

An optimal control approach to particle filtering on Lie groups

Bo Yuan, Qinsheng Zhang, and Yongxin Chen

Abstract— We study the filtering problem over a Lie group that plays an important role in robotics and aerospace applications. We present a new particle filtering algorithm based on stochastic control. In particular, our algorithm is based on a duality between smoothing and optimal control. Leveraging this duality, we reformulate the smoothing problem into an optimal control problem, and by approximately solving it (using, e.g., iLQR) we establish a superior proposal for particle smoothing. Combining it with a suitable designed sliding window mechanism, we obtain a particle filtering algorithm that suffers less from sample degeneracy compared with existing methods. The efficacy of our algorithm is illustrated by a filtering problem over $\text{SO}(3)$ for satellite attitude estimation.

I. INTRODUCTION

The filtering problem to estimate the posterior distribution of the state of a dynamic system from noisy observations plays an essential role in applications in control engineering [1], robotics [2], autonomy [3] etc. For dynamics that has a Euclidean state space, many algorithms such as the extended Kalman filter (EKF), the unscented Kalman filter (UKF), particle filtering, have been proposed for the filtering task. In many applications in robotics and aerospace, the state space of the dynamical system is a manifold instead. The above filtering algorithms cannot directly be extended to the manifold setting due to the differences between the Euclidean spaces and general manifolds.

Designing filters directly evolving on manifolds gives better estimation of posteriors in most cases. For instance, for the agent's pose evolving in the three-dimensional space, the state variable belongs to the special orthogonal group, a Lie group $\text{SO}(3)$. On this manifold, better convergence results from the Invariant Extended Kalman Filters have been shown in [4], [5], [6], [7]. However, these filters are sensitive to the initial state estimation and only locally stable. Alternatively, Lie Group Variational Filters [8], [9] derived from the d'Alembert's Principle are shown to be almost globally asymptotically stable.

We present a particle filtering for dynamical systems on Lie groups [10]. Our algorithm relies on the duality of filtering and stochastic optimal control, which reformulates the smoothing problem into an optimal control problem [11], [12]. If one can find the exact optimal solution for it, then the resulting filtering algorithm can achieve its theoretical maximum effectiveness. In practice, the computation of optimal control problems is demanding, thus we utilize a suboptimal solution instead [13]. The particles in the

This work was supported by NSF under grant 1942523 and 2008513.

Bo Yuan, Q. Zhang and Y. Chen are with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA; {byuan48, qzhang419, yongchen}@gatech.edu

resulting filtering/smoothing algorithm is generated by simulating the dynamics under suboptimal control. Path integral based on the Girsanov theorem [14] and the Kallianpur-Striebel formula [15] is used to compute the weights of these particles. This work extends the Path Integral Particle Filter [16] to the Lie group setting. We show with numerical experiments that our control-based filtering algorithm with the suboptimal control signal is more effective than that with a zero control signal. The latter is effectively the standard Sequential Importance Resampling.

II. NOTATION AND PRELIMINARIES

In this section, we present the notation and background knowledge on Lie groups [17] as well as the duality between smoothing and optimal control.

A. Lie Groups and Lie Algebras

We denote a Lie group by G , a set that is both a group and a d -dimensional differentiable manifold. Let $T_g G$ be the tangent space of G at the element g . The associated Lie algebra \mathfrak{g} is defined as $T_e G$ with e being the identity element.

A vector field V is a smooth mapping by which each element $g \in G$ corresponds to a vector $v \in T_g G$. Let L_g denote the left multiplication by g , i.e., for all $h \in G$, $L_g(h) = gh$. It can be shown that L_g is a diffeomorphism between G and itself. Due to this diffeomorphism, for each vector $a \in \mathfrak{g}$, we can define a vector field V_a that assigns each $g \in G$ to $dL_g(a)$, where $dL_g(a)$ is the shorthand notation of $(g\rho)'(0)$ with $\rho(t)$, a curve on G , satisfying $\rho(0) = e$ and $\rho'(0) = a$. To proceed, we denote the exponential mapping as $\exp: \mathfrak{g} \rightarrow G$. For every $a \in \mathfrak{g}$, $\exp(a) = \gamma(1)$, where $\gamma(t) \in G$ satisfies $\gamma'(t) = dL_{\gamma(t)}(a)$ and $\gamma(0) = e$. As a local diffeomorphism between G and \mathfrak{g} , the exponential mapping plays a crucial role in connecting them. In addition, we point out a Lie algebra \mathfrak{g} , as a vector space, is homeomorphic to \mathbb{R}^d . We denote the corresponding coordinate of $a \in \mathfrak{g}$ by a^\vee .

In many practical applications, Lie groups are assumed to be matrix Lie groups. A matrix Lie group is a closed subgroup of $\text{GL}(n)$, the group consists of all $n \times n$ invertible matrices. In the case of matrix Lie groups, group multiplication corresponds to matrix multiplication, $dL_g(a)$ is the standard matrix multiplication for g and a , i.e., $dL_g(a) = ga$, and the exponential mapping is matrix exponential. In what follows, we adopt \exp to represent both the exponential mapping on Lie groups and matrix exponential, as the exact meaning will be clear from context. Among these matrix Lie groups, the special orthogonal group $\text{SO}(3)$, also named as 3D rotation group, is what we particularly care about. Formally, $\text{SO}(3)$ is the group of all 3×3 real orthogonal matrices with

determinant 1. The Lie algebra of $\text{SO}(3)$, denoted by $\mathfrak{so}(3)$, is all 3×3 skew-symmetric matrices.

B. Smoothing and Stochastic Optimal Control on \mathbb{R}^d

Recently, the duality between smoothing and optimal control inspires a class of filtering algorithms on \mathbb{R}^d [16], [11], [12], [18]. The connection between the two areas is established as follows. Consider a continuous dynamic system evolving on \mathbb{R}^d in which the dynamics are represented by stochastic differential equations

$$\begin{aligned} dX_t &= b(t, X_t)dt + \sigma(t, X_t)dW_t, \quad X_0 \sim \nu_0, \quad (1a) \\ dY_t &= h(t, X_t)dt + \sigma_B dB_t, \quad Y_0 = 0. \quad (1b) \end{aligned}$$

Here, we denote the state variables by $X_t \in \mathbb{R}^d$ and the raw measurements by $Y_t \in \mathbb{R}^m$. The initial distribution of X_t is ν_0 . Moreover, W_t and B_t are two independent multidimensional Wiener processes. The task of the smoothing problem over $[0, T]$ is to estimate the posterior distribution of every X_τ given all the observations $\{Y_t, 0 \leq t \leq T\}$, i.e., $P(X_\tau | Y_t, 0 \leq t \leq T)$ for every $0 \leq \tau \leq T$. Consider another dynamics system with control signal u_t ,

$$d\tilde{X}_t = b(t, \tilde{X}_t)dt + \sigma(t, \tilde{X}_t)(u_t dt + dW_t), \quad \tilde{X}_0 \sim \pi_0, \quad (2)$$

where \tilde{X}_t stands for the state variable under control signal u_t , and \tilde{X}_0 follows π_0 . Note that the control signal relies on both t and \tilde{X}_t . We denote it by u_t for brevity.

Denote the probability measure induced by dynamic (1a) and dynamic (2) by \mathcal{P} and $\tilde{\mathcal{P}}$, respectively. We also let \mathcal{Q}^Y be the posterior measure induced by (1). It has been shown that this smoothing problem can be reformulated to a stochastic optimal control problem by minimizing the KL divergence between $\tilde{\mathcal{P}}$ and \mathcal{Q}^Y [16], [11], [12], [18], i.e.,

$$\min_{\tilde{\mathcal{P}}} \text{KL}(\tilde{\mathcal{P}} \| \mathcal{Q}^Y) := \mathbb{E}_{\tilde{\mathcal{P}}}(\log \frac{d\tilde{\mathcal{P}}}{d\mathcal{Q}^Y}). \quad (3)$$

It is worthwhile to summarize the derivation here, as our central proof is a generalization of it. The prementioned objective functions could be decomposed into two terms,

$$\mathbb{E}_{\tilde{\mathcal{P}}}(\log \frac{d\tilde{\mathcal{P}}}{d\mathcal{Q}^Y}) = \mathbb{E}_{\tilde{\mathcal{P}}}(\log \frac{d\tilde{\mathcal{P}}}{d\mathcal{P}}) - \mathbb{E}_{\tilde{\mathcal{P}}}(\log \frac{d\mathcal{Q}^Y}{d\mathcal{P}}). \quad (4)$$

Under certain regularity conditions, one can adopt Pathwise Kallianpur-Striebel formula (See Proposition 1.4.2. in [15]) which gives

$$\begin{aligned} \frac{d\mathcal{Q}^Y}{d\mathcal{P}} &\propto \exp \left\{ - \int_0^T \frac{1}{2\sigma_B^2} \|h(t, \tilde{X}_t)\|^2 dt \right. \\ &\quad \left. - \int_0^T \frac{1}{\sigma_B^2} Y'_t dh(t, \tilde{X}_t) + \frac{1}{\sigma_B^2} Y'_T h(T, \tilde{X}_T) \right\}. \end{aligned} \quad (5)$$

The other term can be dealt with by Girsanov theorem (See Theorem 8.6.5. in [14]), and one has

$$\log \frac{d\tilde{\mathcal{P}}}{d\mathcal{P}} = \int_0^T \frac{1}{2} \|u_t\|^2 dt + u'_t dW_t + \log \frac{d\nu_0}{d\nu_0}, \quad (6)$$

which infers

$$\text{KL}(\tilde{\mathcal{P}} \| \mathcal{P}) = \mathbb{E} \left\{ \int_0^T \frac{1}{2} \|u_t\|^2 dt \right\} + \text{KL}(\nu_0 \| \nu_0). \quad (7)$$

It follows the objective function of the derived optimal control problem is as follows

$$\begin{aligned} \min_{u, \pi_0} \text{KL}(\nu_0 \| \nu_0) + &\mathbb{E}_{\tilde{\mathcal{P}}} \left\{ \int_0^T \left[\frac{1}{2} \|u_t\|^2 + \frac{1}{2\sigma_B^2} \|h(t, \tilde{X}_t)\|^2 \right] dt \right. \\ &\left. + \int_0^T \frac{1}{\sigma_B^2} Y'_t dh(t, \tilde{X}_t) - \frac{1}{\sigma_B^2} Y'_T h(T, \tilde{X}_T) \right\} \end{aligned} \quad (8)$$

In principle, once the optimal control strategy u_t^* and the optimal initial distribution π_0^* of (8) are found, one can sample trajectories from dynamics (2), and the empirical distribution of these particles forms an approximation of the solution to smoothing problems.

III. PATH INTEGRAL PARTICLE FILTERING ON LIE GROUPS

In this section, we present our main results on an optimal control inspired approach to particle filtering/smoothing over Lie groups.

A. Smoothing as Control Problems for Lie Groups

Consider the following dynamics for filtering on Lie groups

$$dg_t = dL_{g_t}(\xi_t)dt, \quad (9a)$$

$$d\xi_t^\vee = b(t, g_t, \xi_t)dt + \sigma_W dW_t, \quad (9b)$$

$$dY_t = h(t, g_t, \xi_t)dt + \sigma_B dB_t, \quad (9c)$$

$$(g_0, \xi_0^\vee) \sim \nu_0, \quad Y_0 = 0. \quad (9d)$$

Here the velocity variable $\xi_t \in \mathfrak{g}$ stands for elements in Lie algebras, and the rotation variable $g_t \in G$ means elements in Lie groups. The output variable $Y_t \in \mathbb{R}^m$ represents raw observations, and the joint variable (g_t, ξ_t^\vee) stands for the state of the dynamic system. In addition, W_t and B_t represent two multidimensional independent Wiener processes as Section II-B. We also denote the initial distribution of (g_t, ξ_t^\vee) by ν_0 , where ν_0 is a well-defined distribution defined $G \times \mathfrak{g}$. The aim of filtering problem is to approximate the posterior distribution of (g_τ, ξ_τ^\vee) given the observation up to τ , i.e., $P((g_\tau, \xi_\tau^\vee) | Y_t, 0 \leq t \leq \tau)$. Note that the dynamic (9a) implies that, assuming (g_0, ξ_0) is fixed, the state variable $\{(g_\tau, \xi_\tau^\vee), 0 \leq \tau \leq t\}$ is fully determined by $\{\xi_\tau^\vee, 0 \leq \tau \leq t\}$. Hence, loosely speaking, the measure of (g_t, ξ_t^\vee) is the marginal measure of ξ_t^\vee . For simplicity, we assume the parameters of (9) are sufficiently regular so that the solution to this system exists. Actually, the dynamics driven by (9a) and (9b) satisfy the Doléans-Dade and Protter's equations, and the existence conditions of solutions for this type of equations have been studied in [19]. For (9c), we assume h_t is integrable almost surely.

In section II-B, we have shown the smoothing problem on \mathbb{R}^d can be reformulated as a stochastic control problem. In

fact, the same framework can be applied on Lie groups. Let \mathcal{P} be the measure of (g_t, ξ_t^\vee) given by (9a) and (9b). Denote \mathcal{Q}^Y to be its posterior distribution. Moreover, we denote the state of the following Itô process with control signal u_t by $(\tilde{g}_t, \tilde{\xi}_t)$ and the associated measure by $\tilde{\mathcal{P}}$.

$$d\tilde{g}_t = dL_{\tilde{g}_t}(\tilde{\xi}_t)dt, \quad (10a)$$

$$d\tilde{\xi}_t^\vee = b(t, g_t, \tilde{\xi}_t)dt + \sigma_W(u_t dt + dW_t), \quad (10b)$$

$$(\tilde{g}_0, \tilde{\xi}_0) \sim \pi_0. \quad (10c)$$

The next lemma reveals that the Kullback–Leibler divergence of $\tilde{\mathcal{P}}$ and \mathcal{P} of (9) is in the same form as the \mathbb{R}^d case.

Lemma 1: Assume the Novikov's condition holds, i.e., $\mathbb{E}[\exp(\frac{1}{2} \int_0^T \|u_t\|^2 dt)] < \infty$, then

$$\text{KL}(\tilde{\mathcal{P}} \parallel \mathcal{P}) = \mathbb{E} \left\{ \int_0^T \frac{1}{2} \|u_t\|^2 dt \right\} + \text{KL}(\pi_0 \parallel \nu_0). \quad (11)$$

Proof: In the standard Gisanov theorem, an elementary assumption is the dynamic is the state variable is in \mathbb{R}^d , which seems to be violated in our problem. However, the ξ_t^\vee -marginal measure satisfies this assumption and the Novikov's condition holds, so one can compute the Kullback–Leibler divergence of marginals as follows. Let \mathcal{P}_ξ and $\tilde{\mathcal{P}}_\xi$ be the ξ_t^\vee -marginals for \mathcal{P} and $\tilde{\mathcal{P}}$, respectively. Then, assuming the initial values are given we have

$$\log \frac{d\tilde{\mathcal{P}}_\xi}{d\mathcal{P}_\xi} = \int_0^T \frac{1}{2} \|u_t\|^2 dt + u'_t dW_t.$$

As the sample path of $g_{[0,T]}$ is exactly a function of $\xi_{[0,T]}$, it follows

$$\begin{aligned} \log \frac{d\tilde{\mathcal{P}}}{d\mathcal{P}} &= \log \frac{d\tilde{\mathcal{P}}_\xi}{d\mathcal{P}_\xi} + \log \frac{d\pi_0}{d\nu_0} \\ &= \int_0^T \frac{1}{2} \|u_t\|^2 dt + u'_t dW_t + \log \frac{d\pi_0}{d\nu_0}, \end{aligned}$$

which yields the Kullback–Leibler divergence in Lemma 1. Note that the Novikov's condition is already satisfied in most practical applications, so we can assume it is true as most previous literatures do. ■

Lemma 2: Assume h_t is an adapted and continuously differentiable process and the condition, $\int_0^T \|h_s\|^2 ds < \infty$ for any T , holds almost surely. Then we have the Radon–Nikodym derivative,

$$\begin{aligned} \frac{d\mathcal{Q}^Y}{d\mathcal{P}} &\propto \exp \left\{ - \int_0^T \frac{1}{2\sigma_B^2} \|h(t, \tilde{g}_t, \tilde{\xi}_t)\|^2 dt \right. \\ &\quad \left. - \int_0^T \frac{1}{\sigma_B^2} Y'_t dh(t, \tilde{g}_t, \tilde{\xi}_t) + \frac{1}{\sigma_B^2} Y'_T h(T, \tilde{g}_T, \tilde{\xi}_T) \right\}. \end{aligned} \quad (12)$$

Proof: In the original derivation of the Pathwise Kallianpur–Striebel formula (Proposition 1.4.2. in [15]), the dynamics of state variables are assumed to be an adapted càdlàg process in a Polish space. In our setting, the state variable (g_t, ξ_t) follows an continuous diffusion process in the Cartesian product space of g_t and ξ_t , which does satisfy

the prementioned assumption. One can check the assumption that $\int_0^T \|h_s\|^2 ds < \infty$ for any T , holds almost surely, is the main assumption in [15], which yields

$$\begin{aligned} \frac{d\mathcal{Q}^Y}{d\mathcal{P}} &\propto \exp \left\{ - \int_0^T \frac{1}{2\sigma_B^2} \|h(t, \tilde{g}_t, \tilde{\xi}_t)\|^2 dt \right. \\ &\quad \left. + \int_0^T \frac{1}{\sigma_B^2} h'(t, \tilde{g}_t, \tilde{\xi}_t) dY_t \right\}. \end{aligned}$$

It is easy to verify that with the assumption that h_t is adapted and continuously differentiable, the integral $\int h'_t dY_t$ satisfies the stochastic integration by parts formula. Therefore, the Radon–Nikodym derivative in our case is in the same form as the one in [15]. ■

By plugging (11) and (12) into $\text{KL}(\tilde{\mathcal{P}} \parallel \mathcal{Q}^Y)$ in (3), the corresponding stochastic optimal control problem becomes

$$\begin{aligned} \min_{u, \pi_0} \mathbb{E} \left\{ \int_0^T \left[L(t, \tilde{g}_t, \tilde{\xi}_t) + \frac{1}{2} \|u_t\|^2 \right] dt + \Phi(\tilde{g}_T, \tilde{\xi}_T) \right\} \\ + \text{KL}(\pi_0 \parallel \nu_0), \end{aligned} \quad (13)$$

where

$$L(t, x, y) dt = \frac{1}{2\sigma_B^2} \|h(t, x, y)\|^2 dt + \frac{1}{\sigma_B^2} Y'_t dh(t, x, y) \quad (14a)$$

$$\Phi(x, y) = -\frac{1}{\sigma_B^2} Y'_T h(T, x, y). \quad (14b)$$

In [16], the authors point out the optimization problem of this form can be solved separately. The optimal signal u_t^* depends only on t and X_t , thus one can calculate u_t^* initially. It has been shown that u_t^* can be solved in the similar manner for \mathbb{R}^d [20]. After the computation of u_t^* , one can find that the optimal initial distribution π_0^* is the posterior distribution of (g_0, ξ_0^\vee) given all the observations $\{Y_t, 0 \leq t \leq T\}$ which resembles the \mathbb{R}^d setting [16]. Thus, if the solution of (13) exists, the minimum value is 0. For the sake of completeness, we summarize our previous result in the following theorem.

Theorem 3: Assume a solution to (13) exists. and the conditions of Lemma 1 and Lemma 2 hold. For the dynamic system modeled by (9) and (10), denote the related optimal solution to (13) by u_t^* and π_0^* . Then the posterior measure \mathcal{Q}^Y is equal to the measure of dynamics $\tilde{\mathcal{P}}$ driven by u_t^* and π_0^* .

Proof: The form of the corresponding stochastic optimal control problem follows from Lemma 1 and Lemma 2. Since the dynamics $\tilde{\mathcal{P}}$ driven by u_t^* and π_0^* satisfy $\text{KL}(\tilde{\mathcal{P}} \parallel \mathcal{Q}^Y) = 0$, one can recover the posterior measure \mathcal{Q}^Y through sampling from $\tilde{\mathcal{P}}$. ■

B. Algorithms for Smoothing and Filtering Problems

It is worthwhile to point out that the computation cost of u_t^* and π_0^* is demanding. Therefore, we adopt importance sampling with suboptimal solutions. The procedure is essentially the same with the one in [16], as, in Lemma 1, we already show Gisanov theorem is applicable in this case. Consider K trajectories $\{g_i^k, (\xi_i^k)^\vee\}$ generated by

suboptimal conditions and denote the value of $S_u(0, T)$ over the k th trajectory by $S_u^k(0, T)$. Then the posterior distribution of (g_t, ξ_t^\vee) is approximated by $\sum_{k=1}^K w^k \delta_{(g_t^k, (\xi_t^k)^\vee)}$, where

$$w^k \propto \frac{d\nu_0}{d\pi_0}(g_0^k, (\xi_0^k)^\vee) \exp[-S_u^k(0, T)], \quad (15)$$

and

$$S_u(t, s) = \int_t^s \left[L(\tau, \tilde{g}_\tau, \tilde{\xi}_\tau) + \frac{1}{2} \|u_\tau\|^2 \right] d\tau \quad (16)$$

$$+ \int_t^s u'_\tau dW_\tau + \Phi(\tilde{g}_T, \tilde{\xi}_T).$$

In what follows, we divide $[0, T]$ as $0 = t_0 < t_1 < t_2 < \dots < t_N = T$ where for the integer i , $t_i - t_{i-1} = \Delta t$, a fixed time horizon. If Δt is sufficiently small, then the discretization of (9a) can be written with the exponential mapping. We also apply the one-step Euler-Maruyama approximation to (9b), which yields the following discrete simulation step,

$$g_{i+1} = g_i \exp(\Delta t \xi_i) \quad (17a)$$

$$\xi_{i+1}^\vee = \xi_i^\vee + b(i, g_i, \xi_i) \Delta t + \sigma_W(u_i \Delta t + \sqrt{\Delta t} \epsilon_i) \quad (17b)$$

Here, g_0 and ξ_0 are sampled from π_0 , and ϵ_i follows the standard Gaussian distribution. Note that we replace t_i by i for the ease of presentation. For the observation model, using the same approach, we have

$$Y_{i+1} = Y_i + h(i, g_i, \xi_i) \Delta t + \sigma_B \sqrt{\Delta t} \delta_i, \quad (18)$$

with $\delta_i \sim N(0, \mathbb{I})$ being a Gaussian random variable.

Next, consider the discrete version of trajectory cost (15) as follows,

$$S_u(0, T) \approx \sum_{i=0}^{N-1} \left[\frac{1}{2} \|u_i\|^2 + \frac{1}{2\sigma_B^2} \|h(i, \tilde{g}_i, \tilde{\xi}_i)\|^2 \right] \Delta t \quad (19)$$

$$- \sum_{i=0}^{N-1} \frac{1}{\sigma_B^2} h(i, \tilde{g}_i, \tilde{\xi}_i)' (Y_{i+1} - Y_i) + u'_i \sqrt{\Delta t} \epsilon_i,$$

where the first equation follows from integration by parts.

The control-based particle smoothing algorithm on Lie groups is shown in Algorithm 1. The key difference of smoothing and filtering is that filtering refers to a length-varying time window while smoothing corresponds to a fixed time window. A naive but computationally-expensive implementation for filtering problem is to recursively run Algorithm 1 over the increasing time interval. Specifically, for every time point τ , run the smoothing algorithm over $[0, \tau]$ and simulate particles $\{g_\tau^k, \xi_\tau^k\}_{k=1}^K$ to approximate $P((g_\tau, \xi_\tau) | Y_t, 0 \leq t \leq \tau)$. To reduce the computational demand, we combine the sliding window mechanism proposed in [16] to improve the efficiency of filters.

The essential component of sliding window algorithm is to involve a sliding window parameter H . For the j -th time point, if $j \leq H$, then simply run Algorithm 1 over the window $[t_0, t_j]$. Otherwise, solve the smoothing problem

Algorithm 1 Control-based Particle Smoothing on Lie Groups

Input: b, σ_W, h, σ_B : Model formulation and parameters

ν_0 : Initial distribution

N : Total number of time points

K : Total number of particles

Output: $\{g^k, \xi^k\}, \{w^k\}$: Smoothing results

$L, \Phi \leftarrow$ Formulate the corresponding optimal control problem by (14) and (13)

u_t and $\pi_0 \leftarrow$ Find the suboptimal solutions of (13)

for $k \leftarrow 1, \dots, K$ **do**

$\{g_i^k, \xi_i^k\} \leftarrow$ Sample the k th trajectory at time point i following the dynamic (10) with suboptimal conditions

u_t and π_0

$\{w^k\} \leftarrow$ Compute the weights of trajectories by (15)

end for

over $[t_{j-H}, t_j]$ where the initial distribution of t_{j-H} , $\nu_{t_{j-H}}$, is recursively updated by the following formulas. Assume

$$\nu_{t_{j-H-1}} \approx \sum_{k=1}^K w_p^k X_{t_{j-H-1}}^k, \quad (20)$$

then

$$\nu_{t_{j-H}} \approx \sum_{k=1}^K \hat{w}_p^k X_{t_{j-H}}^k. \quad (21)$$

Here $X_{t_{j-H}}^k$ represents the k th particle at t_{j-H} sampled by the smoothing algorithm over $[t_{j-H-1}, t_{j-H}]$, and \hat{w}_p^k is the normalized weights of $w_p^k \exp[-S_u^k(t_{j-H-1}, t_{j-H})]$, where $\exp[-S_u^k(t_1, t_2)]$ stands for the trajectory cost over $[t_1, t_2]$. In the standard particle filtering, a resampling step is performed to reduce weight degeneracy [21]. Therefore, we utilize the resampling step as in [16], which leads to Algorithm 2, the following control-based filtering algorithm on Lie groups for the case $j > H$.

IV. EXPERIMENTS

To demonstrate the efficacy of our algorithm, we consider the following position and velocity estimation problem of a 3D rigid body whose state variable is the tangent bundle of $\text{SO}(3)$. To begin with, the system dynamics are given by

$$dg_t = g_t \xi_t dt \quad (22a)$$

$$d\xi_t^\vee = \mathbb{M}^{-1}(\mathbb{M} \xi_t^\vee \times \xi_t^\vee + \mathbb{H} \sigma u_t) dt + \mathbb{M}^{-1} \mathbb{H} \sigma dW_t \quad (22b)$$

$$(g_0, \xi_0) \sim \pi_0 \quad (22c)$$

In the above dynamics, the variable $g_t \in \text{SO}(3)$ represents the rotation matrix, and the column vector $\xi_t^\vee \in \mathbb{R}^3$ is the body-fixed velocity. We also denote the control signal by $u_t \in \mathbb{R}^3$. Note that for $\text{SO}(3)$, $dL_{g_t}(\xi_t)$ can be simplified to matrix multiplication of g_t and ξ_t . The parameter $\mathbb{M} \in \mathbb{R}^{3 \times 3}$ is the inertia tensor, and $\mathbb{H} \in \mathbb{R}^{3 \times 3}$ represents the control torques. In addition σ is used to modulate the magnitudes of the control signal and noise signal. Here the symbol \times

Algorithm 2 Control-based Particle filtering on Lie Groups

Input: $b, \sigma_W, h, \sigma_B, \nu_0, N, K$
j: Current time point
H: The Length of the sliding window
 $\bar{\gamma}$: Threshold of performing resampling
 $\sum_{k=1}^K w_p^k X_{t_{j-H}}^k$: Prior information at the previous step
Output: $\{X^k\}, \{w^k\}$: Filtering results $\{w_p^k, X_p^k\}$: Prior information for the next sliding window

$L, \Phi \leftarrow$ Formulate the corresponding optimal control problem by (14) and (13) over $[t_{j-H}, t_j]$

for $k \leftarrow 1, \dots, K$ **do**

- $X^k \leftarrow$ Sample the k th trajectory by (17) over $[t_{j-H}, t_j]$ initialized by $X_{t_{j-H}}^k$
- $\exp[-S_u^k(t_{j-H}, t_j)] \leftarrow$ Compute the weight of k th trajectory by (15) over $[t_{j-H}, t_j]$.
- $w^k \propto w_p^k \exp[-S_u^k(t_{j-H}, t_j)]$
- $X_p^k = X_{t_{j-H+1}}^k$
- $\{w_p^k\} \propto \{w^k \exp[-S_u^k(t_{j-H}, t_{j-H+1})]\}$

end for

Calculate the effective ratio γ of the collection of weights $\{w^k\}_{k=1}^K$ by $\gamma = \frac{1}{K \sum_{i=1}^K w_i^2}$.

if $\gamma < \bar{\gamma}$ **then**

- Sample $\{X_p^k\}$ from multinomial($\{X_p^k\}, \{w_p^k\} \exp[-S_u^k(t_{j-H+1}, t_j)]$)
- $\{w_p^k\} \propto \exp[S_u^k(t_{j-H+1}, t_j)]$

end if

stands for the cross product in the three-dimensional vector space. For $\mathfrak{so}(3)$, we define $\vee : \mathfrak{g} \rightarrow \mathbb{R}^3$ by

$$\begin{bmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{bmatrix}^\vee = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}.$$

The deterministic version of the prementioned forward dynamics was initially proposed in [22] and was later adopted in [13] as an underlying dynamics in optimal control.

For the observation model, we assume the sensors have access to both rotation matrix and the velocity. The rotation matrix g_t is measured by the accelerometer whose output is $-g_t' r_g$, where r_g is a unit vector aligned with the gravity. Similarly, we also measure g_t with the magnetometer which gives $g_t' r_b$ with r_b being a unit vector aligned with the magnetic field (See [23] for more discussion of the two sensors). In summary, our observation model is

$$dY_t = h(t, g_t, \xi_t)dt + \sigma_B dB_t, \quad (23)$$

where $h(t, g_t, \xi_t) = (-g_t' r_g, g_t' r_b, \xi_t^\vee) \in \mathbb{R}^9$.

To define an initial distribution π_0 that is easy to sample from, we first generate a sample $x \sim N(0, \Sigma_{6 \times 6})$. Then let the first three components $(x_1, x_2, x_3)'$ be ξ_0^\vee and g_0 is given by $\exp([x_4, x_5, x_6]')$. This approach is consistent with defining the randomness on $\text{SO}(3)$ by latent variables in a vector space and then mapping them to $\text{SO}(3)$ via suitable

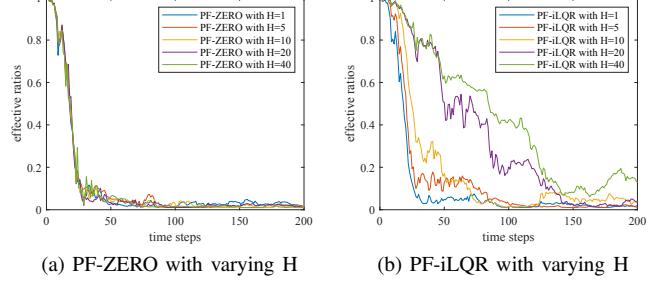


Fig. 1: Comparison of Effective Ratios

mappings [24]. In this way, the joint distribution of (g_0, ξ_0^\vee) is well-defined.

We assume the suboptimal initial distribution ν_0 is always the given one in the prior dynamics. As the control signal u_t is the function of the current state x_t , it is sufficient to solve our original optimal control problem (13) by considering the following deterministic and discrete version,

$$\begin{aligned} \min_{u_t} \sum_{i=0}^{N-1} & \frac{1}{2\sigma_B^2} \|h_i\|_2^2 + \frac{1}{2} \|u_i\|_2^2 - \frac{1}{\sigma_B^2} h_i'(Y_{i+1} - Y_i)/\Delta t \\ & g_{i+1} = g_i \exp(\Delta t \xi_i), \\ & \xi_{i+1}^\vee = \xi_i^\vee + \mathbb{M}^{-1}(\mathbb{M} \xi_t^\vee \times \xi_t^\vee + \mathbb{H} \sigma u_t) \Delta t \\ & + \mathbb{M}^{-1} \mathbb{H} \sigma \sqrt{\Delta t} \epsilon_i. \end{aligned} \quad (24)$$

Here, our ground truth data Y_i is generated by simulation. In the following experiments, we set $\mathbb{M} = \text{diag}([1, 1.11, 1.3])$, $\mathbb{H} = \mathbb{I}_3$, $\sigma = 1$, $\sigma_B = 0.1$, $\Delta t = 0.005$, $N = 200$, $K = 100$, $\bar{\gamma} = 0.1$, $r_g = [0, 0, 1]$, $r_b = [1/\sqrt{2}, 0, 1/\sqrt{2}]$.

The method we adopt to solve this optimal control problem is based on Differential Dynamic Programming over Lie groups (LDDP) proposed in [13]. As our framework does not require the exact optimal control signal, we only utilize the linear approximation of the transition function in this paper, which is the iterative linear quadratic regulator (iLQR). We denote this algorithm by PF-iLQR that means the suboptimal solution of (24) is given by iLQR. As for the baseline algorithm, the suboptimal control signal is set as zero control, which we denote by PF-ZERO. Note that if $H = 1$, then PF-ZERO reduces to the standard Sequential Importance Resampling (SIR) algorithm (See Remark 1 in [16]). We measure the quality of the two algorithms with two criteria: the effective ratio and the error with respect to the simulated ground truth data.

The comparison for the effective ratio is depicted in Figure 1. Clearly, PF-iLQR significantly outperforms PF-ZERO, as the suboptimal signal generated by iLQR would control the evolution of particles to approximate the posterior distribution. We also compare the effects of different H and observe increasing the value of H would in general improve the effective ratios of PF-iLQR. However, very large H may deteriorate the performance. In Table I, PF-iLQR with

$H = 200$ that represents using all available data shows slightly worse results than PF-iLQR with $H = 40$. Similar result has also been observed in [16]. For PF-ZERO, we observe its performance is nearly irrelevant with H .

In addition to the effective ratios, we also test the errors with respect to the simulated data. For the Lie algebra, $\mathfrak{so}(3)$, we compare the mean squared errors. While for $\text{SO}(3)$, we follow the procedure used in [10], which means computing the rotation angle error $\delta_\theta \in [0, 180^\circ]$ by the quaternion representation of rotations. More specifically, Let q^* be the unit quaternion of ground-true rotation and \bar{q} be the weighted average of quaternions. Then $\delta_\theta = 2 \arccos(|\delta q^0|)$ where δq^0 is the first component of $(\bar{q})^{-1} \times q^*$. In each trial, the error is averaged over time, and we repeat each algorithm for 20 trials when H is larger than 1. The number of repeated trials is 50 for $H = 1$ to reduce the impact of randomness. For completeness, we also compare their performance with resampling steps in Table II. Indeed, the improvement with resampling steps is not as significant as the case w/o resampling, as PF-ZERO executes more resampling steps than PF-iLQR. It is worth noting that even in this case, PF-iLQR outperforms PF-ZERO.

TABLE I: Comparison of Errors w/o resampling

	MSE of $\mathfrak{so}(3)$	Errors of $\text{SO}(3)$
PF-ZERO with $H=1$	0.5182	9.0206
PF-ZERO with $H=20$	0.5230	8.4491
PF-iLQR with $H=1$	0.4710	8.8710
PF-iLQR with $H=5$	0.3835	7.4118
PF-iLQR with $H=10$	0.3210	6.9799
PF-iLQR with $H=20$	0.2918	6.4658
PF-iLQR with $H=40$	0.2606	6.3510
PF-iLQR with $H=200$	0.2548	6.3976

TABLE II: Comparison of Errors with resampling

	MSE of $\mathfrak{so}(3)$	Errors of $\text{SO}(3)$
PF-ZERO with $H=1$	0.2934	7.3829
PF-ZERO with $H=20$	0.2801	7.0383
PF-iLQR with $H=1$	0.2924	7.6737
PF-iLQR with $H=5$	0.2721	6.4915
PF-iLQR with $H=10$	0.2799	6.5156
PF-iLQR with $H=20$	0.2648	6.3164
PF-iLQR with $H=40$	0.2544	6.3311
PF-iLQR with $H=200$	0.2533	6.4026

V. CONCLUSIONS

In this paper, we generalized the duality between smoothing and optimal control to Lie groups, and building on this duality we developed a particle filtering algorithm for dynamical systems over Lie groups. Our algorithm uses iLQR on Lie groups to approximately solve the dual optimal control problem. The resulting PF-iLQR algorithm significantly improves the effective ratios and thus the overall performance of particle filtering algorithms.

REFERENCES

- [1] B. D. Anderson and J. B. Moore, Optimal filtering. Courier Corporation, 2012.
- [2] H. Durrant-Whyte and T. Bailey, “Simultaneous localization and mapping: part I,” IEEE robotics & automation magazine, vol. 13, no. 2, pp. 99–110, 2006.
- [3] L. Bugoniu, V. S. Varma, J. Lohéac, A. Codrean, O. Štefan, I.-C. Morărescu, and S. Lasaulce, “Learning control for transmission and navigation with a mobile robot under unknown communication rates,” Control Engineering Practice, vol. 100, p. 104460, 2020.
- [4] T. Zhang, K. Wu, J. Song, S. Huang, and G. Dissanayake, “Convergence and consistency analysis for a 3D invariant-EKF SLAM,” IEEE Robotics and Automation Letters, vol. 2, no. 2, pp. 733–740, 2017.
- [5] N. Y. Ko, G. Song, W. Youn, I. H. Choi, and T. S. Kim, “Improvement of extended Kalman filter using invariant extended Kalman filter,” in 2018 18th International Conference on Control, Automation and Systems (ICCAS). IEEE, 2018, pp. 948–950.
- [6] A. Barrau and S. Bonnabel, “The invariant extended Kalman filter as a stable observer,” IEEE Transactions on Automatic Control, vol. 62, no. 4, pp. 1797–1812, 2016.
- [7] H. A. Hashim, M. Abouheaf, and M. A. Abido, “Geometric stochastic filter with guaranteed performance for autonomous navigation based on IMU and feature sensor fusion,” Control Engineering Practice, vol. 116, p. 104926, 2021.
- [8] M. Izadi and A. K. Sanyal, “Rigid body pose estimation based on the Lagrange-d'Alembert principle,” Automatica, vol. 71, pp. 78–88, 2016.
- [9] A. K. Sanyal, M. Izadi, D. J. Scheeres, G. Misra, and E. Samiei, “Estimation of dynamics of space objects from visual feedback during proximity operations,” in AIAA/AAS Astrodynamics Specialist Conference, 2014, p. 4419.
- [10] C. Zhang, A. Taghvaei, and P. G. Mehta, “Feedback particle filter on Riemannian manifolds and matrix Lie groups,” IEEE Transactions on Automatic Control, vol. 63, no. 8, pp. 2465–2480, 2017.
- [11] S. K. Mitter and N. J. Newton, “A variational approach to nonlinear estimation,” SIAM journal on control and optimization, vol. 42, no. 5, pp. 1813–1833, 2003.
- [12] J. W. Kim and P. G. Mehta, “An optimal control derivation of nonlinear smoothing equations,” in Advances in Dynamics, Optimization and Computation: A volume dedicated to Michael Dellnitz on the occasion of his 60th birthday. Springer, 2020, pp. 295–311.
- [13] G. I. Boutselis and E. Theodorou, “Differential dynamic programming on Lie groups: Derivation, convergence analysis and numerical results,” arXiv preprint arXiv:1809.07883, 2018.
- [14] B. Øksendal, “Stochastic differential equations,” in Stochastic differential equations. Springer, 2003.
- [15] R. Van Handel, “Filtering, stability, and robustness,” Ph.D. dissertation, California Institute of Technology, 2007.
- [16] Q. Zhang, A. Taghvaei, and Y. Chen, “An optimal control approach to particle filtering,” arXiv preprint arXiv:2110.03199, 2021.
- [17] J. Sola, J. Deray, and D. Atchutan, “A micro Lie theory for state estimation in robotics,” arXiv preprint arXiv:1812.01537, 2018.
- [18] E. Todorov, “General duality between optimal control and estimation,” in 2008 47th IEEE Conference on Decision and Control. IEEE, 2008, pp. 4286–4292.
- [19] M. Métivier and J. Pellaumail, “On a stopped Doob’s inequality and general stochastic equations,” The Annals of Probability, pp. 96–114, 1980.
- [20] G. I. Boutselis and E. A. Theodorou, “Path integral control on Lie groups,” in 2018 IEEE Conference on Decision and Control (CDC). IEEE, 2018, pp. 4990–4995.
- [21] A. Doucet, M. Briers, and S. Sénecal, “Efficient block sampling strategies for sequential monte carlo methods,” Journal of Computational and Graphical Statistics, vol. 15, no. 3, pp. 693–711, 2006.
- [22] P. Crouch, “Spacecraft attitude control and stabilization: Applications of geometric control theory to rigid body models,” IEEE Transactions on Automatic Control, vol. 29, no. 4, pp. 321–331, 1984.
- [23] R. Mahony, T. Hamel, and J.-M. Pflimlin, “Nonlinear complementary filters on the special orthogonal group,” IEEE Transactions on automatic control, vol. 53, no. 5, pp. 1203–1218, 2008.
- [24] L. Falorsi, P. de Haan, T. R. Davidson, and P. Forré, “Reparameterizing distributions on Lie groups,” in The 22nd International Conference on Artificial Intelligence and Statistics. PMLR, 2019, pp. 3244–3253.