International Society for Bayesian Analysis, 9^{th} World Meeting,
Hamilton Island, Australia, 2008.

Pooling information across matrix decompositions

Peter Hoff[1*]

[1] University of Washington, Seattle, U.S.A.
[*] hoff@stat.washington.edu

One approach to summarizing relational and other matrix-valued data is with low-rank matrix approximations. For example, the variation among the entries of a symmetric $n \times n$ data matrix Y is often expressed with the eigenvalue decomposition model $Y \sim U \Lambda U^T + E$, where U is an $n \times k$ orthonormal matrix and Λ is a diagonal matrix. In this work we consider pooling information across multiple such data matrices $Y^{(1)}, \ldots, Y^{(p)}$ for situations in which the common cells across matrices $\{y_{i,j}^{(1)}, \ldots, y_{i,j}^{(p)}\}$ represent repeated or multivariate measurements under a common set of conditions $\{i, j\}$. This is accomplished by estimating the parameters in a model for the variability among the orthonormal eigenvector matrices $U^{(1)}, \ldots, U^{(p)}$ of the p data matrices. The model is based on a variation of the matrix Langevin distribution, for which estimation is accomplished primarily with Gibbs sampling. The methodology is applied to the analysis of multivariate relational data where $\{y_{i,j}^{(1)}, \ldots, y_{i,j}^{(p)}\}$ represent multiple dyadic relations between nodes i and j.