New Modeling and Bayes Inference for 3-D Orientations/Rotations and Equivalence Classes of Them

UARS, PARS, and Induced Models

Stephen Vardeman
(With Melissa Bingham, Dan Nordman, Yu Qiu, and Chuanlong Du)

(Supported in part by NSF grant DMS #0502347 EMSW21-RTG)

March 2015

Introduction (Our Original Motivation)

- Quantifying variation in measurements obtained through Electron Backscatter Diffraction (EBSD) — a technique used in studying the microtexture of metals
- Data are (nominally)[1] orientations of cubic crystals at scanned positions on a metal surface

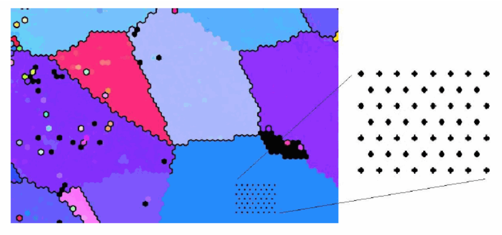

Figure: Grain map from a scan of a nickel specimen, orientation readings every 0.2 μm (clustered somehow).

[1] More on this later.

Introduction (Some Background)

- Orientations (relative to a "world" coordinate system) in 3-d may be represented by 3×3 orthogonal matrices $\mathbf{O}$ with positive determinant (rotation matrices)

- A uniform distribution over the set of such matrices/orientations, say $\boldsymbol{\Omega}$, is the Haar measure, μ

- $\mathbf{O} \sim \mu$ may be generated (among other ways) by
  - generating a 3-d direction as a uniformly distributed unit vector $\mathbf{u}$ (a point on the unit sphere)
  - rotating clockwise about an axis pointed in the $\mathbf{u}$ direction through an (independently chosen) random angle $r \in (-\pi, \pi]$ with pdf proportional to $1 - \cos r$

- Various families of distributions on $\boldsymbol{\Omega}$ have been defined in terms of densities wrt to μ

Introduction (Visualizing)

The problem of modeling (and then inference) for randomness/variation in 3-d orientations might be visualized by picturing where the world coordinate axes end up upon rotation by **O**

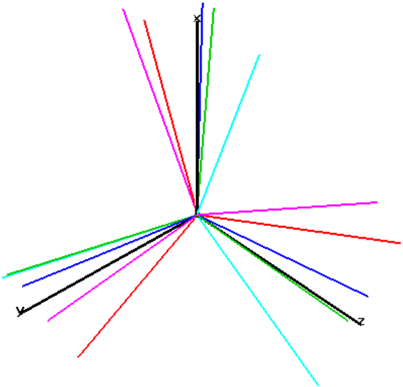

Figure: World axes (black) and several random perturbations.

Introduction (Matrix Fisher)

- The most famous class of distributions on $\mathbf{\Omega}$ is the set of Matrix Fisher von-Mises (Matrix Fisher) distributions with densities wrt to μ

$$a(\mathbf{M})\exp(tr(\mathbf{M}'\mathbf{o})), \mathbf{o} \in \Omega$$

for $\mathbf{M}$ a 3×3 matrix of full rank (and $a(\mathbf{M})$ a normalizing constant) ... this parameterization is not intuitively helpful

- A symmetric version (restriction) of the Matrix Fisher class has densities wrt to μ

$$\frac{\exp(\kappa[tr(\mathbf{S}'\mathbf{o})-1])}{I_0(2\kappa)-I_1(2\kappa)} = b(\kappa)\exp\left(\kappa \cdot tr(\mathbf{S}^T\mathbf{o})\right), \ \mathbf{o} \in \Omega$$

for $\mathbf{S} \in \Omega$, $\kappa \geq 0$, and $b(\kappa)$ a normalizing constant ... here $\mathbf{S}$ serves as a "principal orientation"/center/location of the distribution and $\kappa \geq 0$ controls distribution "spread" ... this is a "two-parameter" family with directly interpretable parameters

Introduction (Symmetric Matrix Fisher)

- As it turns out, a symmetric Matrix Fisher realization with principal orientation **I** may be generated by
  - generating a 3-d direction as a uniformly distributed unit vector **u**
  - rotating clockwise about an axis pointed in the **u** direction through an (independently chosen) random angle $r \in (-\pi, \pi]$ with pdf proportional to

$$(1 - \cos r) \exp(2\kappa \cos r)$$

- If **O** is symmetric Matrix Fisher with principal orientation **I** and **S** $\in \Omega$, then **SO** is Matrix Fisher with the same κ, but principal orientation **S**
- Several other lesser-known symmetric distributions in the literature share the above constructive characterization, but with different pdf's for r
- It's easy enough to find orientation data where symmetry is appropriate, but none of the published models fit

The Uniform-Axis-Random-Spin (UARS) Class

Taking a clue from the characterization of the symmetric Matrix Fisher distributions, a broad class of "location-scale" distributions on orientations has the constructive definition

- $\mathbf{O} \sim \text{UARS}_C(\mathbf{S},\kappa)$ may be generated by
  - generating a 3-d direction as a uniformly distributed unit vector $\mathbf{u}$ (a point on the unit sphere)
  - rotating about an axis pointed in the $\mathbf{u}$ direction clockwise through an (independently chosen) random angle $r \in (-\pi, \pi]$ with pdf $C(\cdot|\kappa)$ symmetric on $(-\pi, \pi]$ for $\kappa \geq 0$ some concentration parameter
  - rotating by $\mathbf{S}$
- By varying C, this covers most (?all?) published symmetric distributions and opens the possibility of considering many new ones
  - A symmetric von-Mises circular distribution choice of C gave us good fits in a materials/measured-crystal-orientation application

Visualizing Two Members of a UARS Family

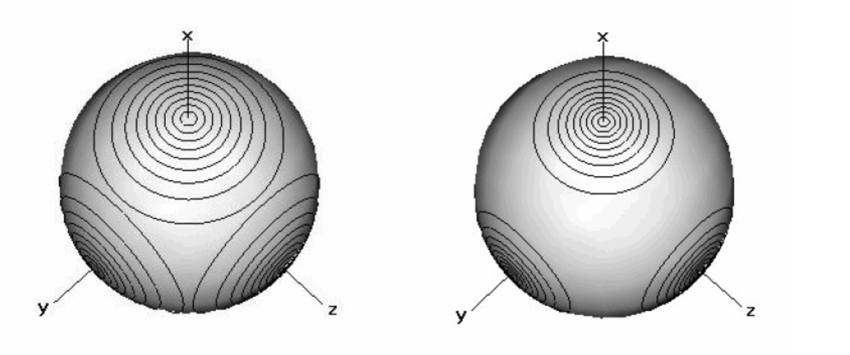

Figure: Probability content contours for rotated axes produced by 2 members of the same UARS family.

Other "Standard" Members of the UARS Class

- Isotropic Gaussian

$$C_{\mathrm{IG}}(r|\kappa) = \frac{1-\cos r}{2\pi} \sum_{m=0}^{\infty} (2m+1) \exp\left[\frac{-m(m+1)}{2\kappa^2}\right] \frac{\sin\left[(m+1/2)r\right]}{\sin(r/2)}$$

- Bunge "Gaussian"

$$C_{\mathrm{Bunge}}(r|\kappa) = \frac{1-\cos r}{2\pi} N(\kappa) \exp\left[-\kappa^2 r^2/2\right]$$

- de la Vallée Poussin (Caley)

$$C_{\mathrm{Poussin}}(r|\kappa) = \frac{1-\cos r}{2\pi} \frac{B(3/2,1/2)}{B(3/2,2\kappa^2+1/2)} \cos^{4\kappa^2}(r/2)$$

- Lorentzian

$$C_{\mathrm{Lorentzian}}(r|\kappa) = \frac{1-\cos r}{2\pi}(1+\lambda) \frac{(1+2\lambda)^2 + 4\lambda(\lambda+1)\cos^2(r/2)}{\left[(1+2\lambda)^2 - 4\lambda(\lambda+1)\cos^2(r/2)\right]^2}$$

Useful (New) Members of the UARS Class

- von-Mises (circular density)

$$C_{\mathrm{vM}}(r|\kappa) = \frac{\exp(\kappa^2 \cos r)}{2\pi I_0(\kappa^2)}$$

(where I_0 denotes the modified Bessel function of order 0)

- wrapped Normal (circular density)

$$C_{\mathrm{wN}}(r|\kappa) = \frac{1}{\kappa^2} + \frac{\kappa}{\sqrt{2\pi}} \int_{-\pi}^{\pi} \sum_{m=-\infty}^{\infty} (2m\pi - r)^2 \exp\left(-(2m\pi - r)^2\kappa^2/2\right) dr$$

- "wrapped Trivariate Normal" (wrapped Maxwell-Boltzman circular density)

$$C_{\mathrm{wTN}}(r|\kappa) = \sum_{m=-\infty}^{\infty} \frac{\kappa^3}{\sqrt{2\pi}}(2m\pi - r)^2 \exp\left[-\kappa^2(2m\pi - r)^2/2\right]$$

(an excellent approximation to the IG density for large κ and easily simulated from)

Distributions of |r|

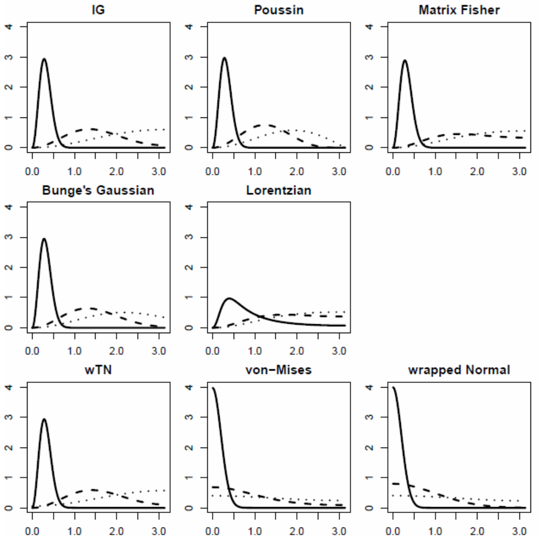

Figure: Densities for $|r|$ for $\kappa = 0.5$ (dotted), 1 (dashed), 5 (solid).

UARS Density

- The UARS$_C$ ($\mathbf{S}$,κ) distribution can be shown to have an R-N derivative wrt to the Haar measure (a density) of the form

$$f_C(\mathbf{o}|\mathbf{S},\kappa) = \frac{4\pi}{3 - tr(\mathbf{S}^T\mathbf{o})} C\left(\arccos[2^{-1}(tr(\mathbf{S}^T\mathbf{o}) - 1)]\Big|\kappa\right), \quad \mathbf{o} \in \Omega$$

  - The UARS class of distributions allows unrestricted choice of symmetric distribution for r... but unless the function $\frac{C(r|\kappa)}{1 - \cos r}$ has a finite limit at $r = 0$, the density will be unbounded at $\mathbf{o} = \mathbf{S}$
  - Standard distributions have a C producing a finite limit
  - We have found the symmetric von-Mises circular distribution version of $C(r|\kappa)$ and the wrapped trivariate normal version of $C(r|\kappa)$ to be useful ... and they are non-zero at $r = 0$

- The density provides a basis for likelihoods and the possibility of inference for UARS ("location" and "scale") parameters

UARS One-Sample Likelihood

For $\mathbf{O}_1, \ldots, \mathbf{O}_n \sim^{iid} \text{UARS}_C(\mathbf{S}, \kappa)$, the one-sample likelihood

$$L_n(\mathbf{S}, \kappa) = \prod_{i=1}^{n} f_C(\mathbf{o}_i | \mathbf{S}, \kappa)$$

serves to guide inference

- For the standard models (where $\dfrac{C(r|\kappa)}{1 - \cos r}$ has a finite limit at $r = 0$) the inference problem is regular (and all standard likelihood inference can be expected to work fine ... the only obvious wrinkle is parameterizing $\mathbf{S}$ for computation)
- For some other useful $C(r|\kappa)$, the problem is not regular ... the likelihood has singularities in $\mathbf{S}$ at every $\mathbf{o}_i$

UARS One-Sample Bayes Inference Generalities

Using a prior for $(\mathbf{S},\kappa)$ of product form $\mu \times J_C$ for μ the Haar measure and J_C the Jeffreys prior for κ for the densities $C(r|\kappa)$ on $(-\pi, \pi]$, it turns out that

- Geometrically interpretable Bayes credible sets are straightforward
- Bayes methods are well-calibrated in terms of producing frequentist performance measures matching fixed Bayes posterior probabilities used in specifying the methods
- Bayes performance in regular cases is at least as good as that of properly calibrated likelihood methods
- Bayes performance in non-regular cases is **strikingly better** (n^{-1} convergence rate in place of $n^{-1/2}$ convergence rate) than that of quasi-likelihood methods

MCMC for UARS One-Sample Bayes Inference

In light of the constructive characterization of UARS variables, it is trivial to implement a Metropolis-Hastings-within-Gibbs MCMC algorithm for any UARS_C one-sample posterior (a proposal for $\mathbf{S}^{j+1}$ is easily generated from a UARS distribution with principal orientation $\mathbf{S}^j$) .... e.g., suppressing dependence upon $\mathbf{o}_1, \mathbf{o}_2, \ldots, \mathbf{o}_n$ let $g_C(\mathbf{S}, \kappa)$ be the product of the UARS_C likelihood and the pdf for J_C on $[0, \infty)$

1. Generate $\mathbf{S}^{j*} \sim \text{UARS}\left(\mathbf{S}^{j-1}, \rho\right)$ as a proposal for $\mathbf{S}^j$
2. Compute $r_{1j} = \dfrac{g_C\left(\mathbf{S}^{j*}, \kappa^{j-1}\right)}{g_C\left(\mathbf{S}^{j-1}, \kappa^{j-1}\right)}$, generate $W_{1j} \sim \text{Ber}(\min(1, r_{1j}))$, and take $\mathbf{S}^j = W_{1j}\mathbf{S}^{j*} + (1 - W_{1j})\,\mathbf{S}^{j-1}$
3. Generate $\log(\kappa^{j*}) \sim \text{N}(\log(\kappa^{j-1}), \sigma^2)$, with κ^{j*} as a candidate for κ^j
4. Compute $r_{2j} = \dfrac{g_C\left(\mathbf{S}^j, \kappa^{j*}\right)\kappa^{j*}}{g_C\left(\mathbf{S}^j, \kappa^{j-1}\right)\kappa^{j-1}}$, generate $W_{2j} \sim \text{Ber}(\min(1, r_{2j}))$, and take $\kappa^j = W_{2j}\kappa^{j*} + (1 - W_{2j})\,\kappa^{j-1}$

UARS One-Sample Bayes Inference (Credible Sets)

For the $\mu \times J_C$ priors, credible intervals for κ based on MCMC iterates are obvious ... the intervals are well-calibrated for all n and generally comparable to (to better than) LRT or qLRT intervals

We have developed (new) *geometrically interpretable* "sets of cones" credible sets for $\mathbf{S}$

- Consider sets consisting of all $\mathbf{S}$ "near" some "center" of the posterior draws
- Measure "nearness" in terms of the maximum (positive) angle between a rotated coordinate axis for the measure of center and the corresponding rotated coordinate axis for $\mathbf{S}$
  - Geometrically, this can be represented by centering cones of fixed angles around rotated coordinate axes for the measure of center
  - In regular cases use an analytic posterior mode as a center ... in non-regular cases use $\mathbf{S}_\mathrm{B}$ maximizing $tr\left(\mathbf{S}'\overline{\mathbf{S}}\right)$ over choice of $\mathbf{S}$ (for $\overline{\mathbf{S}}$ the sample mean of the MCMC iterates)

UARS One-Sample Bayes Inference (Credible Sets)

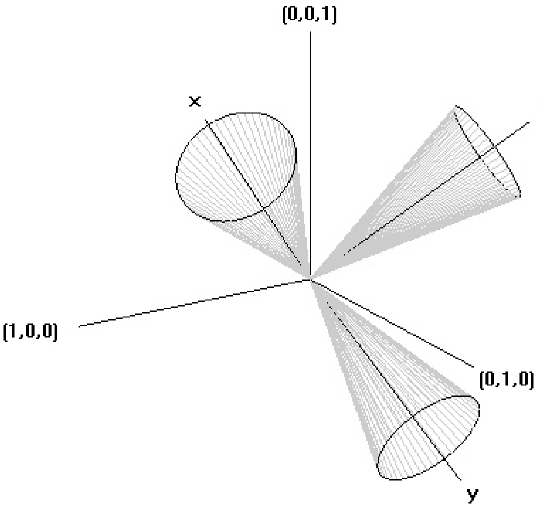

Figure: A "set-of-cones" credible set for **S**.

UARS One-Sample Bayes Inference (Credible Sets)

- Desired credible levels of sets of this form may be obtained by computing the maximum angles between the 3 coordinate axes rotated by $\mathbf{S}_{\mathrm{B}}$ and the corresponding axes for each MCMC iterate and using a desired quantile of these maxima to set a "cone angle around $\mathbf{S}_{\mathrm{B}}$"
- Inversion of likelihood ratio or quasi-likelihood ratio tests about **S** does not seem to produce confidence sets with interpretable geometry
- For the $\mu \times J_C$ priors, these credible sets for **S** are well-calibrated for all n
  - For the regular cases, Bayes sets are comparable to (to slightly smaller than) LRT sets
  - For the non-regular cases, the Bayes sets are **strikingly smaller** than qLRT sets, with Bayes median cone angles decreasing at rate n^{-1}

UARS Random Effects Models

Suppose that for $i = 1, \ldots, r$ and $k = 1, \ldots, m_i$,

$$\mathbf{O}_{ik} = \mathbf{P}_i \mathbf{Q}_{ik} \in \Omega$$

for $\mathbf{P}_i \sim^{iid} \text{UARS}_C(\mathbf{S}, \tau)$ independent of $\mathbf{Q}_{ik} \sim^{iid} \text{UARS}_C(\mathbf{I}, \kappa)$

- r groups with m_i observations in the i^{th} group, where τ represents the size of the between-group variation and κ represents the size of the within-group variation
- Conditioned on $\mathbf{P}_i$ the $\mathbf{O}_{ik}$ are ind $\text{UARS}_C(\mathbf{P}_i, \kappa)$

This is a "3-d orientation data" version of the one-way random effects model

(In the motivating materials example, one might wish to quantify within-grain and between-grain variation in measured orientation ... or for repeat scans, within-single-site variation and between-site variation for a single grain)

UARS Random Effects Model Bayes Inference

- Again (in light of the constructive characterization of UARS variables) it is easy to implement a Metropolis-Hastings-within-Gibbs MCMC algorithm for any UARS_C one-way random effects posterior
  - A proposal for $\mathbf{S}^{j+1}$ is easily generated from a UARS distribution with principal orientation $\mathbf{S}^{j}$
  - A proposal for $\mathbf{P}_i^{j+1}$ is easily generated from a UARS distribution with principal orientation $\mathbf{P}_i^{j}$

- Under the potentially non-informative $\mu \times J_C \times J_C$ prior for $(\mathbf{S}, \tau, \kappa)$ some scattered checks (for both Matrix Fisher and vM-UARS cases) suggest that the correctness of calibration of the one-sample Bayes procedures carries over to one-way random effects models ... *and this methodology appears to be the first available for the problem*

Other Models With UARS Components (Not Yet Addressed)

Regression Models

- For a known parametric function $\mathbf{S}(\boldsymbol{\theta}, \mathbf{x})$ taking values in Ω and data $(\mathbf{x}_i, \mathbf{o}_i)$, $i = 1, \ldots, n$, consider the model $\mathbf{O}_i \sim^{ind} \text{UARS}_C(\mathbf{S}(\boldsymbol{\theta}, \mathbf{x}_i), \kappa)$

Time Series Models

- For $\mathbf{Q}_i \sim^{iid} \text{UARS}_C(\mathbf{S}, \kappa)$ and $\mathbf{S}_0 \in \Omega$ the matrix partial products

$$\mathbf{O}_j = \mathbf{Q}_j \cdots \mathbf{Q}_2 \mathbf{Q}_1 \mathbf{S}_0$$

  form a 3-d orientation "random walk with drift"

  - For $\mathbf{E}_i \sim^{iid} \text{UARS}_C(\mathbf{I}, \tau)$ independent of the $\mathbf{Q}_i$, replacing the $\mathbf{O}_j$ with $\mathbf{E}_j \mathbf{O}_j$ allows for observing the random walk through noise ("state space" modeling)

In such cases, Bayes methods seem "obvious"

Nonsymmetric (PARS) Models With Interpretable Parameters

- One useful way of generalizing the UARS ideas beyond symmetric distributions on Ω is for

$$\mathbf{Q} \sim \text{UARS}_C(\mathbf{S}, \kappa)$$

and $\mathbf{R}$ (independent of $\mathbf{Q}$) representing clockwise rotation about an axis pointed in the direction of the fixed unit vector $\mathbf{V}$ through a random angle $r \sim C\left(r|\tau\right)$ to consider the distribution of

$$\mathbf{O} = \mathbf{RQ}$$

- Distributions so-defined have interpretable parameters $\mathbf{S}, \kappa, \mathbf{V}$, and τ ... the principal orientation and a corresponding concentration of symmetric variation about it, and a **preferred axis** and concentration of additional spin around it controlling the nature and amount of asymmetry ... (mixtures of UARS distributions)

PARS Models

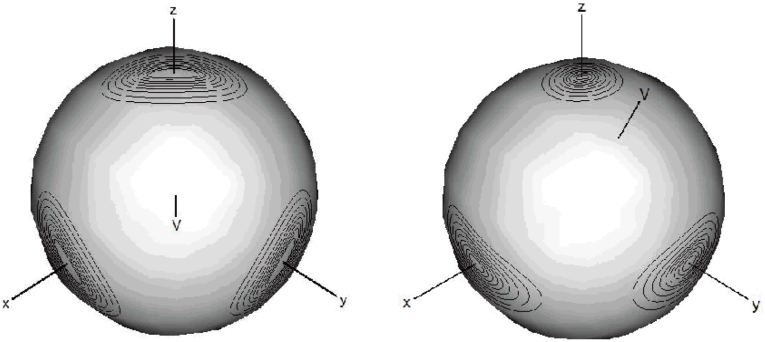

Figure: Contours capturing $10 \cdot i\%$ of the distribution, according to density, for the PARS($\mathbf{I}, 100, \mathbf{V}, 15$) distribution with (a) $\mathbf{V} = (1, 1, 1)/\sqrt{3}$ and (b) $\mathbf{V} = (1, 3, 5)/\sqrt{35}$

PARS Density

- Let $\mathbf{T}(\mathbf{V}, p) \in \boldsymbol{\Omega}$ represent rotation clockwise about an axis $\mathbf{V}$ through angle p
- If p has density $D(p|\tau)$ on $(-\pi, \pi]$ and $\mathbf{O} \sim \text{UARS}_C(\mathbf{S},\kappa)$ then $\mathbf{M} = \mathbf{T}(\mathbf{V}, p)\mathbf{O} \sim \text{PARS}(\mathbf{S}, \kappa, \mathbf{V}, \tau)$ and $(\mathbf{M}, p)$ has joint density

$$g(\mathbf{M}, p|\mathbf{S}, \kappa, \mathbf{V}, \tau) = f(\mathbf{M}|p, \mathbf{S}, \kappa, \mathbf{V}) D(p|\tau)$$

  for

$$f(\mathbf{M}|p, \mathbf{S}, \kappa, \mathbf{V}) = f_C(\mathbf{T}'(\mathbf{V}, p)\mathbf{S}'\mathbf{M}|\mathbf{I}, \kappa)$$

  where $f_C(\mathbf{o}|\mathbf{S}, \kappa)$ is as before the $\text{UARS}(\mathbf{S}, \kappa)$ density
- In principle, one could integrate p out of $g(\mathbf{M}, p|\mathbf{S}, \kappa, \mathbf{V}, \tau)$ to get a density for the $\text{PARS}(\mathbf{S}, \kappa, \mathbf{V}, \tau)$ model
  - It is neither convenient nor necessary to do so in order to enable Bayes inference
  - Rather, one can treat p as an unobserved latent variable in MCMC

Priors, Posterior, and MCMC for 1-Sample PARS Data

- Suppose $\mathbf{m}_1, \ldots, \mathbf{m}_n \sim^{iid} \text{PARS}(\mathbf{S}, \kappa, \mathbf{V}, \tau)$ with corresponding (unobservable) spins $p_1, \ldots, p_n$
- Set independent priors on $\mathbf{S}, \kappa, \mathbf{V}$, and τ

$$\mathbf{S} \sim \mu, \mathbf{V} \text{ uniform on the unit sphere}, \kappa \sim J_C, \text{ and } \tau \sim J_D$$

- The posterior density for the *parameters* $\mathbf{S}$, κ, $\mathbf{V}$, and τ, *and the unobservable spins* $p_1, \ldots, p_n$ is proportional to

$$h_1(\mathbf{S})h_2(\boldsymbol{\kappa})h_3(\tau)h_4(\mathbf{V})\prod_{i=1}^{n} g(\mathbf{m}_i, p_i|\mathbf{S}, \kappa, \mathbf{V}, \tau)$$

- M-H-within-Gibbs can proceed much as before ... proposals for $\mathbf{V}$ are made using Fisher distributions on 3-d unit vectors

Inference

- Inference for $\mathbf{S}, \kappa$, and τ from MCMC iterates is exactly as for UARS models
- Note that PARS$(\mathbf{S}, \kappa, -\mathbf{V}, \tau)$ =PARS$(\mathbf{S}, \kappa, \mathbf{V}, \tau)$ and so credible sets for $\mathbf{V}$ need to be symmetric wrt multiplication by -1
- We make "pairs of cones about a line" credible sets for $\mathbf{V}$
  - For $\mathbf{V}^1, \ldots, \mathbf{V}^N$ the posterior draws, suppose $\hat{\mathbf{V}}$ maximizes $\sum_{i=1}^{N} \left|\hat{\mathbf{V}}'\mathbf{V}^i\right|$
  - A point estimate for $\mathbf{V}$ is the pair $\pm\hat{\mathbf{V}}$
  - For $\gamma^i \in [0, \pi]$ the positive angle between $\pm\hat{\mathbf{V}}$ and $\mathbf{V}^i$ and $\gamma_{.95}$ the .95 quantile of these, cones of constant angle $\gamma_{.95}$ around $\pm\hat{\mathbf{V}}$ bound a 95% credible region for $\mathbf{V}$

Credible Set for the Preferred Axis

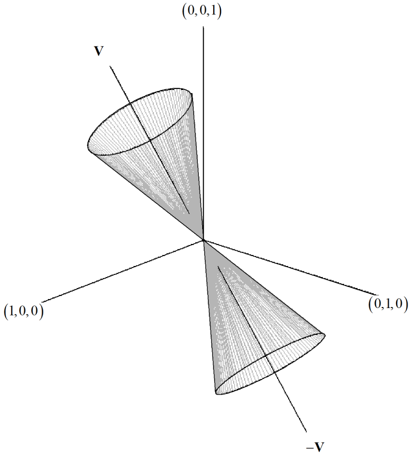

Figure: Credible set for a preferred axis **V**.

Return to the Motivating Application

- EBSD machines return what *appear to be* rotation/orientation matrices (usually in the form of 3 "Euler angles") but (for example for the case of cubic crystals) all that is really known is how corners and edges of the crystal are aligned relative to the world coordinate system (there is no coordinate system on a crystal !!!)
- Sensible definition of "orientation" in this context?
- Probability modeling?
- Inference?

Modeling

- A cubic crystal orientation can be thought of as a child's "jack" (for which there are 24 different possible right-hand coordinate systems, each specifying a different rotation of the world coordinate system)
- For a rotation matrix $\mathbf{O} = (\mathbf{x}, \mathbf{y}, \mathbf{z})$ where $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$ are the 3 columns of $\mathbf{O}$, define the corresponding equivalence class of orientations as

$$
\begin{aligned}
[\mathbf{O}] \;=\; \{ & (\mathbf{x}, \mathbf{y}, \mathbf{z}), (\mathbf{y}, -\mathbf{x}, \mathbf{z}), (-\mathbf{x}, -\mathbf{y}, \mathbf{z}), (-\mathbf{y}, \mathbf{x}, \mathbf{z}), (-\mathbf{z}, \mathbf{y}, \mathbf{x}), \\
& (\mathbf{y}, \mathbf{z}, \mathbf{x}), (\mathbf{z}, -\mathbf{y}, \mathbf{x}), (-\mathbf{y}, -\mathbf{z}, \mathbf{x}), (\mathbf{x}, -\mathbf{z}, \mathbf{y}), (-\mathbf{z}, -\mathbf{x}, \mathbf{y}), \\
& (-\mathbf{x}, \mathbf{z}, \mathbf{y}), (\mathbf{z}, \mathbf{x}, \mathbf{y}), (\mathbf{z}, \mathbf{y}, -\mathbf{x}), (\mathbf{y}, -\mathbf{z}, -\mathbf{x}), (-\mathbf{z}, -\mathbf{y}, -\mathbf{x}), \\
& (-\mathbf{y}, \mathbf{z}, -\mathbf{x}), (\mathbf{x}, \mathbf{z}, -\mathbf{y}), (\mathbf{z}, -\mathbf{x}, -\mathbf{y}), (-\mathbf{x}, -\mathbf{z}, -\mathbf{y}), (-\mathbf{z}, \mathbf{x}, -\mathbf{y}), \\
& (\mathbf{x}, -\mathbf{y}, -\mathbf{z}), (-\mathbf{y}, -\mathbf{x}, -\mathbf{z}), (-\mathbf{x}, \mathbf{y}, -\mathbf{z}), (\mathbf{y}, \mathbf{x}, -\mathbf{z}) \}
\end{aligned}
$$

- A UARS_C $(\mathbf{S}, \kappa)$ model for $\mathbf{O}$ induces a parametric model on $[\mathbf{O}]$'s with parameters $[\mathbf{S}]$ and κ

Bayes Inference for Equivalence Classes of Orientations

- Recognizing real EBSD "data" to be $[\mathbf{O}]_i$'s and extending the methods presented here to this context, we've developed properly calibrated one-sample methods of parametric inference
- κ describes material "texture"
- "Sets of cones" credible sets for $[\mathbf{S}]$ are possible

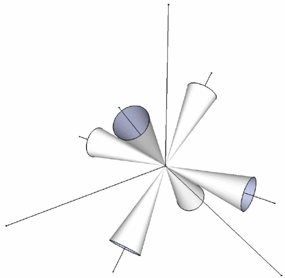

Figure: A credible set for $[\mathbf{S}]$.

Extensions and Other Types of Applications

- Model-based clustering (that takes account of spatial proximity of observations from a single grain) is under development
- Biomechanical modeling and inference (see BNV 2012 *JABES*)
- Kinematics modeling and inference
- Biology (protein folding, for example) modeling and inference

References

1. Bayes one-sample and one-way random effects analyses for 3-d orientations with application to materials science. *Bayesian Analysis*, 2009, Vol. 4, No. 3, pp. 607 - 630, DOI:10.1214/09-BA423. (Bingham, Vardeman, and Nordman)
2. Modeling and inference for measured crystal orientations and a tractable class of symmetric distributions for rotations in 3 dimensions, *Journal of the American Statistical Association*, 2009, Vol. 104, No. 488, pp. 1385-1397, DOI:10.1198/jasa.2009.ap08741. (Bingham, Nordman, and Vardeman)
3. Uniformly hyper-efficient Bayes inference in a class of non-regular problems. *The American Statistician*, 2009, Vol. 63, No. 3, pp. 234-238, DOI:10.1198/tast.2009.08170. (Nordman, Vardeman, and Bingham)
4. Finite sample investigation of likelihood and Bayes inference for the symmetric von Mises-Fisher distribution, *Computational Statistics and Data Analysis*, 2010, Vol. 54, No.5, pp. 1317-1327, DOI:10.1016/j.csda.2009.11.020. (Bingham, Nordman, and Vardeman)

References

5. Bayes inference for a new class of nonsymmetric distributions for 3-dimensional rotations. *Journal of Agricultural, Biological, and Environmental Statistics*, 2012, Vol. 17, No. 4, pp. 527-543, DOI: 10.1007/s13253-012-0107-9. (Bingham, Nordman, and Vardeman)
6. One-sample Bayes inference for existing symmetric distributions on 3-d rotations. *Computational Statistics and Data Analysis*, 2014, Vol. 71, pp. 520-529, DOI:10.1016/j.csda.2013.02.004. (Qiu, Nordman, and Vardeman)
7. A wrapped trivariate normal distribution for 3-D rotations and Bayes inference. *Statistica Sinica*, 2014, Vol. 24, No. 2, pp. 897-917, DOI:10.5705/ss.2011.235. (Qiu, Nordman, and Vardeman)
8. One-sample Bayes inference for a new family of distributions on equivalence classes of 3-D orientations with applications to materials science. To appear in *Technometrics*. (Du, Nordman, and Vardeman)

Thanks

Thanks for your kind attention!

Questions/comments?

Simple UARS Distributional Facts

The constructive definition of the UARS class enables many simple and potentially useful properties to be identified ... among these are:

- If $\mathbf{O} \sim \text{UARS}(\mathbf{I}, \kappa)$, then $\mathbf{O}^T \sim \text{UARS}(\mathbf{I}, \kappa)$
- If $\mathbf{O} \sim \text{UARS}(\mathbf{I}, \kappa)$, then $\mathbf{S} \cdot \mathbf{O} \cdot \mathbf{S}^T \sim \text{UARS}(\mathbf{I}, \kappa)$ for any 3×3 rotation matrix $\mathbf{S}$
- If $\mathbf{O} \sim \text{UARS}(\mathbf{I}, \kappa)$, then $\mathbf{S} \cdot \mathbf{O}$ and $\mathbf{O} \cdot \mathbf{S} \sim \text{UARS}(\mathbf{S}, \kappa)$
- If $\mathbf{O} \sim \text{UARS}(\mathbf{S}, \kappa)$, then $\mathbf{O}^T \sim \text{UARS}(\mathbf{S}^T, \kappa)$
- If $\mathbf{O} \sim \text{UARS}(\mathbf{S}, \kappa)$, then $\mathbf{S}^T \cdot \mathbf{O}$ and $\mathbf{O} \cdot \mathbf{S}^T \sim \text{UARS}(\mathbf{I}, \kappa)$
- Suppose $\mathbf{O} \sim \text{UARS}(\mathbf{I}, \kappa)$ and $\mathbf{O} = (\mathbf{X}\ \mathbf{Y}\ \mathbf{Z})$, where $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ are the three columns of $\mathbf{O}$. Let $P_{\mathbf{X}}$ be the spherical distribution of $\mathbf{X}$ about $(1, 0, 0)^T$, $P_{\mathbf{Y}}$ be the spherical distribution of $\mathbf{Y}$ about $(0, 1, 0)^T$, and $P_{\mathbf{Z}}$ be the spherical distribution of $\mathbf{Z}$ about $(0, 0, 1)^T$. Then $P_{\mathbf{X}} = P_{\mathbf{Y}} = P_{\mathbf{Z}}$.

UARS One-Sample Likelihood and Quasi-Likelihood Inference Generalities

Extrapolating from our experience with the (regular) Matrix Fisher case and the (non-regular) vM-UARS case (where $C(r|\kappa)$ is symmetric von Mises on $(-\pi,\pi])$ it seems likely that

- For **regular versions** of the UARS class, LRT's and Wald tests will work as one would expect, and limiting results will be relevant for fairly small sample sizes (convergence to limiting distributions seemingly faster for LRT statistics than for Wald statistics)

- For **non-regular versions** of the UARS class, one may operate with a quasi-likelihood

$$Q_n(\mathbf{S},\kappa) = \prod_{i=1}^{n} \left(\frac{3 - tr(\mathbf{S}^T \mathbf{o}_i)}{2} \right) f_C(\mathbf{o}_i|\mathbf{S},\kappa)$$

  and get limiting results (for quasi-LRT and quasi-Wald test statistics) qualitatively like those for regular cases

Angle-Axis Form

$$\mathbf{T}(\mathbf{U}, r) = \mathbf{U}\mathbf{U}' + (\mathbf{I}_{3\times 3} - \mathbf{U}\mathbf{U}')\cos r + \begin{pmatrix} 0 & -u_3 & u_2 \\ u_3 & 0 & -u_1 \\ -u_2 & u_1 & 0 \end{pmatrix} \sin r \in \Omega$$