t e^{A(t-s)}F \, \mathrm{d}W_t, \qquad t \ge 0. \tag{5}$$

Since $X_0 \sim \mathcal{N}(\mu_0, K_0)$, $\{X_t\}_{t\ge0}$ is a *Gaussian process* governed by the linear SDE $\mathrm{d}X_t = AX_t \, \mathrm{d}t + F \, \mathrm{d}W_t$. Since the observation model $\mathrm{d}Y_t = CX_t \, \mathrm{d}t + G \, \mathrm{d}V_t$ is also linear, the state-observation process $\{(X_t, Y_t)\}_{t\ge0}$ is also Gaussian. This fact is immensely helpful because, just as in the discrete-time case, the conditional mean $\widehat{X}_t = \mathbf{E}[X_t|\mathcal{F}^Y_t]$ turns out to be a *linear functional* of $\{Y_s : 0 \le s \le t\}$. Let us spell this out more precisely.

Let us first center the processes X and Y:

$$\tilde{X}_t := X_t - \mathbf{E}[X_t], \qquad \tilde{Y}_t := \int_0^t C\tilde{X}_t \, \mathrm{d}t + \int_0^t G \, \mathrm{d}V_t. \tag{6}$$

Then evidently $\mathcal{F}^Y_t = \mathcal{F}^{\tilde{Y}}_t$ for each t and $(\tilde{X}_t, \tilde{Y}_t)$ is a zero-mean Gaussian process. Fix some t. For each $k = 1, 2, \ldots$, consider the σ-algebra

$$\mathcal{G}_k := \sigma\left(\tilde{Y}_{j2^{-k}t} - \tilde{Y}_{(j-1)2^{-k}t} : j = 1, 2, \ldots, 2^k\right) \tag{7}$$

and consider the conditional mean $\mathbf{E}[\tilde{X}_t|\mathcal{G}_k]$. Moreover, the random vector

$$\tilde{Y}^k := \left(\tilde{Y}_{2^{-k}t}, \tilde{Y}_{2\cdot2^{-k}t} - \tilde{Y}_{2^{-k}t}, \ldots, \tilde{Y}_t - \tilde{Y}_{(1-2^{-k})t}\right), \tag{8}$$

is Gaussian, and therefore

$$\mathbf{E}[\tilde{X}_t|\tilde{Y}^k] = \mathrm{Cov}(\tilde{X}_t, \tilde{Y}^k) \cdot \mathrm{Cov}(\tilde{Y}^k)^{-1}\tilde{Y}^k \tag{9}$$

(recall that both $\tilde{X}_t$ and $\tilde{Y}_t$ are zero-mean). In particular, there exist deterministic matrices H^k_j, $j = 1, 2, \ldots, 2^k$, such that

$$\begin{aligned}\mathbf{E}[\tilde{X}_t|\mathcal{G}_k] &= H^k_1\tilde{Y}_{2^{-k}t} + \sum_{j=1}^{2^k-1} H^k_{j+1}(\tilde{Y}_{(j+1)2^{-k}t} - \tilde{Y}_{j2^{-k}t}) && (10)\\ &= \int_0^t H^k(t,s) \, \mathrm{d}\tilde{Y}_s, && (11)\end{aligned}$$

where $s \mapsto H^k(t,s)$ is the deterministic piecewise constant matrix-valued function

$$H^k(t,s) := H^k_j, \qquad t \in [(j-1)2^{-k}t, j2^{-k}t). \tag{12}$$

Now, since $\{Y_t\}_{t \ge 0}$ is a diffusion process and thus has continuous sample paths, and since the times $j2^{-k}t$, $k = 1,2,\ldots$, $j = 1,2,\ldots,2^k$ are dense in the interval $[0,t]$, the σ-algebra $\mathcal{F}^{\tilde{Y}}_t$ is the smallest σ-algebra containing $\mathcal{G}_1, \mathcal{G}_2, \ldots$. Using this and the result known as Lévy's upward theorem, it can be shown that $\mathbf{E}[\tilde{X}_t|\mathcal{G}_k] \to \mathbf{E}[X_t|\mathcal{F}^Y_t]$ in mean square as $k \to \infty$. The upshot of all this is that, for each t, there exists a deterministic function $s \mapsto H(t,s)$ with values in $\mathbb{R}^{n \times p}$, such that

$$\mathbf{E}[X_t|\mathcal{F}^Y_t] = \mathbf{E}[X_0] + \int_0^t H(t,s)C(X_s - \mathbf{E}[X_s])\,\mathrm{d}s + \int_0^t H(t,s)G\,\mathrm{d}V_s, \tag{13}$$

such that $\int_0^t \|H(t,s)\|^2\,\mathrm{d}s < \infty$ (here, $\|\cdot\|$ is the Frobenius norm). Thus, $\widehat{X}_t = \mathbf{E}[X_t|\mathcal{F}^Y_t]$ is, indeed, a linear functional of $\{Y_s : 0 \le s \le t\}$.

Now the plan of attack is clear: Among all linear functionals of $\{Y_s : 0 \le s \le t\}$ of the form given by the right-hand side of (13), $\mathbf{E}[X_t|\mathcal{F}^Y_t]$ is the one that minimizes the mean squared error. We will now show that the problem of finding $H(t,s)$ can be cast as a deterministic LQR problem, and $H(t,s)$ will be the optimal control.

To that end, let us fix the horizon $T < \infty$. Given any deterministic function $H : [0,T] \to \mathbb{R}^{n \times p}$ with $\int_0^T \|H(t)\|^2\,\mathrm{d}t < \infty$, consider the Itô process

$$L^H_t := \mathbf{E}[X_t] + \int_0^t H(s)C(X_s - \mathbf{E}[X_s])\,\mathrm{d}s + \int_0^t H(s)G\,\mathrm{d}V_s, \tag{14}$$

which is $\{\mathcal{F}^Y_t\}$-adapted (why?). Given a vector $v \in \mathbb{R}^n$, consider the cost

$$J_v(H) := \mathbf{E}[(v^\top(X_T - L^H_T))^2]. \tag{15}$$

Then, for any $v \in \mathbb{R}^n$, by the properties of the conditional expectation,

$$\mathbf{E}[v^\top X_T|\mathcal{F}^Y_T] = v^\top \mathbf{E}[X_T|\mathcal{F}^Y_T] = J_v(H^*), \tag{16}$$

where $H^* : [0,T] \to \mathbb{R}^{n \times p}$ is the function that minimizes $J_v(H)$ for every v. Remarkably, as we will now show, the problem of finding H^* is actually a deterministic LQR problem!

To that end, for any $t \in [0,T]$ consider the matrix-valued ODE

$$\frac{\mathrm{d}}{\mathrm{d}s}\Psi^H_{s,t} = -\Psi^H_{s,t}A + H(s)C, \qquad \Psi^H_{t,t} = I_n. \tag{17}$$

Consider the process $U_t := \Psi^H_{t,T}(X_t - \mathbf{E}[X_t])$. Then $U_T = X_T - \mathbf{E}[X_T]$, and Itô's lemma gives

$$X_T - \mathbf{E}[X_T] = \Psi^H_{0,T}(X_0 - \mathbf{E}[X_0]) + \int_0^T H(s)C(X_s - \mathbf{E}[X_s])\,\mathrm{d}s + \int_0^T \Psi^H_{s,T}F\,\mathrm{d}W_s. \tag{18}$$

Consequently,

$$\begin{aligned} X_T - L_T^H &= X_T - \mathbf{E}[X_T] + \mathbf{E}[X_T] - L_t^H \qquad (19)\\ &= \Psi^H_{0,T}(X_0 - \mathbf{E}[X_0]) + \int_0^T H(s)C(X_s - \mathbf{E}[X_s])\,\mathrm{d}s + \int_0^T \Psi^H_{s,T} F\,\mathrm{d}W_s \\ &\qquad - \int_0^t H(s)C(X_s - \mathbf{E}[X_s])\,\mathrm{d}s - \int_0^t H(s)G\,\mathrm{d}V_s \qquad (20)\\ &= \Psi^G_{0,T}(X_0 - \mathbf{E}[X_0]) + \int_0^T \Psi^H_{s,T}F\,\mathrm{d}W_s - \int_0^t H(s)G\,\mathrm{d}V_t. \qquad (21) \end{aligned}$$

The whole point of these manipulations is that the resulting expression for $X_T - L_T^H$ does not involve any time integrals, only the initial condition and two Itô integrals with respect to two independent Wiener processes. Now, for any deterministic vector $v \in \mathbb{R}^n$,

$$v^\top (X_T - L_T^H) = v^\top \Psi^H_{0,T}(X_0 - \mathbf{E}[X_0]) + \int_0^T v^\top \Psi^H_{s,T} F\,\mathrm{d}W_s - \int_0^t v^\top H(s) G\,\mathrm{d}V_s, \qquad (22)$$

and we can compute the expected value of $(v^\top(X_T - L_T^H))^2$ by the Itô isometry:

$$\begin{aligned} J_v(H) &= \mathbf{E}[(v^\top(X_T - L_T^H))^2] \qquad (23)\\ &= \int_0^T \left\{ \|F^\top (\Psi^H_{t,T})^\top v\|^2 + \|G^\top H(t)^\top v\|^2 \right\} \mathrm{d}t + v^\top \Psi^H_{0,T} K_0 (\Psi^H_{0,T})^\top v. \qquad (24) \end{aligned}$$

Now let $\xi_t := (\Psi^H_{T-t,T})^\top v$ and $\alpha_t := H(T-t)^\top v$. Using these definitions and (17), we obtain the following deterministic ODE:

$$\frac{\mathrm{d}}{\mathrm{d}t}\xi_t = A^\top \xi_t - C^\top \alpha_t, \qquad \xi_0 = v \qquad (25)$$

and we can now express the cost $J_v(H)$ in terms of α as

$$\begin{aligned} J_v(H) &= J(\alpha) \\ &= \int_0^T \left\{ \xi_t^\top F F^\top \xi_t + \alpha_t^\top G G^\top \alpha_t \right\} \mathrm{d}t + \xi_T^\top K_0 \xi_T. \qquad (26) \end{aligned}$$

This is a deterministic LQR problem in disguise: the $\mathbb{R}^n$-valued state ξ_t with initial condition $\xi_0 = v$ is controlled by means of the $\mathbb{R}^p$-valued control α over the finite horizon T, with running cost $c(\xi, \alpha) = \xi^\top F F^\top \xi + \alpha^\top G G^\top \alpha$ and terminal cost $c_T(\xi) = \xi^\top K_0 \xi$, where $GG^\top \succ 0$ by our assumption on G. In particular, we know that the optimal control $\alpha^* = \{\alpha^*_t\}$ achieves

$$J(\alpha^*) = \min_\alpha J(\alpha) = v^\top \tilde{K}_0 v, \qquad (27)$$

where $\tilde{K}_0$ is obtained by solving the Riccati differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{K}_t = -\tilde{K}_t A^\top - A\tilde{K}_t + \tilde{K}_t C^\top (GG^\top)^{-1} C \tilde{K}_t - FF^\top, \qquad \tilde{K}_T = K_0 \qquad (28)$$

and

$$\alpha^*_t = -(GG^\top)^{-1} C \tilde{K}_t \xi_t. \tag{29}$$

Defining $K_t := \tilde{K}_{T-t}$, we obtain the ODE

$$\frac{\mathrm{d}}{\mathrm{d}t} K_t = K_t A^\top + A K_t - K_t C^\top (GG^\top)^{-1} C K_t - FF^\top, \tag{30}$$

which is exactly the ODE (3b) for the error covariance $\mathbf{E}[(X_t - \widehat{X}_t)(X_t - \widehat{X}_t)^\top | \mathcal{F}^Y_t]$. It is straightforward but tedious to fill in the remaining details.

The proof that the innovations process $\mathrm{d}\bar{V}_t$ is a Wiener process uses a Girsanov-type martingale argument, and is omitted.

14.1.2 The controlled case

Now we will show that, just as in the discrete-time case, the structure of the linear filter is exactly the same for any sufficiently regular $\{\mathcal{F}^Y_t\}$-adapted control process U. This is a consequence of the linear structure of the problem. First we recall that, by Lemma 14.1, the state process is given by

$$\begin{aligned} X^U_t &= e^{At} X_0 + \int_0^t e^{A(t-s)} B U_s \,\mathrm{d}s + \int_0^t e^{A(t-s)} \,\mathrm{d}F W_s & (31)\\ &= X^0_t + \int_0^t e^{A(t-s)} B U_s \,\mathrm{d}s, & (32) \end{aligned}$$

where we have denoted by X^U the controlled state process and by X^0 the uncontrolled process (the one with $U_t \equiv 0$ for all t). If we denote by Y^U the observation process under the control U, i.e.,

$$Y^U_t = \int_0^t C X^U_s \,\mathrm{d}s + \int_0^t G \,\mathrm{d}V_s \tag{33}$$

and define the σ-algebra $\mathcal{F}^{Y,U}_t := \sigma(Y^U_t : 0 \le s \le t)$, then the second term in (32) is $\{\mathcal{F}^{Y,U}_t\}$-adapted. Thus, by linearity of (conditional) expectation,

$$\begin{aligned} \widehat{X}^U_t &:= \mathbf{E}[X^U_t | \mathcal{F}^{Y,U}_t] & (34)\\ &= \mathbf{E}[X^0_t | \mathcal{F}^{Y,U}_t] + \int_0^t e^{A(t-s)} B U_s \,\mathrm{d}s. & (35) \end{aligned}$$

The following result is now trivial:

Theorem 14.2 (Kalman–Bucy, controlled case) *Suppose that the control process $U = \{U_t\}_{t \ge 0}$ is such that the σ-fields $\mathcal{F}^{Y,U}_t$ and $\mathcal{F}^{Y,0}_t$ are equal for all $t < \infty$, and $\mathbf{E}[\|X^U_t - X^0_t\|] < \infty$ for all t. Then Theorem 14.1 holds.*

Proof: If $\mathcal{F}^{Y,U}_t = \mathcal{F}^Y_t$ for all t, then $\mathbf{E}[X^0_t | \mathcal{F}^{Y,U}_t] = \mathbf{E}[X^0_t | \mathcal{F}^Y_t] = \widehat{X}^0_t$, so

$$\widehat{X}_t = \widehat{X}^0_t + \int_0^t e^{A(t-s)} B U_s \,\mathrm{d}s \tag{36}$$

is easily seen to be given by (3a), with K_t evolving according to (3b). The claim that $\mathbf{E}[(X_t^U - \widehat{X}_t^U)(X_t^U - \widehat{X}_t^U)^\top | \mathcal{F}_t^{Y,U}] = K_t$ follows from two facts: $X_t^U - \widehat{X}_t^U = X_t^0 - \widehat{X}_t^0$ and $\mathcal{F}_t^{Y,U} = \mathcal{F}_t^Y$. Finally, since $X_t^U - \widehat{X}_t^U = X_t^0 - \widehat{X}_t^0$, it follows that the innovations process $\bar{V}_t^U$ is the same as $\bar{V}_t^0 = \bar{V}_t$. ■

The conditions of the theorem hold, for example, when $U_t = g(X_t, t)$ for some function $g : \mathbb{R}^n \times [0, \infty) \to \mathbb{R}^m$ which is Lipschitz in x uniformly in t, i.e., there exists some constant $0 < C < \infty$, such that

$$\|g(x,t) - g(x',t)\| \le C\|x - x'\| \tag{37}$$

for all $x, x' \in \mathbb{R}^n$, $t \ge 0$.

14.2 LQG control with partial observations

We now consdier the control problem (2) and wish to minimize the expected cost

$$J_T(U) := \mathbf{E}\left[\int_0^T \left\{(X_t^U)^\top Q X_t^U + U_t^\top R U_t\right\} \mathrm{d}t + (X_T^U)^\top Q_T X_T^U\right] \tag{38}$$

over all $\mathbb{R}^m$-valued $\{\mathcal{F}_t^{Y,U}\}$-adapted controls that satisfy the conditions of Theorem 14.2. As usual Q and Q_T are symmetric positive semidefinite $n \times n$ matrices and R is a symmetric positive definite $m \times m$ matrix. We have the following result

Theorem 14.3 (continuous-time LQG problem, partial observations) *Let N_t denote the solution of the Riccati equation*

$$\frac{\mathrm{d}}{\mathrm{d}t} N_t = A N_t + N_t A^\top - N_t C^\top (G G^\top)^{-1} C N_t + F F^\top \tag{39}$$

where N_0 is the covariance matrix of X_0, and let M_t denote the solution of the time-reversed Riccati equation

$$\frac{\mathrm{d}}{\mathrm{d}t} M_t + A^\top M_t + M_t A - M_t B R^{-1} B^\top M_t + Q = 0 \tag{40}$$

with the terminal condition $M_T = Q_T$. Then an optimal strategy for the LQG problem (2) is given by the estimated state feedback law $U_t^* = -R^{-1} B^\top M_t \widehat{X}_t$, where $\widehat{X}_t$ evolves according to

$$\mathrm{d}\widehat{X}_t = (A - BR^{-1}B^\top M_t)\widehat{X}_t \,\mathrm{d}t + N_t (G^{-1}C)^\top G^{-1}(\mathrm{d}Y_t^{U^*} - C\widehat{X}_t \,\mathrm{d}t), \qquad \widehat{X}_0 = \mathbf{E}[X_0] \tag{41}$$

and $\widehat{X}_t = \widehat{X}_t^{U^*}$, $N_t = K_t$ are the optimal state estimate and error covariance for $X_t^{U^*}$.

Proof: We proceed as in the discrete-time case. For any control U satisfying our regularity assumptions, the Kalman-Bucy theorem holds and

$$\mathbf{E}[(X_t^U)^\top Q X_t^U] = \mathbf{E}[(\widehat{X}_t^U)^\top Q \widehat{X}_t^U] + \mathbf{E}[(X_t^U - \widehat{X}_t^U)^\top Q (X_t^U - \widehat{X}_t^U)] \tag{42}$$

$$= \mathrm{tr}[Q K_t], \tag{43}$$

where K_t is the Kalman filter error covariance, which does not depend on U (or on the observations). Then

$$J_T(U) = J'_T(U) + \int_0^T \operatorname{tr}[QK_t]\,\mathrm{d}t + \operatorname{tr}[Q_T K_T], \tag{44}$$

where

$$J'_T(U) = \mathbf{E}\left[\int_0^T \left\{(\widehat{X}^U_t)^\top Q \widehat{X}^U_t + U_t^\top R U_t\right\}\mathrm{d}t + (\widehat{X}^U_T)^\top Q_T \widehat{X}^U_T\right]. \tag{45}$$

Thus, any admissible control U^* that minimizes $J'_T(U)$ also minimizes $J_T(U)$. Now, by Theorem 14.1, the estimated state $\widehat{X}^U_t$ satisfies the SDE

$$\mathrm{d}\widehat{X}^U_t = A\widehat{X}^U_t\,\mathrm{d}t + BU_t\,\mathrm{d}t + K_t(G^{-1}C)^\top\,\mathrm{d}\bar{V}_t, \tag{46}$$

where $\bar{V}$ is a Wiener process. This, together with the cost $J'_T(\cdot)$, defines a *fully observed* LQR problem, where the optimal control is given by $U^*_t = g^*(\widehat{X}^{U^*}_t, t)$ with $g^*(x,t) := -R^{-1}B^\top M_t x$. Since $g(x,t)$ is uniformly Lipschitz on $\mathbb{R}^n \times [0,T]$, and therefore admissible. ■

14.3 Nonlinear filtering in continuous time

Now we will consider a more general problem: Let $X = \{X_t\}_{t\ge 0}$ be a continuous-time Markov process with state space $\mathcal{X}$ and generator $\mathcal{A}$, so that, in particular, for any bounded measurable function $f : \mathcal{X} \to \mathbb{R}$,

$$M^f_t := f(X_t) - f(X_0) - \int_0^t \mathcal{A}f(X_s)\,\mathrm{d}s \tag{47}$$

is an $\mathcal{F}^X_t$-martingale: for any $0 \le s \le t$, $\mathbf{E}[M^f_t - M^f_s|\mathcal{F}^X_s] = M^f_s$. For example, X could be a finite-state process with $\mathcal{A}f = \Lambda f$, where Λ is the transition intensities matrix of X, or it could be a diffusion process with $\mathcal{A}f = b^\top \nabla f + \frac{1}{2}\operatorname{tr}[\sigma\sigma^\top \nabla^2 f]$, where b and σ are the drift and the diffusion coefficients of X. We now assume that X is not available for observation, but instead we observe a noisy version

$$Y_t = \int_0^t h(X_s)\,\mathrm{d}s + V_t, \tag{48}$$

where $h : \mathcal{X} \to \mathbb{R}$ is a given function and $V = \{V_t\}_{t\ge 0}$ is a standard one-dimensional Wiener process. (We could just as well consider the case of vector-valued observations, but that would only mean extra notation.) The *nonlinear* (or *Bayesian*) *filtering problem* is about computing the posterior distribution $\pi_t(\cdot) := \mathbf{P}[X_t \in \cdot|\mathcal{F}^Y_t]$, for each $t \ge 0$. More precisely, given a bounded measurable function $f : \mathcal{X} \to \mathbb{R}$, we wish to compute the conditional mean $\pi_t(f) := \mathbf{E}[f(X_t)|\mathcal{F}^Y_t]$.

In the preceding section, we have solved this problem for the special case when X was a diffusion process governed by the Itô SDE $\mathrm{d}X_t = AX_t\,\mathrm{d}t + F\,\mathrm{d}W_t$ and where Y was a noisy linear functional of X with $h(x) = Cx$. For that particular problem, we were fortunate in that a direct solution was possible by exploiting the linear Gaussian structure of the problem together with the fact that, in that instance, the filtering distribution was Gaussian and we just had to compute its mean and covariance. Now we will consider the general case.

We begin by examining what looks like a fairly boring instance of the problem, i.e., the one with $h(x) \equiv 0$ for all x. Then the observation process Y is just the Wiener process V, and it is independent of X. In othe words, if we denote by μ_X the probability law of the signal X, by μ_Y the probability law of Y, and by $\mathbf{Q}$ the joint probability law of the pair process (X, Y), then $\mathbf{Q} = \mu_X \otimes \mu_Y$, and moreover μ_Y is just the probability law of the Wiener process. On the one hand, under $\mathbf{Q}$, the observation process is completely uninformative about the signal. On the other hand, the computation of π_t is rather simple. More generally, if f is any bounded measurable function of the paths $X_{[0,T]}$ and $Y_{[0,T]}$ for some $T > 0$, then

$$\mathbf{E}_{\mathbf{Q}}[f(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T] = \int f(x_{[0,T]}, Y_{[0,T]})\mu_X(\mathrm{d}x_{[0,T]}), \tag{49}$$

i.e., we just marginalize out the signal. Now let $\mathbf{P}$ denote the probability law of (X, Y) for a given $h : \mathcal{X} \to \mathbb{R}$ and, for each T, let $\mathbf{P}_T$ and $\mathbf{Q}_T$ denote the corresponding marginal laws of $(X_{[0,T]}, Y_{[0,T]})$. We begin by stating the following key fact: for any T, there exists a function Z_T of $(X_{[0,T]}, Y_{[0,T]})$, such that

$$\mathbf{E}_{\mathbf{P}}[f(X_{[0,T]}, Y_{[0,T]})] = \int Z_T(x_{[0,T]}, y_{[0,T]})f(x_{[0,T]}, y_{[0,T]})\mathbf{Q}(\mathrm{d}x_{[0,T]}, \mathrm{d}y_{[0,T]}). \tag{50}$$

In other words, Z_T is the *density* (or *Radon–Nikodym derivative*) $\mathrm{d}\mathbf{P}_T/\,\mathrm{d}\mathbf{Q}_T$ of $\mathbf{P}_T$ with respect to $\mathbf{Q}_T$. Moreover, it has the explicit form

$$\frac{\mathrm{d}\mathbf{P}_T}{\mathrm{d}\mathbf{Q}_T} = Z_T(X_{[0,T]}, Y_{[0,T]}) = \exp\left(\int_0^T h(X_t)\,\mathrm{d}Y_t - \frac{1}{2}\int_0^T |h(X_t)|^2\,\mathrm{d}t\right), \tag{51}$$

where the stochastic integral $\int_0^T h(X_t)\,\mathrm{d}Y_t$ is defined in the same way as before as a mean-square limit of sums of the form

$$\sum_{i=0}^{n-1} h(X_{t_i})(Y_{t_{i+1}} - Y_{t_i}) \tag{52}$$

as the step size $\max_i(t_{i+1} - t_i)$ of the partition $0 = t_0 < t_1 < \ldots < t_{n-1} < t_n = T$ goes to zero. The formula (51) is a very special case of a powerful general result known as *Girsanov's theorem*, and can be found in any textbook on stochastic calculus. To give at least a formal justification for it, fix a partition $0 = t_0 < t_1 < \ldots < t_{n-1} < t_n = T$ and consider the samples $(X_{t_0}, Y_{t_0}), (X_{t_1}, Y_{t_1}), \ldots, (X_{t_n}, Y_{t_n})$. If $\max_i(t_{i+1} - t_i)$ is small, then, under $\mathbf{P}$, the conditional law of Y_{t_i} given $(X_{t_0}, Y_{t_0}), \ldots, (X_{t_{i-1}}, Y_{t_{i-1}})$ is approximately Gaussian with mean $Y_{t_{i-1}} + h(X_{t_{i-1}})(t_i - t_{i-1})$ and variance $t_i - t_{i-1}$; on the other hand, under $\mathbf{Q}$, the conditional law of Y_{t_i} given $(X_{t_0}, Y_{t_0}), \ldots, (X_{t_{i-1}}, Y_{t_{i-1}})$ is Gaussian with mean

$Y_{t_{i-1}}$ and variance $t_i - t_{i-1}$. Therefore, the likelihood ratio $\mathrm{d}\mathbf{P}_T/\,\mathrm{d}\mathbf{Q}_T$ can be approximated by

$$\begin{aligned}
&\prod_{i=1}^n \frac{\exp\left(-\frac{1}{2(t_i-t_{i-1})}(Y_{t_i} - Y_{t_{i-1}} - h(X_{t_{i-1}})(t_i - t_{i-1}))^2\right)}{\exp\left(-\frac{1}{2(t_i-t_{i-1})}(Y_{t_i} - Y_{t_{i-1}})^2\right)} \\
&\qquad = \prod_{i=1}^n \exp\left(-\frac{1}{2(t_i - t_{i-1})}(h^2(X_{t_{i-1}})(t_i - t_{i-1}) - 2(t_i - t_{i-1}))h(X_{t_{i-1}})(Y_{t_i} - Y_{t_{i-1}}))\right) \\
&\qquad \exp\left\{\sum_{i=1}^n h(X_{t_{i-1}})(Y_{t_i} - Y_{t_{i-1}}) - \frac{1}{2}\sum_{i=1}^n |h(X_{t_{i-1}})|^2(t_i - t_{i-1})\right\}, \qquad (53)
\end{aligned}$$

and we obtain (51) as we take $\max_i(t_i - t_{i-1}) \to 0$. Now that we have Z_T, we can compute the conditional expectation $\mathbf{E}_\mathbf{P}[f(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T]$ as follows:

$$\mathbf{E}_\mathbf{P}[f(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T] = \frac{\mathbf{E}_\mathbf{Q}[f(X_{[0,T]}, Y_{[0,T]})Z_T(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T]}{\mathbf{E}_\mathbf{Q}[Z_T(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T]}, \qquad (54)$$

where, since X and Y are independent under $\mathbf{Q}$,

$$\mathbf{E}_\mathbf{Q}[f(X_{[0,T]}, Y_{[0,T]})Z_T(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T] = \int f(x_{[0,T]}, Y_{[0,T]})Z_T(x_{[0,T]}, Y_{[0,T]})\mu_X(\mathrm{d}x_{[0,T]}) \qquad (55)$$

and

$$\mathbf{E}_\mathbf{Q}[Z_T(X_{[0,T]}, Y_{[0,T]})|\mathcal{F}^Y_T] = \int Z_T(x_{[0,T]}, Y_{[0,T]})\mu_X(\mathrm{d}x_{[0,T]}). \qquad (56)$$

Eq. (54) is known as the *Kallianpur–Striebel formula*. We can now write down an explicit expression for the conditional mean $\pi_t(f) := \mathbf{E}_\mathbf{P}[f(X_t)|\mathcal{F}^Y_t]$. If we define the *unnormalized filter* as

$$\sigma_t(f) := \int f(x_t)Z_t(x_{[0,t]}, Y_{[0,t]})\mu_X(\mathrm{d}x_{[0,T]}), \qquad (57)$$

then it follows from the Kallianpur–Striebel formula that

$$\pi_t(f) = \frac{\sigma_t(f)}{\sigma_t(1)}, \qquad (58)$$

where the denominator is obtained by computing the unnofmrlized filtering mean (57) of the constant function $f(x) = 1$.

Now, in the discrete-time case, we had shown that the computation of both the unnormalized and the normalized filtering distributions could be done via recursive updates $\sigma_t \to \sigma_{t+1}$ and $\pi_t \to \pi_{t+1}$. Our next order of business is to show that the continuous-time analogue of this is a pair of stochastic differential equations. The equation for σ_t is known as the *Duncan–Mortensen–Zakai equation* (often referred to as just the Zakai equation), while the one for π_t is known as the *Kushner–Stratonovich equation.*

We will start with the former. The first thing we need is an expression for $f(X_t)Z_t$. To that end, we will need the Itô calculus version of the product rule: for any two Itô processes M and N,

$$\mathrm{d}(M_t N_t) = M_t\,\mathrm{d}N_t + N_t\,\mathrm{d}M_t + \mathrm{d}M_t\,\mathrm{d}N_t, \tag{59}$$

which can usually be simplified further using the rules $\mathrm{d}t \cdot \mathrm{d}W_t = \mathrm{d}t \cdot \mathrm{d}t = 0$ and $\mathrm{d}W_t \cdot \mathrm{d}W_t = \mathrm{d}t$. From (47), we have

$$\mathrm{d}f(X_t) = \mathcal{A}f(X_t)\,\mathrm{d}t + \mathrm{d}M^f_t, \tag{60}$$

and an application of Itô's lemma gives

$$\mathrm{d}Z_t = Z_t h(X_t)\,\mathrm{d}Y_t. \tag{61}$$

Then, using (59), we obtain

$$\begin{aligned} \mathrm{d}(f(X_t)Z_t) &= f(X_t)\,\mathrm{d}Z_t + Z_t\,\mathrm{d}f(X_t) + \mathrm{d}Z_t \cdot \mathrm{d}f(X_t) & (62)\\ &= h(X_t)f(X_t)Z_t\,\mathrm{d}Y_t + \mathcal{A}f(X_t)Z_t\,\mathrm{d}t + Z_t\,\mathrm{d}M^f_t, & (63) \end{aligned}$$

where the term $\mathrm{d}Z_t \cdot \mathrm{d}M^f_t$ is zero since $Y_t M^f_t$ is a martingale. It then follows that

$$\begin{aligned} f(X_t)Z_t &= f(X_0)Z_0 + \int_0^t Z_s f(X_s)h(X_s)\,\mathrm{d}Y_s + \int_0^t Z_s \mathcal{A}f(X_s)\,\mathrm{d}s + \int_0^t Z_s\,\mathrm{d}M^f_s \\ &= f(X_0) + \int_0^t Z_s f(X_s)h(X_s)\,\mathrm{d}Y_s + \int_0^t Z_s \mathcal{A}f(X_s)\,\mathrm{d}s + \int_0^t Z_s\,\mathrm{d}M^f_s, & (64) \end{aligned}$$

where we have used the fact that $Z_0 = 1$. Now we take expectation of both sides of the above equation with respect to $\mathbf{P}$ to get

$$\sigma_t(f) = \nu(f) + \int_0^t \sigma_s(\mathcal{A}f)\,\mathrm{d}s + \int_0^t \sigma_s(hf)\,\mathrm{d}Y_s, \tag{65}$$

where $\nu(\cdot) := \mathbf{P}[X_0 \in \cdot]$, $(hf)(x) = h(x)f(x)$, and where we have used the fact that, since M^f is a martingale, the expected value of the stochastic integral $\int_0^t Z_s\,\mathrm{d}M^f_s$ is zero. Eq. (65) is the Zakai equation for the unnormalized filter. To obtain the Kushner–Stratonovich equation, we apply Itô's lemma to the Kallianpur–Striebel formula $\pi_t(f) = \sigma_t(f)/\sigma_t(1)$. The calculation is straightforward if a bit tedious, and it gives

$$\pi_t(f) = \nu(f) + \int_0^t \pi_s(\mathcal{A}f)\,\mathrm{d}s + \int_0^t \left\{\pi_s(hf) - \pi_s(h)\pi_s(f)\right\}(\mathrm{d}Y_s - \pi_s(h)\,\mathrm{d}s). \tag{66}$$