

A WRAPPED TRIVARIATE NORMAL DISTRIBUTION AND BAYES INFERENCE FOR 3-D ROTATIONS

Yu Qiu, Daniel J. Nordman and Stephen B. Vardeman

Iowa State University

Abstract: For modeling orientation data represented as 3×3 rotation matrices, we develop a wrapped trivariate normal distribution (wTND) under which random rotations have simple geometric construction as symmetric errors about a mean. While of interest in its own right, the wTND also provides simple and effective approximations to the isotropic Gaussian distribution on rotations, with some advantages over approximations based on other commonly used models. We develop non-informative Bayes inference for the wTND via Markov Chain Monte Carlo methods that allow straightforward computations in a model where maximum likelihood is undefined. Credible regions for model parameters (including a fixed 3×3 mean rotation) are shown to possess good frequentist coverage properties. We illustrate the model and inference method with orientation data collected in texture analysis from materials science.

Key words and phrases: CLT, credible set of cones, isotropic Gaussian distribution, MCMC, UARS model.

1. Introduction

Three-dimensional orientation data are of interest in such fields as human kinematics, vectorcardiography, structural geology, robotics and materials science (cf., Downs (1972); Chang (1998); Matthies, Muller, and Vinel (1988); Rancourt, Rivest, and Asselin (2000); Stavdahl et al. (2005); Bingham, Nordman, and Vardeman (2009a)); see Mardia and Jupp (2000, Sec. 13.2.1) for an introduction. With such data, each observation is represented by a 3×3 rotation matrix in $SO(3)$, the set of orthogonal matrices with determinant 1, and typically denotes the orientation of some object after rotating its reference frame in \mathbb{R}^3 away from some “world” reference frame. For clarity in what follows, we refer to a probability model for a random 3×3 rotation matrix as a *rotational distribution*.

In many applications involving orientation data, the rotational distributions used are symmetric or *isotropic* (having central or rotationally invariant densities) about a central rotation in $SO(3)$, and intended to model the variability in orientation data as due to directionally symmetric random perturbations about

an underlying mean rotation parameter. Using such rotational distributions is akin to using errors ε , symmetrically distributed around 0 in a standard location model $Y = \mu + \varepsilon$ for real-valued data. In the statistical literature, the oldest and most common distribution on $SO(3)$ of this form is the isotropic version of the Matrix Fisher distribution (cf., Downs (1972); Khatri and Mardia (1977)). Other such models include the isotropic Cayley distribution (León, Massé, and Rivest (2006, Sec. 5.2)), Bunge's Gaussian distribution (Bunge (1982)), the Lorentzian distribution (Matthies (1982)), the de la Vallée Poussin distribution (Schaeben (1997)) and the isotropic Gaussian distribution on $SO(3)$ (cf., Nikolayev and Savyolova (1997)). All of these belong to a general class of isotropic distributions on $SO(3)$, referred to here as "uniform-angle-random-spin" (UARS) distributions, that have intuitive interpretation as random "rotational errors," as well as a simple geometric construction in terms of Euler's axis-angle representation of rotations; see, for example, Bingham, Nordman, and Vardeman (2009a) and Hielscher, Schaeben, and Siemes (2010).

Our purposes in this manuscript are two-fold. First, we wish to clarify the isotropic Gaussian distribution (IGD) on $SO(3)$, a rotational distribution from texture analysis that is not widely appreciated in the statistical literature. An appealing property of this distribution for modeling orientations is its position as a type of "normal" distribution for rotations, by serving as the distributional limit for compositions of large numbers of independent, small random rotations. However, since its proposal (Savyolova (1984); Matthies, Muller, and Vinel (1988)), the IGD on $SO(3)$ has been criticized as having no motivation through a meaningful central limit theorem (CLT) argument with rotations (cf., Schaeben (1992)). We point out that there is indeed a simple, rigorous argument showing that the IGD on $SO(3)$ does have a CLT-related motivation for modeling orientation data, giving it the same kind of justification as is usually provided for the normal and log-normal distributions in other statistical modeling applications.

Our second and main aim here is to develop a new family of isotropic distributions on $SO(3)$, referred to as the wrapped trivariate normal distribution (wTND) family. These rotational distributions are motivated by a CLT in \mathbb{R}^3 , rather than a CLT in $SO(3)$ directly, along with an exponential mapping of \mathbb{R}^3 onto $SO(3)$. One major motivation for the wTND is that it has a fairly simple distributional form for statistical inference, unlike the IGD on $SO(3)$ that has a rather complicated density (as do several other UARS models shown in Section 2.1, e.g., Bunge's Gaussian, Lorentzian). Additionally, the wTND turns out to closely approximate the IGD on $SO(3)$ in many practical situations, more so than many other commonly used isotropic models for rotations. Such approximations are useful, not only because the IGD on $SO(3)$ has CLT motivations, but also because *any* highly concentrated UARS distribution for rotational errors with a continuously differentiable density will follow the IGD on $SO(3)$

(see Section 2.2), and so the wTND model can become a tractable substitute for highly concentrated orientations. Implicit assumptions involving trivariate normal distributions have also appeared for approximating highly concentrated Matrix Fisher distributions (Rancourt, Rivest, and Asselin (2000)) that we more formally and directly characterize. We then describe one-sample Bayes inference for the wTND, using non-informative priors on the two parameters of the distribution: one (location) parameter is a fixed mean rotation $\mathbf{S} \in SO(3)$, and the other parameter $\kappa \in (0, \infty)$ controls the variability of random rotations from the wTND. We use Bayes inference because the approach is computationally straightforward and well-defined, unlike maximum likelihood. Simulations also indicate that, with non-informative priors, the resulting Bayes credible regions have excellent frequentist properties.

The rest of the manuscript is organized as follows. Section 2 describes the UARS-framework for isotropic distributions on $SO(3)$ and provides a CLT motivation for the IGD on $SO(3)$. Section 3 provides the wTND, along with some simulation studies indicating the effectiveness of its approximation to the IGD on $SO(3)$ compared to some competing approximations. We outline one-sample Bayes inference for the wTND in Section 4 and examine the procedure through simulation in Section 5. Section 6 illustrates an application of the wTND to orientation data collected in texture analysis, and Section 7 provides concluding remarks. An on-line Supplementary Appendix contains additional results.

2. Preliminaries: UARS models and the IGD on $SO(3)$

2.1. The UARS class: rotationally symmetric models on $SO(3)$

“Uniform-angle-random-spin” (UARS) distributions for random rotations can be described using a stochastic version of Euler’s angle-axis representation for rotations. For $\mathbf{v} = (v_1, v_2, v_3)^T \in \mathbb{R}^3$, define a mapping

$$\mathbf{A}(\mathbf{v}) = \begin{pmatrix} 0 & -v_3 & v_2 \\ -v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{pmatrix}$$

of \mathbb{R}^3 to the space $so(3)$ of real-valued skew-symmetric 3×3 matrices, and define the matrix exponential

$$\exp(\mathbf{B}) = \sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{B}^k$$

for $\mathbf{B} \in so(3)$. Then,

$$\exp(\mathbf{A}(\mathbf{v})) = (\cos \|\mathbf{v}\|) \mathbf{I}_3 + \frac{\sin \|\mathbf{v}\|}{\|\mathbf{v}\|} \mathbf{A}(\mathbf{v}) + \frac{1 - \cos \|\mathbf{v}\|}{\|\mathbf{v}\|^2} \mathbf{v} \mathbf{v}^T$$

represents a rotation of the identity matrix \mathbf{I}_3 , the set of standard coordinate vectors in \mathbb{R}^3 , by an angle of $\|\mathbf{v}\|$ about a vector or signed axis $\mathbf{v} \in \mathbb{R}^3$ following the right-hand rule (cf., Mardia and Jupp (2000, p.287)). Then, letting $\mathbf{u} = (u_1, u_2, u_3)^T$ be uniformly distributed over the \mathbb{R}^3 unit sphere S^2 and, independently, letting r denote a random draw from an *angular distribution* on $(-\pi, \pi]$ having a symmetric density $g(\cdot|\kappa)$ around zero whose spread is controlled by the concentration parameter $\kappa > 0$, a random UARS rotation with mean direction \mathbf{I}_3 is given by

$$\mathbf{M}(r, \mathbf{u}) \equiv \exp(\mathbf{A}(r\mathbf{u})) = (\cos r)\mathbf{I}_3 + (\sin r)\mathbf{A}(\mathbf{u}) + (1 - \cos r)\mathbf{u}\mathbf{u}^T, \quad (2.1)$$

a rotation by the random angle r about the random vector $\mathbf{u} \in \mathbb{R}^3$. This construction is then used to define a UARS distribution with (fixed) mean rotation $\mathbf{S} \in SO(3)$ by $\mathbf{O} = \mathbf{S} \cdot \mathbf{M}(r, \mathbf{u})$ (or equivalently $\mathbf{M}(r, \mathbf{u}) \cdot \mathbf{S}$), representing a directionally symmetric perturbation of \mathbf{S} . We refer to the rotational distribution of \mathbf{O} as a UARS model with mean $\mathbf{S} \in SO(3)$ and angular density $g(\cdot|\kappa)$.

As an important feature of UARS models, each rotational distribution in the UARS class is completely characterized by some angular distribution in the definition (2.1), and all of the previously mentioned common families of isotropic distributions on $SO(3)$ correspond to different choices of angular densities $g(\cdot|\kappa)$ defined on $(-\pi, \pi]$; these are listed in Table 1 (e.g., isotropic Matrix Fisher, Cayley, Bunge’s Gaussian, and IGD). Given an angular density $g(\cdot|\kappa)$ on $(-\pi, \pi]$ and mean parameter $\mathbf{S} \in SO(3)$, a UARS rotation \mathbf{O} has a corresponding density on $SO(3)$ given by

$$f(\mathbf{O}|\mathbf{S}, \kappa) = \frac{4\pi}{3 - \text{tr}(\mathbf{S}^T \mathbf{O})} g(\arccos[2^{-1}(\text{tr}(\mathbf{S}^T \mathbf{O}) - 1)]|\kappa), \quad \mathbf{O} \in SO(3), \quad (2.2)$$

with respect to the uniform distribution on $SO(3)$ which provides a dominating measure on $SO(3)$; taking $g(r) = [1 - \cos(r)]/[2\pi]$, $r \in (-\pi, \pi]$ and $\mathbf{S} = \mathbf{I}_3$ in (2.2) gives the density $f(\mathbf{O}) = 1$ of the uniform distribution on $SO(3)$ (Miles (1965)).

We thank referees for suggesting other generalizations and characterizations of UARS distributions. If a random variable t has a density $\tilde{g}(t)$ (with respect to the Lebesgue measure on \mathbb{R}) and, independently, \mathbf{u} is uniformly distributed on S^2 , the unit sphere in \mathbb{R}^3 , then $\mathbf{S} \exp(\mathbf{A}(t\mathbf{u}))$ is UARS-distributed with mean rotation \mathbf{S} and the density of the wrapped angle $r = t(\text{mod}2\pi)$ is

$$g(r) = \sum_{m=-\infty}^{\infty} \tilde{g}(r + m2\pi), \quad r \in (-\pi, \pi], \quad (2.3)$$

on the unit circle S^1 ; this is relevant for the wrapped trivariate normal distribution (wTND) described in Section 3.1. Any random orientation \mathbf{O} having a

Table 1. Angular density functions, with indicated concentration parameters, for the random angle $r \in (-\pi, \pi]$ defining common UARS models (2.1) on $SO(3)$. (Below I_i denotes the modified Bessel function of order i , $C(\cdot)$ denotes a normalizing constant, and $\lambda = \lambda(\kappa_L) = \kappa_L/2 - 0.5 + 2/(\kappa_L + 2)^2$ puts the Lorentzian distribution on roughly the same scale as the others, though the Lorentzian shape differs from the others for large concentrations κ_L .)

| Model | Angular Density |
|--|---|
| ^a Isotropic Cayley (κ_C) or ^b de la Vallée Poussin | $\frac{1 - \cos r}{2\pi} \frac{\sqrt{\pi} \Gamma(2\kappa_C^2 + 2) (1 + \cos r)^{2\kappa_C^2}}{2^{2\kappa_C^2} \Gamma(2\kappa_C^2 + 1/2)}$ |
| ^c Isotropic Matrix Fisher (κ_F) | $\frac{1 - \cos r}{2\pi} \frac{\exp(\kappa_F^2 \cos r)}{I_0(\kappa_F^2) - I_1(\kappa_F^2)}$ |
| ^d Bunge's Gaussian (κ_{BG}) | $\frac{1 - \cos r}{2\pi} C(\kappa_{BG}) \exp[-\kappa_{BG}^2 \frac{r^2}{2}]$ |
| ^e Lorentzian ($\lambda = \lambda(\kappa_L)$) | $\frac{1 - \cos r}{2\pi} (1 + \lambda) \frac{(1 + 2\lambda)^2 + 4\lambda(\lambda + 1) \cos^2(r/2)}{[(1 + 2\lambda)^2 - 4\lambda(\lambda + 1) \cos^2(r/2)]^2}$ |
| ^f Isotropic Gaussian (κ_{IG}) | $\frac{1 - \cos r}{2\pi} \sum_{m=0}^{\infty} (2m + 1) \exp[-\frac{m(m + 1)}{2\kappa_{IG}^2}] \frac{\sin[(m + 1/2)r]}{\sin(r/2)}$ |
| Wrapped Trivariate Normal (κ) | $\sum_{m=-\infty}^{\infty} \frac{\kappa^3}{\sqrt{2\pi}} (2m\pi - r)^2 \exp[-\frac{\kappa^2(2m\pi - r)^2}{2}]$ |

^aLeón, Massé, and Rivest (2006, Sec. 3.2).

^bSchaeben (1997).

^cDowns (1972); Khatri and Mardia (1977); Matthies, Muller, and Vinel (1988); León, Massé, and Rivest (2006, Sec. 3.2).

^dBunge (1982); Matthies, Muller, and Vinel (1988); Bucharova and Savyolova (1993).

^eMatthies (1982); Matthies, Muller, and Vinel (1988).

^fSavyolova (1984); Matthies, Muller, and Vinel (1988); Nikolayev and Savyolova (1997); Borovkov and Savyolova (2007).

density (2.2) on $SO(3)$ with respect to the uniform distribution that depends on \mathbf{O} only through a function $h[tr(\mathbf{O})]$, has a UARS distribution with mean rotation \mathbf{I}_3 and an angle r with Lebesgue density $h[1 + 2 \cos r](1 - \cos r)/(2\pi)$ on $r \in (-\pi, \pi]$; the density for \mathbf{O} is zonal/central on $SO(3)$ in that $h[tr(\mathbf{O})] = h[tr(\mathbf{O}_1^T \mathbf{O} \mathbf{O}_1)]$ for $\mathbf{O}, \mathbf{O}_1 \in SO(3)$. If \mathbf{v} has an isotropic (rotation-invariant or spherically symmetric) distribution on \mathbb{R}^3 , then $\exp(\mathbf{A}(\mathbf{v}))$ has a UARS distribution and any UARS distribution can be obtained this way. Finally, as unit quaternions (vectors on S^3 , the unit sphere in \mathbb{R}^4) can be equivalently used to represent rotations, if $\mathbf{u} = (u_1, u_2, u_3)^T$ and r denote the random Euler axis-angle in the UARS formulation (2.1), then the Cayley-Klein map $\rho(\mathbf{w}) = \mathbf{I}_3 + 2w_1\mathbf{A}((w_2, w_3, w_4)^T) + 2\mathbf{A}((w_2, w_3, w_4)^T)^2$ of the random quaternion $\mathbf{w} =$

$(w_1, w_2, w_3, w_4)^T = (u_1 \sin(r/2), u_2 \sin(r/2), u_3 \sin(r/2), \cos(r/2))^T$ has a UARS distribution on $SO(3)$ with mean rotation \mathbf{I}_3 , and all distributions on unit quaternions which are rotationally symmetric about $(0, 0, 0, 1)^T$ induce UARS distributions on $SO(3)$ through this mapping. For more on UARS distributions and map-induced distributions on $SO(3)$ via the exponential map on $SO(3)$ or Cayley-Klein map on S^3 , see Prentice (1986), Schaeben and Nikolayev (1998), Mardia and Jupp (2000, Chap. 13.2), León, Massé, and Rivest (2006, Sec. 3.2), Bingham, Nordman, and Vardeman (2009a), and Hielscher, Schaeben, and Siemes (2010).

2.2. A CLT motivation for the IGD on $SO(3)$

In critiquing several UARS models on $SO(3)$ used in texture analysis, including Bunge’s Gaussian, the Lorentzian and the isotropic Matrix Fisher distributions (cf., Table 1), Matthies, Muller, and Vinel (1988, p.85) argued that it may be physically plausible to imagine crystal orientations observed in materials as built from composition of small, independent rotations in the texture development and therefore reasonable to motivate a “normal” distribution for orientation data by a CLT for rotations. Those authors informally provided a density on $SO(3)$ for the limit distribution of rotational compositions, and Savyolova (1984) derived the same density by characterizing a “normally” distributed rotation as having an infinitely divisible distribution. This density corresponds to the isotropic Gaussian distribution (IGD) on $SO(3)$ (see Table 1), that has been further studied and generalized by Nikolayev and Savyolova (1997). Schaeben (1992) and Schaeben and Nikolayev (1998, Sec. 5) criticized the work of Matthies, Muller, and Vinel (1988), arguing that no physically meaningful CLT argument for rotations could motivate the IGD as “normal” on $SO(3)$ and that no CLT analog exists for compositions in $SO(3)$ under assumptions similar to those for the CLT in Euclidean spaces. But this is untrue, as seen in Proposition 1, which straightforwardly combines a CLT result of Parthasarathy (1964) on $SO(3)$ with a triangular array of UARS-distributed rotations (see the Supplementary Material for details).

Proposition 1. *Suppose $r_{1,n}, \dots, r_{n,n}$ are iid draws from a symmetric distribution on $(-\pi, \pi]$ with variance $\sigma^2 > 0$ and, independently, let $\mathbf{u}_{1,n}, \dots, \mathbf{u}_{n,n}$ be iid vectors, uniformly distributed on S^2 . Fix $\mathbf{S} \in SO(3)$ and define UARS rotations $\mathbf{O}_{1,n}, \dots, \mathbf{O}_{n,n}$ by forming $\mathbf{O}_{i,n}$ with angle $r_{i,n}/\sqrt{n}$ and axis $\mathbf{u}_{i,n}$ in (2.1). Then, the composition $\mathbf{O}^{(n)} = \mathbf{S} \prod_{i=1}^n \mathbf{O}_{i,n}$ converges in distribution to an isotropic Gaussian distribution on $SO(3)$ as $n \rightarrow \infty$, a UARS model with mean rotation \mathbf{S} and the angular density in Table 1 having concentration parameter $\kappa_N = \sqrt{3}/\sigma$.*

Hence, the IGD on $SO(3)$ does indeed have a CLT-motivation as the limit of many “small” i.i.d. physical rotations in 3-D, supporting the argument of

Matthies, Muller, and Vinel (1988). Beyond Proposition 1, the composition of independent UARS-distributed rotations is *always* UARS-distributed, the UARS class of distributions is closed under composition (as proven in the Supplementary Material). But more is true for the IGD on $SO(3)$, because the convolution of independent rotations with an IGD has again an IGD (cf., Nikolayev and Savyolova (1997, Thm. 3)). In these ways, the IGD on $SO(3)$ does behave like the “normal” distributions associated with any Euclidean space.

The IGD on $SO(3)$ is again characterized by an angular density (cf., Table 1) in (2.1) which is not particularly tractable. Other angular densities with analytically simpler forms, like those associated with Bunge’s Gaussian and the isotropic Matrix Fisher distributions on $SO(3)$, have been suggested as approximations for the IGD (cf., Nikolayev and Savyolova (1997)). But these have also been criticized as having shortcomings (cf., Matthies, Muller, and Vinel (1988); Bucharova and Savyolova (1993)). For example, the normalizing constant in the angular density for Bunge’s Gaussian distribution is not expressible in a closed form (cf., Table 1), and the Matrix Fisher-based approximation is not good except for very large concentrations (see Section 3.2 and Figure 2). This motivates us to consider a wrapped trivariate normal distribution (wTND) to provide a simple distributional approximation to the IGD on $SO(3)$.

Before leaving this section, we add that the criticisms of Schaeben (1992) and Schaeben and Nikolayev (1998) largely concerned an operational definition of a “normal” distribution on $SO(3)$. As with the normal distribution on \mathbb{R} , various characterizations of “normality” exist for rotations. For example, Bunge’s (1982) Gaussian distribution is an analog of the real-valued normal distribution in terms of being a solution to a heat equation on manifolds (cf., Bucharova and Savyolova (1993)). Schaeben (1992) commented that, statistically speaking, the isotropic Matrix Fisher distribution could be argued to be “normal” due to its matrix density representation (2.2),

$$f(\mathbf{O}|\mathbf{S}, \kappa) = \frac{\exp[\kappa_F^2 2^{-1}(\text{tr}(\mathbf{S}^T \mathbf{O}) - 1)]}{I_0(\kappa_F^2) - I_1(\kappa_F^2)}, \quad \mathbf{O} \in SO(3),$$

which has an exponential form, decaying away from its mode \mathbf{S} (as does a normal distribution in Euclidean space). On the other hand, the Matrix Fisher rotational distribution is itself not the limit distribution of small rotational compositions and we do not know if this family is closed under convolutions (see the Supplementary Material for evidence that it is not closed). As Schaeben and Nikolayev (1998) also noted, the isotropic Matrix Fisher distribution closely matches the IGD on $SO(3)$ for highly concentrated orientations, which has close connections to results for directional data on the \mathbb{R}^p unit sphere S^{p-1} (cf., Roberts and Ursell (1960); Hartman and Watson (1974); Kent (1978); Mardia and Jupp (2000,

p.173)). While true, for *any* concentrated UARS distribution with continuously differentiable density and mean rotation \mathbf{I}_3 , a Taylor expansion of the density about \mathbf{I}_3 shows that it is close to that of an IGD on $SO(3)$. And compared to the Matrix Fisher and other UARS models, the wTND of the next section can provide closer approximations to the IGD on $SO(3)$ for a wider range of concentration parameters.

3. Wrapped Trivariate Normal Distributions on $SO(3)$

3.1. Definition and motivation

Suppose \mathbf{x} has a trivariate normal distribution $N(\mathbf{0}_3, \kappa^{-2}\mathbf{I}_3)$ on \mathbb{R}^3 with component variance $\kappa^{-2} > 0$. Then, by wrapping \mathbb{R}^3 onto $SO(3)$, as in Section 2.1, $\mathbf{S} \exp[\mathbf{A}(\mathbf{x})]$ defines a wrapped trivariate normal distribution (wTND) on orientations with (fixed) mean rotation $\mathbf{S} \in SO(3)$ and concentration parameter $\kappa > 0$. For a random variable b independent of \mathbf{x} with $P(b = 1) = P(b = -1) = 1/2$, one may decompose $\mathbf{x} = t\mathbf{u}$ in terms of independent $t = b\|\mathbf{x}\|$ and $\mathbf{u} = b\mathbf{x}/\|\mathbf{x}\|$, where \mathbf{u} is uniformly distributed on S^2 , to see that the wTND is a UARS model with an angle-axis construction (2.1) defined by \mathbf{u} and $r = t(\bmod 2\pi)$. As $\kappa^2 t^2$ has a chi-square distribution with 3 degrees of freedom, it follows from (2.3) that a random “spin” or angle $r \in (-\pi, \pi]$ has a (Lebesgue) density

$$g_{wTN}(r|\kappa) = \frac{\kappa^3}{\sqrt{2\pi}} \sum_{m=-\infty}^{\infty} (2m\pi - r)^2 \exp\left(-\frac{\kappa^2(2m\pi - r)^2}{2}\right) \quad (3.1)$$

corresponding to a wrapped (symmetrized) Maxwell-Boltzmann distribution with concentration parameter $\kappa > 0$ (up to scaling, the Maxwell-Boltzmann distribution is that of the square-root of a χ_3^2 variable and appears in modeling particle speeds in statistical mechanics, cf., Peckham and McNaught (1992)). The wTND then has a particularly direct and simple path to simulation that can be attractive for modelers. One may either simulate and wrap independent $N(0, \kappa^{-2})$ values, or simulate a random angle from the wTN angular density (3.1) for use in (2.1) via $r = (-1)^b |\kappa^{-1} w^{1/2} - \pi \lfloor \kappa^{-1} w^{1/2} / \pi \rfloor|$ with χ_3^2 random variable w and an independent Bernoulli variable b (0 or 1 with equal probabilities). The wrapped kernel in (3.1) also closely resembles that of the wrapped normal density (with standard deviation $\delta > 0$) on $(-\pi, \pi]$,

$$g(r|\delta) = \frac{1}{\sqrt{2\pi\delta}} \sum_{m=-\infty}^{\infty} \exp\left(-\frac{(2m\pi - r)^2}{2\delta^2}\right), \quad r \in (-\pi, \pi]$$

which is a commonly used angular distribution for modeling for 2-D rotations.

The wTND has a CLT-related motivation because sums of iid small variance quantities in \mathbb{R}^3 lead to trivariate normal distributions in Euclidean space which

can then be wrapped onto $SO(3)$. In particular, taking $\mathbf{s}_n = \sum_{i=1}^n r_{i,n} \mathbf{u}_{i,n}$ with iid $\mathbf{u}_{1,n}, \dots, \mathbf{u}_{n,n}$ uniformly-distributed on S^2 and iid random angles $r_{1,n}, \dots, r_{n,n}$ from a common distribution on $(-\pi, \pi]$ having mean zero and variance σ^2/n for some $\sigma > 0$, $\exp[\mathbf{A}(\mathbf{s}_n)]$ converges to wTND on $SO(3)$ with mean rotation \mathbf{I}_3 and $\kappa = \sqrt{3}/\sigma$ by the usual CLT in \mathbb{R}^3 . Because $\exp[\mathbf{A}(\mathbf{s}_n)] \approx \prod_{i=1}^n \mathbf{M}(r_{i,n}, \mathbf{u}_{i,n})$ for small rotations in (2.1), the wTND also approximates the IGD on $SO(3)$ as the limit of a large number of compositions of “small” independent random rotations (cf., Proposition 1), a phenomenon which is next investigated through simulation.

3.2. Comparisons of the wTND to other UARS models

One expects the wTND (and other UARS models) to be close to the IGD on $SO(3)$ for sufficiently large concentrations κ . To gain some rough idea of how large κ must be for effective approximations, in Figure 1 we plot the cumulative distribution functions of $|r|$ for a random angle $r \in (-\pi, \pi]$ from the symmetric angular density (3.1) of the wTND as well as the angular density from the IGD on $SO(3)$ with concentrations $\kappa = 3, 2, 1, 0.5$ (cf., Table 1), and we compare these against the true sampling distribution of the absolute angle $|r_n|$ resulting from the composition of n iid rotation matrices $\mathbf{M}(r_{i,n}, \mathbf{u}_{i,n})$ with the angles $r_{i,n}$ having $\text{uniform}(-3\kappa^{-1}n^{-1/2}, 3\kappa^{-1}n^{-1/2})$ distributions for $n = 4, 10$. The comparisons show that, at least when $\kappa \geq 2$, the wTND effectively approximates the IGD’s angular distribution, which is in turn a good approximation to the real angular distribution that describes the composition. The central limit convergence of products to a IGD limit appears to be remarkably fast, suggesting potentially wide-spread applications for the family (and good approximations to it) where observed physical orientations are plausibly modeled as derived from multiple small random perturbations of a basic orientation.

Common rotational models belong to the UARS class and so can be described in terms of their angular densities, listed in Table 1 (an alternative description of UARS distributions through related densities is described in the Supplementary Material). Where necessary, we have reparameterized the densities from their most common forms so that all parameters κ are non-negative and control the concentrations of the distributions in a similar manner. Except for the Lorentzian case, the angular densities for the models in Table 1 nearly match that of the IGD if the parameters κ are large enough. In Figure 2, we also compare the (absolute) angular densities from Table 1 for the isotropic Cayley, a de la Vallée Poussin distribution, Matrix Fisher, Bunge’s Gaussian, and wTN models to that of the IGD on $SO(3)$ for $\kappa = 10, 5, 2, 1$. We can see that the angular density from the wTND approximates the angular density of the IGD much better than the Cayley and Matrix Fisher-distributions, and at least as well as Bunge’s Gaussian

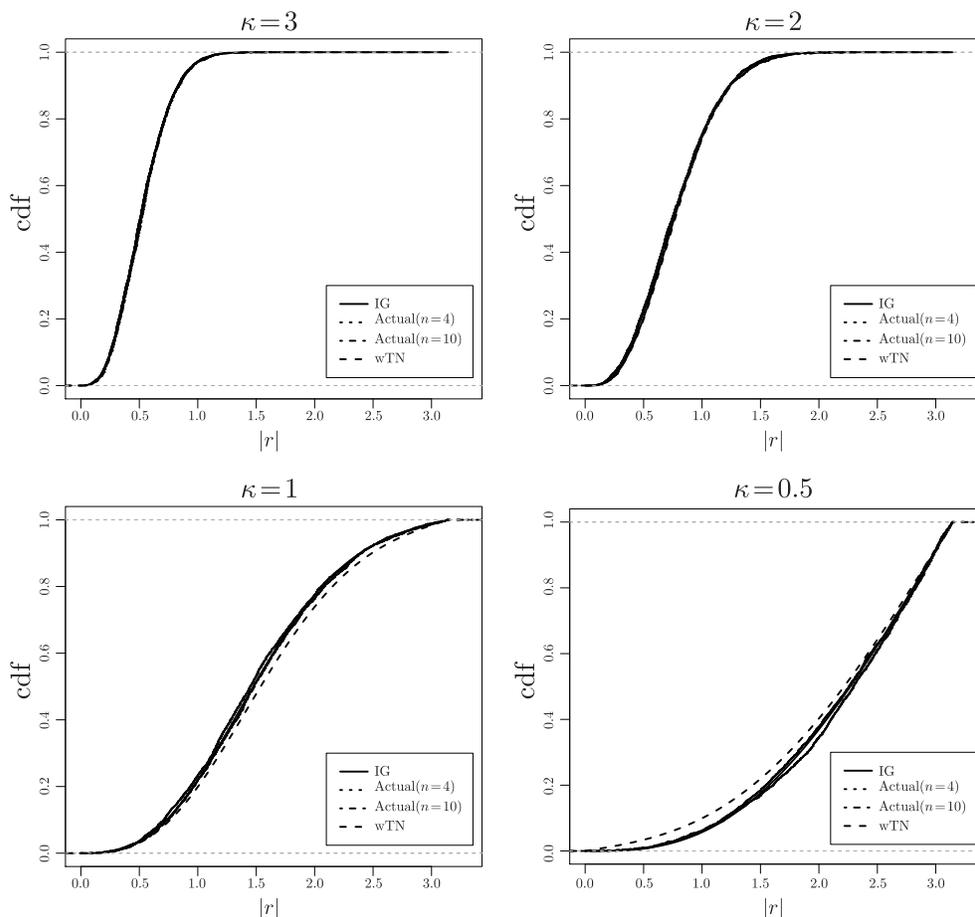


Figure 1. Cumulative distribution functions (cdf) of $|r|$ for angles $r \in (-\pi, \pi]$ from the symmetric angular densities associated with isotropic Gaussian (IG) and wrapped trivariate normal (wTN) distributions on $SO(3)$ with different concentration parameters $\kappa > 0$. Also provided are the cdfs of the absolute angle $|r_n|$ (approximated from 100,000 simulations) as determined by the product of $n = 4, 10$ independent UARS-distributed rotation matrices (each having uniform $(-3\kappa^{-1}n^{-1/2}, 3\kappa^{-1}n^{-1/2})$ angular distributions).

distribution when κ is small (though, as indicated in Table 1, the angular density from the wTND has a closed form while the normalizing constant of Bunge's Gaussian distribution has to be numerically determined for each concentration parameter κ in Figure 2).

4. One-sample Bayes Methods for wTND on $SO(3)$

From (3.1), we obtain the density (with respect to the uniform distribution)

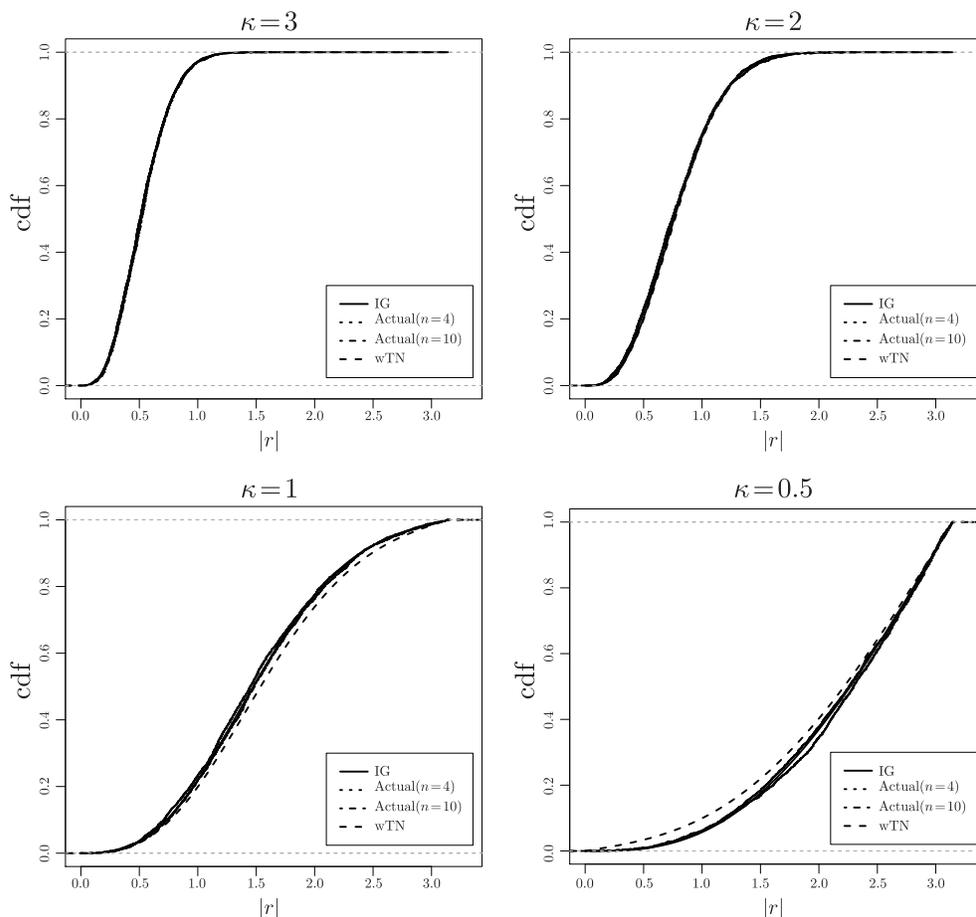


Figure 2. Densities for $|r|$ when $r \in (-\pi, \pi]$ follows the symmetric angular density associated with the isotropic Gaussian (IG), wrapped trivariate normal (wTN), isotropic Matrix Fisher, isotropic Cayley, and Bunge’s Gaussian rotational distributions.

for wTND on $SO(3)$ as

$$f(\mathbf{O}|\mathbf{S}, \kappa) = \frac{4\pi}{3 - \text{tr}(\mathbf{S}^T \mathbf{O})} g_{wTN}(\arccos[2^{-1}(\text{tr}(\mathbf{S}^T \mathbf{O}) - 1)]|\kappa), \quad \mathbf{O} \in SO(3) \tag{4.1}$$

from (2.2). The density (4.1) has a singularity at $\mathbf{O} = \mathbf{S}$. (The other models represented in Table 1 do not have such singularities due to the term $1 - \cos r$ in their densities for r .) However, this does not prevent us from developing useful Bayes inference, where maximum likelihood estimation would technically be undefined. In fact, due to the non-regularity of the likelihood function, the convergence rate of Bayes procedures for estimating \mathbf{S} can be *super-efficient* and

observably so in realistic sample sizes for small κ , as we will illustrate with simulations in Section 5. At the same time, for sufficiently large concentration parameters κ , the wTND can also behave “regularly” whereby the numerator of its matrix density (4.1) decays to zero rapidly enough to effectively cancel out the singularity; the simulations of Section 5 will also clarify this behavior.

For Bayes inference, we would like to identify potentially non-informative prior distributions for the parameters \mathbf{S} and κ of the wTND(\mathbf{S}, κ), so that the resulting credible regions have good frequentist coverage properties. To this end, we use a prior selection approach as in Bingham, Vardeman, and Nordman (2009b); Bingham, Nordman, and Vardeman (2009c). As a prior for the mean rotation parameter \mathbf{S} , we use the uniform distribution on $SO(3)$ having density $p(\mathbf{S}) = 1, \mathbf{S} \in SO(3)$. For the concentration parameter κ , we use the Jeffreys prior for the angular density. Here it is slightly more convenient for discussion and plotting purposes to consider the corresponding prior for the spread parameter $\eta = -\log \kappa$, which has density

$$J(\eta) = \exp(-\eta)\sqrt{\mathcal{I}(\exp(-\eta))}, \quad \eta \in (-\infty, \infty)$$

for

$$\begin{aligned} \mathcal{I}(\kappa) &= E \left[\left(\frac{d}{d\kappa} \log g_{wTN}(r|\kappa) \right)^2 \right] \\ &= -\frac{9}{\kappa^2} + \frac{\kappa^5}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \frac{(\sum_{m=-\infty}^{\infty} (2m\pi - r)^4 \exp(-\kappa^2(2m\pi - r)^2/2))^2}{\sum_{m=-\infty}^{\infty} (2m\pi - r)^2 \exp(-\kappa^2(2m\pi - r)^2/2)} dr. \end{aligned}$$

While this density does not have a closed form, $J(\eta)$ can be evaluated numerically and we display this (improper) Jeffreys prior density in Figure 3. Because $J(\eta) \rightarrow 0$ as $\eta \rightarrow \infty$ and $J(\eta) \rightarrow \sqrt{6}$ as $\eta \rightarrow -\infty$, to determine $J(\eta)$ numerically in simulations, we use $J(\eta) \approx \sqrt{6}$ when $\eta < -0.5$, $J(\eta) \approx 0$ when $\eta > 2$ and, for $-0.5 \leq \eta \leq 2$, we fit a cubic spline to approximate $J(\eta)$ after calculating the density at grid points $-0.5 + 2.5/1000 \cdot i, i = 0, 1, \dots, 1000$. From Figure 3, we see that the Jeffreys prior density, perhaps surprisingly, is not a simple monotone or unimodal function; it has turning points around $\eta = 0.5$ and $\eta = 0.85$. The non-monotonicity affects the behavior of samples simulated from the posterior distribution, especially for small sample sizes.

For a random sample $\mathbf{O}_i, i = 1, \dots, n$, from the wTND, the likelihood function for (\mathbf{S}, η) is

$$L(\mathbf{S}, \eta) \propto \frac{\prod_{i=1}^n g_{wTN}(\arccos[2^{-1}(tr(\mathbf{S}^T \mathbf{O}_i) - 1)] | \exp(-\eta))}{\prod_{i=1}^n (3 - tr(\mathbf{S}^T \mathbf{O}_i))},$$

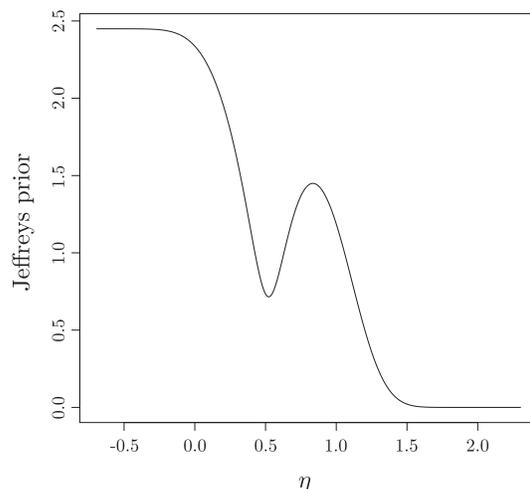


Figure 3. (Improper) Jeffreys prior density for $\eta = -\log \kappa$.

by (4.1). Multiplying by prior densities $p(\mathbf{S})$ and $J(\eta)$ gives a posterior density $h(\mathbf{S}, \eta)$ for (\mathbf{S}, η) proportional to

$$\left(\frac{\prod_{i=1}^n g_{wTN}(\arccos[2^{-1}(tr(\mathbf{S}^T \mathbf{O}_i) - 1)] | \exp(-\eta))}{\prod_{i=1}^n (3 - tr(\mathbf{S}^T \mathbf{O}_i))} \right) J(\eta).$$

We sample a sequence (\mathbf{S}^j, η^j) from the posterior distribution using a Metropolis-Hastings-within-Gibbs (MHG) algorithm with variables $\mathbf{O}_1, \dots, \mathbf{O}_n \in SO(3)$ and the starting values \mathbf{S}^0, η^0 , as follows:

1. As a proposal for \mathbf{S}^j , generate \mathbf{S}^{j*} from the isotropic Matrix Fisher distribution with location parameter \mathbf{S}^{j-1} and concentration κ_F , see the angular density in Table 1 in (2.1). Here κ_F is a tuning parameter.
2. Compute $r_j^{(1)} = h(\mathbf{S}^{j*}, \eta^{j-1})/h(\mathbf{S}^{j-1}, \eta^{j-1})$ and generate $w_j^{(1)} \sim \text{Bernoulli}(\min(1, r_j^{(1)}))$. Take $\mathbf{S}^j = w_j^{(1)}\mathbf{S}^{j*} + (1 - w_j^{(1)})\mathbf{S}^{j-1}$.
3. Generate normal $\eta^{j*} \sim N(\eta^{j-1}, \gamma^2)$. Here γ is a tuning parameter.
4. Compute $r_j^{(2)} = h(\mathbf{S}^j, \eta^{j*})/h(\mathbf{S}^j, \eta^{j-1})$ and generate $w_j^{(2)} \sim \text{Bernoulli}(\min(1, r_j^{(2)}))$. Take $\eta^j = w_j^{(2)}\eta^{j*} + (1 - w_j^{(2)})\eta^{j-1}$.

Section 5 describes a simulation study of one-sample Bayes inference for the wTND using this algorithm and we explain how posterior draws can be used to

Table 2. Values of tuning parameters $\kappa_F = \sqrt{2\rho}$ and γ expressed in terms of (ρ, γ) .

| | $n = 10$ | | $n = 30$ | | $n = 100$ | | $n = 300$ | | $n = 1000$ | |
|-----------------|----------|----------|----------|----------|-----------|----------|-----------|----------|------------|----------|
| | ρ | γ | ρ | γ | ρ | γ | ρ | γ | ρ | γ |
| $\eta = 1.3$ | 5 | 0.5 | 50 | 0.5 | 1000 | 0.4 | 5000 | 0.3 | 200000 | 0.25 |
| $\eta = 0.85$ | 1 | 0.7 | 10 | 0.5 | 200 | 0.2 | 1500 | 0.15 | 4000 | 0.08 |
| $\eta = 0.5$ | 0.5 | 0.7 | 0.7 | 0.5 | 2 | 0.3 | 4 | 0.2 | 20 | 0.1 |
| $\eta = 0$ | 2 | 0.7 | 4 | 0.5 | 15 | 0.13 | 50 | 0.08 | 200 | 0.05 |
| $\eta = -0.347$ | 5 | 0.4 | 20 | 0.23 | 50 | 0.13 | 150 | 0.07 | 500 | 0.04 |
| $\eta = -1.151$ | 33 | 0.4 | 100 | 0.23 | 350 | 0.13 | 800 | 0.07 | 3000 | 0.04 |
| $\eta = -1.844$ | 150 | 0.4 | 300 | 0.23 | 1200 | 0.13 | 4000 | 0.07 | 12000 | 0.04 |
| $\eta = -3.454$ | 4000 | 0.4 | 10000 | 0.23 | 35000 | 0.13 | 80000 | 0.07 | 300000 | 0.04 |

construct credible regions for $\mathbf{S} \in SO(3)$ and $\eta \in \mathbb{R}$.

5. Bayes Credible Regions and Coverage Accuracy

We conducted a simulation study for several different combinations (n, η) . In generating rotation data from the wTND($\mathbf{S}, \kappa = \exp[-\eta]$), we set the true mean rotation \mathbf{S} to be \mathbf{I}_3 , as the choice of \mathbf{S} is irrelevant (cf., Bingham, Nordman, and Vardeman (2009a)). The values for the spread parameter η were $-3.454, -1.844, -1.151, -0.347, 0, 0.5, 0.85, 1.3$ and sample sizes were $n = 10, 30, 100, 300, 1,000$. For each combination (n, η) , we simulated 4,000 data sets and, with each, we generated $N = 100,000$ samples from the posterior distribution using the MHG algorithm after a 25,000 iteration burn-in period. After inspecting several different starting values and finding the simulation results to be insensitive to this choice, we chose starting values for \mathbf{S}^0 and η^0 in the simulation study to be the true parameters. The tuning parameters κ_F and γ listed in Table 2 were chosen to keep the Metropolis-Hastings jumping rates between 30% and 40%.

For the purpose of analysis, a 95% credible level was used. Two types of credible intervals for η were obtained from the posterior sampling, equal-tail (ET) intervals and shortest length (SL) intervals. Credible regions for \mathbf{S} were constructed using the method of “credible sets of cones” described by Bingham, Nordman, and Vardeman (2009c). Thus if $\mathbf{S}^1, \dots, \mathbf{S}^N$ denote the posterior samples, we define a Bayes point estimate \mathbf{S}_B of the mean rotation as the maximizer of $\sum_{j=1}^N \text{tr}(\mathbf{S}_B^T \mathbf{S}^j)$, the Bayes estimator under a squared error loss function $\text{tr}[(\mathbf{S}_B - \mathbf{S})(\mathbf{S}_B - \mathbf{S})^T]$, and then define a credible region by a “set of cones” of angle a around each column vector in \mathbf{S}_B , where a is the 95th percentile of $\{a_1, \dots, a_N\}$ and each a_j represents the maximum arccosine value (between 0 and π) of the diagonal elements of $\mathbf{S}_B^T \mathbf{S}^j$. Hence, a region for \mathbf{S} can be graphically illustrated as in Figure 4 and the size of the region is defined in terms of the angle between the centers (columns of \mathbf{S}_B) and edges of the cones.

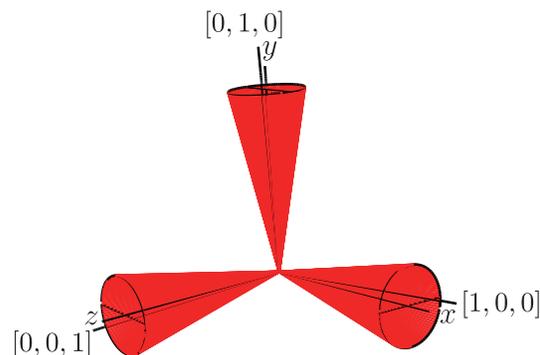


Figure 4. A 95% credible region for the parameter \mathbf{S} with x, y and z representing the orientation, column vectors, of the Bayes point estimate $\mathbf{S}_B = [x\ y\ z]$ for \mathbf{S} .

After finding the credible regions for \mathbf{S} and η for each of the 4,000 data sets at each (n, η) , we determined whether the regions for \mathbf{S} and η contained the true values. This provided the coverage rates for \mathbf{S} and η for the (n, η) combinations in Table 3. For both \mathbf{S} and η , the frequentist coverage rates of Bayes regions are consistent with their credible levels and as sample size increased, the coverage rates converged to the nominal ones. This indicates that the current Bayes approach is effective for obtaining good frequentist coverage accuracy.

We also considered median sizes for the 4,000 credible regions for \mathbf{S} and η , where we used the cone angle to characterize the size of a region for \mathbf{S} and computed lengths of both ET and SL credible intervals for η . Results are summarized in Tables 4 and 5. From Table 4, the two methods for obtaining 95% intervals for η produce similar results for all combinations (n, η) . For a fixed η , as sample size n increases, the intervals become narrower, as expected. However, for fixed n , the width of interval for η is not strictly monotone decreasing in η due to the effect of the prior shape and, in particular, for $\eta < 0$, concentration parameter $\kappa > 1$, the width does not change as η decreases.

With credible regions for \mathbf{S} , as seen in Table 5, the median cone angle decreases as n increases for any fixed η . However, for each $\eta > 0$, this convergence rate (found by regressing the log of median angle over the log of n for $n = 100, 300, 1,000$) is approximately $O(1/n)$ due to the non-regularity of the likelihood (cf., a circular data case in Nordman, Vardeman, and Bingham (2009)). For $\eta \leq 0$, the empirical convergence rate is approximately $O(1/\sqrt{n})$. This is consistent with our claim that for a large concentration parameter κ , the wTND effectively approximates the IGD on $SO(3)$ (having regular behavior). Thus, for large κ , there is effectively no wrapping involved in the angular density (3.1) from the wTND and the only real contribution to the summation (3.1) is the $m = 0$ component,

Table 3. Coverage rates (%) for \mathbf{S} and η using 95% Bayes credible regions for different combinations of (n, η) ; credible regions for η characterized here are ET intervals (with SL intervals performing similarly).

| (n, η) | \mathbf{S} | η | (n, η) | \mathbf{S} | η |
|--------------|--------------|--------|----------------|--------------|--------|
| (10, 1.3) | 95.5 | 95.6 | (10, -0.347) | 96.0 | 95.9 |
| (30, 1.3) | 95.9 | 93.9 | (30, -0.347) | 95.4 | 94.6 |
| (100, 1.3) | 92.9 | 94.9 | (100, -0.347) | 95.0 | 96.0 |
| (300, 1.3) | 95.2 | 95.4 | (300, -0.347) | 96.3 | 96.2 |
| (1000, 1.3) | 95.2 | 94.9 | (1000, -0.347) | 95.2 | 95.0 |
| (10, 0.85) | 93.2 | 98.6 | (10, -1.151) | 95.5 | 94.6 |
| (30, 0.85) | 95.9 | 98.2 | (30, -1.151) | 93.2 | 95.5 |
| (100, 0.85) | 95.1 | 95.5 | (100, -1.151) | 93.1 | 96.0 |
| (300, 0.85) | 94.9 | 94.5 | (300, -1.151) | 94.2 | 94.1 |
| (1000, 0.85) | 95.0 | 95.1 | (1000, -1.151) | 95.0 | 94.6 |
| (10, 0.5) | 93.5 | 96.4 | (10, -1.844) | 97.3 | 95.9 |
| (30, 0.5) | 96.7 | 95.6 | (30, -1.844) | 93.6 | 93.6 |
| (100, 0.5) | 95.3 | 95.2 | (100, -1.844) | 94.3 | 94.3 |
| (300, 0.5) | 94.7 | 94.9 | (300, -1.844) | 95.3 | 94.0 |
| (1000, 0.5) | 95.1 | 95.3 | (1000, -1.844) | 95.0 | 95.4 |
| (10, 0) | 95.9 | 94.4 | (10, -3.454) | 94.7 | 95.4 |
| (30, 0) | 95.5 | 95.9 | (30, -3.454) | 95.0 | 95.5 |
| (100, 0) | 95.5 | 92.7 | (100, -3.454) | 93.1 | 95.0 |
| (300, 0) | 94.7 | 93.9 | (300, -3.454) | 95.4 | 95.1 |
| (1000, 0) | 95.2 | 95.0 | (1000, -3.454) | 95.0 | 94.9 |

for which the r^2 term there essentially behaves like $1 - \cos r$. Intuitively, this $1 - \cos r$ factor allows the wTN density to “look like” the angular densities in Table 1 corresponding to regular rotational distributions. Even for large concentrations, the wTN model is non-regular (due to the spikes in (4.1)) but this aspect is not practically “seen” at even fairly large sample sizes.

6. An Application to Orientation Data from EBSD

Here we make use of part of a data set collected in the study of Bingham, Lograsso, and Laabs (2010). That paper provides details of an electron backscatter diffraction (EBSD) experiment done to measure crystal orientations in a nickel specimen. Fourteen repeat scans were made on a 2-D rectangular grid on the specimen’s planar surface, at over 4,000 sites per scan. We use data from a particular 4×28 sub-grid and a single scan. The EBSD measurement device returned an orientation matrix (in terms of 3 Euler angles) at each location, and we consider the characterization of variation in orientations across the grid.

We used the Bayes methods in Bingham, Nordman, and Vardeman (2009c) and here to fit both isotropic Matrix Fisher and wTN models to the 112 observed

Table 4. Median width of 95% Bayes credible intervals for η with different combinations of (n, η) and both equal-tail (ET) and shortest-length (SL) intervals.

| (n, η) | ET Width | SL Width | (n, η) | ET Width | SL Width |
|--------------|----------|----------|----------------|----------|----------|
| (10,1.3) | 1.429 | 1.405 | (10, -0.347) | 0.559 | 0.554 |
| (30,1.3) | 0.574 | 0.566 | (30, -0.347) | 0.300 | 0.298 |
| (100,1.3) | 0.432 | 0.428 | (100, -0.347) | 0.161 | 0.160 |
| (300,1.3) | 0.351 | 0.348 | (300, -0.347) | 0.093 | 0.092 |
| (1000,1.3) | 0.282 | 0.276 | (1000, -0.347) | 0.051 | 0.051 |
| (10, 0.85) | 1.275 | 1.252 | (10, -1.151) | 0.543 | 0.538 |
| (30, 0.85) | 0.895 | 0.870 | (30, -1.151) | 0.299 | 0.297 |
| (100, 0.85) | 0.330 | 0.322 | (100, -1.151) | 0.161 | 0.160 |
| (300, 0.85) | 0.161 | 0.160 | (300, -1.151) | 0.093 | 0.092 |
| (1000, 0.85) | 0.086 | 0.086 | (1000, -1.151) | 0.051 | 0.051 |
| (10, 0.5) | 1.210 | 1.192 | (10, -1.844) | 0.543 | 0.538 |
| (30, 0.5) | 0.787 | 0.781 | (30, -1.844) | 0.298 | 0.297 |
| (100, 0.5) | 0.470 | 0.469 | (100, -1.844) | 0.161 | 0.160 |
| (300, 0.5) | 0.313 | 0.307 | (300, -1.844) | 0.092 | 0.092 |
| (1000, 0.5) | 0.157 | 0.153 | (1000, -1.844) | 0.051 | 0.050 |
| (10, 0) | 1.370 | 1.322 | (10, -3.454) | 0.542 | 0.538 |
| (30, 0) | 0.358 | 0.347 | (30, -3.454) | 0.298 | 0.297 |
| (100, 0) | 0.171 | 0.170 | (100, -3.454) | 0.161 | 0.160 |
| (300, 0) | 0.098 | 0.097 | (300, -3.454) | 0.092 | 0.092 |
| (1000, 0) | 0.053 | 0.053 | (1000, -3.454) | 0.051 | 0.050 |

orientations. Although the computations involved were much more complicated, we also fit the IGD to the data by maximum likelihood. Estimated concentration parameters for these fits were, respectively,

$$\hat{\kappa}_F = 1.365, \quad \hat{\kappa} = 0.974, \quad \text{and} \quad \hat{\kappa}_{IG} = 0.932.$$

In texture analysis, the absolute value $|r|$ of the random spin $r \in (-\pi, \pi]$ in a UARS rotation (2.1) is often referred to as a misorientation angle, the smallest (non-negative) angle in an axis-angle representation needed to align a rotation (2.1) back to a standard reference frame \mathbf{I}_3 , cf., Randle (2003); note $|r|$ has a density on $[0, \pi]$ that twice the angular density (for r) listed in Table 1 for the isotropic Matrix Fisher distribution, the wTND, and the IGD. For each of these models, Figure 5 plots the fitted cumulative distribution function for the misorientation angle $|r|$. These are plotted against the empirical distribution $\{\widehat{|r|}_{ij} : i = 1, \dots, 4; j = 1, \dots, 28\}$ of misorientation angles, computed as $\widehat{|r|}_{ij} = \arccos\{[tr(\hat{\mathbf{S}}^T \mathbf{O}_{ij}) - 1]/2\}$ using a non-parametric “moment” estimator $\hat{\mathbf{S}}$ of the mean rotation for de-trending, defined as the maximizer of $\sum_{i,j} tr(\mathbf{S}^T \mathbf{O}_{ij})$; this estimation of misorientation angles uses the fact that a UARS orientation $\mathbf{O} = \mathbf{S} \cdot \mathbf{M}(r, \mathbf{u})$ satisfies $tr(\mathbf{S}^T \mathbf{O}) = tr(\mathbf{M}(r, \mathbf{u})) = 1 + 2 \cos |r|$, from (2.1).

Table 5. Median cone angle (in radians) of Bayes credible sets for \mathbf{S} with different combinations (n, η) and the apparent moderate sample size convergence rate of the median angles for fixed η .

| (n, η) | Angle | Apparent Convergence Rate | (n, η) | Angle | Apparent Convergence Rate |
|--------------|-------|------------------------------|----------------|-------|------------------------------|
| (10, 1.3) | 1.492 | $n^{-1.008}$ | (10, -0.347) | 0.667 | $n^{-0.505}$ |
| (30, 1.3) | 0.346 | | (30, -0.347) | 0.390 | |
| (100, 1.3) | 0.102 | | (100, -0.347) | 0.211 | |
| (300, 1.3) | 0.031 | | (300, -0.347) | 0.122 | |
| (1000, 1.3) | 0.010 | | (1000, -0.347) | 0.066 | |
| (10, 0.85) | 1.542 | $n^{-1.042}$ | (10, -1.151) | 0.288 | $n^{-0.497}$ |
| (30, 0.85) | 1.503 | | (30, -1.151) | 0.161 | |
| (100, 0.85) | 0.243 | | (100, -1.151) | 0.088 | |
| (300, 0.85) | 0.068 | | (300, -1.151) | 0.050 | |
| (1000, 0.85) | 0.022 | | (1000, -1.151) | 0.028 | |
| (10, 0.5) | 1.546 | $n^{-0.920}$ | (10, -1.844) | 0.141 | $n^{-0.499}$ |
| (30, 0.5) | 1.525 | | (30, -1.844) | 0.080 | |
| (100, 0.5) | 0.831 | | (100, -1.844) | 0.044 | |
| (300, 0.5) | 0.153 | | (300, -1.844) | 0.028 | |
| (1000, 0.5) | 0.098 | | (1000, -1.844) | 0.014 | |
| (10, 0) | 1.485 | $n^{-0.651}$ | (10, -3.454) | 0.028 | $n^{-0.476}$ |
| (30, 0) | 0.725 | | (30, -3.454) | 0.016 | |
| (100, 0) | 0.371 | | (100, -3.454) | 0.009 | |
| (300, 0) | 0.123 | | (300, -3.454) | 0.005 | |
| (1000, 0) | 0.082 | | (1000, -3.454) | 0.003 | |

The plot suggests that the fitted IGD and wTND are essentially identical, and do a better job of describing the “texture” of the nickel specimen in terms of variability in crystal orientations across this grid of locations than does the fitted Matrix Fisher model. The methods of this paper further establish that 95% cones for the mean rotation \mathbf{S} in the wTND have angle 22.86° and that 95% limits for κ are 0.895 and 1.053. This illustrates the utility of the wTN model and our method of non-informative Bayes inference.

7. Conclusion

We have provided a physical framework to motivate the isotropic Gaussian distribution (IGD) on $SO(3)$ as the limit distribution of a composition of large number of small, independent rotational errors (specifically, rotationally symmetric errors from the uniform-angle-random-spin (UARS) class of rotational distributions).

In part because the IGD has a complicated distributional form, we have developed a new UARS model as the wrapped trivariate normal distribution

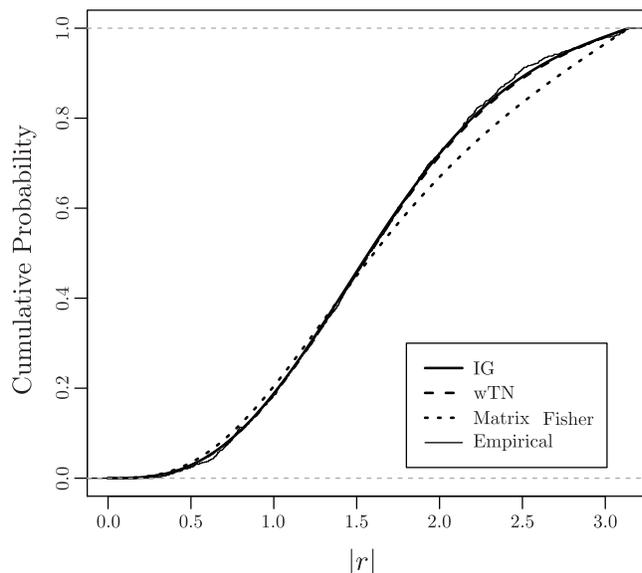


Figure 5. For the EBSD nickel data, plot of the empirical distribution of estimated misorientation angles and the cumulative distribution for $|r|$ in fitted models for the isotropic Matrix Fisher distribution, the wTND, and the IGD on $SO(3)$.

(wTND); it is tractable and provides natural approximations for the limit behavior of the composition of many small independent rotations. We have demonstrated the straightforward implementation and effectiveness of non-informative Bayes inference for these distributions.

There remains the question of conducting inference for the IGD directly. This remains a topic of future research, but we believe that the appropriate Bayes approach may offer a practical solution. As with the wTND, Bayes inference for the IGD is suggested purely on computational grounds, but more analytical work is required to develop non-informative priors with this model. In general, however, we expect the basic prescription of “product of uniform prior on \mathbf{S} and Jeffreys prior on κ ” plus “MHG sampling to approximate posteriors” to be reasonable for essentially any one-sample UARS model. As building blocks for more complicated models, UARS families, including the wTND, and generalizations of the one-sample Bayes analyses have their place in regression, time series, spatial, and other kinds of statistical modeling and inference.

Acknowledgements

We thank two referees and an associate editor for detailed and constructive comments which improved the manuscript. The associate editor and one

referee in particular contributed very substantially to a comprehensive revised manuscript.

References

- Bingham, M. A., Lograsso, B. K. and Laabs, F. C. (2010). A statistical analysis of the variation in measured crystal orientations obtained through electron backscatter diffraction. *Ultramicroscopy* **110**, 1312-1319.
- Bingham, M. A., Nordman, D. J. and Vardeman, S. B. (2009a). Modeling and inference for measured crystal orientations and a tractable class of symmetric distributions for rotations in three dimensions. *J. Amer. Statist. Assoc.* **104**, 1385-1397.
- Bingham, M. A., Vardeman, S. B. and Nordman, D. J. (2009b). Bayes one-sample and one-way random effects analyses for 3-D orientations with application to materials science. *Bayesian Anal.* **4**, 607-630.
- Bingham, M. A., Nordman, D. J. and Vardeman, S. B. (2009c). Finite-sample investigation of likelihood and Bayes inference for the symmetric von Mises-Fisher distribution. *Comput. Statist. Data Anal.* **54**, 1317-1327.
- Borovkov, M. and Savyolova, T. (2007). The computational approaches to calculate normal distributions on the rotation group. *J. Appl. Cryst.* **40**, 449-455.
- Bucharova, T. I. and Savyolova, T. I. (1993). Application of normal distributions on $SO(3)$ and S^n for orientation distribution function approximation. *Textures Microstruct.* **21**, 161-176.
- Bunge, H. J. (1982). *Texture Analysis in Material Science*. Butterworth, London.
- Chang, T. (1998). Estimating the relative rotation of two tectonic plates from boundary crossings. *J. Amer. Statist. Assoc.* **83**, 1178-1183.
- Downs, T. D. (1972). Orientation Statistics. *Biometrika* **59**, 665-676.
- Hartman, P. and Watson, G. S. (1974). Normal distribution functions on spheres and the modified Bessel functions. *Annal. Probab.* **2**, 593-607.
- Hielscher, R., Schaeben, H., and Siemes, H. (2010). Orientation distribution within a single hematite crystal. *Math. Geosci.* **42**, 359-375.
- Kent, J. (1978). Limiting behaviour of the von Mises-Fisher distribution. *Math. Proc. Cambridge Philos. Soc.* **84**, 531-536.
- Khatri, C. G. and Mardia, K. V. (1977). The Von Mises-Fisher matrix distribution in orientation statistics. *J. Roy. Statist. Soc. Ser. B*, **39**, 95-106.
- León, C. A., Massé, J.-C., and Rivest, L.-P. (2006). A statistical model for random rotations. *J. Multivariate Anal.* **97**, 412-430.
- Mardia, K.V. and Jupp, P.E. (2000). *Directional Statistics*. Wiley, New York.
- Matthies, S. (1982). Form effects in the description of the orientation distribution function (ODF) of texturized materials by model components. *Physica Status Solidi (b)* **112**, 705-716.
- Matthies, S., Muller, J., and Vinel, G. W. (1988). On the normal distribution in the orientation space. *Textures Microstruct.* **10**, 77-96.
- Miles, R. E. (1965). On random rotations in \mathbb{R}^3 . *Biometrika* **52**, 636-639.
- Nikolayev, D. I. and Savyolova, T. I. (1997). Normal distribution on the rotation group $SO(3)$. *Textures Microstruct.* **29**, 201-233.

- Nordman, D. J., Vardeman, S. B., and Bingham, M. A. (2009). Uniformly hyper-efficient Bayes inference in a class of non-regular problems. *Amer. Statist.* **63**, 234-238.
- Parthasarathy, K. P. (1964). The Central limit theorem for the rotation group. *Theory Probab. Appl.* **9**, 273-282.
- Peckham, G. D. and McNaught, I. J. (1992). Applications of the Maxwell-Boltzmann distribution. *J. Chem. Educ.* **69**, 554-558.
- Prentice, M. J. (1986). Orientation statistics without parametric assumptions. *J. Roy. Statist. Soc. Ser. B* **48**, 214-222.
- Randle, V. (2003). *Microtexture Determination and its Applications*. Maney for The Institute of Materials, Minerals and Mining, London.
- Rancourt, D., Rivest, L. P., and Asselin, J. (2000). Using orientation statistics to investigate variations in human kinematics. *J. Roy. Statist. Soc. Ser. C* **49**, 81-94.
- Roberts, P. H. and Ursell, H. D. (1960). Random walk on a sphere and on a Riemannian manifold. *Philos. Trans. R. Soc. Lond. A* **252**, 317-356.
- Savyolova, T. I. (1984). Distribution functions of grains with respect to polycrystal orientations and its Gaussian approximations. *Ind. Lab.* **50**, 468-474.
- Schaeben, H. (1992). "Normal" orientation distributions. *Textures Microstruct.* **19**, 197-202.
- Schaeben, H. (1997). The de la Vallée Poussin standard orientation density function *Textures and Microstructures* **33**, 365-373.
- Schaeben, H. and Nikolayev, D. I. (1998). The central limit theorem in texture component fit methods. *Acta Appl. Math.* **53**, 59-87.
- Stavdahl, O., Bondhus, A. K., Pettersen, K. Y. and Malvig, K. E. (2005). Optimal statistical operators for 3-dimensional rotational data: Geometric interpretations and application to prosthesis kinematics. *Robotica* **3**, 283-292.

Department of Statistics, Iowa State University, Ames, IA, US 50011, USA.

E-mail: yuqiu1982@gmail.com

Department of Statistics, Iowa State University, Ames, IA, US 50011, USA.

E-mail: dnordman@iastate.edu

Departments of Statistics and Industrial and Manufacturing Systems Engineering, Iowa State University, Ames, IA, US 50011.

E-mail: vardeman@iastate.edu

(Received September 2011; accepted April 2013)