Fast approximate Bayesian inference of HIV indicators

PhD Student Presentations & Networking Event, Alan Turing Institute

Adam Howes

Imperial College London

May 2023

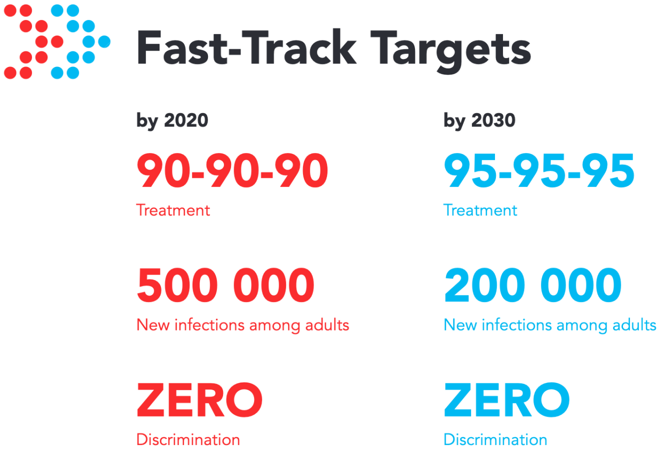

Figure 1: Ambitious targets required to end the AIDS epidemic as a public health threat by 2030.

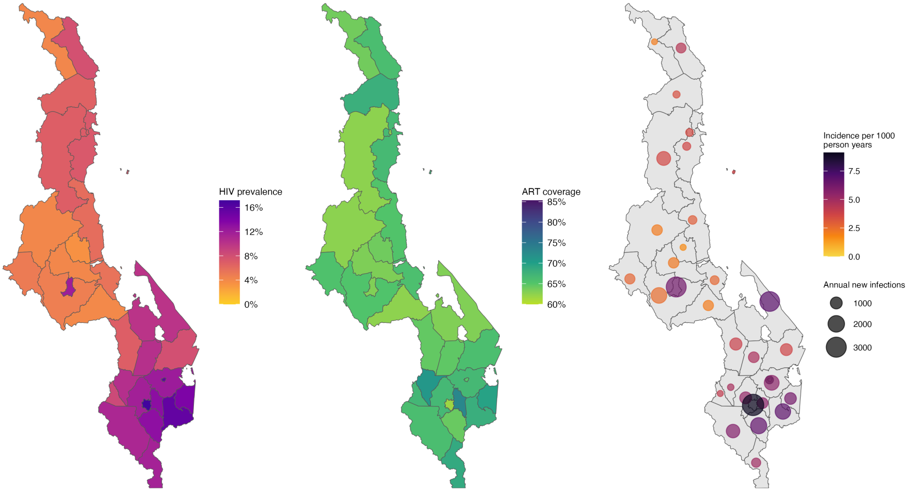

Figure 2: More effective intervention based on granular estimates of HIV indicators. One size does not fit all!

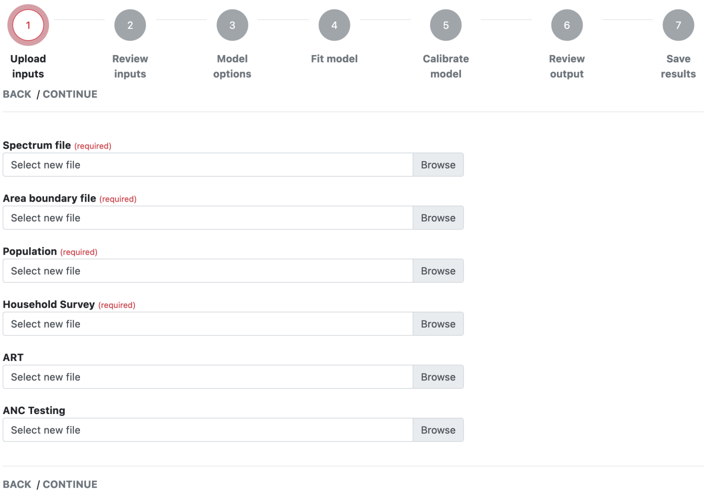

Figure 3: Generation of estimates by country teams strengthens data quality, use and ownership. User interface from https://naomi.unaids.org/.

Want

Fast approximate Bayesian inference for a complex, spatiotemporal, evidence synthesis model

Strategy

1. Marginal Laplace approximation
2. Adaptive Gauss-Hermite quadrature
3. Principal components analysis

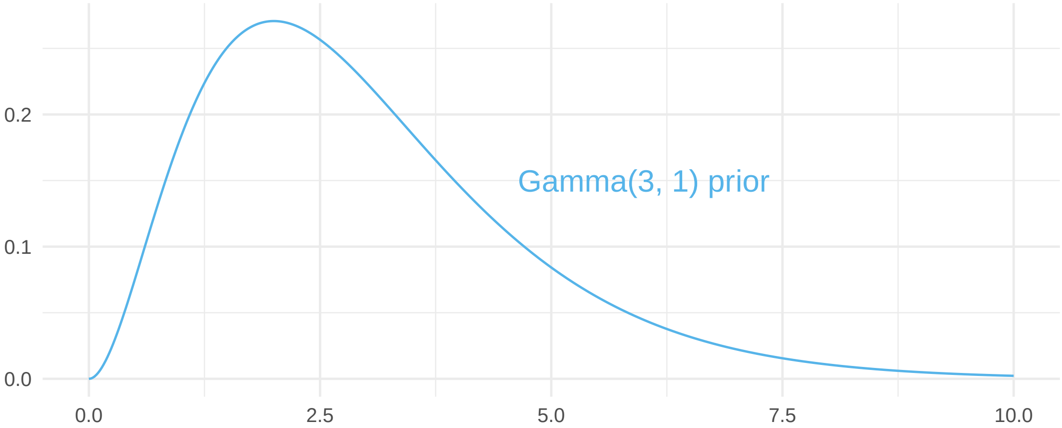

Figure 4: A Gamma prior with $a = 3$ and $b = 1$.

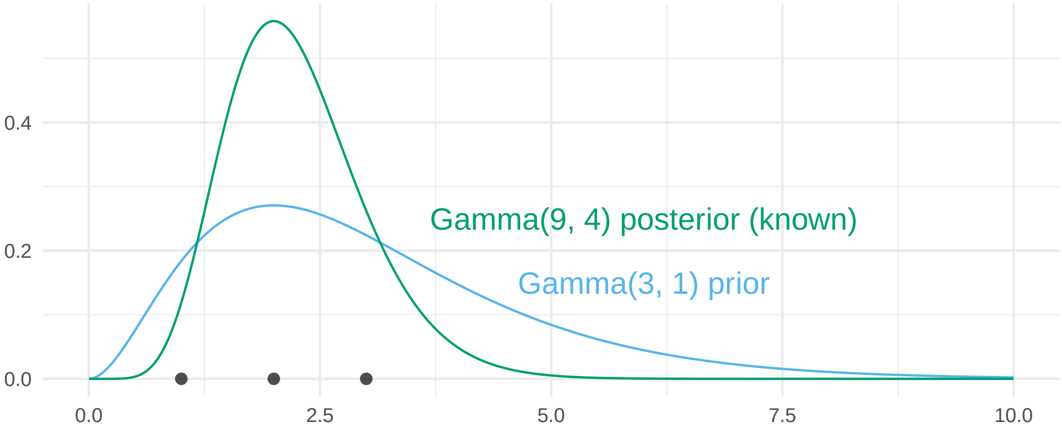

Figure 5: Draw 3 points from Poisson(3), then compute the posterior.

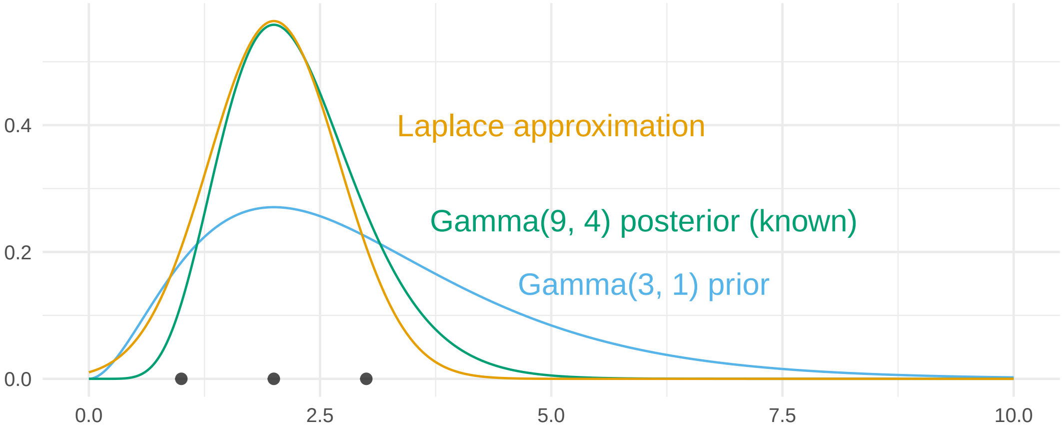

Figure 6: The Laplace approximation in this case is good near the mode but not in the tails.

- For Malawi, the model has 491 parameters
- 467 have a joint Gaussian prior: call them the latent field x
- 24 are not Gaussian: call them hyperparameters θ
- Use the Laplace approximation only for the latent field marginal posterior!

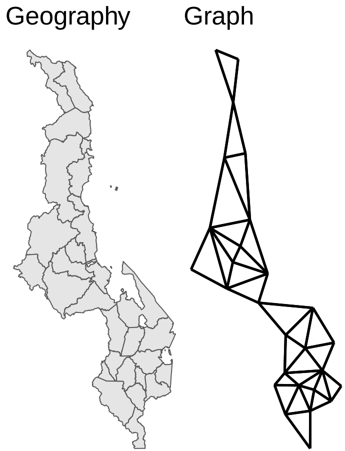

Figure 7: Spatial random effects $\phi_i \mid \phi_{-i} \sim \mathcal{N}\left(\frac{1}{n_{\delta i}} \sum_{j:j\sim i} \phi_j, \frac{1}{n_{\delta i}\tau_\phi}\right)$ are included in the latent field. We assume that neighbouring districts are similar: first law of geography.

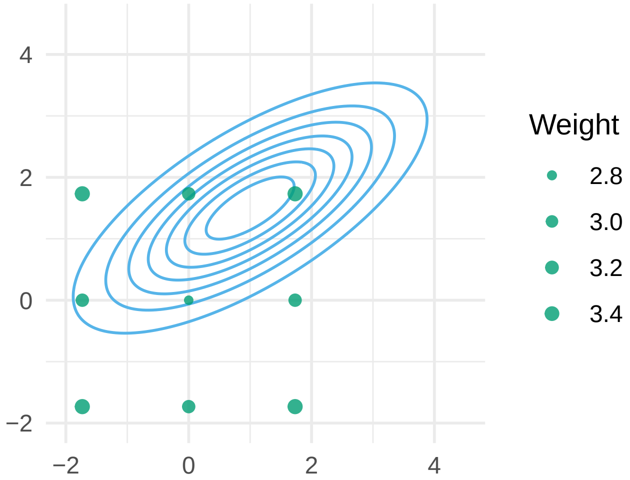

Figure 8: Unadapted Gauss-Hermite points in two dimensions with $k = 3$.

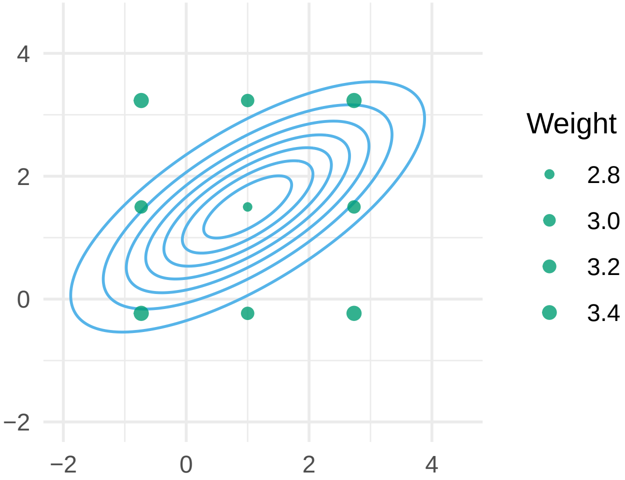

Figure 9: Add the mean $z + \hat{\theta}$.

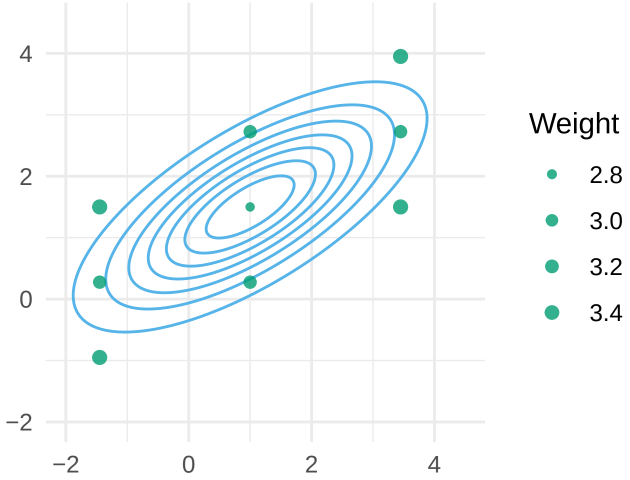

Figure 10: First option: rotate by the lower Cholesky $Lz + \hat{\theta}$.

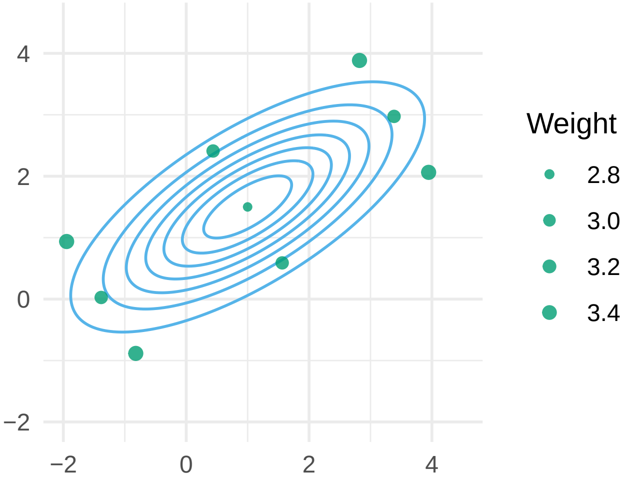

Figure 11: Second option: rotate using the eigendecomposition $E\Lambda^{1/2}z + \hat{\theta}$.

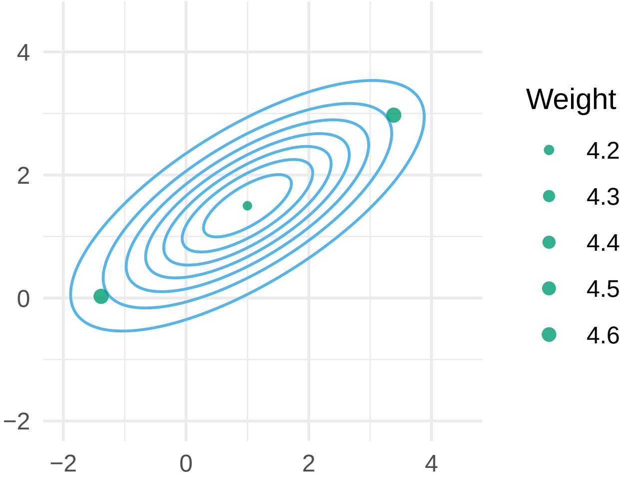

Figure 12: Now keeping only the first principal component, $s = 1$.

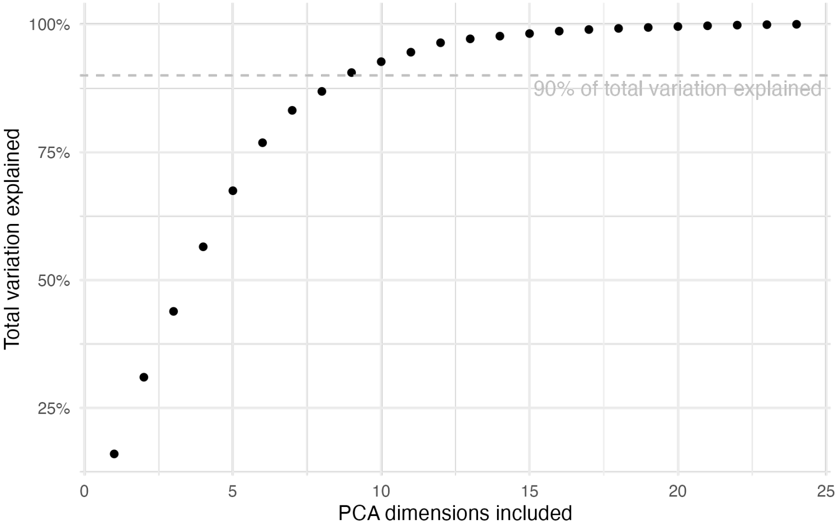

Figure 13: Scree plot suggests 10 or so dimensions is enough. We use $s = 8$ to avoid long computation times.

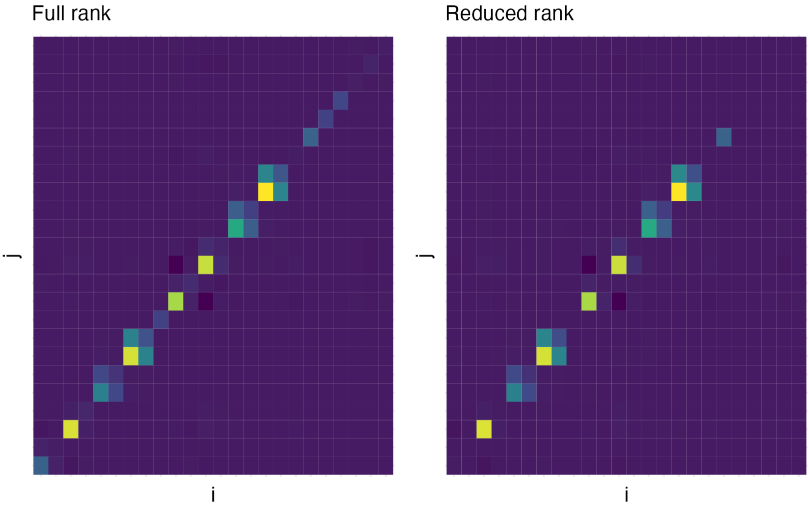

Figure 14: With 8 dimenions, the hyperparameter covariance matrix is accurately reproduced.

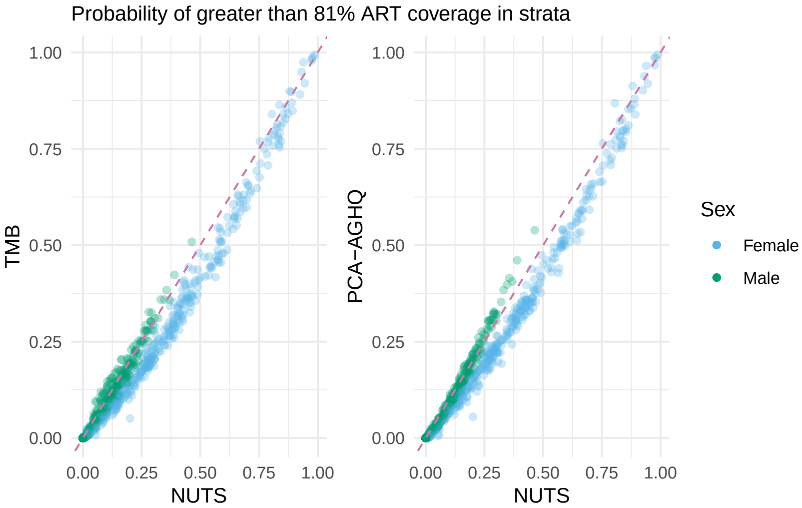

Figure 15: Reduced RMSE by 10 percent, but still a work in progress! Also seeing small improvements using Kolmogorov-Smirnov tests, Pareto-smoothed importance sampling, maximum mean discrepancy.

Thanks for listening!

- For more about Naomi, see Eaton et al. (2021)
- Joint work with Alex Stringer (Waterloo), Seth Flaxman (Oxford), and Jeff Eaton (Harvard, Imperial)
- For more about me, see athowes.github.io/about

MRC Centre for Global Infectious Disease Analysis
Imperial College London
MACHINE LEARNING & GLOBAL HEALTH NETWORK

References I

Eaton, Jeffrey W., Laura Dwyer-Lindgren, Steve Gutreuter, Megan O'Driscoll, Oliver Stevens, Sumali Bajaj, Rob Ashton, et al. 2021. "Naomi: A New Modelling Tool for Estimating HIV Epidemic Indicators at the District Level in Sub-Saharan Africa." *Journal of the International AIDS Society* 24 (S5): e25788.