Proximal Recursion for the Wonham Filter

Abhishek Halder[1], and Tryphon T. Georgiou[2]

***Abstract*— This paper contributes to the emerging viewpoint that governing equations for state estimation, conditioned on the history of noisy measurements, can be viewed as gradient flow on the manifold of joint probability density functions with respect to suitable metrics. Herein, we focus on the Wonham filter where the prior dynamics is given by a continuous time Markov chain on a finite state space; the measurement model includes noisy observation of the (possibly nonlinear function of) state. We establish that the posterior flow given by the Wonham filter can be viewed as the small time-step limit of proximal recursions of certain functionals on the probability simplex. The results of this paper extend our earlier work where similar proximal recursions were derived for the Kalman-Bucy filter.**

I. Introduction

We consider the problem of estimating the state of a continuous time Markov chain $X(t)$ on finite state space $\Omega = \{a_1, \ldots, a_m\}$ with $m \times m$ transition[1] rate matrix $\boldsymbol{Q}$, i.e.,

$$Q_{ij} \geq 0, \text{ for } i \neq j, \quad \text{and} \quad Q_{ii} = -\sum_{\substack{j=1\\ j\neq i}}^{m} Q_{ij} < 0.$$

In words, matrix $\boldsymbol{Q}$ has non-negative off-diagonal and negative diagonal elements such that each row sum equals zero. To ease notation, we hereafter write $X(t) \sim \text{Markov}\,(\boldsymbol{Q})$. Suppose that one observes the process $Z(t)$ governed by the Itô stochastic differential equation (SDE)

$$\mathrm{d}Z(t) = h\left(X(t)\right)\,\mathrm{d}t + \sigma_V(t)\,\mathrm{d}V(t), \tag{1}$$

where $h(\cdot)$ is a deterministic injective function of the state, $\sigma_V(t)$ is continuously differentiable and bounded away from zero for all $t \geq 0$, and the standard Wiener process $V(t)$ is independent of the process $X(t)$. One typically refers to $h(\cdot)$ as the sensing or measurement model, and $V(t)$ as the measurement noise. Given the history of noisy observation $\{Z(s), 0 \leq s \leq t\}$, the objective of the estimation problem is to compute the *conditional probability* of the state $X(t)$, i.e., to compute the *posterior probabilities*

$$\pi_i^+(t) := \mathbb{P}\big\{X(t) = a_i \,|\, Z(s), 0 \leq s \leq t\big\},\; i = 1, \ldots, m. \tag{2}$$

Let the initial occupation probability (row) vector be $\boldsymbol{\pi}_0$ satisfying $\boldsymbol{\pi}_0 \geq 0$ elementwise, and $\boldsymbol{\pi}_0 \mathbb{1} = 1$, where $\mathbb{1}$ denotes a column vector of ones. The time evolution of the *prior distribution* $\boldsymbol{\pi}^-(t) := \{\pi_1^-(t), \ldots, \pi_m^-(t)\}$ is governed by the ordinary differential equation (ODE)

$$\dot{\boldsymbol{\pi}}^-(t) = \boldsymbol{\pi}^-(t)\boldsymbol{Q}, \qquad \boldsymbol{\pi}^-(0) = \boldsymbol{\pi}_0. \tag{3}$$

In other words, (3) gives the *unconditional probabilities* of the state $X(t)$, i.e., $\pi_i^-(t) = \mathbb{P}(X(t) = a_i)$, $i = 1, \ldots, m$.

In [1], Wonham showed that for the state-observation model given by

$$X(t) \sim \text{Markov}\,(\boldsymbol{Q})\,, \tag{4a}$$
$$\mathrm{d}Z(t) = h\left(X(t)\right)\,\mathrm{d}t + \sigma_V(t)\,\mathrm{d}V(t), \tag{4b}$$

the posterior probability $\boldsymbol{\pi}^+(t) := \{\pi_1^+(t), \ldots, \pi_m^+(t)\}$ evolves according to the Itô SDE

$$\mathrm{d}\boldsymbol{\pi}^+(t) = \boldsymbol{\pi}^+(t)\boldsymbol{Q}\,\mathrm{d}t + \frac{1}{\left(\sigma_V(t)\right)^2}\boldsymbol{\pi}^+(t)\left(\boldsymbol{H} - \widehat{h}(t)\boldsymbol{I}\right) \times \left(\mathrm{d}Z(t) - \widehat{h}(t)\mathrm{d}t\right), \tag{5}$$

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

$$\boldsymbol{H} := \mathrm{diag}\left(h(a_1), \ldots, h(a_m)\right), \quad \widehat{h}(t) := \sum_{i=1}^{m} h(a_i)\pi_i^+(t). \tag{6}$$

The vector SDE (5) has since been known as the Wonham filtering equation that allows computing the conditional probabilities of the state. Reference [1, eqn. (21)] derived (5) with $h(\cdot)$ as the identity map; the form (5) has appeared in the literature since then – for recent references see e.g., [2, eqn. (2)] and [3, eqn. (5)]. Having obtained $\boldsymbol{\pi}^+$ from (5), assuming that the points in Ω are elements of a linear space, one can compute the optimal (in the minimum mean squared error sense) state estimate given by the conditional expectation

$$\widehat{X}(t) := \mathbb{E}\left[X(t) \mid Z(s), 0 \leq s \leq t\right] = \sum_{i=1}^{m} a_i \pi_i^+(t). \tag{7}$$

The purpose of this paper is to give new variational interpretation of the flow $\boldsymbol{\pi}^+(t)$ governed by (5). Specifically, we seek a gradient flow description for the evolution of the posterior or conditional probability on the standard simplex

$$\Delta^{m-1} := \{\boldsymbol{\pi} \in \mathbb{R}^m_{\geq 0} \mid \boldsymbol{\pi}\mathbb{1} = 1\}. \tag{8}$$

Such interpretations were uncovered in [6] for nonlinear filtering with zero prior dynamics and, more generally, in our recent works [4], [5] for the Kalman-Bucy filter. Results of such flavor are not only fundamental in systems-theoretic context, but may also be transformative in computation since they open up the possibility to solve the filtering equations

[1]Department of Applied Mathematics, University of California, Santa Cruz, CA 95064, USA, ahalder@ucsc.edu

[2]Department of Mechanical and Aerospace Engineering, University of California, Irvine, CA 92617, USA, tryphon@uci.edu

[1]We suppose that the Markov chain is time-homogeneous, i.e., the associated transition probability matrix is $\exp(\boldsymbol{Q}t)$ for all $t \geq 0$.

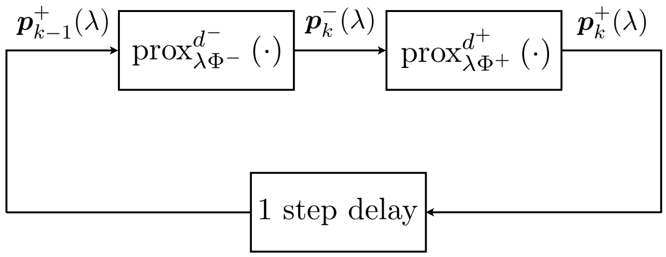

Fig. 1: The block diagram showing the prior and posterior computation as sequential proximal updates given by (10) for $k \in \mathbb{N}$.

via proximal algorithms [9], [10]. This is pursued in [11], [12] for fast computation of the prior joint probability density functions without spatial discretization.

This paper is structured as follows. In Section II, we outline the main ideas for gradient flow formulation via proximal recursion. The recursions for computing the posterior in the Wonham filter are derived in Section III, followed by the same for the computing the prior in Section IV. Numerical examples are given in Section V to illustrate the scope of the proposed framework. Section VI concludes the paper.

Notations: We use $\circ$ to denote function composition, and $\langle\cdot,\cdot\rangle$ to denote the standard Euclidean inner product. The notation $\nabla_{\boldsymbol{x}}$ stands for standard Euclidean gradient w.r.t. vector $\boldsymbol{x}$. Furthermore, $\odot$ denotes elementwise multiplication. The notation $\exp(\cdot)$ with vector argument means elementwise exponential, and the same with matrix argument denotes the exponential matrix.

II. Main Idea

We adopt a metric viewpoint of gradient flow that approximates the flow of probability distribution $\boldsymbol{\pi}(t)$ starting from a given initial condition $\boldsymbol{\pi}_0 := \boldsymbol{\pi}(0)$, as the small time-step limit of a variational recursion in the form

$$\boldsymbol{p}_k(\lambda) = \underset{\boldsymbol{p}}{\arg\inf}\, \frac{1}{2} d^2\left(\boldsymbol{p}, \boldsymbol{p}_{k-1}\right) \;+\; \lambda\Phi(\boldsymbol{p}), \tag{9}$$

where $\boldsymbol{p}_0 \equiv \boldsymbol{\pi}_0$, $k \in \mathbb{N}$, and λ is the step-size. Here, $d(\cdot,\cdot)$ is a distance functional between two probability distributions, and the functional $\Phi(\cdot)$ depends on the generator of the flow $\boldsymbol{\pi}(t)$. In particular, the functionals $d(\cdot,\cdot)$ and $\Phi(\cdot)$ are to be chosen such that $\boldsymbol{p}_k(h) \to \boldsymbol{\pi}(t = k\lambda)$ as $\lambda \downarrow 0$.

The recursion (9) is reminiscent of the Euclidean setting, where the gradient flow for the ODE $\dot{\boldsymbol{x}} = -\nabla_{\boldsymbol{x}}\Phi(\boldsymbol{x})$ can be approximated via a recursion of the form (9) with $d(\cdot,\cdot)$ as the Euclidean distance metric, and $\Phi(\cdot,\cdot)$ as in the argument generating the vector field. In the optimization literature, the operators associated with such recursions are termed as *Moreau-Yosida proximal operators* [7]–[10], denoted as

$$\operatorname{prox}_{\lambda\Phi}^{d},$$

which reads, proximal operator for functional $\lambda\Phi$ w.r.t. d. Likewise, we use the same notation for the right-hand-side of (9) in our more general setting. This allows interpreting the discrete time-stepping as steepest descent of the functional Φ w.r.t. distance d. Proximal operators have also been used in general Hilbert spaces [13], and in the space of probability density functions [4], [5], [11], [12], [14], [15]. The idea of applying proximal recursion in the space of probability measures appeared first in [14]; see also [15].

In the filtering context, we think of the computation of prior followed by that of posterior, as composition of respective proximal operators. Denoting the approximate prior and posterior probability vectors for the proximal recursions associated with Wonham filter as $\boldsymbol{p}_k^-$ and $\boldsymbol{p}_k^+$, respectively, we write

$$\begin{aligned}\boldsymbol{p}_k^-(\lambda) &= \operatorname{prox}_{\lambda\Phi^-}^{d^-}\left(\boldsymbol{p}_{k-1}^+\right)\\ &= \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf}\, \frac{1}{2}\left(d^-\left(\boldsymbol{p},\boldsymbol{p}_{k-1}^+\right)\right)^2 + \lambda\Phi^-(\boldsymbol{p}),\end{aligned} \tag{10a}$$

$$\begin{aligned}\boldsymbol{p}_k^+(\lambda) &= \operatorname{prox}_{\lambda\Phi^+}^{d^+}\left(\boldsymbol{p}_k^-\right)\\ &= \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf}\, \frac{1}{2}\left(d^+\left(\boldsymbol{p},\boldsymbol{p}_k^-\right)\right)^2 + \lambda\Phi^+(\boldsymbol{p}),\end{aligned} \tag{10b}$$

where $k \in \mathbb{N}$, $\lambda > 0$ is the step-size, and $(d^\pm, \Phi^\pm)$ are to be determined functional pairs guaranteeing $\boldsymbol{p}_k^+(\lambda) \to \boldsymbol{\pi}^+(t = k\lambda)$ as $\lambda \downarrow 0$, wherein $\boldsymbol{\pi}^+(t)$ solves (5). In other words, $(d^\pm, \Phi^\pm)$ are to be designed such that the composite map $\operatorname{prox}_{\lambda\Phi^+}^{d^+} \circ \operatorname{prox}_{\lambda\Phi^-}^{d^-}$ approximates the flow of (5) in small time-step limit (see Fig. 1).

Next, we focus on the problem of designing the pair (d^+, Φ^+) in (10b).

III. Proximal Recursion for the Posterior

We first derive a proximal recursion of the form (10b) for the posterior update in the special case $\boldsymbol{Q} \equiv \boldsymbol{0}$ (Section III-A). The proof for the same recovers the explicit stochastic integral formula given in [1, eqn. (5)]. We then show that the same proximal recursion applies for the general $\boldsymbol{Q} \neq \boldsymbol{0}$ case (Section III-B).

A. The Case of Zero Prior Dynamics

As in [1, Section 2], we start with the simple case when the state X, instead of being a Markov chain, is a random variable taking values in $\Omega = \{a_1, \ldots, a_m\}$ with a (known) prior probability distribution $\boldsymbol{p}_0 \in \Delta^{m-1}$ at $t = 0$.

For $k \in \mathbb{N}$, let $t_{k-1} := (k-1)\lambda$ where λ is the step size, and let $\{Z_{k-1}\}_{k\in\mathbb{N}}$ be the sequence of samples of the process $Z(t)$ at $\{t_{k-1}\}_{k\in\mathbb{N}}$. Introducing $Y_{k-1} := (Z_k - Z_{k-1})/\lambda$, we consider the functional

$$\Phi^+(\boldsymbol{p}) := \frac{1}{2\left(\sigma_V(t_{k-1})\right)^2}\mathbb{E}\left[\left(Y_{k-1} - h\right)^2\right], \tag{11}$$

where the expectation operator $\mathbb{E}[\cdot]$ is taken w.r.t. the probability vector $\boldsymbol{p} \in \Delta^{m-1}$. The following result shows that with Φ^+ as in (11), the functional $\frac{1}{2}(d^+)^2$ in (10b) can be taken as the Kullback-Leibler divergence D_{KL}, given by

$$D_{\mathrm{KL}}(\boldsymbol{\alpha} \parallel \boldsymbol{\beta}) := \sum_{i=1}^{m}\alpha_i \log(\alpha_i/\beta_i), \text{ for } \boldsymbol{\alpha}, \boldsymbol{\beta} \in \Delta^{m-1}. \tag{12}$$

In other words, (10b) can be viewed as an entropic proximal mapping [16], [17].

Theorem 1. *Let $\Phi^+(\boldsymbol{p})$ be as in (11), and consider the proximal recursion*

$$\boldsymbol{p}_k^+(\lambda) = \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf}\; D_{\mathrm{KL}}\left(\boldsymbol{p} \parallel \boldsymbol{p}_k^-\right) + \Phi^+(\boldsymbol{p}),\; k\in\mathbb{N}, \quad (13)$$

with initial condition $\boldsymbol{p}_0 \in \Delta^{m-1}$. Let $\boldsymbol{\pi}^+(t)$ be the flow generated by (5) with $\boldsymbol{Q}\equiv\boldsymbol{0}$ and initial condition $\boldsymbol{\pi}^+(0)\equiv\boldsymbol{p}_0$. Then $\boldsymbol{p}_k^+(\lambda)\to\boldsymbol{\pi}^+(t=k\lambda)$ as $\lambda\downarrow 0$.

Proof. See Appendix A. ■

B. The General Case

In the following, we formally state and prove that in the general case $\boldsymbol{Q}\neq\boldsymbol{0}$, the recursion (13) still applies for the posterior computation. Compared to the proof of Theorem 1, the proof now will differ since the map $\boldsymbol{p}_{k-1}^+\mapsto\boldsymbol{p}_k^-$ is no longer identity, and one does not have an analytical solution for the SDE (5) for $\boldsymbol{Q}\neq\boldsymbol{0}$, in general.

Theorem 2. *Let $\Phi^+(\boldsymbol{p})$ be as in (11), and consider the proximal recursion*

$$\boldsymbol{p}_k^+(\lambda) = \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf}\; D_{\mathrm{KL}}\left(\boldsymbol{p} \parallel \boldsymbol{p}_k^-\right) + \Phi^+(\boldsymbol{p}),\; k\in\mathbb{N}, \quad (14)$$

with initial condition $\boldsymbol{p}_0 \in \Delta^{m-1}$. Let $\boldsymbol{\pi}^+(t)$ be the flow generated by (5) with initial condition $\boldsymbol{\pi}^+(0)\equiv\boldsymbol{p}_0$. Then $\boldsymbol{p}_k^+(\lambda)\to\boldsymbol{\pi}^+(t=k\lambda)$ as $\lambda\downarrow 0$.

Proof. See Appendix B. ■

IV. Proximal Recursion for the Prior

We now derive a proximal recursion of the form (10a) for the prior update. We assume that the Markov chain $X(t)$ is irreducible. Then, $X(t)$ has a unique stationary distribution vector $\boldsymbol{\pi}_\infty$ which is the limit $\lim_{t\to\infty}\boldsymbol{\pi}^-(t)$ and is elementwise positive. We further assume that the Markov chain is reversible, i.e., that the detailed balance condition

$$(\boldsymbol{\pi}_\infty)_i\boldsymbol{Q}(i,j) = (\boldsymbol{\pi}_\infty)_j\boldsymbol{Q}(j,i) \quad (15)$$

holds. Here $(\cdot)_i$ denotes the ith entry of.

Starting from a known probability distribution

$$\boldsymbol{\pi}^-(t_{k-1}) = \boldsymbol{p}_{k-1}^+(\lambda),\;\text{ for } t_{k-1}=(k-1)\lambda,$$

the evolution of the prior $\boldsymbol{\pi}^-(t)$ of $X(t)$ is governed by the prior dynamics (3). Thus,

$$\boldsymbol{\pi}^-(\underbrace{t_{k-1}+\lambda}_{t_k=k\lambda}) = \boldsymbol{p}_{k-1}^+(\lambda)\exp\left(\lambda\boldsymbol{Q}\right),$$

equivalently, $\boldsymbol{p}_{k-1}^+(\lambda) = \boldsymbol{\pi}^-(t_k)\exp\left(-\lambda\boldsymbol{Q}\right)$. Hence,

$$\boldsymbol{p}_{k-1}^+(\lambda) = \boldsymbol{\pi}^-(t_k)(\boldsymbol{I}-\lambda\boldsymbol{Q}) + o(\lambda), \quad (16)$$

where $o(\lambda)$ signifies the "order of".

As a consequence of (15), the transition rate matrix $\boldsymbol{Q}$ defines a symmetric operator when considered with respect to the inner product

$$\langle\boldsymbol{p},\boldsymbol{q}\rangle_{\boldsymbol{\pi}_\infty} := \sum_i\frac{(\boldsymbol{p})_i(\boldsymbol{q})_i}{(\boldsymbol{\pi}_\infty)_i}.$$

Indeed, if $\boldsymbol{D}_{\boldsymbol{\pi}_\infty}$ denotes the diagonal matrix formed with the entries of $\boldsymbol{\pi}_\infty$, then in matrix notation, (15) becomes

$$\boldsymbol{D}_{\boldsymbol{\pi}_\infty}\boldsymbol{Q} = \boldsymbol{Q}^T\boldsymbol{D}_{\boldsymbol{\pi}_\infty}, \quad (17)$$

and

$$\begin{aligned}\langle\boldsymbol{pQ},\boldsymbol{q}\rangle_{\boldsymbol{\pi}_\infty} &= \boldsymbol{pQD}_{\boldsymbol{\pi}_\infty}^{-1}\boldsymbol{q}^T\\ &= \boldsymbol{pD}_{\boldsymbol{\pi}_\infty}^{-1}\boldsymbol{Q}^T\boldsymbol{q}^T\\ &= \langle\boldsymbol{p},\boldsymbol{qQ}\rangle_{\boldsymbol{\pi}_\infty},\end{aligned}$$

where, in the above, ^{T} denoting matrix/vector transposition. We can now express $\boldsymbol{p}_k^-(\lambda)$ as a solution to a proximal recursion. Define the quadratic form

$$\Phi^-(\boldsymbol{p}) = -\frac{1}{2}\langle\boldsymbol{pQ},\boldsymbol{p}\rangle_{\boldsymbol{\pi}_\infty}, \quad (18)$$

and denote

$$\|\boldsymbol{p}\|_{\boldsymbol{\pi}_\infty}^2 = \langle\boldsymbol{p},\boldsymbol{p}\rangle_{\boldsymbol{\pi}_\infty}.$$

Theorem 3. *Let $\Phi^-(\boldsymbol{p})$ be as in (18). The λ-approximate prior satisfies the following proximal recursion*

$$\boldsymbol{p}_k^-(\lambda) = \underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf}\;\frac{1}{2}\|\boldsymbol{p}-\boldsymbol{p}_{k-1}^+\|_{\boldsymbol{\pi}_\infty}^2 + \lambda\Phi^-(\boldsymbol{p}). \quad (19)$$

Proof. The stationarity condition for (19) becomes

$$\begin{aligned}\boldsymbol{0} &= \frac{\partial}{\partial\boldsymbol{p}}\left(\frac{1}{2}\|\boldsymbol{p}-\boldsymbol{p}_{k-1}^+\|_{\boldsymbol{\pi}_\infty}^2 + \lambda\Phi^-(\boldsymbol{p})\right)\\ &= (\boldsymbol{p}-\boldsymbol{p}_{k-1}^+)\boldsymbol{D}_{\boldsymbol{\pi}_\infty}^{-1} - \lambda\boldsymbol{p}\frac{1}{2}\left(\boldsymbol{QD}_{\boldsymbol{\pi}_\infty}^{-1}+\boldsymbol{D}_{\boldsymbol{\pi}_\infty}^{-1}\boldsymbol{Q}^T\right)\\ &= (\boldsymbol{p}-\boldsymbol{p}_{k-1}^+)\boldsymbol{D}_{\boldsymbol{\pi}_\infty}^{-1} - \lambda\boldsymbol{pQD}_{\boldsymbol{\pi}_\infty}^{-1},\end{aligned}$$

since $\boldsymbol{D}_{\boldsymbol{\pi}_\infty}^{-1}\boldsymbol{Q}^T = \boldsymbol{QD}_{\boldsymbol{\pi}_\infty}^{-1}$ from (17). Here, by $\boldsymbol{0}$ we denote the zero vector of compatible dimensions. Thus,

$$\boldsymbol{p}_{k-1}^+(\lambda) = \boldsymbol{p}(\boldsymbol{I}-\lambda\boldsymbol{Q}).$$

For sufficiently small λ, the matrix $\boldsymbol{I}-\lambda\boldsymbol{Q}$ is invertible and the unique minimizer $\boldsymbol{p}$ has positive entries. Moreover, since $\boldsymbol{Q}\mathbb{1}=\boldsymbol{0}$, $\boldsymbol{p}\mathbb{1}=1$ and hence $\boldsymbol{p}\in\Delta^{m-1}$, thus, we set

$$\boldsymbol{p}_k^-(\lambda) = \boldsymbol{p}_{k-1}^+(\lambda)(\boldsymbol{I}-\lambda\boldsymbol{Q})^{-1}.$$

Comparing with (16) we see that

$$\boldsymbol{p}_k^-(\lambda) = \boldsymbol{\pi}^-(k\lambda)+o(\lambda),$$

which is our desired result. ■

V. Numerical Examples

To illustrate the proximal recursions proposed in Sections III and IV, we now give two numerical examples. Both of our examples concern estimating the state of a 3-state continuous time Markov chain, the first one being reversible while the second is not.

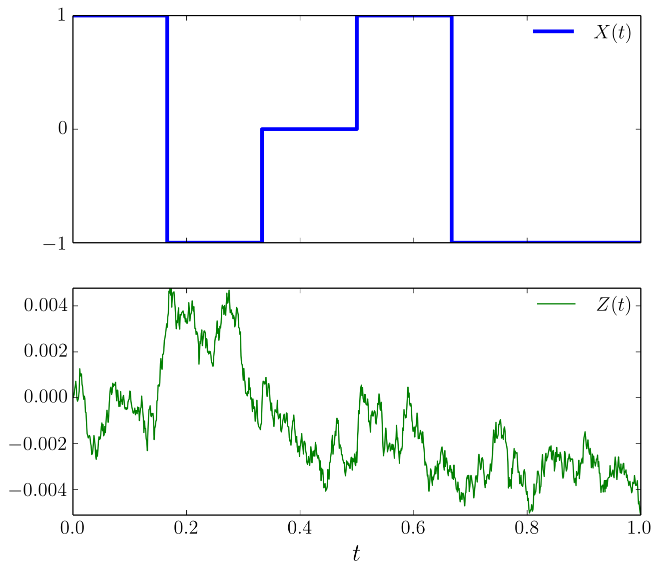

(a) A sample path of the state $X(t)$ (*top*) and of the observation process $Z(t)$ (*bottom*) shown for Example 1 in Section V.

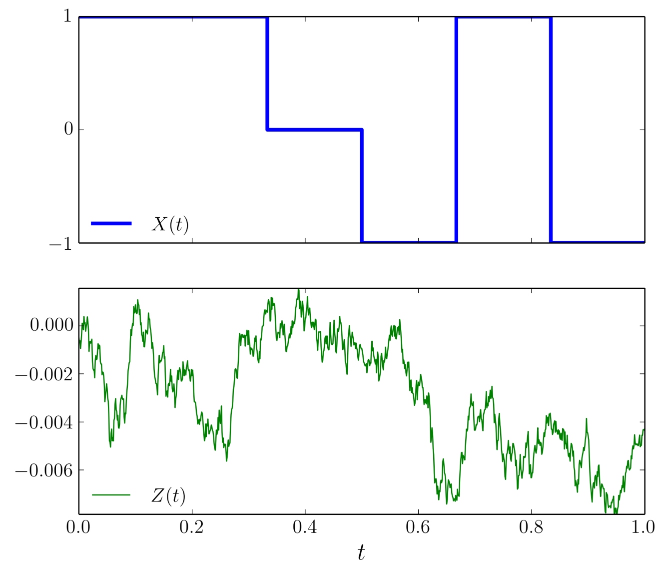

(a) A sample path of the state $X(t)$ (*top*) and of the observation process $Z(t)$ (*bottom*) shown for Example 2 in Section V.

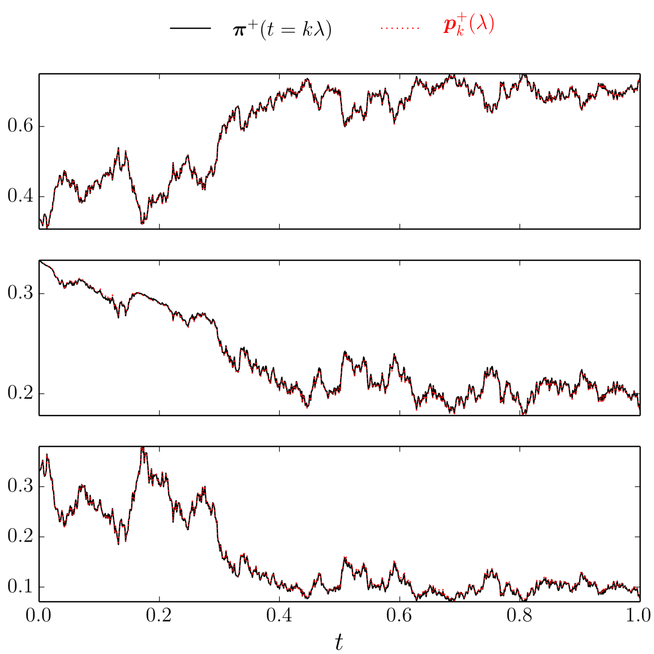

(b) Starting from the initial occupation probability vector $(1/3, 1/3, 1/3)$, shown above are sample paths for the first (*in top*), the second (*in middle*), and the third (*in bottom*) component of the true (*black, solid*) and approximate (*red, dashed*) posterior probability vectors for Example 1 in Section V.

Fig. 2: Simulation results for Example 1 in Section V.

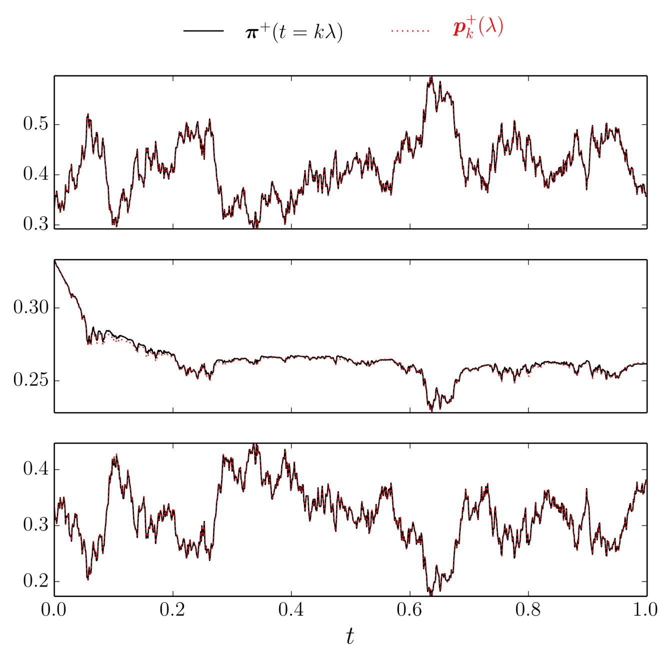

(b) Starting from the initial occupation probability vector $(1/3, 1/3, 1/3)$, shown above are sample paths for the first (*in top*), the second (*in middle*), and the third (*in bottom*) component of the true (*black, solid*) and approximate (*red, dashed*) posterior probability vectors for Example 2 in Section V.

Fig. 3: Simulation results for Example 2 in Section V.

A. Example 1

We consider estimating the state of an irreducible Markov chain $X(t)$ taking values in $\{-1, 0, 1\}$ with rate matrix

$$\boldsymbol{Q} = \begin{bmatrix} -1 & 1/2 & 1/2 \\ 2 & -2 & 0 \\ 3 & 0 & -3 \end{bmatrix}. \tag{20}$$

It is easy to verify that the stationary distribution $\boldsymbol{\pi}_\infty = (12/17, 3/17, 2/17)$, and that the reversibility condition (15) holds. In (4b), we set $h(X(t)) = 0.01X(t)$, and $\sigma_V = 0.01$.

Fig. 2a shows a sample path for $X(t)$, and the same for $Z(t)$. In Fig. 2b, we compare the time evolution of the components of the associated posterior probability vector $\boldsymbol{\pi}^+(t)$ from the Wonham filter (*in black, solid*), with the same of the approximator $\boldsymbol{p}_k^+(\lambda)$ (*in red, dashed*) computed via the proposed proximal recursion framework (Fig. 1) with step-size $\lambda = 10^{-3}$ and $t \in [0, 1]$. We only show here the result for a fixed initial condition $\boldsymbol{\pi}_0 = (1/3, 1/3, 1/3)$; the trends are similar for other initial conditions. Computing $\boldsymbol{\pi}^+(t)$ in the Wonham filter entails numerically solving the system of coupled nonlinear SDEs (5), done here via the Euler-Maruyama method. In contrast, computing $\boldsymbol{p}_k^+(\lambda)$ entails recursive evaluation of (33) and (24). In our numerical experiments, the latter was observed to enjoy about an order of magnitude computational speed up. Fig. 2b shows that the respective posterior sample paths match, as predicted.

B. Example 2

Next we consider a Markov chain $X(t)$ taking values in $\{-1, 0, 1\}$ with rate matrix

$$\boldsymbol{Q} = \begin{bmatrix} -5 & 3 & 2 \\ 4 & -10 & 6 \\ 3 & 4 & -7 \end{bmatrix}, \tag{21}$$

and $h(\cdot)$, σ_V, λ as before. Here again, for $t \in [0, 1]$, we compare the posterior sample paths computed from the Wonham filter (5) with its approximator computed via the proposed proximal recursion approach. Notice that for prior computation, although one does not have a metric version of the variational formula (10a), the approximation (33) still applies for small λ. Thus, $\boldsymbol{p}_k^+(\lambda)$ can still be computed by recursive evaluation of (33) and (24).

Fig. 3a shows a sample path for $X(t)$, and the same for $Z(t)$. In Fig. 3b, we show the corresponding $\boldsymbol{\pi}^+(t)$ from the Wonham filter (*in black, solid*), and $\boldsymbol{p}_k^+(\lambda)$ (*in red, dashed*) from the proximal recursion for this case, starting from the initial condition $\boldsymbol{\pi}_0 = (1/3, 1/3, 1/3)$. The respective sample paths are in agreement, as expected.

VI. Conclusions

This purpose of this paper is to expand on the list of examples where the governing equations for state estimation, conditioned on the history of noisy measurements, can be expressed as gradient flow on a space of probability density functions with respect to a suitable metric. This viewpoint promises a new class of estimation algorithms, taking advantage of implementation of flows via recursive application of proximal projections. Prior work elucidated the case of the Kalman-Bucy filter, and therefore, in the present work we sought to extend the paradigm to case of the Wonham filter. The latter estimates the state of a continuous time Markov chain on a finite state space based on noisy observations. In this paper we have established that the posterior flow that is provided by the Wonham filter can be expressed as the small time-step limit of proximal recursions of certain functionals on the probability simplex. Our preliminary numerical experiments reported here hint at possible computational advantages of the proposed approach, especially for large Markov chains. This will be systematically investigated in our future work.

Acknowledgement

Partial support was provided by NSF under grants 1665031, 1807664 and 1839441, and by AFOSR under grant FA9550-17-1-0435.

Appendix

A. Proof of Theorem 1

With the stated choices for the pair (d^+, Φ^+), we are led to the proximal map of the form

$$\underset{\boldsymbol{p}\in\Delta^{m-1}}{\arg\inf}\; D_{\mathrm{KL}}\left(\boldsymbol{p} \parallel \boldsymbol{p}_k^-\right) + \langle \boldsymbol{c}_{k-1}, \boldsymbol{p}\rangle, \quad k \in \mathbb{N}, \tag{22}$$

where the i-th component of the row vector $\boldsymbol{c}_{k-1}$ is

$$\boldsymbol{c}_{k-1}(i) := \frac{\lambda}{2\left(\sigma_V(t_{k-1})\right)^2}\left(Y_{k-1} - h(a_i)\right)^2,\; i = 1, \ldots, m. \tag{23}$$

The objective in (22) is strictly convex since D_{KL} is strictly convex in $\boldsymbol{p}$, and the other summand is linear in $\boldsymbol{p}$. Therefore, the $\arg\inf$ in (22), which we denote by $\boldsymbol{p}_k^+$, is unique. By direct calculation (setting the gradient of Lagrangian w.r.t. $\boldsymbol{p}$ to zero, and enforcing the constraint $\boldsymbol{p}\mathbb{1} = 1$), we get

$$\boldsymbol{p}_k^+ = \boldsymbol{p}_k^- \odot \exp\left(-\boldsymbol{c}_{k-1}\right) / \left(\left(\boldsymbol{p}_k^- \odot \exp(-\boldsymbol{c}_{k-1})\right)\mathbb{1}\right). \tag{24}$$

Since in this case, we have no prior dynamics, therefore $\boldsymbol{p}_k^- \equiv \boldsymbol{p}_{k-1}^+$ for all $k \in \mathbb{N}$. Hence, we can rewrite (24) in terms of the prior probability distribution $\boldsymbol{p}_0$ as

$$\boldsymbol{p}_k^+ = \boldsymbol{p}_0 \odot \exp(-\boldsymbol{\gamma}_{k-1}) / \left((\boldsymbol{p}_0 \odot \exp(-\boldsymbol{\gamma}_{k-1}))\mathbb{1}\right), \tag{25}$$

where the $1 \times m$ vector

$$\boldsymbol{\gamma}_{k-1} := \sum_{r=1}^{k} \boldsymbol{c}_{r-1}. \tag{26}$$

Noting that (25) is simply normalization (i.e., Kullback-Leibler projection onto probability simplex) of the vector $\boldsymbol{p}_0 \odot \exp(-\boldsymbol{\gamma}_{k-1})$, we now unpack $\boldsymbol{\gamma}_{k-1}$ as function of λ and the sampled process $\{Z_{k-1}\}_{k\in\mathbb{N}}$.

Because $h(\cdot)$ is injective, the random variable $\widetilde{X} := h(X)$ takes values in $\{\widetilde{a}_1, \ldots, \widetilde{a}_m\}$, and $\widetilde{X} = \widetilde{a}_i := h(a_i)$ with probability $\boldsymbol{p}_0(i)$. Combining (23) and (26), we then have

$$\boldsymbol{\gamma}_{k-1}(i) = \frac{\lambda}{2}\sum_{r=1}^{k}\frac{\left(Y_{r-1} - \widetilde{a}_i\right)^2}{\left(\sigma_V(t_{r-1})\right)^2}, \quad i = 1, \ldots, m. \tag{27}$$

Expanding the square in the numerator of each summand in (27), substituting $Y_{r-1} = (Z_r - Z_{r-1})/\lambda$ for $r = 1, \dots, k$, and rearranging yields

$$\boldsymbol{\gamma}_{k-1}(i) = \underbrace{\frac{1}{2}\sum_{r=1}^{k}\frac{(Z_r - Z_{r-1})^2}{\lambda\,(\sigma_V(t_{r-1}))^2}}_{\text{term 1}} - \underbrace{\widetilde{a}_i\sum_{r=1}^{k}\frac{Z_r - Z_{r-1}}{(\sigma_V(t_{r-1}))^2}}_{\text{term 2}} + \underbrace{\frac{1}{2}\widetilde{a}_i^2\sum_{r=1}^{k}\frac{\lambda}{(\sigma_V(t_{r-1}))^2}}_{\text{term 3}}. \tag{28}$$

To simplify the term 1 indicated in (28), we use the Euler-Maruyama update for (1) given by

$$Z_r = Z_{r-1} + h(a_i)\lambda + \sigma_V(t_{r-1})\,(V_r - V_{r-1}) + O(\lambda^2), \tag{29}$$

where $r = 1, \dots, k$, and the increments $(V_r - V_{r-1})$ are i.i.d. zero mean normal random variables with variance λ. From (29), we get

$$(Z_r - Z_{r-1})^2 = (\sigma_V(t_{r-1}))^2\,\lambda, \tag{30}$$

where we used[2] $\lambda^2 = 0$, $(V_r - V_{r-1})^2 = \lambda$, and $(V_r - V_{r-1})\lambda = 0$. Consequently, term 1 in (28) equals $k/2$.

Combining (25), (28), and (30), we thus obtain

$$\boldsymbol{p}_k^+(i) = \frac{\boldsymbol{p}_0(i)\exp(-k/2 + \text{term 2} - \text{term 3})}{\sum_{i=1}^{m}\boldsymbol{p}_0(i)\exp(-k/2 + \text{term 2} - \text{term 3})}. \tag{31}$$

Passing (31) to the limit $\lambda \downarrow 0$, recalling that $\widetilde{a}_i = h(a_i)$, and that the time interval $[0, t]$ was divided into sub-intervals with breakpoints $t_0 \equiv 0, t_1, \dots, t_k \equiv t$, we arrive at

$$\lim_{\lambda\downarrow 0}\boldsymbol{p}_k^+(i) = \frac{\boldsymbol{p}_0(i)\exp\Big(h(a_i)\int_0^t\frac{\mathrm{d}Z(s)}{(\sigma_V(s))^2} - \frac{1}{2}(h(a_i))^2\int_0^t\frac{\mathrm{d}s}{(\sigma_V(s))^2}\Big)}{\sum_{i=1}^{m}\boldsymbol{p}_0(i)\exp\Big(h(a_i)\int_0^t\frac{\mathrm{d}Z(s)}{(\sigma_V(s))^2} - \frac{1}{2}(h(a_i))^2\int_0^t\frac{\mathrm{d}s}{(\sigma_V(s))^2}\Big)}. \tag{32}$$

The right-hand-side of (32) is exactly the solution of the SDE (5) for $\boldsymbol{Q} \equiv \boldsymbol{0}$ with initial condition $\boldsymbol{\pi}^+(0) = \boldsymbol{p}_0$; see [1, Appendix 2] for a proof. Therefore, we conclude that $\lim_{\lambda\downarrow 0}\boldsymbol{p}_k^+ = \boldsymbol{\pi}^+(t = k\lambda)$, as desired. ∎

B. Proof of Theorem 2

We start by noting that the development in Appendix A up to expression (24) still applies. Also, since $h(\cdot)$ is injective, the process $h(X(t)) \sim \text{Markov}(\boldsymbol{Q})$ takes values in $\{h(a_1), \dots, h(a_m)\}$.

For $X(t) \sim \text{Markov}(\boldsymbol{Q})$, the map $\boldsymbol{p}_{k-1}^+ \mapsto \boldsymbol{p}_k^-$, $k \in \mathbb{N}$, corresponding to the Euler discretization of (3) is

$$\boldsymbol{p}_k^- = \boldsymbol{p}_{k-1}^+(\boldsymbol{I} + \lambda\boldsymbol{Q}) + O(\lambda^2). \tag{33}$$

[2]These are equivalent to the well-known relations in stochastic calculus: $(\mathrm{d}t)^2 = 0$, $(\mathrm{d}V)^2 = \mathrm{d}t$, and $\mathrm{d}V\mathrm{d}t = 0$.

Let $\Delta Z_{k-1} := Z_k - Z_{k-1}$. Substituting $Y_{k-1} = \Delta Z_{k-1}/\lambda$ in (23), expanding the square, and using (30), we have

$$\exp(-\boldsymbol{c}_{k-1}(i)) = \exp(-1/2)\times\exp\left(\frac{h(a_i)\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}\right)\times \exp\left(-\frac{\lambda(h(a_i))^2}{2\,(\sigma_V(t_{k-1}))^2}\right), \tag{34}$$

for $i = 1, \dots, m$.

Up to first order, the second exponential factor in (34) approximates as $1 + h(a_i)\Delta Z_{k-1}/(\sigma_V(t_{k-1}))^2$. For the third exponential factor in (34), notice that since $\lambda \equiv \mathrm{d}t$ for $\lambda \downarrow 0$, therefore from (4b), we have $h^2(a_i) = (\mathrm{d}Z - \sigma_V\mathrm{d}V)^2/\lambda^2 = 0$, wherein we used $(\mathrm{d}Z)^2 = \sigma_V^2\lambda$, $(\mathrm{d}V)^2 = \lambda$, and $\lambda\mathrm{d}V = 0$. Putting these together, (34) yields

$$\exp(-\boldsymbol{c}_{k-1}(i)) \approx \exp(-1/2)\left(1 + \frac{h(a_i)\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}\right). \tag{35}$$

Substituting for $\boldsymbol{p}_k^-$ and $\exp(-\boldsymbol{c}_{k-1})$ in (24) from (33) and (35), respectively, we get

$$\boldsymbol{p}_k^+(i) = \nu(i)/\delta, \quad i = 1, \dots, m, \tag{36}$$

where

$$\nu(i) := \left(\boldsymbol{p}_{k-1}^+(i) + \lambda\sum_{j=1}^{m}\boldsymbol{p}_{k-1}^+(j)\boldsymbol{Q}(i,j)\right)\left(1 + \frac{h(a_i)\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}\right), \tag{37a}$$

$$\delta := \sum_{i=1}^{m}\nu(i). \tag{37b}$$

From (29), we observe that $\lambda\Delta Z_{k-1} = h(a_i)\lambda^2 + \sigma_V(t_{k-1})\lambda(V_k - V_{k-1}) = 0$, as both λ^2 and $\lambda(V_k - V_{k-1})$ are zero. This allows us to simplify (37a) as

$$\nu(i) = \boldsymbol{p}_{k-1}^+(i) + \frac{\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}h(a_i)\boldsymbol{p}_{k-1}^+(i) + \lambda\sum_{j=1}^{m}\boldsymbol{p}_{k-1}^+(j)\boldsymbol{Q}(i,j) + O(\lambda^2), \tag{38}$$

and consequently, (37b) reduces[3] to

$$\delta = \sum_{i=1}^{m}\nu(i) = 1 + \frac{\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}\widehat{h}(t_{k-1}). \tag{39}$$

Combining (36), (38) and (39), we obtain

$$\boldsymbol{p}_k^+(i) - \boldsymbol{p}_{k-1}^+(i) = \left(\frac{\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}\boldsymbol{p}_{k-1}^+\left(h(a_i) - \widehat{h}(t_{k-1})\right) + \lambda\sum_{j=1}^{m}\boldsymbol{p}_{k-1}^+(j)\boldsymbol{Q}(i,j)\right)\times\left(1 + \frac{\Delta Z_{k-1}}{(\sigma_V(t_{k-1}))^2}\widehat{h}(t_{k-1})\right)^{-1}. \tag{40}$$

Up to first order, the second factor in (40) approximates as $1 - \Delta Z_{k-1}\widehat{h}(t_{k-1})/(\sigma_V(t_{k-1}))^2$. Using this approximation

[3]We use here that $\sum_{i=1}^{m}\boldsymbol{Q}(i,j) = 0$ for all $j = 1, \dots, m$.

together with $(\Delta Z_{k-1})^2 = \lambda(\sigma_V(t_{k-1}))^2$ (from (30)), and that $\lambda\Delta Z_{k-1} = 0$ (as before), (40) simplifies as

$$\boldsymbol{p}_k^+(i) - \boldsymbol{p}_{k-1}^+(i) = \lambda \sum_{j=1}^{m} \boldsymbol{p}_{k-1}^+(j)\boldsymbol{Q}(i,j) + \frac{\boldsymbol{p}_{k-1}^+(i)}{(\sigma_V(t_{k-1}))^2} \times \left(h(a_i) - \widehat{h}(t_{k-1})\right)\left(\Delta Z_{k-1} - \widehat{h}(t_{k-1})\lambda\right), \quad (41)$$

which is exactly the first order (Euler-Maruyama) discretization of the SDE (5). Specifically, in the limit $\lambda \downarrow 0$, (41) reduces to (5). Hence the statement. ■

REFERENCES

[1] W.M. Wonham, "Some applications of stochastic differential equations to optimal nonlinear filtering". *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, Vol. 2, No. 3, pp. 347–369, 1964.

[2] Q. Zhang, Y.G. George, and J.B. Moore, "Two-time-scale approximation for Wonham filters". *IEEE Transactions on Information Theory*, Vol. 53, No. 5, pp. 1706–1715, 2007.

[3] T. Yang, P.G. Mehta, and S.P. Meyn, "Feedback particle filter for a continuous-time Markov chain". *IEEE Transactions on Automatic Control*, Vol. 61, No. 2, pp. 556–561, 2016.

[4] A. Halder, and T.T. Georgiou, "Gradient flows in uncertainty propagation and filtering of linear Gaussian systems". *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 3081–3088, 2017.

[5] ——, "Gradient flows in filtering and Fisher-Rao geometry". *2018 Annual American Control Conference (ACC)*, pp. 4281–4286, 2018.

[6] R.S. Laugesen, P.G. Mehta, S.P. Meyn, and M. Raginsky, "Poisson's equation in nonlinear filtering". *SIAM Journal on Control and Optimization*, Vol. 53, No. 1, pp. 501–525, 2015.

[7] J-J. Moreau, "Proximité et dualité dans un espace hilbertien". *Bulletin de la Société mathématique de France*, Vol. 93, pp. 273–299, 1965.

[8] K. Yosida, *Functional Analysis*, Spinger-Verlag, 1964.

[9] R.T. Rockafeller, "Monotone operators and the proximal point algorithm". *SIAM Journal on Control and Optimization*, Vol. 14, No. 5, pp. 877–898, 1976.

[10] N. Parikh, and S. Boyd, "Proximal algorithms". *Foundations and Trends® in Optimization*, Vol. 1, No. 3, pp. 127–239, 2014.

[11] K.F. Caluya, and A. Halder, "Proximal recursion for solving the Fokker-Planck equation". *2019 Annual American Control Conference (ACC)*, accepted, preprint arXiv:1809.10844, 2019.
web: `https://arxiv.org/pdf/1809.10844.pdf`

[12] ——, "Gradient flow algorithms for density propagation in stochastic systems". preprint, 2019.
web: `https://www.abhishekhalder.org/proxFPK.pdf`

[13] H.H. Bauschke, and P.L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics, Springer; 2011.

[14] R. Jordan, D. Kinderlehrer, and F. Otto, "The variational formulation of the Fokker–Planck equation". *SIAM Journal on Mathematical Analysis*, Vol. 29, No. 1, pp. 1–17, 1998.

[15] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.

[16] M. Teboulle, "Entropic proximal mappings with applications to nonlinear programming". *Mathematics of Operations Research*, Vol. 17, No. 3, pp. 670–690, 1992.

[17] Y. Censor, and S.A. Zenios, "Proximal minimization algorithm with D-functions". *Journal of Optimization Theory and Applications*, Vol. 73, No. 3, pp. 451–464, 1992.