JOHN ORMEROD, **University of Sydney**

Particle Based Collapsed Variational Approximation for Bayesian Linear Model Averaging

While Bayesian model averaging has several desirable properties, it is computationally expensive unless the number of models to be averaged over is small. Typically the number of models to be averaged grows exponentially in the number of covariates and some form of approximation is required. In this paper we explore a novel particle based collapsed variational approximation for Bayesian model averaging. The resulting objective function can be optimized in a highly parallel manner. We explore several different prior specifications which lead to Bayes factors with closed forms. We show empirically that our approach is fast and effective for moderately large problems on several simulated and publicly available datasets, particularly when parallel computing resources are available. An R package is available implementing our approach. (Joint work with Mark Greenaway).