

Exploring Applications of Bingham Distribution for Characterizing Uncertainty over $SO(3)$

16-833 Final Report

Ratnesh Madaan [ratneshm], Sudharshan Suresh [sudhars1], Ankita Kalra [akalra1]

I. INTRODUCTION

In this project, we explore the potential of the Bingham distribution in localization and mapping approaches. Namely, this includes (i) the formulation of uncertainty over $SO(3)$ via the Quaternion Bingham Filter (QBF) (ii) an investigation into surface normal segmentation with a Bingham Mixture Model, akin to previous work in global scene segmentation [1].

Rotational data is frequently encountered in robotics, however most existing methods make a Gaussian assumption on uncertainty. The goal of this project is to explore a suitable alternate assumption on noise - the Bingham distribution. To this effect, we introduce the quaternion Bingham process model, which can be further studied in [2]. Given 3-D rotational data, we investigate a filtering method where the dynamics/measurements are affected by random Bingham distributed noise. This represents an ideal combination of elements - while the quaternions provide a minimal representation of orientations sans degeneracy, the QBF allows for very precise tracking in the presence of large disturbances. Such a method has been shown to have lower tracking error than the Extended Kalman Filter (EKF) for dynamic states.

Section IV discusses the theory and implementation of such a filter, and evaluates it on a synthetic dataset we collect that possesses the rotational characteristics we so desire.

Secondly, we wish to investigate the segmentation of surface normals using a Bingham Mixture Model to capture global scene segmentation, inspired by [1]. While [1] introduced a Dirichlet Process over the von-Mises-Fischer distribution for global segmentation, intuitively a Bingham Mixture Model is a much simpler way to accomplish the same. Modeling surface normals as quaternions, and segmenting based on directionality seems like an ideal use case for the Bingham distribution.

Section V describes surface normal segmentation performed in [1], and attempts to perform Bingham clustering of normals obtained in Kaess' [3] planar SLAM pipeline. Potentially, this could lead to the representation of a Bingham distribution over S^3 to model the uncertainty over the quaternion representation of planes introduced in [3]. In the iSAM based factor graph formulation of planar SLAM, this boils down to a change in the measurement model of planes from a Gaussian in the tangent space of S^3 at the linearization point to a Bingham over S^3 itself.

While our results are not significant, they do enforce the validity and scope of the Bingham distribution in robotics. As section II discusses, this niche body of work has grown in significance in recent years. We constantly revisited and re-framed our goals in this project due to implementational difficulties, however our learnings and takeaways have been rather significant.

In general [4], [5] are great resources to understand the Bingham distribution and mixture models over them.

II. RELATED WORK

A. Filter-based approaches

The Bingham distribution stems from the work of Bingham [6], where he parameterizes the distribution with a concentration matrix Z and orientation matrix M . We maintain the same naming conventions in this report. He also elucidates on a numerical example of the same from the domain of geology.

Glover et al. [2] uses two version (first-order and second order) of a QBF to track a ping-pong ball. He further establishes the lower tracking error of the QBF has compared to the EKF. Kurz et al. [7] presents a recursive implementation of the Bingham filter in two dimensions, and evaluates it on simulated data.

The implementational tools in MATLAB/C++ stem from the work of the above two authors -

- 1) The Bingham Statistics Library (from Glover) provides a framework for implementation of Bingham distributions on unit spheres \mathbb{S}^1 , \mathbb{S}^2 , and \mathbb{S}^3 .
- 2) libDirectional, from Kurz et al. [8], is a library for directional statistics as well as recursive estimation on directional manifolds.

Kaess et al. [3] proposed a planar SLAM formulation which models planes as quaternions. It also establishes a minimal representation for plane parameters. We briefly discuss the application of Bingham statistics in such a planar SLAM formulation in section VI.

B. Scene segmentation

Kaess's approach to plane segmentation borrows from Holz et al. [9]. Surface normals are extracted by taking the cross-product of vectors tangential to the local surface. Points are clustered in normal space to get a set of planes, which are merged if they show similar local surface normals. An alternate method of plane segmentation is presented by Straub et al. [10]. While they propose a von-Mises-Fisher distribution, they actually approximate it with a Bingham distribution in their calculation of posterior of global segmentation, as shown in Fig of [11]. This is a more mathematically sound way to cluster the data, readily generalizable to higher dimensions. This segmentation is then incorporated into semi-dense SLAM.

III. THEORY

A. The Bingham Distribution

The Bingham Distribution is an antipodally-symmetric probability distribution, by taking into account the facts that the mapping from quaternion to the rotation space is two to one and that antipodal quaternions $(q, -q)$ represent the same 3-D rotation/orientation on a unit hypersphere.

It is derived from a zero-mean Gaussian on R^{d+1} , conditioned to lie on the surface of the unit hypersphere S^d . Thus, the Bingham's probability density function (PDF) is proportional to the PDF of a zero-mean Gaussian, but it differs from the Gaussian in its normalization constant. It is intuitive to use it to represent uncertainty over the rotation space along with unit quaternions. One could also consider a hemisphere of S^3 and use another distribution like Von Mises-Fisher, for instance, to model the rotation space; however it is convenient to consider the whole hypersphere coupled with a Bingham and not bother with the discontinuities at the equator. Here, it is worth mentioning that the VMF distribution is also a maximum entropy distribution over a hypersphere, and it could be a viable avenue for modeling rotations.

Mathematically, the Bingham distribution is represented as:

$$f(\mathbf{x}; \Lambda, V) = \frac{1}{F(\Lambda)} \exp\left\{\sum_{i=1}^d \lambda_i (\mathbf{v}_i^T \mathbf{x})^2\right\} \quad (1)$$

$$= \frac{1}{F(\Lambda)} \exp\{\mathbf{x}^T V \Lambda V^T \mathbf{x}\} \quad (2)$$

where \mathbf{x} is a unit vector on the hypersphere S^d , Λ is a $d \times d$ diagonal matrix of concentration parameters $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d (\leq 0)$, V is the matrix of the eigenvectors (v_1, v_2, \dots, v_d) of the distribution, and F is the normalization constant. The $(d+1)^{th}$ eigenvalue λ_{d+1} (and its corresponding eigenvector v_{d+1}) are omitted by adding $-\lambda_{d+1}$ to all eigenvalues, without affecting the distribution.

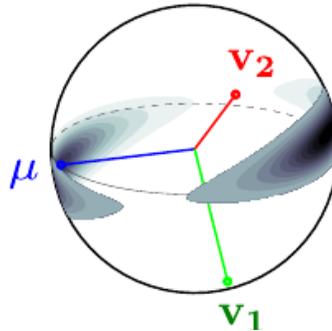


Fig. 1: Bingham distribution on S^2 , as taken from [4]. We see the mode and two orthogonal vectors in this graphic, along with the concentration densities on the hypersphere.

In figure 1, we see an example of a Bingham distribution in 3-D on the sphere S^2 . In this example, $\lambda_1 < \lambda_2 < 0$, and so the distribution is more spread out in the direction of \mathbf{v}_2 than it is along \mathbf{v}_1 . Unlike the Gaussian distribution (which has one mode), we have already seen that the Bingham distribution has at least two modes, M and $-pt$. In fact, the Bingham can also have infinitely many modes!

If one or more of the λ_i 's in equation 1 is zero, then any $x \in S^d$ in the subspace spanned by μ and the \mathbf{v}_i 's with $\lambda_i = 0$ will also be a mode of the distribution. As an extreme example, if all of the λ_i 's are zero, then every \mathbf{x} is equally likely, which means that the distribution is uniform. If only some of the λ_i 's are zero, then the distribution will have rings of equal probability around the hypersphere, as shown in figure 2. For the ones who are familiar with the multivariate Gaussian distribution will recognize the exponent of equation 2 as the exponent of Gaussian in information form, where the information (inverse covariance) matrix has

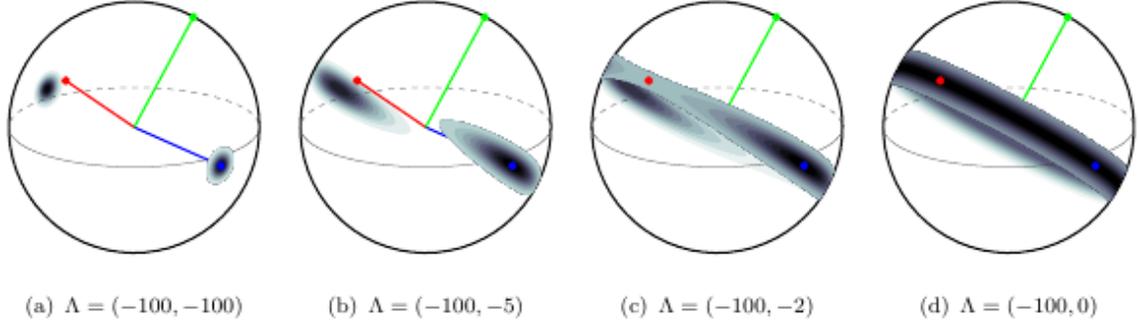


Fig. 2: How the Bingham distribution appears when you vary one of the concentration parameters to 0. [4]

been decomposed into its eigenvector and eigenvalue matrices, V and Λ , and where the Bingham exponent is missing a factor of $-\frac{1}{2}$, which is moved into the Bingham's normalization constant. Also missing are the $(d+1)$ 'st eigenvector and eigenvalue, \mathbf{v}_{d+1} and λ_{d+1} . This is because one can add a constant, ϵ , to all of the original $(d+1)$ λ 's and obtain the same PDF (up to a proportionality constant, which will be compensated for in the normalization term, F):

$$\begin{aligned} e^{\mathbf{x}^T V(\Lambda + \epsilon I) V^T \mathbf{x}} &= e^{\mathbf{x}^T V \Lambda V^T \mathbf{x} + \mathbf{x}^T V \epsilon I V^T \mathbf{x}} = e^{\mathbf{x}^T V \Lambda V^T \mathbf{x} + \epsilon \mathbf{x}^T \mathbf{x}} \\ &= e^{\mathbf{x}^T V \Lambda V^T \mathbf{x}} \cdot e^\epsilon \end{aligned}$$

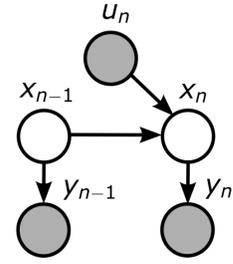
B. Quaternions

There are various ways to represent $SO(3)$ including Euler angles, rotation matrices, axis-angle representation and unit quaternions, and all of them have their pros and cons. However, we choose quaternions as they are: (a) not as ambiguous as Euler Angles apart from antipodal symmetry, (b) do not suffer from the Gimbal lock problem, (c) easier to compose as compared to the axis-angle representation, and (d) allow for faster computation. For our case, an even stronger motivation for using quaternions is that the Bingham distribution captures the antipodal topology of the quaternion space perfectly. To top that, the Bingham is a member of the exponential family and is the maximum entropy distribution on the hypersphere [4]. There has been a recent surge of interest in and the revival of Bingham due to the Glover's work [4] and they provide a promising model for rotational uncertainty.

IV. THE QUATERNION BINGHAM FILTER

A. Description

An in-depth treatment of the quaternion Bingham filter (QBF) can be found in [2]. Here we will expand the theory around only the discrete-time first-order quaternion Bingham filter.



$$\begin{aligned} \mathbf{X}_n &= \mathbf{W}_n \circ \mathbf{u}_n \circ \mathbf{X}_{n-1} \\ \mathbf{y}_n &= \mathbf{Z}_n \circ \mathbf{X}_n \end{aligned}$$

Fig. 3: The first-order discrete quaternion Bingham process depicted in graphical model form, from [2]

In the case of a *first order discrete time quaternion Bingham process*, the state at time n is a unit quaternion \mathbf{x}_n . Further, the control input at that timestep, \mathbf{u}_n , is a unit quaternion as well. One can obtain the next state \mathbf{x}_{n+1} as the previous state, rotated by the control input. The observations at each timestep \mathbf{y}_n are the states corrupted with system disturbance. In our case the noise comes from an independent Bingham distribution.

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^2 . The control inputs \mathbf{u}_n are analogous to rotational accelerations. Process noise \mathbf{w}_n enters the system in the velocity dynamics.

For notational clarity, $\mathbf{x}_n = \mathbf{w}_n \circ \mathbf{u}_n \circ \mathbf{x}_{n-1}$ represents composition of rotations. The state equations, with the graphical model for the first order process can be seen in Figure 3. [2] also elucidates on the second-order filter, which we will not include in our implementation.

The QBF is recursive in nature, similar to the Bayes filter that we have studied in class. Its state distribution is projected into the Bingham family after each time step. The derivation for the first-order filter bears resemblance to that of the Kalman filter. We have an initial state distribution on x_0 , $B_{x0} = \text{Bingham}(Z_0, M_0)$. We wish

to compute the posterior $f(x_n|u_1, y_1 \dots u_n, y_n)$ given the observations $y_1 \dots y_n$ and control inputs $u_1 \dots u_n$. In the manner familiar to us, we proceed to apply the Markov assumption and Bayes' rule -

$$\begin{aligned} B_{x_n} &= f(x_n|u_1, y_1, \dots, u_n, y_n) \\ &\propto f(y_n|x_n)f(x_n|u_1, y_1, \dots, u_{n-1}, y_{n-1}, u_n) \\ &\propto f(y_n|x_n) \int_{x_{n-1}} f(x_n|x_{n-1}, u_n) B_{x_{n-1}}(x_{n-1}) \\ &\propto f(y_n|x_n)(f_{w_n} \circ u_n \circ B_{x_{n-1}})(x_n) \end{aligned}$$

To evaluate the first term $f(y_n|x_n)$, we use the equation $y_n = z_n \circ x_n$. This is nothing but a Bingham defined by -

$$y_n|x_n \approx \text{Bingham}(y_n; Z_o, M_o \circ x_n)$$

B. Experiments and Results

Synthetic data generation: We evaluate our filter on quaternions generated in simulation. We generate controls for a fixed frame to gradually precess, based on an input spin rate and precession rate. We selected such a pattern, as it covers a wide range of rotations - suited well for application of a Bingham Filter. The pattern of the ground-truth synthetic data can be seen in Figure 4. The video of the same can be viewed [here](#).

The initial state of the filter is set with the below parameters -

$$M = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \quad Z = \begin{bmatrix} -1 \\ -1 \\ -1 \\ 0 \end{bmatrix}$$

We then add Bingham distributed system noise B_{sys} , with parameters -

$$M_{\text{sys}} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \quad Z_{\text{sys}} = 10^4 \times \begin{bmatrix} -5 \\ -5 \\ -5 \\ 0 \end{bmatrix}$$

We perform a stochastic sampling from this distribution via Kent's method [12] to give us the noisy terms.

We model the known system disturbance into our QBF, and predict and plot the rotation at each time step. We have generated similar plots for our filter results, which shows accurate tracking. In Figure 5, we see the imperfect plot generated by the raw noisy values. The recovered plot resembles the ground truth. A video of the same can be seen [here](#). Our Bingham filter was created with the libDirectional library for MATLAB [8].

V. BINGHAM NORMAL SEGMENTATION

Inspired by the results of Straub et al.'s work on Direction-Aware SLAM [1] and in formulating and capturing a generative model of surface normals in indoor scenes [11], we wanted to investigate the effectiveness of using a mixture model over Bingham distributions to do similar tasks. Straub uses a Dirichlet Process over the von Mises Fischer distribution in S^2 for both tasks.

Antipodal surface normals are different and we wouldn't want to use a Bingham over S^2 to capture such data. However, we can augment surface normal data and move it to S^3 by making pure quaternions (real part is zero) out of them.

We tried two methods for segmenting normals with Bingham mixture models. A Bingham Mixture Model (BMM) is akin to a Gaussian Mixture Model, and is nothing but a weighted sum of Bingham, $p(x, BMM, \alpha) = \sum_{i=1}^k \alpha_i p(x, B_i)$ [5]. Intuitively, a BMM over S^2 is not ideal for fitting a mixture model over surface normals, as antipodal data points over S^2 are different normals.

We made pure quaternions from the surface normal data obtained from frames of Kaess' staircase dataset from [3] by simply using zeros as the real part of the quaternion, and using the vector part as the surface normal itself. Fitting a Bingham over S^3 however didn't give ideal structures as can be seen in Figure 6. We then tried fitting a BMM over S^2 over the raw surface normal data, the results of which can be seen in Figure 7 and [this video](#).

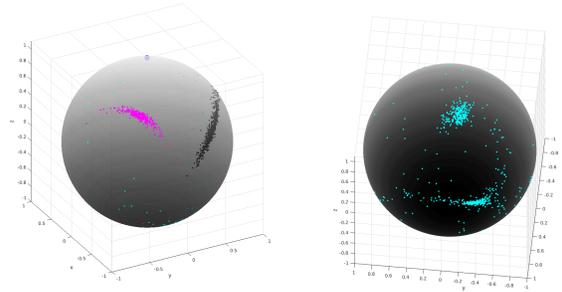


Fig. 6: (a) Actual distribution of normals on a unit sphere for frame 10 of the stairs dataset (b) Bingham clustering for the same frame.

[5] introduced a simple RANSAC based method to fit BMMs to data, where outliers are defined by a probability less than the reciprocal of the surface area

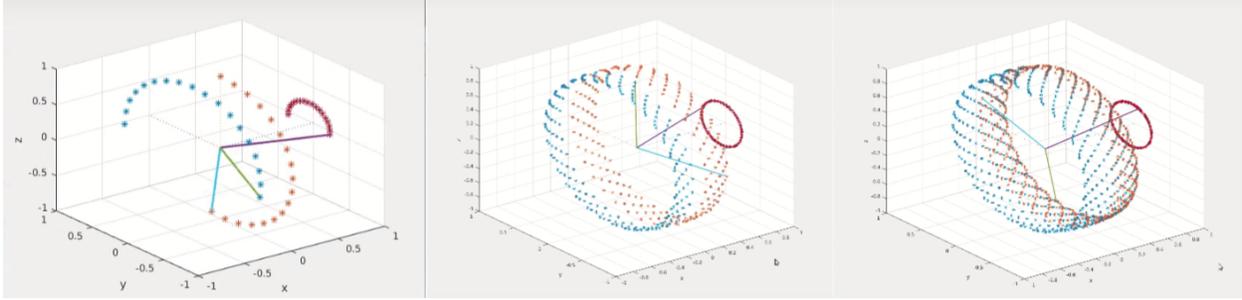


Fig. 4: Synthetic data, with axis points plotted over time. This is `quat_0.001_0.3_0.03.csv` in the data folder.

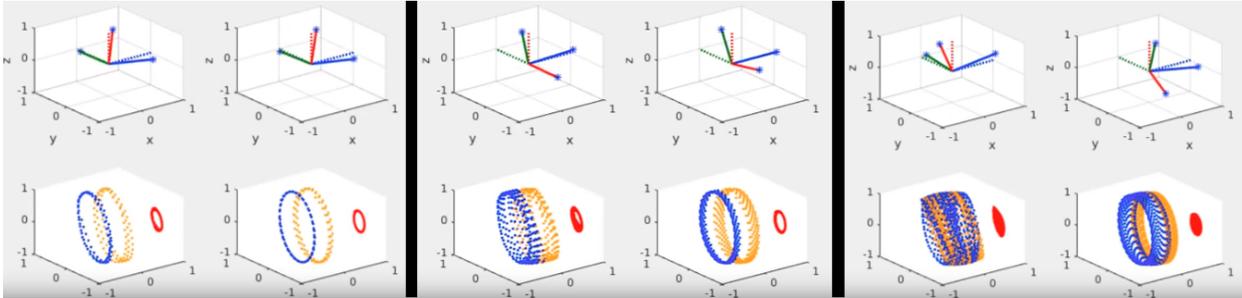


Fig. 5: Three images that show the results of the QBF. Each image consists of 4 subplots - TL: noisy quaternion visualized. TR: filtered quaternion. BL: axes points over time for noisy quaternion BR: axes points over time for QBF result

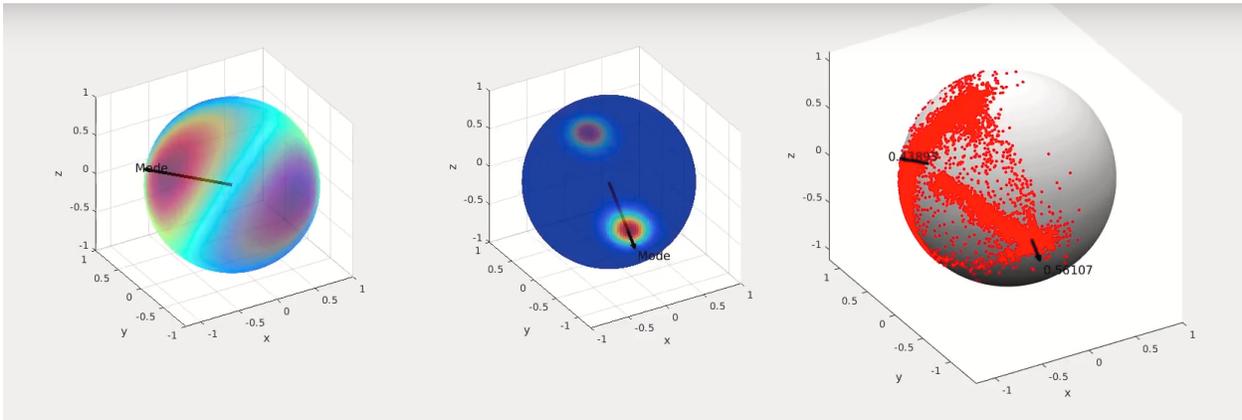


Fig. 7: Bingham clustering results for S^2 . The frame being reference, with the segmented planes in Kaess' pipeline are shown in figure 8

of the corresponding hypersphere. To fit a BMM, we transcribed Glover's C implementation of the BMM-SAC algorithm to MATLAB and use it in conjunction with `libDirectional`.

A. Reasons for Negative Results

1) *BMM over S^3* : To fit the S^3 BMM, we call the `bingham_cluster()` function from the file `bingham.c` in Glover's `libbingham` repo. The BMM-SAC implementation in the above has a couple of hyperparameters: the minimum number of points in

one component, and the minimum pdf of a quaternion sample to be assigned to a specific component. It could be that we did not tune these parameters enough to get good results. Another reason could be that the BMM-SAC is lacking in itself, and the community needs to come up with a smarter algorithm, perhaps something on the lines of Expectation Maximization for fitting Gaussian Mixture Models.

2) *BMM over S^2* : Using a BMM over S^2 is flawed by design, as there is no reason to have antipodal

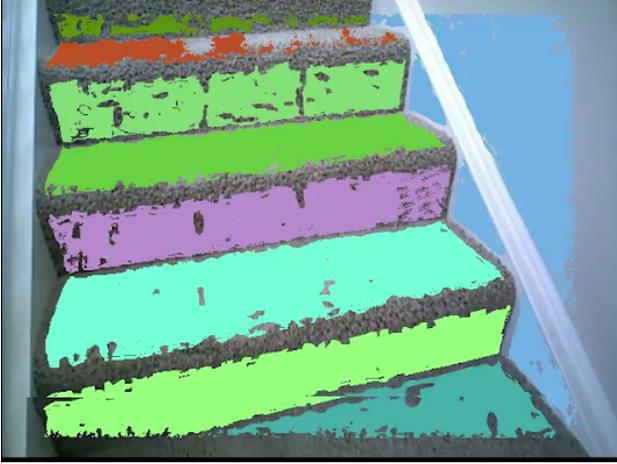


Fig. 8: Plane segmentation in Kaess’ pipeline from frame 500. From the image we can see two major dominant planes, along with a third (light blue) one. Its bingham clustering is shown in figure 7.

symmetry in S^2 - normals pointing opposite to each other most probably belong to opposite surfaces (unless the respective points at which the normals were evaluated were really close in R^3 which is not the case here). We however still proceeded with implementation of the S^2 BMM, driven by curiosity. As can be seen Fig 4 and the video linked in the previous paragraph, the modes of the component do somewhat agree with the distribution of the raw data - the two components of the Bingham can be seen as representing the perpendicular walls in the data. However, we are missing the 3rd mode of the dataset, which corresponds to the floor.

For both S^2 and S^3 , the data itself is noisy, which could be another reason for the failure of fitting the correct mixture model in the absence of better algorithms for the same.

VI. APPLICATION IN PLANAR SLAM

This final section discusses the application of the Bingham distribution in Kaess’ SLAM formulation [3]. Planes inherently have only 3 degrees of freedom - orientation can be modeled by two angles α and β , and orthogonal distance to origin by d . However Kaess [3] explains that this minimal representation has singularities like Euler angles, which would lead to problems in optimization. He proposes to model the plane with a homogeneous vector $\pi = (\pi_1, \pi_2, \pi_3, \pi_4) \in \mathbb{P}^3$. Then, a point $p = (p_1, p_2, p_3, p_4) \in \mathbb{P}^3$ -

$$\pi_1 p_1 + \pi_2 p_2 + \pi_3 p_3 + \pi_4 p_4 = 0$$

The above representation can be mapped to the standard plane equation in R^3 [3]. We define the normal

vector, $n = (\pi_1, \pi_2, \pi_3)^T / \sqrt{\pi_1^2 + \pi_2^2 + \pi_3^2}$ and distance from origin, $d = -\pi_4 / \sqrt{\pi_1^2 + \pi_2^2 + \pi_3^2}$. Then, the point $p^{xyz} = (p_1/p_4, p_2/p_4, p_3/p_4)^T$ lies on the plane if:

$$n^T p^{xyz} = d$$

[3] establishes a minimal representation using the above over-parameterized homogeneous representation by simply normalizing it such that π lies on the unit hypersphere in \mathbb{R}^4 : $\pi' = \pi / \|\pi\| \in S^3$. For optimization, a minimal representation of $SO(3)$, its lie-algebra is used, which is essentially the set of skew symmetric matrices over the 3-vector axis-angle representation of rotations.

In [3]’s formulation, the uncertainty of the plane z_{π_x} measured from pose x is modeled by a zero-mean Gaussian with a 3×3 covariance matrix Σ in the tangent space:

$$\pi_x = T_{gx}^{-T} \pi \oplus v, \quad v \sim \mathcal{N}(0, \Sigma)$$

Here T_{gx} is the transformation matrix between the global frame and frame corresponding to pose x (plane measurements are relative). This means that the probability of a plane measurement $\hat{\pi}$ given the actual observation z_{π_x} at a pose x is:

$$p(\hat{\pi} | z_{\pi_x}) = \frac{1}{\sqrt{(2\pi)^3 |\Sigma|}} \exp\left(-\frac{1}{2} \|h(T_{gx}, \hat{\pi}) \ominus z_{\pi_x}\|_{\Sigma}^2\right)$$

Following from [4], we define a base Bingham distribution $B_0(\Lambda_0, V_0)$, which has its mode at the identity quaternion:

$$V_0 = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \Lambda_0 = \text{diag}(\lambda_1, \lambda_2, \lambda_3)$$

Instead of having the mean at the sensor observation, as one would do with a Gaussian, we need to pre-rotate B_0 accordingly with the observed planar measurement quaternion, z_{π_x} . [4], [2] define the pre-rotation of a Bingham with a quaternion, which essentially is doing quaternion pre-multiplication of each of the 3 eigenquaternions of B_0 (columns of V_0), which we denote by $V_p = z_{\pi_x} \circ V_0$. Now if v_q is a sample from this rotated Bingham distribution, the predicted measurement, π_x is obtained by simply pre-rotating z_{π_x} :

$$\pi_x = v_q z_{\pi_x} \quad v_q \sim \text{Bingham}(\Lambda_0, V_p = z_{\pi_x} \circ V_0)$$

We can define the probability of a plane measurement $\hat{\pi}$ given the actual observation z_{π_x} at a pose x as:

$$p(\hat{\pi} | z_{\pi_x}) = \frac{1}{F(\Lambda_0)} \exp\left(\hat{\pi}^T (z_{\pi_x} \circ V_0) \Lambda_0 (z_{\pi_x} \circ V_0) \hat{\pi}\right)$$

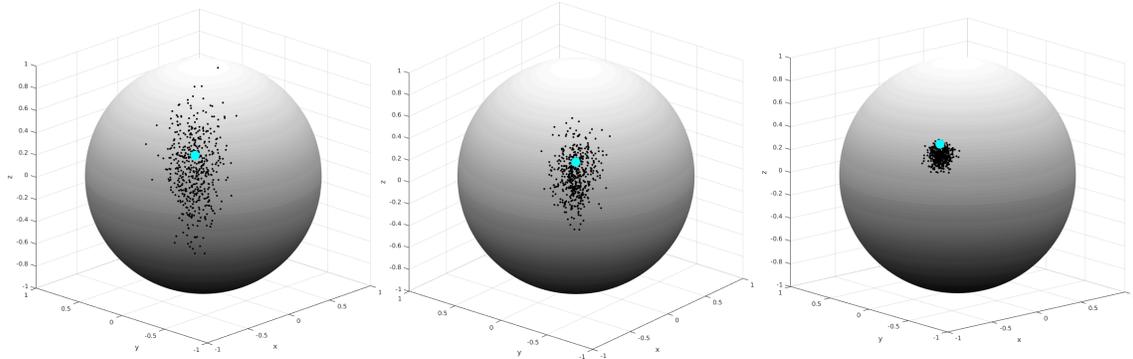


Fig. 9: Possible sensor models with concentration parameters, $\Lambda =$ (a) $(-30, -30, -600)$ (b) $(-60, -60, -900)$ (c) $(-800, -800, -900)$, and eigen-quaternion matrix = V_0 (defined in text). These plots are obtained by rotating the point in cyan by 1000 quaternion samples from respective bingham distribution about the origin (center of the sphere). As the cyan point is on the sphere, the resulting point will also lie on the sphere. According to [2]’s [library](#), the concentration params are defined in the range from -900 to 0. The less negative a concentration parameter, the more is the uncertainty about the corresponding eigen-quaternion.

We visualize possible sensor models in Fig 9. To visualize multiple quaternions, we pick an arbitrary point on a 3-D sphere (shown in cyan), and rotate it about the origin with 1000 quaternions sampled from the Bingham sensor model and plot the resulting points (which would still lie on the sphere) in black. This visualization is inspired by the EGI plots of [13]. The cyan point can be thought of as the unit normal vector of the planar observation (with the distance to origin stripped off).

VII. CONCLUSION, FUTURE WORK, LEARNINGS

The Bingham distribution has a great use case in the interesting problem of modeling uncertainty over the manifold of rotations in SLAM. Recent work in segmentation and filtering has buoyed interest in such representations in the field of robotics.

Future work include the testing of our QBF on real-world data. Some possible options come to mind, such as initially testing it out on a challenging sequence from the EuRoC dataset [14]. We believe that the filter can find a very good use case in state estimation for drones, and running it onboard would be the ideal end goal.

Upon correcting the issues in our normal clustering (section V), we wish to perform plane segmentation and compare our results with that of Kaess et al. [3]. This will allow us to confirm if a distribution on a hypersphere is the best way of characterizing planes.

There were significant learnings from the project -

- Making the QBF meshed right into the course, where the EKF was handled in detail

- It was very interesting trying to formulate uncertainty in a manner different from the standard Gaussian methods.
- In order to come up with an alternate formulation for the planar SLAM problem, we gained a great understanding in Kaess’ work on infinite planar SLAM [3].
- We also stumbled onto a lot of very impressive work from Julian Straub, one of our main motivations for going forward with this project.
- Our [reading list](#) lists papers from all the above authors.

REFERENCES

- [1] J. Straub, R. Cabezas, J. Leonard, and J. W. Fisher III, “Direction-aware semi-dense slam,” *arXiv preprint arXiv:1709.05774*, 2017.
- [2] J. Glover and L. P. Kaelbling, “Tracking 3-d rotations with the quaternion bingham filter,” 2013.
- [3] M. Kaess, “Simultaneous localization and mapping with infinite planes,” in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 4605–4611, IEEE, 2015.
- [4] J. M. Glover, *The quaternion Bingham distribution, 3D object detection, and dynamic manipulation*. PhD thesis, Massachusetts Institute of Technology, 2014.
- [5] J. Glover, “Monte carlo pose estimation with quaternion kernels and the bingham distribution,” 2012.
- [6] C. Bingham, “An antipodally symmetric distribution on the sphere,” *The Annals of Statistics*, pp. 1201–1225, 1974.
- [7] G. Kurz, I. Gilitschenski, S. Julier, and U. D. Hanebeck, “Recursive estimation of orientation based on the bingham distribution,” in *Information Fusion (FUSION), 2013 16th International Conference on*, pp. 1487–1494, IEEE, 2013.
- [8] G. Kurz, I. Gilitschenski, F. Pfaff, and L. Drude, “libdirectional,” 2015.
- [9] D. Holz, S. Holzer, R. B. Rusu, and S. Behnke, “Real-time plane segmentation using rgb-d cameras,” in *Robot Soccer World Cup*, pp. 306–317, Springer, 2011.
- [10] J. Straub, J. Chang, O. Freifeld, and J. Fisher III, “A dirichlet process mixture model for spherical data,” in *Artificial Intelligence and Statistics*, pp. 930–938, 2015.

- [11] J. Straub, G. Rosman, O. Freifeld, J. J. Leonard, and J. W. Fisher, "A mixture of manhattan frames: Beyond the manhattan world," in *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pp. 3770–3777, IEEE, 2014.
- [12] J. T. Kent, A. M. Ganeiber, and K. V. Mardia, "A new method to simulate the bingham and related distributions in directional data analysis with applications," *arXiv preprint arXiv:1310.8110*, 2013.
- [13] S. Riedel, Z.-C. Marton, and S. Kriegel, "Multi-view orientation estimation using bingham mixture models," in *Automation, Quality and Testing, Robotics (AQTR), 2016 IEEE International Conference on*, pp. 1–6, IEEE, 2016.
- [14] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The euroc micro aerial vehicle datasets," *The International Journal of Robotics Research*, 2016.