Well-Tempered MCMC Simulations for Population Pharmacokinetic Models

Frederic Y. Bois[1], Nan-Hung Hsieh[2], Wang Gao[3], Weihsueh A. Chiu[2], Brad Reisfeld[4]
[1] Certara UK Limited, Simcyp Division, Sheffield, UK;   [2] Texas A&M University, USA;
[3] Université de Technologie de Compiegne, France;   [4] Colorado State University, USA.

Abstract

We improve the simulated tempering Markov Chain Monte Carlo (TMCMC) algorithm of Geyer & Thompson (1995) to automatize it. We then reap its advantages: faster convergence; sampling from complex multi-modal distributions; estimation of normalization constants giving access to inference about model structure and model choice.

Background

A full Bayesian treatment of complex population PK/PD models gives access to powerful inference, including on model structure. TMCMC can sample efficiently from sharp multi-modal posteriors, providing useful identifiability insight for model simplification; it has assured convergence with a single simulated chain if infinite temperature is reached; it can be used to compute accurate Bayes factors for formal model comparison and choice.

Methods

Adequate scales of temperatures are obtained through a Robbins-Munro process and adaptive optimization. Optimized grids of perks and associated pseudo-priors were used to perform thermodynamic integration, bridging the joint prior and the posterior distributions, for two stylized case studies and two realistic population pharmacokinetic inference problems: One for theophylline (Trembath 1980), with a small pharmacokinetic model; the other one for acetaminophen with a large PBPK model (Zurlinden 2017, Hsieh 2018). For each, we use standard Bayesian multilevel models.

Results

Tempered MCMC sampling creates a smooth distribution landscape easier to explore and sample from.

Efficient mixing of the chain requires setting perk weights (pseudo-priors) iteratively.

**Automatic tempering tested on a simple 1-compartment model**

Results

In the case of a sharply bimodal regression model, the two high posterior density region (in a three dimensional space) were accurately identified, where other algorithms, such as MCMC or Hamiltonian MCMC failed. The two peaks lie on different planes and are marginally separated by 600 times their width.


Sharp bimodal distribution of the slope in a pathologic regression.

In the case of a Gaussian mean, we get a close estimate of the normalization constant needed to compute Bayes factors.

Posterior histograms of a Gaussian mean at perk 0 (left) and perk 1 (right, after *normalization*), with the exact densities (red lines).


The theophylline population PK model with inter- and intra-individual variability is clearly favored compared to a model without intra-individual variability, based on the Bayes factor (equal to 10^{+9}).
The posterior parameter estimates are obtained 8 times faster and are close to those of standard MCMC, except that TMCMC identified a secondary mode of the posterior of a population variance.

Automatic tempering tested on a simple 1-compartment model

For the acetaminophen PBPK model, the posterior distributions of the most sensitive parameters of the acetaminophen PBPK are close than those obtained with MCMC, but obtained 20 times faster due to the improved mixing of the chain mixing and assured convergence.
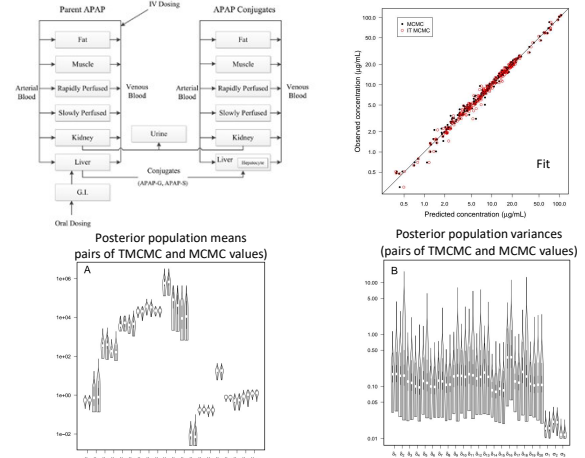
Posterior inference on the acetaminophen PBPK model

Conclusions

- TMCMC has a number of advantages over other MCMC algorithms, if it is implemented efficiently. Our case studies show a full range of its capabilities, which are useful for pharmacokinetic or pharmacodynamic modeling, but also apply generally.

References

- Geyer CJ, Thompson EA (1995) Journal of the American Statistical Association 90:909.
- Trembath PW, Boobis SW (1980) British Journal of Clinical Pharmacology 9:365.
- Hsieh N-H, Reisfeld B, Bois FY, Chiu WA (2018) Frontiers in Pharmacology 9, doi:10.3389/fphar.2018.00588
- Zurlinden TJ, Reisfeld B (2017) European Journal of Drug Metabolism and Pharmacokinetics 42:143, doi:10.1007/s13318-016-0329-2