Tracking the spin on a ping pong ball with the quaternion Bingham filter

Tracking the Spin on a Ping Pong Ball with the Quaternion Bingham Filter

Jared Glover and Leslie Pack Kaelbling

***Abstract*— A deterministic method for sequential estimation of 3-D rotations is presented. The Bingham distribution is used to represent uncertainty directly on the unit quaternion hypersphere. Quaternions avoid the degeneracies of other 3-D orientation representations, while the Bingham distribution allows tracking of large-error (high-entropy) rotational distributions. Experimental comparison to a leading EKF-based filtering approach on both synthetic signals and a ball-tracking dataset shows that the Quaternion Bingham Filter (QBF) has lower tracking error than the EKF, particularly when the state is highly dynamic. We present two versions of the QBF– suitable for tracking the state of first- and second-order rotating dynamical systems.**

I. Introduction

As any fan of this high-speed sport knows, table tennis is a game of spin. Because of the high-friction soft rubber surfaces of modern ping pong paddles, ball spin—which can reach 150 rotations per second—plays an enormous role in determining the trajectory of the ball after being hit. A top-spin ball tends to fly off of your paddle up into the air, while an under-spin ball will drop down into the net. Due to air resistance, spin can also change the in-flight trajectories of balls as well, a phenomenon known as the "Magnus effect" which leads to the well-known "curveball" in the sport of baseball. Ping pong players go to great length to disguise the spins they put on the ball, and much of the training of table tennis professionals goes into the technique and strategy of handling different types of spins. Swings which impart different types of spin are given different names, like "loop", "chop", "push", and "flip".

Several robots have been programmed to play ping pong over the years [1], [16], [20], [17]. However, only the earliest of these systems (by Russell Andersson in the late 1980s) made any attempt to track the spin on the ball. Because it used only indirect measurements of the spin (via the Magnus effect), the spin estimates were extremely noisy [1]. Fortunately for Andersson, he used only low-friction wooden paddles (with no rubber surface), so the effects of spin were minimized, and the robot was able to hold its own against novice players at moderate speeds and spins.

In this paper we present an approach to track the spin on the ping pong ball from direct measurements of the ball's orientation with a high-speed camera (Figure 2).

This work was supported in part by the NSF under Grant No. 1117325. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. We also gratefully acknowledge support from ONR MURI grant N00014-09-1-1051, from AFOSR grant FA2386-10-1-4135 and from the Singapore Ministry of Education under a grant to the Singapore-MIT International Design Center.

Jared Glover and Leslie Pack Kaelbling are with the Massachusetts Institute of Technology. Email: {jglov,lpk}@mit.edu

Fig. 1. Our ping pong robot uses a 7-dof Barrett WAM arm with a ping pong paddle rigidly attached at the wrist.

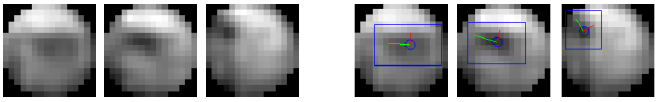

Fig. 2. Detecting the logo on a ping pong ball in high-speed images can be a tricky task, due to the small size of the ball (40mm) and the motion blur. (Left) Original images. (Right) Ball orientation estimates.

3-D rotational data occurs in many disciplines, from geology to robotics to physics. Yet modern statistical inference techniques are seldom applied to such data sets, due to the complex topology of 3-D rotation space, and the well-known aliasing problems caused by orientations "wrapping around" back to zero. Many probability distributions exist for modeling uncertainty on rotational data, yet difficulties often arise in the mechanics of complex inference tasks. The technical contribution of this paper is to explore one distribution—the Bingham—which is particularly well-suited for inference, and to derive some common operations on Bingham distributions as a reference for future researchers.

We present a new deterministic method—the Quaternion Bingham Filter (QBF)—for approximate recursive inference in quaternion Bingham processes. The quaternion Bingham process is a type of dynamic Bayesian network (DBN) on 3-D rotational data, where both process dynamics and measurements are perturbed by random Bingham rotations. The QBF uses the Bingham distribution to represent state uncertainty on the unit quaternion hypersphere. Quaternions avoid the degeneracies of other 3-D orientation representations, while the Bingham enables accurate tracking of high-noise signals. Performing exact inference on quaternion Bingham processes requires the composition of Bingham distributions, which results in a non-Bingham density. Therefore, we approximate the resulting composed distribution as

a Bingham using the method of moments, in order to keep the filtered state distribution in the Bingham family.

We compare the quaternion Bingham filter to a previous approach to tracking rotations based on the Extended Kalman Filter (EKF) and find that the QBF has lower tracking error than the EKF, particularly when the state is highly dynamic. We evaluate the performance of the QBF on both synthetic rotational process signals and on a real dataset containing 3-D orientation estimates of a spinning ping-pong ball tracked in high-speed video. We also derive the true probability density function (PDF) for the composition of two Bingham distributions, and report the empirical error of the moment-matching composition approximation for various distributions. We begin by introducing the Bingham distribution and presenting operations on it that will be needed by the filter. Then, we derive the first- and second-order quaternion Bingham processes and the QBF for estimating their state. We conclude with experiments on artificial and real data.

II. Distributions on rotations

The problem of how to represent a probability distribution on the space of rotations in three dimensions has been a subject of considerable study. Representing the distribution directly in the space of Euler angles is difficult because of singularities in the space when two of the angles are aligned (known as gimbal lock). A more appropriate space for representing distributions on rotations is the space of unit quaternions: a rotation becomes a point on the 4-dimensional unit hypersphere, $\mathbb{S}^3$. This space lacks singularities, but has the difficulty that the representation is not unique: both $\mathbf{q}$ and $-\mathbf{q}$ represent the same rotation. Putting a Gaussian distribution directly in quaternion space does not respect the underlying topology of 3-D rotations; however, this approach has been the basis of tracking methods based on approximations of the Kalman filter [13], [4], [15], [5], [10], [11]. A more appropriate method is to represent distributions in an $\mathbb{R}^3$ space that is tangent to the quaternion hypersphere at the mean of the distribution [6]; but such a tangent-space approach will be unable to effectively represent distributions that have large variances. In many perceptual problems, it may be possible to make observations that provide significant information about only one or two dimensions, yielding high-variance estimates. For this reason, we use the Bingham distribution.

The Bingham distribution is commonly used as a distribution on 3-D rotations as unit quaternions [2], [9], [18]. Its density function (PDF) is given by

$$f(\mathbf{x};\Lambda,V) = \frac{1}{F}\exp\{\sum_{i=1}^{d}\lambda_i(\mathbf{v_i}^T\mathbf{x})^2\} \tag{1}$$

where $\mathbf{x}$ is a unit vector on the surface of the sphere $\mathbb{S}^d \subset \mathbb{R}^{d+1}$, F is a normalization constant, Λ is a vector of non-positive (≤ 0) concentration parameters, and the columns $\mathbf{v_i}$ of the $(d+1)\times d$ matrix V are orthogonal unit vectors.

The Bingham distribution is the *maximum entropy* distribution on the hypersphere which matches the sample inertia matrix $E[\mathbf{x}\mathbf{x}^T]$ [14]. Therefore, it may be better suited to representing random process noise on the hypersphere than some other distributions, such as (projected) tangent-space Gaussians. Binghams are also quite flexible, since a concentration parameter, λ_i, of zero indicates that the distribution is completely uniform in the direction of $\mathbf{v_i}$. They are therefore very useful in tracking problems where there is high, anisotropic noise. For example, to track the ping pong ball based on detections of its logo, the position of the logo can often be detected much more reliably than its orientation, so the axis (from the center of the ball to the logo) of the ball's 3-D orientation estimate will have less uncertainty than the angle.

III. Operations on Bingham distributions

In order to implement the quaternion Bingham filter, we need to be able to perform several operations on Bingham distributions. To our knowledge, all of these operations, except for computing the normalization constant, are new contributions of this paper. More operations (including calculation of KL-divergence and sampling methods) are presented in the accompanying tech report [7].

The Normalization constant. The primary difficulty with using the Bingham distribution in practice lies in computing the normalization constant, F. Since the distribution must integrate to one over its domain ($\mathbb{S}^d$), we can write the normalization constant as

$$F(\Lambda) = \int_{x\in\mathbb{S}^d}\exp\{\sum_{i=1}^{d}\lambda_i(\mathbf{v_i}^T\mathbf{x})^2\} = |\mathbb{S}^d|\cdot{}_1F_1(\frac{1}{2};\frac{d+1}{2};\Lambda) \tag{2}$$

where ${}_1F_1()$ is a hyper-geometric function of matrix argument [2]. Evaluating ${}_1F_1()$ is expensive, so we precompute a lookup table of F-values over a discrete grid of Λ's, and use tri-linear interpolation to quickly estimate normalizing constants on the fly.

Product of Bingham PDFs. The correction step of the filter requires multiplying PDFs. The product of two Bingham PDFs is given by adding their exponents:

$$\begin{aligned} f(x;&\Lambda_1,V_1)f(x;\Lambda_2,V_2) \\ &= \frac{1}{F_1F_2}\exp\{\mathbf{x}^T(\sum_{i=1}^{d}\lambda_{1i}\mathbf{v_{1i}}\mathbf{v_{1i}}^T + \lambda_{2i}\mathbf{v_{2i}}\mathbf{v_{2i}}^T)\mathbf{x}\} \\ &= \frac{1}{F_1F_2}\exp\{\mathbf{x}^T(C_1+C_2)\mathbf{x}\} \end{aligned} \tag{3}$$

After computing the sum $C = C_1 + C_2$ in the exponent of equation 3, we transform the PDF to standard form by computing the eigenvectors and eigenvalues of C, and then subtracting off the lowest magnitude eigenvalue from each spectral component, so that only the eigenvectors corresponding to the largest d eigenvalues (in magnitude) are kept, and $\lambda_1 \leq \cdots \leq \lambda_d \leq 0$ (as in equation 1).

Rotation by a fixed quaternion. To find the effect of the control on the predictive distribution, we rotate the prior by $\mathbf{u}$. Given $\mathbf{q} \sim \text{Bingham}(\Lambda, V)$, $\mathbf{u} \in \mathbb{S}^3$, and $\mathbf{s} = \mathbf{u}\circ\mathbf{q}$, then $\mathbf{s} \sim \text{Bingham}(\Lambda, \mathbf{u}\circ V)$, where $\mathbf{u}\circ V \triangleq [\mathbf{u}\circ\mathbf{v_1}, \mathbf{u}\circ\mathbf{v_2}, \mathbf{u}\circ\mathbf{v_3}]$. In other words, $\mathbf{s}$ is distributed according to

a Bingham whose orthogonal direction vectors have been rotated (on the left) by $\mathbf{u}$. Similarly, if $\mathbf{s} = \mathbf{q} \circ \mathbf{u}$ then $\mathbf{s} \sim \text{Bingham}(\Lambda, V \circ \mathbf{u})$.

Proof for $\mathbf{s} = \mathbf{u} \circ \mathbf{q}$: Since unit quaternion rotation is invertible and volume-preserving, we have

$$f_{\mathbf{s}}(\mathbf{s}) = f_{\mathbf{q}}(\mathbf{u^{-1}} \circ \mathbf{s}) = \frac{1}{F} \exp\{\sum_{i=1}^{d} \lambda_i (\mathbf{v_i}^T (\mathbf{u^{-1}} \circ \mathbf{s}))^2\}$$
$$= \frac{1}{F} \exp\{\sum_{i=1}^{d} \lambda_i ((\mathbf{u} \circ \mathbf{v_i})^T \mathbf{s})^2\} \quad .$$

Quaternion inversion. Given $\mathbf{q} \sim \text{Bingham}(\Lambda, V)$ and $\mathbf{s} = \mathbf{q^{-1}}$, then $\mathbf{s} \sim \text{Bingham}(\Lambda, JV)$, where J is the quaternion inversion matrix, $J = \begin{bmatrix} 1 & & & \\ & -1 & & \\ & & -1 & \\ & & & -1 \end{bmatrix}$. The proof follows the same logic as in the previous section.

Composition of quaternion Binghams. Implementing the QBF requires the computation of the PDF of the composition of two independent Bingham random variables. Letting $\mathbf{q} \sim \text{Bingham}(\Lambda, V)$ and $\mathbf{r} \sim \text{Bingham}(\Sigma, W)$, we wish to find the PDF of $\mathbf{s} = \mathbf{q} \circ \mathbf{r}$.

The **true distribution** is the convolution, in $\mathbb{S}^3$, of the PDFs of the component distributions[1].

$$f_{true}(\mathbf{s}) = \int_{\mathbf{q} \in \mathbb{S}^3} f(\mathbf{s}|\mathbf{q}) f(\mathbf{q})$$
$$= \frac{{}_1F_1(\frac{1}{2}; \frac{4}{2}; C(\mathbf{s}))}{|\mathbb{S}^3| \cdot {}_1F_1(\frac{1}{2}; \frac{4}{2}; \Lambda) {}_1F_1(\frac{1}{2}; \frac{4}{2}; \Sigma)}, \qquad (4)$$

where $C(\mathbf{s}) = \sum_{i=1}^{3} \Big(\sigma_i (\mathbf{s} \circ \mathbf{w_i}^{-1})(\mathbf{s} \circ \mathbf{w_i}^{-1})^T + \lambda_i \mathbf{v_i} \mathbf{v_i}^T \Big)$.

To **approximate** the PDF of $\mathbf{s}$ with a Bingham density, $f_B(\mathbf{s}) = f_B(\mathbf{q} \circ \mathbf{r})$, it is sufficient to compute the second moments of $\mathbf{q} \circ \mathbf{r}$, since the inertia matrix, $E[\mathbf{s}\mathbf{s}^T]$ is the sufficient statistic for the Bingham. This moment-matching approach is equivalent to the variational approach, where f_B is found by minimizing the KL divergence from f_B to f_{true}.

Noting that $\mathbf{q} \circ \mathbf{r}$ can be written as $(\mathbf{q}^T H_1^T \mathbf{r}, \quad \mathbf{q}^T H_2^T \mathbf{r}, \quad \mathbf{q}^T H_3^T \mathbf{r}, \quad \mathbf{q}^T H_4^T \mathbf{r})$, where $H_1 = \begin{bmatrix} 1&0&0&0 \\ 0&1&0&0 \\ 0&0&1&0 \\ 0&0&0&1 \end{bmatrix}$, $H_2 = \begin{bmatrix} 0&-1&0&0 \\ 1&0&0&0 \\ 0&0&0&1 \\ 0&0&-1&0 \end{bmatrix}$, $H_3 = \begin{bmatrix} 0&0&-1&0 \\ 0&0&0&-1 \\ 1&0&0&0 \\ 0&1&0&0 \end{bmatrix}$, and $H_4 = \begin{bmatrix} 0&0&0&-1 \\ 0&0&1&0 \\ 0&-1&0&0 \\ 1&0&0&0 \end{bmatrix}$, we find that

$$E[s_i s_j] = E[\mathbf{q}^T H_i^T \mathbf{r}\mathbf{r}^T H_j \mathbf{q}] \quad ,$$

which has 16 terms of the form $\pm r_a r_b q_c q_d$, where $a, b, c, d \in \{1, 2, 3, 4\}$. Since $\mathbf{q}$ and $\mathbf{r}$ are independent, $E[\pm r_a r_b q_c q_d] = \pm E[r_a r_b] E[q_c q_d]$. Therefore (by linearity of expectation), every entry in the matrix $E[\mathbf{s}\mathbf{s}^T]$ is a quadratic function of elements of the inertia matrices of $\mathbf{q}$ and $\mathbf{r}$, which can be easily computed given the Bingham normalization constants and their partial derivatives. The entire closed form for $E[\mathbf{s}\mathbf{s}^T]$ is given in the accompanying tech report [7].

Estimating the error of approximation. To estimate the error in the Bingham approximation to the composition of two quaternion Bingham distributions, $B_1 \circ B_2$, we approximate the KL divergence from f_B to f_{true} using a finite element approximation on the quaternion hypersphere

$$D_{KL}(f_B \parallel f_{true}) = \int_{\mathbf{x} \in \mathbb{S}^3} f_B(\mathbf{x}) \log \frac{f_B(\mathbf{x})}{f_{true}(\mathbf{x})}$$
$$\approx \sum_{\mathbf{x} \in \mathcal{F}(\mathbb{S}^3)} f_B(\mathbf{x}) \log \frac{f_B(\mathbf{x})}{f_{true}(\mathbf{x})} \cdot \Delta\mathbf{x}$$

where $\mathcal{F}(\mathbb{S}^d)$ and $\{\Delta\mathbf{x}\}$ are the points and volumes of the finite-element approximation to $\mathbb{S}^3$, based on a recursive tetrahedral-octahedral subdivision method [19].

Entropy. The entropy of a Bingham distribution with PDF f is given by:

$$h(f) \quad = \quad -\int_{x \in \mathbb{S}^d} f(x) \log f(x) \quad = \quad \log F - \Lambda \cdot \frac{\nabla F}{F} \quad . \qquad (5)$$

The proof is given in the accompanying tech report [7]. Since both the normalization constant, F, and its gradient with respect to Λ, ∇F, are stored in a lookup table, the entropy is trivial to approximate via interpolation, and can be used on the fly without any numerical integration over hyperspheres.

First-order

$\mathbf{x}_n = \mathbf{w}_n \circ \mathbf{u}_n \circ \mathbf{x}_{n-1}$
$\mathbf{y}_n = \mathbf{z}_n \circ \mathbf{x}_n$

Second-order

$\mathbf{v}_n = \mathbf{w}_n \circ \mathbf{u}_n \circ \mathbf{v}_{n-1}$
$\mathbf{x}_n = \mathbf{v}_{n-1} \circ \mathbf{x}_{n-1}$
$\mathbf{y}_n = \mathbf{z}_n \circ \mathbf{x}_n$

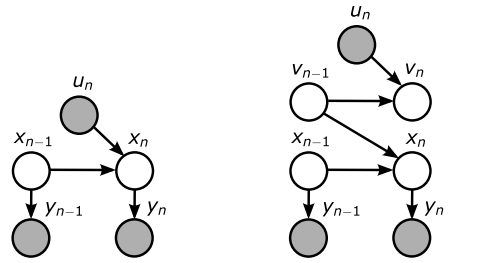

Fig. 3. Process and graphical models for the discrete quaternion Bingham process.

IV. Discrete-time quaternion Bingham process

The first-order discrete-time quaternion Bingham process has, as its state, $\mathbf{x}_n$, a unit quaternion representing the orientation of interest at time n. The system's behavior is conditioned on control inputs $\mathbf{u}_n$, which are also unit quaternions. The new orientation is the old orientation rotated by the control input and then by independent noise $\mathbf{w_n} \sim \text{Bingham}(\Lambda_p, V_p)$. Note that "$\circ$" denotes quaternion multiplication, which corresponds to composition of rotations for unit quaternions. ($\mathbf{q} \circ \mathbf{r}$ means "rotate by $\mathbf{r}$ and then by $\mathbf{q}$".)

The second-order quaternion Bingham process has state $(\mathbf{x}_n, \mathbf{v}_n)$, where $\mathbf{x}_n$ represents orientation and the quaternion $\mathbf{v}_n$ represents discrete rotational velocity at time n. The control inputs $\mathbf{u}_n$ are analogous to rotational accelerations. Process noise $\mathbf{w}_n$ enters the system in the velocity dynamics.

In both the first-order and second-order systems, observations $\mathbf{y}_n$ are given by the orientation $\mathbf{x}_n$ corrupted by independent Bingham noise $\mathbf{z_n} \sim \text{Bingham}(\Lambda_o, V_o)$. One

[1] See the accompanying tech report for a full derivation [7].

common choice for V_p and V_o is $\begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$, which means that the mode is the quaternion identity, $(1, 0, 0, 0)$. (Any V matrix whose top row contains all zeros will have this mode.) Figure 3 shows the process and graphical models for the discrete quaternion Bingham process.

V. Discrete quaternion Bingham filter

The state of a discrete-time quaternion Bingham process can be estimated using a discrete-time quaternion Bingham filter, which is a recursive estimator similar in structure to a Kalman filter. Unlike the Kalman filter, however, the QBF is approximate, in the sense that the tracked state distribution is projected to be in the Bingham family after every time step. The second-order QBF will also require an assumption of independence between $\mathbf{x}_n$ and $\mathbf{v}_n$, given all the data up to time n. Both the first-order and second-order QBFs are examples of *assumed density filtering*—a well-supported approximate inference method in the DBN literature [3]. We will start by deriving the first-order QBF, which follows the Kalman filter derivation quite closely. Note that the following derivations rely on several properties of Bingham distributions which were detailed in section III.

First-order QBF. Given a distribution over the initial state $\mathbf{x}_0$, $B_{\mathbf{x}_0} \sim \text{Bingham}(\Lambda_0, V_0)$, and an action-observation sequence $\mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n$, the goal is to compute the posterior distribution $f(\mathbf{x}_n \mid \mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n)$. We can use Bayes' rule and the Markov property to decompose this distribution as follows:

$$\begin{aligned} B_{\mathbf{x}_n} &= f(\mathbf{x}_n|\mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n) \\ &\propto f(\mathbf{y}_n|\mathbf{x}_n) f(\mathbf{x}_n|\mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_{n-1}, \mathbf{y}_{n-1}, \mathbf{u}_n) \\ &= f(\mathbf{y}_n|\mathbf{x}_n) \int_{\mathbf{x}_{n-1}} f(\mathbf{x}_n \mid \mathbf{x}_{n-1}, \mathbf{u}_n) B_{\mathbf{x}_{n-1}}(\mathbf{x_{n-1}}) \\ &= f(\mathbf{y}_n|\mathbf{x}_n)(f_{\mathbf{w}_n} \circ \mathbf{u}_n \circ B_{\mathbf{x}_{n-1}})(\mathbf{x}_n) \ . \end{aligned}$$

where $f_{\mathbf{w}_n} \circ \mathbf{u}_n \circ B_{\mathbf{x}_{n-1}}$ means rotate $B_{\mathbf{x}_{n-1}}$ by $\mathbf{u}_n$ and then convolve with the process noise distribution, $f_{\mathbf{w}_n}$. For the first term, $f(\mathbf{y}_n|\mathbf{x}_n)$, recall that the observation process is $\mathbf{y}_n = \mathbf{z}_n \circ \mathbf{x}_n$, so $\mathbf{y}_n|\mathbf{x}_n \sim Bingham(\mathbf{y_n}; \Lambda_o, V_o \circ \mathbf{x_n})$, where we used the result from section III for rotation of a Bingham by a fixed quaternion. Thus the distribution for $\mathbf{y}_n|\mathbf{x}_n$ is

$$f(\mathbf{y}_n|\mathbf{x}_n) = \frac{1}{F_o} \exp \sum_{i=1}^{3} \lambda_{oi} (\mathbf{y_n}^T (\mathbf{v_{oi}} \circ \mathbf{x_n}))^2 \ .$$

Now we can rewrite $\mathbf{y_n}^T(\mathbf{v_{oi}} \circ \mathbf{x_n})$ as $(\mathbf{v_{oi}^{-1}} \circ \mathbf{y_n})^T \mathbf{x_n}$, so that $f(\mathbf{y}_n|\mathbf{x}_n)$ is a Bingham density on $\mathbf{x_n}$, $\text{Bingham}(\mathbf{x_n}; \Lambda_o, V_o^{-1} \circ \mathbf{y_n})$. Thus, computing $B_{\mathbf{x}_n}$ reduces to multiplying two Bingham PDFs on $\mathbf{x}_n$, which is given in section III.

Second-order QBF. Given a factorized distribution over the initial state $f(\mathbf{x}_0, \mathbf{v}_0) = B_{\mathbf{x}_0} B_{\mathbf{v}_0}$ and an action-observation sequence $\mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n$, the goal is to compute the joint posterior distribution $f(\mathbf{x}_n, \mathbf{v}_n \mid \mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n)$. However, the joint distribution on $\mathbf{x}_n$ and $\mathbf{v}_n$ is too difficult to represent, so we instead compute the marginal posteriors over $\mathbf{x}_n$ and $\mathbf{v}_n$ separately, and approximate the joint posterior as the product of the marginals. The marginal posterior on $\mathbf{x}_n$ is

$$\begin{aligned} B_{\mathbf{x}_n} &= f(\mathbf{x}_n|\mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n) \\ &\propto f(\mathbf{y}_n|\mathbf{x}_n) \int_{\mathbf{x}_{n-1}} f(\mathbf{x}_n|\mathbf{x}_{n-1}) B_{\mathbf{x}_{n-1}}(\mathbf{x}_{n-1}) \\ &= f(\mathbf{y}_n|\mathbf{x}_n)(B_{\mathbf{v}_{n-1}} \circ B_{\mathbf{x}_{n-1}})(\mathbf{x}_n) \end{aligned}$$

since we assume $\mathbf{x}_{n-1}$ and $\mathbf{v}_{n-1}$ are independent given all the data up to time $n-1$.

Similarly, the marginal posterior on $\mathbf{v}_n$ is

$$\begin{aligned} B_{\mathbf{v}_n} &= f(\mathbf{v}_n|\mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_n, \mathbf{y}_n) \\ &\propto \int_{\mathbf{v}_{n-1}} f(\mathbf{v}_n \mid \mathbf{v}_{n-1}, \mathbf{u}_n) B_{\mathbf{v}_{n-1}}(\mathbf{v}_{n-1}) \\ &\quad \cdot \int_{\mathbf{x}_n} f(\mathbf{y}_n|\mathbf{x}_n) f(\mathbf{x}_n|\mathbf{v}_{n-1}, \mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_{n-1}, \mathbf{y}_{n-1}). \end{aligned}$$

Once again, $f(\mathbf{y}_n|\mathbf{x}_n)$ can be written as a Bingham density on $\mathbf{x}_n$, $\text{Bingham}(\mathbf{x}_n; \Lambda_o, V_o^{-1} \circ \mathbf{y}_n)$. Next, note that $\mathbf{x}_n = \mathbf{v}_{n-1} \circ \mathbf{x}_{n-1}$ so that $f(\mathbf{x}_n|\mathbf{v}_{n-1}, \mathbf{u}_1, \mathbf{y}_1, \ldots, \mathbf{u}_{n-1}, \mathbf{y}_{n-1}) = B_{\mathbf{x}_{n-1}}(\mathbf{v}_{n-1}^{-1} \circ \mathbf{x}_n)$, which can also be re-written as a Bingham on $\mathbf{x}_n$. Now letting $\mathbf{x}_{n-1} \sim \text{Bingham}(\Sigma, W)$, and since the product of two Bingham PDFs is Bingham, the integral over $\mathbf{x}_n$ becomes proportional to a Bingham normalization constant, ${}_1F_1(\frac{1}{2}; \frac{4}{2}; C(\mathbf{v}_{n-1}))$, where

$$C(\mathbf{v}_{n-1}) = \sum_{i=1}^{3} \Big(\sigma_i (\mathbf{v}_{n-1} \circ \mathbf{w}_i)(\mathbf{v}_{n-1} \circ \mathbf{w}_i)^T + \lambda_{oi} (\mathbf{v}_{oi}^{-1} \circ \mathbf{y}_n)(\mathbf{v}_{oi}^{-1} \circ \mathbf{y}_n)^T \Big) \ .$$

Comparing $C(\mathbf{v}_{n-1})$ with equation 4 in section III we find that ${}_1F_1(\frac{1}{2}; \frac{4}{2}; C(\mathbf{v}_{n-1})) \propto (f_{\mathbf{y}_n|\mathbf{x}_n} \circ B_{\mathbf{x}_{n-1}}^{-1})(\mathbf{v}_{n-1})$. Thus,

$$\begin{aligned} B_{\mathbf{v}_n} &\propto \int_{\mathbf{v}_{n-1}} f(\mathbf{v}_n \mid \mathbf{v}_{n-1}, \mathbf{u}_n) B_{\mathbf{v}_{n-1}}(\mathbf{v}_{n-1}) \\ &\qquad \cdot (f_{\mathbf{y}_n|\mathbf{x}_n} \circ B_{\mathbf{x}_{n-1}}^{-1})(\mathbf{v}_{n-1}) \\ &= (f_{\mathbf{w}_n} \circ \mathbf{u}_n \circ (B_{\mathbf{v}_{n-1}} \cdot (f_{\mathbf{y}_n|\mathbf{x}_n} \circ B_{\mathbf{x}_{n-1}}^{-1})))(\mathbf{v}_n) \end{aligned}$$

In other words, to update the belief on $\mathbf{v}_n$, we first convolve the inverse belief on $\mathbf{x}_{n-1}$ with the measurement distribution, then multiply by the belief on $\mathbf{v}_{n-1}$, rotate by the control input $\mathbf{u}_n$, and convolve by the noise distribution, $f_{\mathbf{w}_n}$. In the next section, each of these operations on Bingham distributions will be explained in detail.

A. Extensions

For some applications (such as tracking a ping pong ball through a bounce), the process and observation models of the QBFs described above may be somewhat restrictive. Several extensions are possible, as we outline here.

Sampling-based methods / Tracking the ball through a bounce. The QBF can handle arbitrary process and control models by sampling from the current state distribution, applying the process/control function to each sample, and then fitting new Bingham distributions to the resulting post-process/control samples. This is precisely the method we use

in our experiments to track the ping pong ball through a bounce on the table. The process model is taken from Andersson's ball physics derivation [1], and is a complex function of the ball's tracked angular and translational velocities,

$$\mathbf{w_f} = \mathbf{w_i} + \frac{3\mu}{2r}(\hat{v}_{ry}, -\hat{v}_{rx}, 0)v_{iz}(1+\epsilon),$$

where $\mathbf{w_i}$ and $\mathbf{w_f}$ are the initial and final (pre- and post-bounce) angular velocity vectors, μ and ϵ are the coefficients of friction and restitution, r is the ball's radius, v_{iz} is the ball's initial translational z-velocity, and $\hat{\mathbf{v}}_\mathbf{r} = \mathbf{v_r}/\|\mathbf{v_r}\|$, where $\mathbf{v_r} = (v_{iy}+w_{ix}r, v_{ix}-w_{iy}r, 0)$ is the relative velocity of the surface of the ball with respect to the table.

Quaternion exponentiation / Continuous-time filters. Quaternion exponentiation for unit quaternions is akin to scaling in Euclidean space. If $\mathbf{q}$ represents a 3-D rotation of angle θ about the axis v, then $\mathbf{q}^a$ is a rotation of $a\theta$ about v. This operation would be needed to handle a continuous-time update in the second-order Bingham filter, since the orientation needs to be rotated by some fraction of the spin quaternion at each (time-varying) time step.

It is possible to incorporate quaternion exponentiation in all parts of the model via a moment-matching method for Bingham exponentiation (akin to the moment-matching method for Bingam composition). However, an additional Taylor-approximation is needed to approximate the second moments of the exponentiated Bingham as a function of the second and higher even moments of the original distribution. (The odd moments of a Bingham are always zero due to symmetry.)

VI. Experimental Results

We compare the quaternion Bingham filter against an extended Kalman filter (EKF) approach in quaternion space [13], where process and observation noise are generated by Gaussians in $\mathbb{R}^4$, the measurement function normalizes the quaternion state (to project it onto the unit hypersphere), and the state estimate is renormalized after every update. We chose the EKF both due to its popularity and because LaViola reports in [13] that it has similar (slightly better) accuracy to the unscented Kalman filter (UKF) in several real tracking experiments. We adapted two versions of the EKF (for first-order and second-order systems) from LaViola's EKF implementation by changing from a continuous to a discrete time prediction update. We also mapped QBF (Bingham) noise parameters to EKF (Gaussian) noise parameters by empirically matching second moments from the Bingham to the projected Gaussian—i.e., the Gaussian after it has been projected onto the unit hypersphere.

Synthetic Data. To test the first-order quaternion Bingham filter, we generated several synthetic signals by simulating a quaternion Bingham process, where the (velocity) controls were generated so that the nominal process state (before noise) would follow a sine wave pattern on each angle in Euler angle space. We chose this control pattern in order to cover a large area of 3-D rotation space with varying rotational velocities. Two examples of synthetic signals along with quaternion Bingham filter output are shown in figure 4. Their observation parameters were $\Lambda_o = (-50, -50, -50)$, which gives moderate, isotropic observation noise, and $\Lambda_o = (-10, -10, -1)$, which yields moderately high noise in the first two directions, and near-uniform noise in the third direction. We estimated the composition approximation error (KL-divergence) for 9 of these signals, with both isotropic and nonisotropic noise models, from all combinations of (Λ_p, Λ_o) in $\{(-50,-50,-50), (-200,-200,-200), (-10,-10,-1)\}$. The mean composition error was .0012, while the max was .0197, which occurred when Λ_p and Λ_o were both $(-10,-10,-1)$.

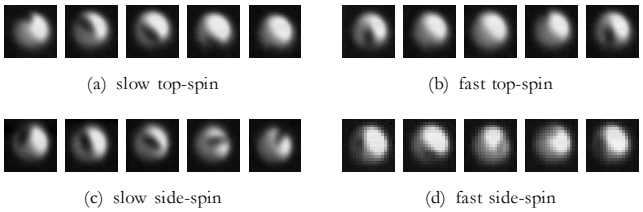

(a) slow top-spin (b) fast top-spin

(c) slow side-spin (d) fast side-spin

Fig. 5. Example image sequences from the spinning ping-pong ball dataset. In addition to lighting variations and low image resolution, high spin rates make this dataset extremely challenging for orientation tracking algorithms. Also, because the cameras were facing top-down towards the table, tracking side-spin relies on correctly estimating the orientation of the elliptical marking in the image, and is therefore much harder than tracking top-spin or under-spin.

For the EKF comparison, we wanted to give the EKF the best chance to succeed, so we generated the data from a projected Gaussian process, with process and observation noise generated according to a projected Gaussian (in order to match the EKF dynamics model) rather than from Bingham distributions. We ran the first-order QBF and EKF on 270 synthetic projected Gaussian process signals (each with 1000 time steps) with different amounts of process and observation noise, and found the QBF to be more accurate than the EKF on 268/270 trials. The mean angular change in 3-D orientation between time steps were 7, 9, and 18 degrees for process noise parameters -400, -200, and -50, respectively (where -400 means $\Lambda_p = (-400, -400 - 400)$, etc.).

The most extreme cases involved anisotropic observation noise, with an average improvement over the EKF mean error rate of 40-50%. The combination of high process noise and low observation noise also causes trouble for the EKF. Table I summarizes the results.

Spinning ping-pong ball dataset To test the second-order QBF, we collected a dataset of high-speed videos of 73 spinning ping-pong balls in flight (Figure 5). On each ball we drew a solid black ellipse over the ball's logo to allow the high-speed (200fps) vision system to estimate the ball's orientation by finding the position and orientation of the logo[2]. However, an ellipse was only drawn on one side of each ball, so the ball's orientation could only be estimated when the logo was visible in the image. Also, since ellipses are symmetric, each logo detection has two possible

[2]Detecting the actual logo on the ball, without darkening it with a marker, would require improvements to our camera setup.

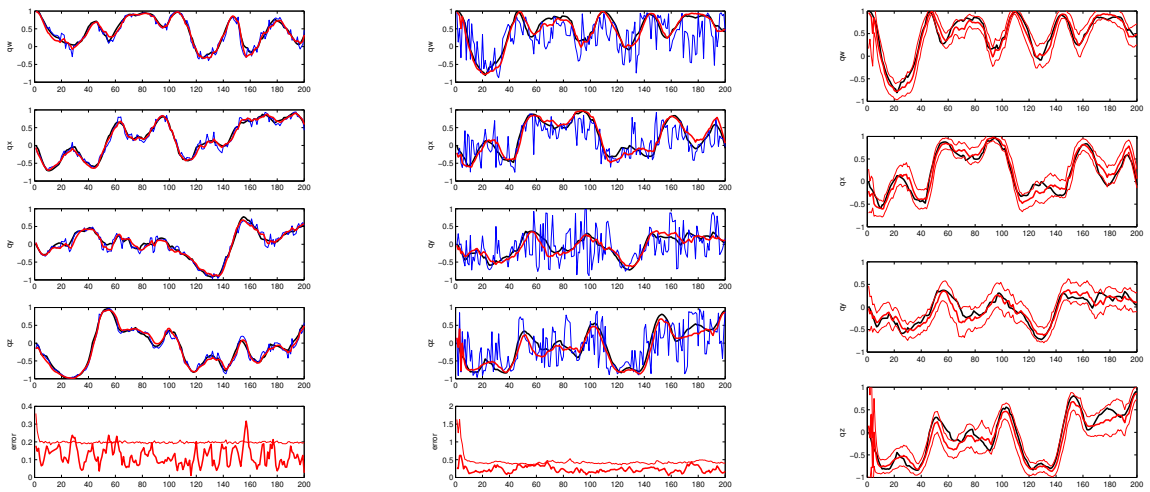

(a) $\Lambda_o = (-50, -50, -50)$ (b) $\Lambda_o = (-10, -10, -1)$ (c) $\Lambda_o = (-10, -10, -1)$

Fig. 4. Two simulated runs with the quaternion Bingham filter—(b) and (c) are different plots of the same simulation. In all figures, the thick black line is the true process signal, generated with isotropic process noise $\Lambda_p = (-400, -400, -400)$. The thin blue lines in (a) and (b) are the observation signal, and the thick red line is the filtered output. Rows 1-4 in each figure show the 4 quaternion dimensions (q_w, q_x, q_y, q_z). Row 5 in (a) and (b) shows the error between filter output and true state (thick red line), together with the QBF sample error 90% confidence bounds (thin red line). Marginal sample 90% confidence bounds are also shown in the thin red lines in (c).

observation noise	-400, -400, -10			-400			-50		
process noise	-50	-200	-400	-50	-200	-400	-50	-200	-400
% improvement	**37.3**	**45.6**	**54.3**	**19.9**	**3.33**	**1.52**	**3.42**	**0.72**	**0.47**
$\pm$	(5.1)	(5.0)	(6.4)	(1.8)	(0.53)	(0.44)	(0.88)	(0.40)	(0.27)

TABLE I

PROJECTED GAUSSIAN PROCESS SIMULATIONS. AVERAGE % MEAN ERROR DECREASE FOR QBF OVER EKF.

orientation interpretations[3]. The balls were spinning at 25-50 revolutions per second (which equates to a 45-90 degree orientation change per frame), making the filtering problem extremely challenging due to aliasing effects. We used a ball gun to shoot the balls with consistent spin and speed, at 4 different spin settings (topspin, underspin, left-sidespin, and right-sidespin) and 3 different speed settings (slow, medium, fast), for a total of 12 different spin types. Although we initially collected videos of 107 ball trajectories, the logo could only be reliably found in 73 of them; the remaining 34 videos were discarded. Although not our current focus, adding more cameras, adding markings to the ball, and improving logo detections would allow the ball's orientation and spin to be tracked on a larger percentage of such videos. To establish an estimate of ground truth, we then manually labeled each ball image with the position and orientation of the logo (when visible), from which we recovered the ball orientation (up to symmetry). We then used least-squares non-linear regression to smooth out our (noisy) manual labels by finding the constant rotation, $\hat{\mathbf{s}}$, which best fit the labeled orientations for each trajectory[4].

[3]We disambiguated between the two possible ball orientation observations by picking the observation with highest likelihood under the current QBF belief.

[4]Due to air resistance and random perturbations, the spin was not really constant throughout each trajectory. But for the short duration of our experiments (40 frames), the constant spin approximation was sufficient.

To run the second-order QBF on this data, we initialized the QBF with a uniform orientation distribution $B_{\mathbf{x}_0}$ and a low concentration ($\Lambda = (-3, -3 - 3)$) spin distribution $B_{\mathbf{v}_0}$ centered on the identity rotation, $(1, 0, 0, 0)$. In other words, we provided no information about the ball's initial orientation, and an extremely weak bias towards slower spins. We also treated the "no-logo-found" (a.k.a. "dark side") observations as a very noisy observation of the logo in the center of the back side of the ball at an arbitrary orientation, with $\Lambda_o = (-3.6, -3.6, 0)$[5]. When the logo was detected, we used $\Lambda_o = (-40, -40, -10)$ for the observation noise. A process noise with $\Lambda_p = (-400, -400, -400)$ was used throughout, to account for small perturbations to spin.

Results of running the second-order QBF (QBF-2) are shown in figure 6. We compared the second-order QBF to the second-order EKF (EKF-2) and also to the first-order QBF and EKF (QBF-1 and EKF-1), which were given the difference between subsequent orientation observations as their observations of spin. The solid, thin, blue line in each plot marked "oracle prior" shows results from running QBF-2 with a best-case-scenario prior, centered on the average ground truth spin for that spin type, with $\Lambda = (-10, -10, -10)$. We show mean orientation and spin errors (to regressed ground truth), and also spin classification accuracy using the MAP estimate of spin type (out of 12)

[5]We got this Λ_o by fitting a Bingham to all possible dark side orientations.

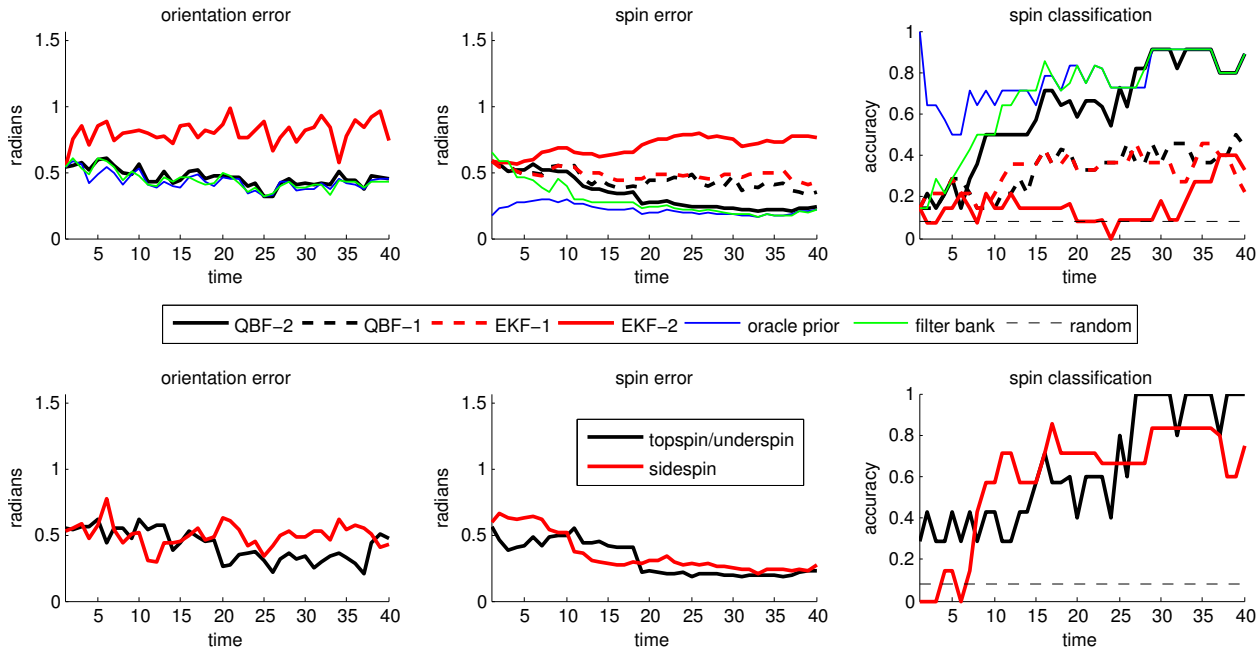

Fig. 6. Spinning ping-pong ball tracking results. Top row: comparison of QBF-2 (with and without an oracle-given prior) to QBF-1, EKF-1, EKF-2, and random guessing (for spin classification); QBF-1 and EKF-1 do not show up in the orientation error graph because they only tracked spin. Note that QBF-2 quickly converges to the oracle error and classification rates. Bottom row: QBF-2 results broken down into top-spin/under-spin vs. side-spin. As mentioned earlier, the side-spin data is harder to track due to the chosen camera placement and ball markings for this experiment.

given the current spin belief[6]. The results clearly show that QBF-2 does the best job of identifying and tracking the ball rotations on this extremely challenging dataset, achieving a classification rate of 91% after just 30 video frames, and a mean spin (quaternion) error of 0.17 radians (10 degrees), with an average of 6.1 degrees of logo axis error and 6.8 degrees of logo angle error. In contrast, the EKF-2 does not significantly outperform random guessing, due to the extremely large observation noise and spin rates in this dataset. In the middle of the pack are QBF-1 and EKF-1, which converge much more slowly since they use the raw observations (rather than the smoothed orientation signal used by QBF-2) to estimate ball spin. Finally, to address the aliasing problem, we ran a set of 12 QBFs in parallel, each with a different spin prior mode (one for each spin type), with $\Lambda = (-10, -10, -10)$. At each time step, the filter was selected with the highest total data likelihood. Results of this "filter bank" approach are shown in the solid, thin, green line in figure 6.

Tracking through the bounce. We also used the sampling-based method outlined in section V-A to track the ball through the bounce for 5 topspin/right-sidespin and 5 underspin/left-sidespin ball trajectories, and found that incorporating the bounce model as a sample-based process update in the quaternion Bingham filter (rather than restarting the filter after the bounce) resulted in a significant reduction in tracking error post-bounce (Figure 7).

In Figure 8 we show an example of the output of the second-order QBF we used to track the orientation and

[6]Spin was classified into one of the 12 spin types by taking the average ground truth spin for each spin type and choosing the one with the highest likelihood with respect to the current spin belief.

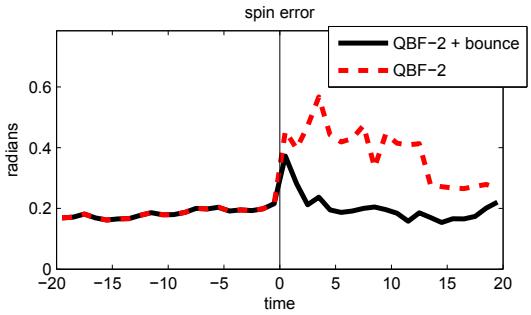

Fig. 7. Average QBF spin tracking error as it tracks the spin through the ball's bounce on the table. Post-bounce errors are significantly lower with the sample-based bounce tracking method (solid black line) outlined section V-A.

spin on the ball through one of the underpin/left-sidespin trajectories. In the first image frame, no logo is detected, so the orientation distribution is initialized to the "dark side" of the ball, and the spin distribution is close to a uniform distribution. After a few more frames, the filter has an accurate estimate of both the orientation and spin of the ball, with fairly high concentration parameters (low-uncertainty) in its Bingham distributions. After the bounce, the sample-based process update correctly updates the orientation and spin on the ball, and the filter maintains correct, high-concentration distributions.

VII. Conclusion

For many control and vision applications, the state of a dynamic process involving 3-D orientations and spins must be estimated over time, given noisy observations. Previ-

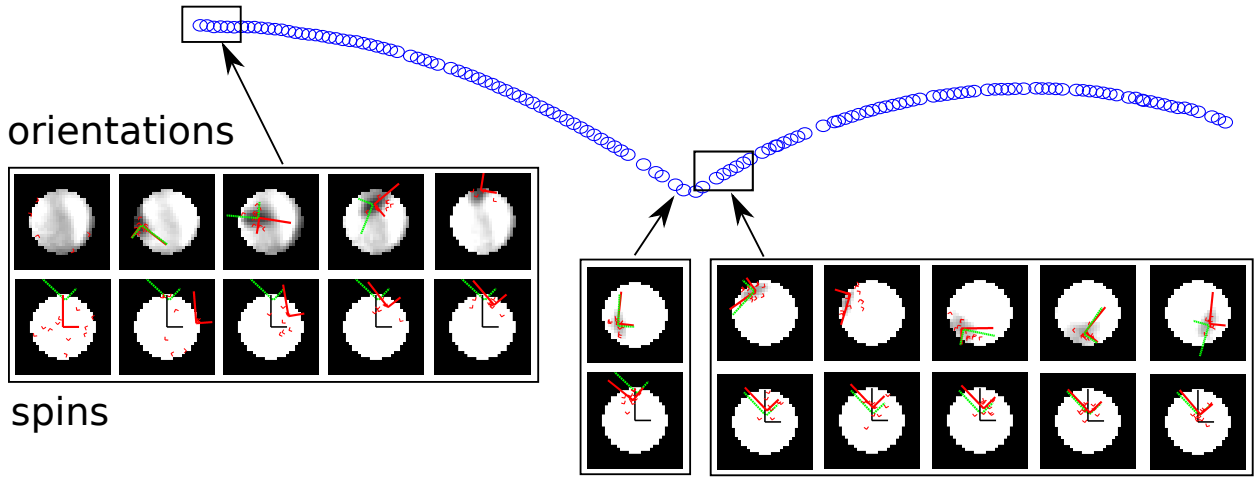

Fig. 8. An example ball trajectory (underspin + left-sidespin) and the state of the QBF as it tracks the ball's orientation and spin through the bounce. In the top row of ball images, the big red (solid-line) axis is the mode of the QBF's orientation distribution, and the small red axes are random samples from the orientation distribution. The big green (dashed-line) axis is the detected ball orientation in that image. In the bottom row, the big red (solid-line) axis is the mode of the QBF's spin distribution, and the small red axes are random samples from the spin distribution. The big green (dashed-line) axis is the ground truth spin, and the black axis in the center is the identity (no-spin), for reference.

ously, such estimation was limited to slow-moving signals with low-noise observations, where linear approximations to 3-D rotation space were adequate. The contribution of our approach is that the quaternion Bingham filter encodes uncertainty directly on the unit quaternion hypersphere, using a distribution—the Bingham—with nice mathematical properties enabling efficient approximate inference, with no restrictions on the magnitude of process dynamics or observation noise. Because of the compact nature of 3-D rotation space and the flexibility of the Bingham distribution, we can use the QBF not only for tracking but also for identification of signals, by starting the QBF with an extremely unbiased prior, a feat which previously could only be matched by more computationally-intensive algorithms, such as discrete Bayesian filters or particle filters.

References

[1] Russell L. Andersson. *A robot ping-pong player: experiment in real-time intelligent control*. MIT Press, Cambridge, MA, USA, 1988.

[2] Christopher Bingham. An antipodally symmetric distribution on the sphere. *The Annals of Statistics*, 2(6):1201–1225, November 1974.

[3] Xavier Boyen and Daphne Koller. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, UAI'98, page 3342, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

[4] Yee-Jin Cheon and Jong-Hwan Kim. Unscented filtering in a unit quaternion space for spacecraft attitude estimation. In *Industrial Electronics, 2007. ISIE 2007. IEEE International Symposium on*, pages 66–71, 2007.

[5] D. Choukroun, I.Y. Bar-Itzhack, and Y. Oshman. Novel quaternion Kalman filter. *Aerospace and Electronic Systems, IEEE Transactions on*, 42(1):174–190, 2006.

[6] Wendelin Feiten, Pradeep Atwal, Robert Eidenberger, and Thilo Grundmann. 6D pose uncertainty in robotic perception. In *Advances in Robotics Research*, pages 89–98. Springer Berlin Heidelberg, 2009.

[7] Jared Glover and Leslie Pack Kaelbling. Tracking 3-d rotations with the quaternion bingham filter. *Technical Report MIT-CSAIL-TR-2013-005*, 2013.

[8] Jared Glover and Sanja Popovic. Bingham procrustean alignment for object detection in clutter. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013.

[9] Jared Glover, Radu Rusu, and Gary Bradski. Monte carlo pose estimation with quaternion kernels and the bingham distribution. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2011.

[10] A. Kim and M.F. Golnaraghi. A quaternion-based orientation estimation algorithm using an inertial measurement unit. In *Position Location and Navigation Symposium, 2004. PLANS 2004*, pages 268–272, 2004.

[11] E. Kraft. A quaternion-based unscented Kalman filter for orientation tracking. In *Information Fusion, 2003. Proceedings of the Sixth International Conference of*, volume 1, pages 47–54, 2003.

[12] Gerhard Kurz, Igor Gilitschenski, Simon Julier, and Uwe D Hanebeck. Recursive estimation of orientation based on the bingham distribution. In *Information Fusion (FUSION), 2013 16th International Conference on*, pages 1487–1494. IEEE, 2013.

[13] J.J. LaViola. A comparison of unscented and extended Kalman filtering for estimating quaternion motion. In *American Control Conference, 2003. Proceedings of the 2003*, volume 3, pages 2435–2440 vol.3, 2003.

[14] K. V. Mardia. Characterizations of directional distributions. In *Statistical Distributions in Scientific Work*, volume 3, pages 365–385. D. Reidel Publishing Company, Dordrecht, 1975.

[15] J.L. Marins, Xiaoping Yun, E.R. Bachmann, R.B. McGhee, and M.J. Zyda. An extended Kalman filter for quaternion-based orientation estimation using MARG sensors. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 4, pages 2003–2011 vol.4, 2001.

[16] Fumio Miyazaki, Michiya Matsushima, and Masahiro Takeuchi. Learning to dynamically manipulate: A table tennis robot controls a ball and rallies with a human being. In *Advances in Robot Control*, pages 317–341. 2006.

[17] K. Muelling, J. Kober, O. Kroemer, and J. Peters. Learning to select and generalize striking movements in robot table tennis. (3):263–279, 2013.

[18] Stephen R Niezgoda and Jared Glover. Unsupervised learning for efficient texture estimation from limited discrete orientation data. *Metallurgical and Materials Transactions A*, pages 1–15, 2013.

[19] S. Schaefer, J. Hakenberg, and J. Warren. Smooth subdivision of tetrahedral meshes. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 147–154, Nice, France, 2004. ACM.

[20] Yichao Sun, Rong Xiong, Qiuguo Zhu, Jun Wu, and Jian Chu. Balance motion generation for a humanoid robot playing table tennis. In *2011 11th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 19–25. IEEE, October 2011.