

Riemannian Laplace Approximation with the Fisher Metric

Hanlin Yu¹, Marcelo Hartmann¹, Bernardo Williams¹, Mark Girolami²
and Arto Klami¹

1 University of Helsinki

2 University of Cambridge and The Alan Turing Institute

Abstract:

Laplace's method approximates a target density with a Gaussian distribution at its mode. It is computationally efficient and asymptotically exact for Bayesian inference due to the Bernstein-von Mises theorem, but for complex targets and finite-data posteriors it is often too crude an approximation. A recent generalization of the Laplace Approximation transforms the Gaussian approximation according to a chosen Riemannian geometry providing a richer approximation family, while still retaining computational efficiency. However, as shown here, its properties depend heavily on the chosen metric, indeed the metric adopted in previous work results in approximations that are overly narrow as well as being biased even at the limit of infinite data. We correct this shortcoming by developing the approximation family further, deriving two alternative variants that are exact at the limit of infinite data, extending the theoretical analysis of the method, and demonstrating practical improvements in a range of experiments.