Proximal Recursion for the Wonham Filter

Abhishek Halder (UC Santa Cruz)

&

Tryphon T. Georgiou (UC Irvine)

CDC, Nice, December 2019

State Estimation for Continuous Time Markov Chain

- $X(t) \sim \text{Markov}\left(\boldsymbol{Q}\right)$ on some finite state space $\Omega = \{a_1, \ldots, a_m\}$.

- The $m \times m$ transition rate matrix $\boldsymbol{Q}$ satisfies $Q_{ij} \geq 0$ for $i \neq j$, $Q_{ii} = -\sum_{j \neq i} Q_{ij} < 0$.

- Assume: the Markov chain is time homogeneous, i.e., the transition probability matrix is $\exp\left(t\boldsymbol{Q}\right)$, $\forall t \geq 0$.

- Given initial occupation probability row vector $\boldsymbol{\pi}_0 \in \Delta^{m-1}$ (standard simplex in $\mathbb{R}^m$)

The Nonlinear Estimation Problem

Dynamics:

state model: $X(t) \sim \text{Markov}\,(\boldsymbol{Q})\,, \quad \boldsymbol{\pi}_0 \in \Delta^{m-1}$

observation model: $\mathrm{d}Z(t) = h\,(X(t))\,\mathrm{d}t + \sigma_V(t)\,\mathrm{d}t$

- $h(\cdot)$ is deterministic injective function of state.
- $\sigma_V(t) \in C^1$, bounded away from zero for all $t \geq 0$.
- Standard Wiener process $V(t)$ is indep. of $X(t)$.

The Nonlinear Estimation Problem

Dynamics:

state model: $X(t) \sim \text{Markov}\,(Q)\,, \quad \boldsymbol{\pi}_0 \in \Delta^{m-1}$

observation model: $\mathrm{d}Z(t) = h\,(X(t))\,\mathrm{d}t \,+\, \sigma_V(t)\,\mathrm{d}t$

- $h(\cdot)$ is deterministic injective function of state.
- $\sigma_V(t) \in C^1$, bounded away from zero for all $t \geq 0$.
- Standard Wiener process $V(t)$ is indep. of $X(t)$.

Compute posterior probabilities (MMSE estimates):

$\pi_i^+(t) := \mathbb{P}\{X(t) = a_i \mid Z(s), 0 \leq s \leq t\}$, $i = 1, \ldots, m$.

Exact Solution: Wonham Filter (1964-65)

J.SIAM CONTROL
Ser. A, Vol. 2, No. 3
Printed in U.S.A., 1965

SOME APPLICATIONS OF STOCHASTIC DIFFERENTIAL EQUATIONS TO OPTIMAL NONLINEAR FILTERING*

W. M. WONHAM†

Posterior prob. $\boldsymbol{\pi}^+(t) := \{\pi_1^+(t), \ldots, \pi_m^+(t)\}$ solves:

$$\mathrm{d}\boldsymbol{\pi}^+(t) = \boldsymbol{\pi}^+(t)\boldsymbol{Q}\,\mathrm{d}t + \frac{1}{(\sigma_V(t))^2}\boldsymbol{\pi}^+(t)\left(\boldsymbol{H} - \widehat{h}(t)\boldsymbol{I}\right) \times \left(\mathrm{d}Z(t) - \widehat{h}(t)\mathrm{d}t\right)$$

with initial condition $\boldsymbol{\pi}^+(0) = \boldsymbol{\pi}_0$.

$$\boldsymbol{H} := \mathrm{diag}\left(h(a_1), \ldots, h(a_m)\right), \quad \widehat{h}(t) := \sum_{i=1}^{m} h(a_i)\pi_i^+(t).$$

The Present Paper

New variational interpretation of the flow $\boldsymbol{\pi}^{+}(t)$

The Present Paper

New variational interpretation of the flow $\boldsymbol{\pi}^{+}(t)$

Main idea: stochastic flow $\sim$ proximal recursion

Construct gradient descent of a stochastic functional Φ:

$$\boldsymbol{p}_k(\lambda) = \underbrace{\underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf} \frac{1}{2} d^2\left(\boldsymbol{p}, \boldsymbol{p}_{k-1}\right) + \lambda\Phi(\boldsymbol{p})}_{\text{prox}^{d}_{\lambda\Phi}(\boldsymbol{p}_{k-1})}, \boldsymbol{p}_0 \equiv \boldsymbol{\pi}_0, k \in \mathbb{N}$$

λ is the step-size

$d(\cdot,\cdot)$ is a distance functional between prob. vectors

$\Phi(\cdot)$ depends on the generator of the flow $\boldsymbol{\pi}(t)$

Stochastic Flow $\sim$ Proximal Recursion

$$\boldsymbol{p}_k(\lambda) = \underbrace{\underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf} \frac{1}{2} d^2\left(\boldsymbol{p}, \boldsymbol{p}_{k-1}\right) + \lambda\Phi(\boldsymbol{p})}_{\text{prox}^d_{\lambda\Phi}(\boldsymbol{p}_{k-1})}, \boldsymbol{p}_0 \equiv \boldsymbol{\pi}_0, k \in \mathbb{N}$$

Design (d, Φ) such that $\boldsymbol{p}_k(\lambda) \to \boldsymbol{\pi}(t = k\lambda)$ as $\lambda \downarrow 0$ a.s.

This is gradient descent of Φ w.r.t. distance d

Familiar in $\mathbb{R}^n$: Grad Descent $\leftrightsquigarrow$ Prox

$$\mathbf{x}_k = \mathbf{x}_{k-1} - \lambda \nabla \phi(\mathbf{x}_{k-1})$$

$$\Updownarrow$$

$$\mathbf{x}_k = \operatorname{prox}_{\lambda\phi}^{\|\cdot\|_2}(\mathbf{x}_{k-1})$$

$$:= \operatorname*{argmin}_{\mathbf{x}\in\mathbb{R}^n} \left\{ \tfrac{1}{2}\|\mathbf{x} - \mathbf{x}_{k-1}\|_2^2 + \lambda\phi(\mathbf{x}) \right\}$$

Familiar in $\mathbb{R}^n$: Grad Descent $\leftrightsquigarrow$ Prox

$$\mathbf{x}_k = \mathbf{x}_{k-1} - \lambda \nabla \phi(\mathbf{x}_{k-1})$$

$$\Updownarrow$$

$$\mathbf{x}_k = \operatorname{prox}_{\lambda\phi}^{\|\cdot\|_2}(\mathbf{x}_{k-1})$$

$$:= \operatorname*{argmin}_{\mathbf{x}\in\mathbb{R}^n} \left\{\tfrac{1}{2}\|\mathbf{x} - \mathbf{x}_{k-1}\|_2^2 + \lambda\phi(\mathbf{x})\right\}$$

This is nice because

- argmin of $\phi \equiv$ fixed point of prox. operator
- prox. is smooth even when ϕ is not
- reveals metric structure of gradient descent

Back to the Estimation Problem

Idea: posterior flow $\sim$ composition of prox. operators

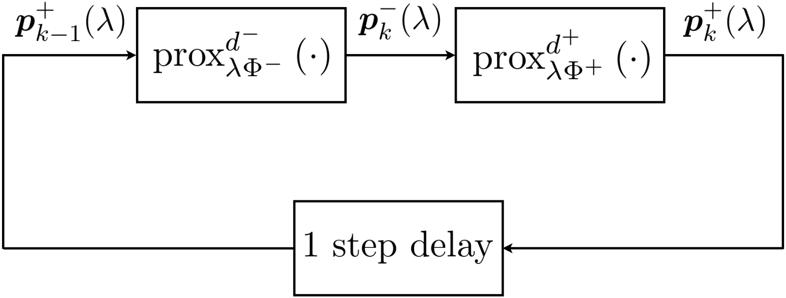

$(\boldsymbol{p}_k^{-}, \boldsymbol{p}_k^{+})$ = (approx. prior, approx. posterior)

solves Wonham SDE

Design $(d^{\pm}, \Phi^{\pm})$ s.t. $\boldsymbol{p}_k^{+}(\lambda) \rightarrow \boldsymbol{\pi}^{+}(t = k\lambda)$ as $\lambda \downarrow 0$ a.s.

Main Results

Proximal recursion for the posterior

Theorem

Let $t_{k-1} := (k-1)\lambda$, $k \in \mathbb{N}$, and $Z_{k-1} := Z(t = t_{k-1})$.

Also, let $Y_{k-1} := (Z_k - Z_{k-1})/\lambda$.

Kullback-Leibler divergence

Then, $\frac{1}{2}(d^+)^2 = D_{\mathrm{KL}}\left(\boldsymbol{p} \parallel \boldsymbol{p}_k^-\right) := \sum_{i=1}^{m} p_i \log\left(\frac{p(i)}{p_k^-(i)}\right)$,

and $\Phi^+(\boldsymbol{p}) = \frac{1}{2(\sigma_V(t_{k-1}))^2}\mathbb{E}_{\boldsymbol{p}}\left[(Y_{k-1} - h)^2\right]$.

Main Results

Proximal recursion for the prior

Theorem

unique stationary prob. $\pi_\infty \in \text{interior}\left(\Delta^{m-1}\right)$

detailed balance: $\pi_\infty(i) Q_{ij} = \pi_\infty(j) Q_{ji}$

Assume $X(t)$ is irreducible and reversible.

Def. inner product $\langle \boldsymbol{p}, \boldsymbol{q} \rangle_{\pi_\infty} := \sum_i \frac{p(i)q(i)}{\pi_\infty(i)}$, $\boldsymbol{p}, \boldsymbol{q} \in \Delta^{m-1}$.

Then, $d^- = \|\boldsymbol{p} - \boldsymbol{p}_{k-1}^+\|_{\pi_\infty}$, and $\Phi^-(\boldsymbol{p}) = -\frac{1}{2}\langle \boldsymbol{p}\boldsymbol{Q}, \boldsymbol{p} \rangle_{\pi_\infty}$.

Other inner products work too: $\langle \boldsymbol{p}, \boldsymbol{q} \rangle_{\pi_\infty} := \sum_i \boldsymbol{p}(i)\boldsymbol{q}(i)\boldsymbol{\pi}_\infty(i)$

If not reversible, then $\boldsymbol{p}_k^-(\lambda) = \boldsymbol{p}_{k-1}^+(\lambda)(\boldsymbol{I} - \lambda \boldsymbol{Q})^{-1} + o(\lambda)$

Quick Recap

$$\begin{aligned} \boldsymbol{p}_k^-(\lambda) &= \operatorname{prox}_{\lambda\Phi^-}^{d^-}\left(\boldsymbol{p}_{k-1}^+\right) && \textbf{[prior update]} \\ &= \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf} \quad \frac{1}{2}\|\boldsymbol{p}-\boldsymbol{p}_{k-1}^+\|_{\boldsymbol{\pi}_\infty}^2 - \frac{\lambda}{2}\langle \boldsymbol{p}\boldsymbol{Q},\boldsymbol{p}\rangle_{\boldsymbol{\pi}_\infty} \end{aligned}$$

$$\begin{aligned} \boldsymbol{p}_k^+(\lambda) &= \operatorname{prox}_{\lambda\Phi^+}^{d^+}\left(\boldsymbol{p}_k^-\right) && \textbf{[posterior update]} \\ &= \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf} \quad D_{\mathrm{KL}}\left(\boldsymbol{p} \parallel \boldsymbol{p}_k^-\right) + \frac{\lambda}{2\left(\sigma_V(t_{k-1})\right)^2}\mathbb{E}_{\boldsymbol{p}}\left[\left(Y_{k-1}-h\right)^2\right] \end{aligned}$$

Numerical Results

Example 1:

$X(t)$ reversible on $\Omega = \{-1, 0, 1\}$, $h(X(t)) = 0.01X(t)$,

$$\text{rate matrix } \boldsymbol{Q} = \begin{bmatrix} -1 & 1/2 & 1/2 \\ 2 & -2 & 0 \\ 3 & 0 & -3 \end{bmatrix}, \sigma_V = 0.01.$$

Example 2:

$X(t)$ non-reversible on $\Omega = \{-1, 0, 1\}$, $h(X(t)) = 0.01X(t)$,

$$\text{rate matrix } \boldsymbol{Q} = \begin{bmatrix} -5 & 3 & 2 \\ 4 & -10 & 6 \\ 3 & 4 & -7 \end{bmatrix}, \sigma_V = 0.01.$$

Numerical Results: Example 1

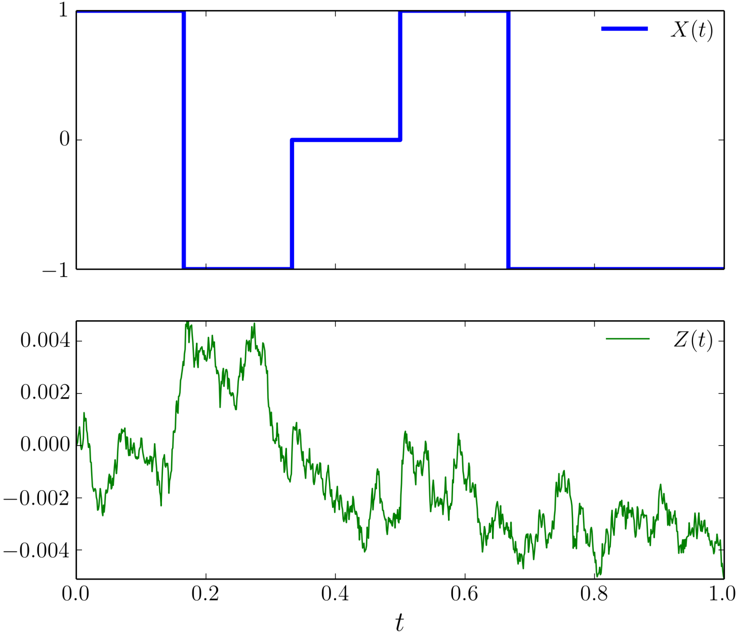

Numerical Results: Example 1 (contd.)

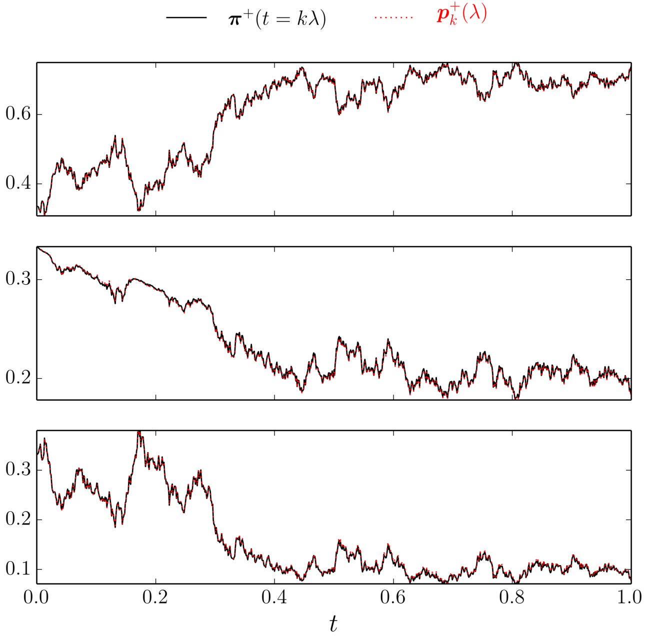

Numerical Results: Example 2

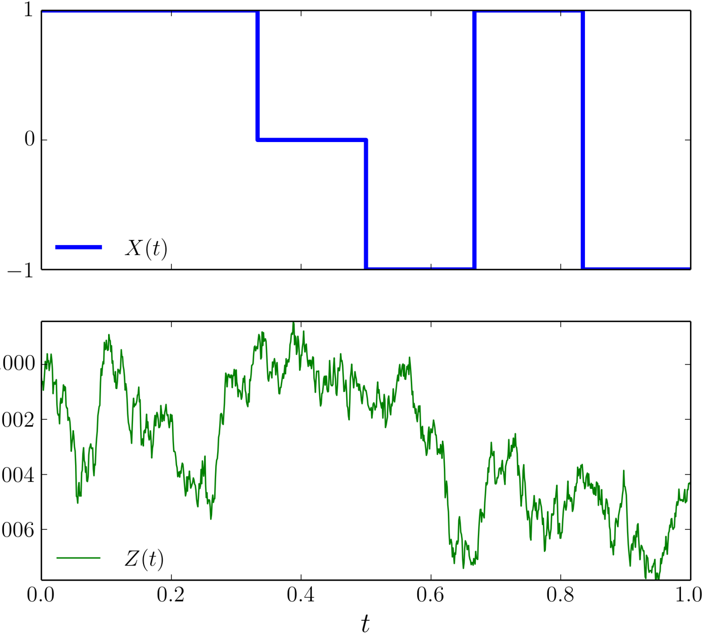

Numerical Results: Example 2 (contd.)

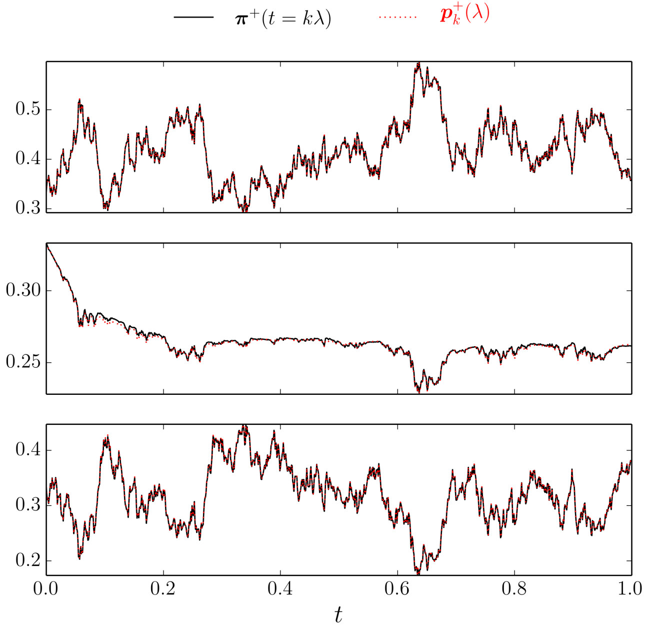

Summary

- General idea: nonlinear filtering as gradient descent
- This work: recovers Wonham filter as composition of prox. operators
- Our prior work: recovered Kalman-Bucy filter (CDC 2017, ACC 2018) as composition of prox. operators
- **Future work:** computation for nonlinear filtering

Thank You