A Bayesian Spatial Model for Predicting the Location of Head Impacts

Michael T. Lawson

PhD Student, UNC-Chapel Hill Dept. of Biostatistics

August 2, 2016

Outline

Acknowledgments

- Daniel Hernandez-Stumpfhauser (UNC-Chapel Hill, Dept. of Biostatistics)
- Amy Herring (UNC-Chapel Hill, Dept. of Biostatistics)
- Gunter Siegmund (MEA Forensic; University of British Columbia)
- Steve Marshall, Jason Mihalik, Kevin Guskiewicz (UNC-Chapel Hill, Matthew Gfeller Center)
- NIEHS grant T32ES007018

Overview of Motivating Experiment

- Overall goal of experiment: Use the output of an imperfect helmet-based accelerometer device to predict the true location of head impacts.
- A helmet with the test device is fitted around a sensor-filled headform attached to a neckform. The researcher sets a kinetic striking device up in prespecified locations, then hits the helmet at prespecified speeds.
- There are 12 impact locations (figure on next slide) and 5 speeds per location, each replicated several times.
- The headform provides gold standard measurements; the device's output is known to be flawed (see Siegmund et al., *Annals of Biomedical Engineering* 2016).

Figure 1: 12 Impact Locations

Device Output

- The location of a head impact is of particular interest for studying the biomechanics of head impacts, and may be useful for diagnosing and treating disease in the future.
- The helmet device outputs the direction (unit vector in 3D-space) and magnitude (scalar) of the impact's peak linear acceleration (PLA).
- However, the magnitude and direction come from *different* vectors. As such, we cannot work with the two as a single object, meaning we must predict direction on the unit vector scale.
- Note: we can represent any unit vector in 3-space as a pair of angles (θ, ϕ) using the conventions of spherical coordinates; we will use this in our plots.

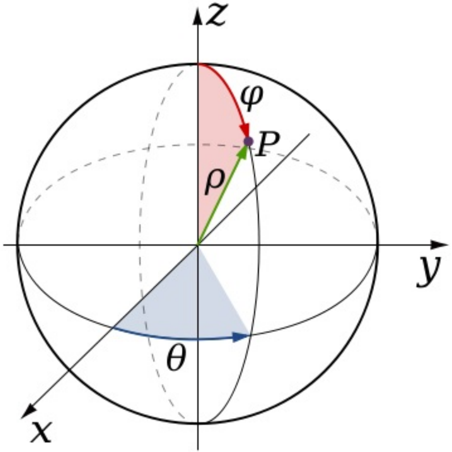

Figure 2: Spherical Coordinates

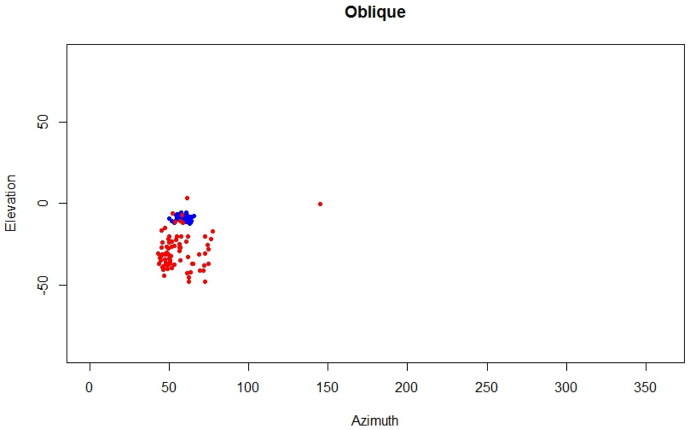

Figure 3: Device's Imperfect Output (Oblique)

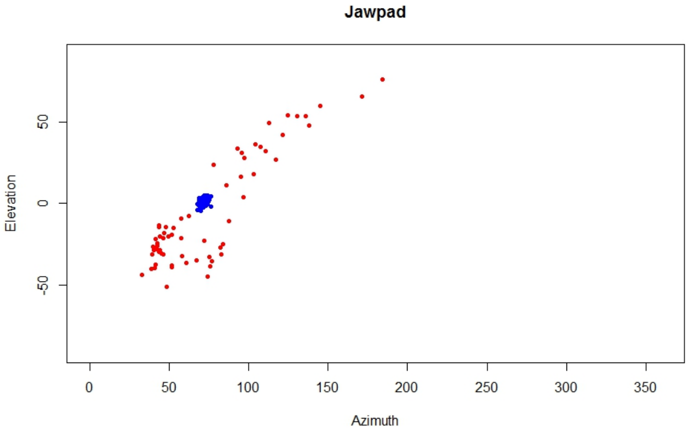

Figure 4: Device's Imperfect Output (Jawpad)

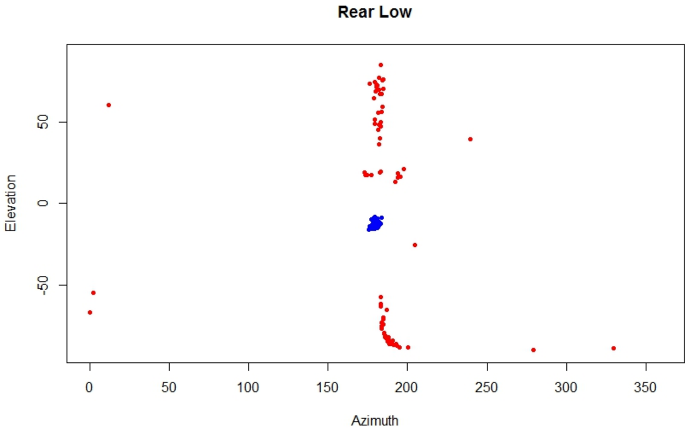

Figure 5: Device's Imperfect Output (Rear Low)

Data Structure

- Our predictor and outcome are both unit vectors in 3-dimensional space. Let U come from the gold standard headform, and let V come from the device.
- Let i index the impact location, $i = 1, ..., I$, and j index observations at each impact location, $j = 1, ..., n_i$.
- Our data consist of $N = \sum_{i=1}^{I} n_i$ pairs of unit vectors, (U_{ij}, V_{ij}).

Modeling Unit Vectors

- To model unit vectors, we use the projected normal distribution.
- If $X \sim \mathcal{N}_3(\mu, \Sigma)$, then $\dfrac{X}{\|X\|} \sim \mathcal{PN}_3(\mu, \Sigma)$.
- The projected normal distribution has a problem with identifiability: if $X/\|X\| = U$, then $rX/\|rX\| = U$ for any $r > 0$. To combat this issue, we "anchor" the covariance matrix Σ by setting its bottom-rightmost element equal to 1.
- Note that if U is any projected normal unit vector and r is the nonnegative length of U, we can compose $rU = X$, where X is Gaussian. We take advantage of this in model building.

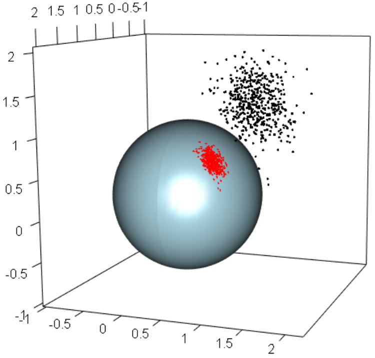

Figure 6: Visualizing $\mathcal{PN}_3$

Likelihood

Let the observed pairs of unit vectors (U_{ij}, V_{ij}) have the pair of latent (unobserved) lengths (r_{ij}, ρ_{ij}) that makes the pair bivariate Gaussian, $(r_{ij}U_{ij}, \rho_{ij}V_{ij}) = (X_{ij}, Y_{ij})$. Then

$$
\begin{aligned}
p(r_{ij}, U_{ij}, &\rho_{ij}, V_{ij} | \text{location} = i, \mu_{Xi}, \mu_{Yi}, \beta_i, \Sigma_{Xi}) \\
=& r_{ij}^2 \mathcal{N}_3(x_{ij}; \mu_{Xi} + \beta_i y_{ij}, \Sigma_{Xi}) \rho_{ij}^2 \mathcal{N}_3(y_{ij}; \mu_{Yi}, I_{3\times 3}) \\
=& r_{ij}^2 \rho_{ij}^2 \mathcal{N}_6 \left(\begin{pmatrix} x_{ij} \\ y_{ij} \end{pmatrix}; \begin{pmatrix} \mu_{Xi} + \beta_i \mu_{Yi} \\ \mu_{Yi} \end{pmatrix}, \begin{pmatrix} \Sigma_{Xi} + \beta_i \beta_i^T & \beta_i \\ \beta_i^T & I_{3\times 3} \end{pmatrix} \right)
\end{aligned}
$$

β_i is a diagonal 3×3 matrix that quantifies the linear dependence between X and Y in location i.

Parameterization of Covariance

In general, we will parameterize Σ_{Xi} as

$$\Sigma_{Xi} = \begin{pmatrix} \sigma^2_{Xi1} + \phi^T_{Xi1}\Lambda_{Xi}\phi_{Xi1} & \phi^T_{Xi}\Lambda_{Xi} \\ \Lambda_{Xi}\phi_{Xi1} & \Lambda_{Xi} \end{pmatrix}$$

where the 2×2 matrix Λ_{Xi} is

$$\Lambda_{Xi} = \begin{pmatrix} \sigma^2_{Xi2} + \phi^2_{Xi2} & \phi_{Xi2} \\ \phi_{Xi2} & 1 \end{pmatrix}.$$

Some Notation

- Note that each μ_{Xi}, $i = 1, ..., I$, is a 3×1 vector. Take the first component of each μ_{Xi} and form $\mu_X^{(1)} = (\mu_{X1}^{(1)}, \mu_{X2}^{(1)}, ..., \mu_{XI}^{(1)})$.
- Similarly form $\mu_X^{(2)}$, $\mu_X^{(3)}$, $\mu_Y^{(1)}$, $\mu_Y^{(2)}$, and $\mu_Y^{(3)}$.
- Construct $\phi_X^{(1)}$, $\phi_X^{(2)}$, $\phi_X^{(3)}$, $\beta^{(1)}$, $\beta^{(2)}$, and $\beta^{(3)}$ in the same way.
- Let $\frac{1}{\sigma^2_{Xi1}} = (z_{Xi}^{(1)})^2$ and $\frac{1}{\sigma^2_{Xi2}} = (z_{Xi}^{(2)})^2$. Define $z_X^{(1)}$ and $z_X^{(2)}$ analogously.
- Each impact location i from the experiment has a fixed location on the sphere. Let η_i denote this location.

Priors

For $s = 1, 2, 3$ and $t = 1, 2$,

$$\mu_X^{(s)} \sim GP(0, K_{\mu_X}(\eta, \eta'))$$
$$\mu_Y^{(s)} \sim GP(0, K_{\mu_Y}(\eta, \eta'))$$
$$\beta^{(s)} \sim GP(0, K_{\beta}(\eta, \eta'))$$
$$\phi_X^{(s)} \sim GP(0, K_{\phi_X}(\eta, \eta'))$$
$$z_X^{(t)} \sim GP(0, K_{z_X}(\eta, \eta'))$$

where K is a squared exponential covariance function. Note that similar quantities (e.g. $\mu_X^{(1)}$ and $\mu_X^{(2)}$) have Gaussian process priors that have identical parameters but are independent.

Fitting

- We can use slice sampling to sample the latent lengths r and ρ.
- Conditional on the latent lengths, the full conditionals for the μ_X, μ_Y, β, and ϕ_X are available in closed form.
- The full conditionals for the z_X are not available in closed form.
- As such, we employ Hamiltonian Monte Carlo sampling embedded within a Gibbs sampler.

Prediction

- One complication: the experiment does not touch the right hemisphere of the helmet. There are 8 locations in the left hemisphere and 4 directly down the center of the helmet.
- Symmetry assumption: if $M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$,

$$p(U|\mu_{Xi,right}, \Sigma_{Xi,right}) = \mathcal{PN}_3(U; M\mu_{Xi,left}, M\Sigma_{Xi,left}M^T)$$

- To obtain a predictive distribution on the sphere, $p(U_{ij,new}|V_{ij,new}, data)$. To do this, we need a prior distribution on the impact locations.
- For this analysis, we reflected the 6 left-hemisphere locations to the right hemisphere, then placed a uniform discrete prior on the 16 resulting locations.

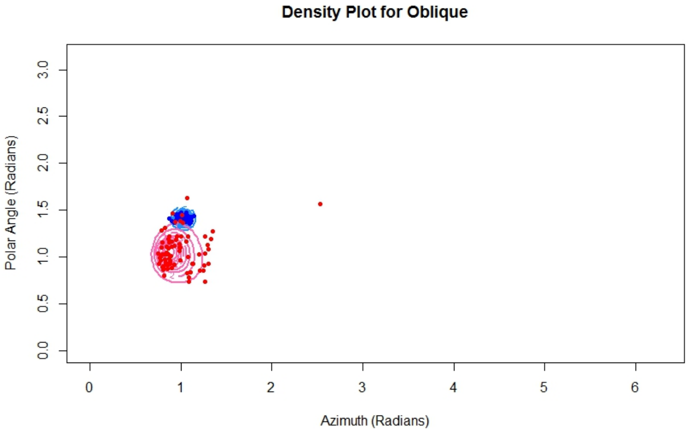

Figure 7: Regression Density Estimation (Oblique)

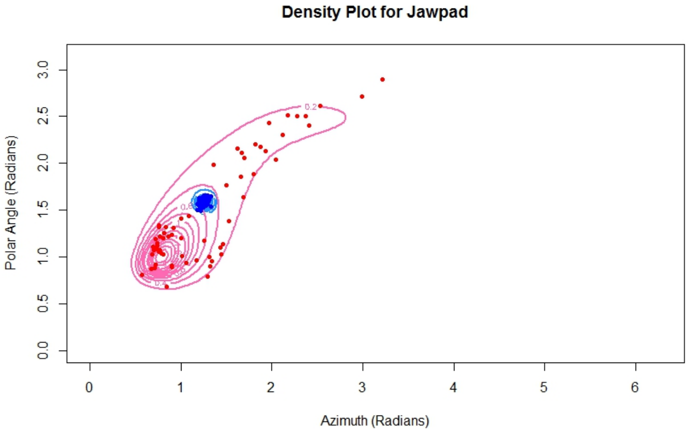

Figure 8: Regression Density Estimation (Jawpad)

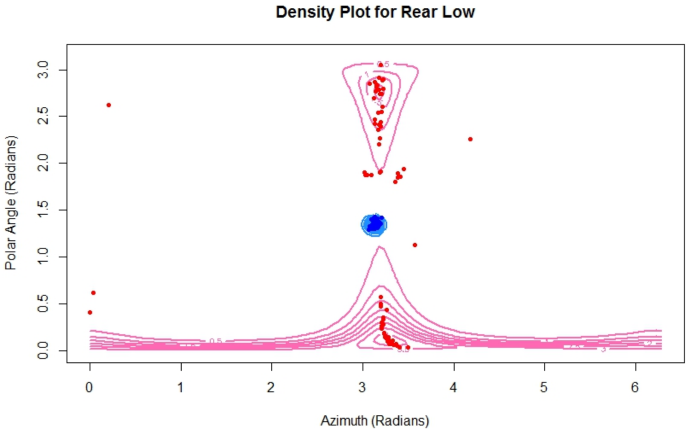

Figure 9: Regression Density Estimation (Rear Low)

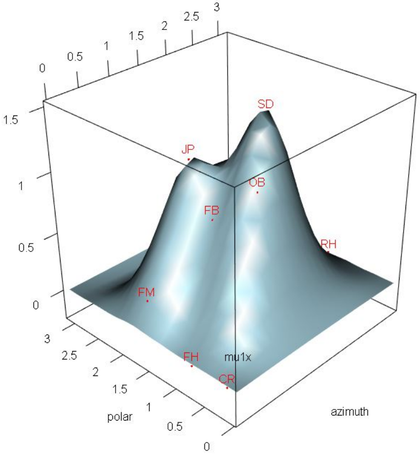

Figure 10: Regression Surface: μ_{X1}

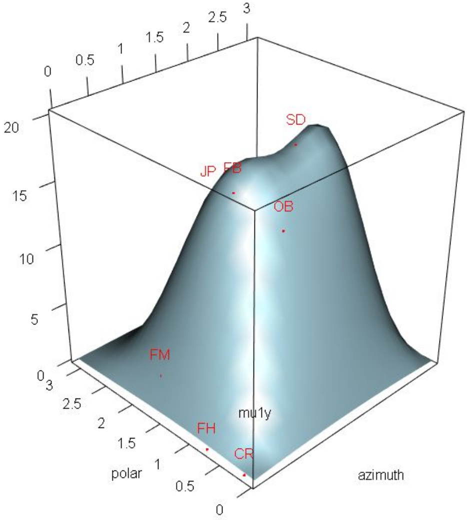

Figure 11: Regression Surface: μ_{Y1}

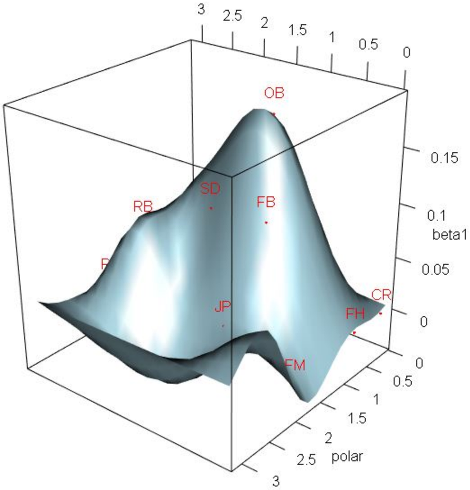

Figure 12: Regression Surface: β_1

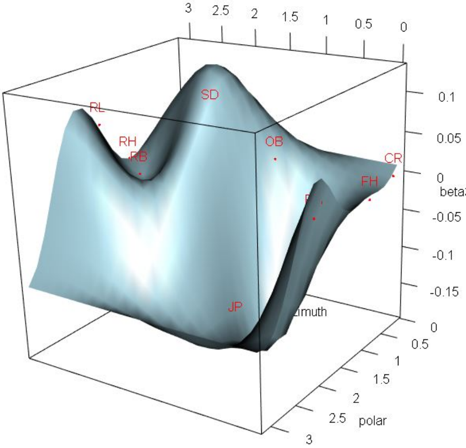

Figure 13: Regression Surface: β_3

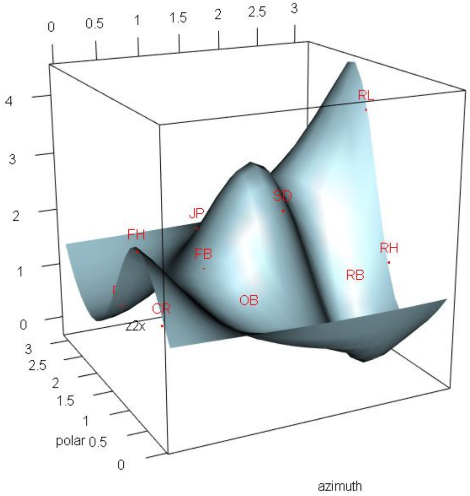

Figure 14: Regression Surface: $z_X^{(1)}$

Next Steps

- Updates to the experiment
  - Both hemispheres
  - Many “small” vs. few “large” locations
  - Target locations to regions of high uncertainty
- Updates to the model
  - Add covariates
- Adapt to real-world data
  - Calibrate using observed impact location distribution
  - Employ in real-time