The Best Performing Monte Carlo Estimator? A likelihood Framework for Compromising between Statistical Efficiency and Computational Efficiency

Xiao-Li Meng, Harvard University

In presenting a Monte Carlo estimator, it is a rather common phenomenon that as soon as the speaker finishes advertising how efficient statistically a proposed estimator is compared to other estimators in the literature, someone in the audience would pose that “nasty” question: “But how about the CPU time?” How to formulate the trade-off between statistical efficiency and computational efficiency has been the most challenging theoretical problem in statistical computing, even though in practice we routinely make ad hoc judgments. The first part of this talk demonstrates that there is a constructive likelihood theory that helps to address this trade-off problem, where the statistical efficiency is captured by a group invariance constraint imposed on the model parameter, which is the baseline measure. [The baseline measure is the model parameter because fundamentally Monte Carlo estimation is about approximating a computationally intractable measure (e.g., Lebesgue measure) by a tractable finite counting measure for which we can enumerate all its possible states (Kong et al, 2003, JRSSB).] The increased statistical efficiency however is achieved at the computational expense of the additional function evaluations needed for computing the maximum likelihood estimator. Hence the size (and the structure) of the group is a suitable indicator of computational efficiency. The second part of this talk demonstrates the practicality of this trade-off formulation in the context of warp bridge sampling (Meng and Schilling, 2002), where the quest for the optimal warp transformation can be adequately addressed as a problem of estimating optimal group constraints (among a class of groups) on the baseline measure. Our key emphasis is that between the two extremes, namely the analytical integration which corresponds to using full group constraints and (the current) vanilla Monte Carlo integration which corresponds to using no/trivial group constraint, there are many more choices of group constraints that can sensibly capture various trade-offs between statistical and computational efficiency.
(Based on joint work with Zhiqiang Tan and David Jones.)