Optimal Control Variates for MCMC

Anand Radhakrishnan, Sean Meyn

Department of Electrical Engineering, University of Florida

Objective:

Minimize asymptotic variance of MCMC estimates using control variates

Introduction

In many applications, we need to evaluate expectations that are difficult to compute:

$$\eta = \int c(x)\pi(x)dx$$

$c\colon \mathbb{R}^\ell \to \mathbb{R}$ is a measurable function, π is a probability density in $\mathbb{R}^\ell$.

Markov-Chain Monte Carlo (MCMC) methods provide numerical algorithms to obtain empirical estimates:

$$\eta(t) = \frac{1}{t}\int_0^t c(\Phi(s))\, ds$$

Φ is a Markov process with steady state distribution π.

Rate of convergence captured by Central Limit Theorem, with asymptotic variance

$$\gamma^2 = \lim_{t\to\infty} \mathsf{E}\left[\left(\frac{1}{\sqrt{t}}\int_0^t (c(\Phi(s)) - \eta) ds\right)^2\right]$$

Langevin Diffusion

Langevin diffusion is an ℓ-dimensional gradient flow with “noise”:

$$d\Phi(t) = -\nabla\mathcal{U}(\Phi(t))\, dt + \sqrt{2}\, dW(t)$$

$\boldsymbol{W} = \{W(t) : t \geq 0\}$ is a standard Brownian motion on $\mathbb{R}^\ell$

Under general conditions, this diffusion is reversible, with unique invariant density, $\pi = e^{-\mathcal{U}+\Lambda}$, where Λ is a normalizing constant.

Differential generator of the Langevin diffusion,

$$\mathcal{D}f := -\nabla\mathcal{U}\cdot f + \Delta f, \qquad f \in C^2$$

Poisson’s equation for the Langevin diffusion:

$$\mathcal{D}h = -\tilde{c}, \qquad \tilde{c} = c - \eta$$

Representation for γ^2 in terms of the solution to Poisson’s equation h:

$$\gamma^2 = 2\langle h\,,\, \tilde{c}\rangle_\pi$$

For Langevin diffusion, on integration by parts, we have

$$\boxed{\gamma^2 = 2\langle h\,,\, \tilde{c}\rangle_\pi = -2\langle h\,,\, \mathcal{D}h\rangle_\pi = 2\langle \nabla h\,,\, \nabla h\rangle_\pi = 2\|\nabla h\|_\pi^2}$$

Control variates and TD-learning

A **control variate** n has zero mean with respect to the density π, i.e. $\int n(x)\pi(x)dx = 0$.

Eg: $n = \mathcal{D}f$, for a “nice” f, i.e. $f \in L_1(\pi)$ and $f \in C^2$.

We look at an improved estimator for η,

$$\eta_\theta(t) := \frac{1}{t}\int_0^t c_\theta(\Phi(s))ds, \qquad \text{where } c_\theta := c + \underbrace{\mathcal{D}h_\theta}_{\text{control variate}}$$

$$h_\theta := \sum_{i=1}^d \theta_i\psi_i, \qquad \text{where } \theta \in \mathbb{R}^d \text{ and } \psi_i \text{ is the } i^{th} \text{basis function.}$$

Ideally, we prefer $\mathcal{D}h$, but h is difficult to solve. Approximating h is hard and is the subject of TD-learning:

$$\min_\theta \|h - h_\theta\|_\pi^2$$

Our goal is to find the “optimal” control variate that minimizes asymptotic variance, i.e.

$$\min_\theta \gamma_\theta^2 := \min_\theta \|\nabla h - \nabla h_\theta\|_\pi^2$$

which is just as hard in general. However for Langevin, we have a simple solution.

Algorithm

$$\theta^* := M_\ell^{-1}b, \qquad M_{ij} := \sum_{k=1}^{d}\langle \partial_k\psi_i\,,\, \partial_k\psi_j\rangle_\pi, \qquad \underbrace{b_i := \langle \psi_i\,,\, \tilde{c}\rangle_\pi}_{\text{Exact solution}}.$$

$$\theta(t) = M^{-1}(t)b(t), \qquad M(t) = M(0) + \int_0^t \nabla\psi(\Phi(s))\,\nabla\psi^T(\Phi(s))ds, \qquad \underbrace{b(t) = \int_0^t \psi(\Phi(s))\,\tilde{c}(\Phi(s))ds.}_{\text{Monte Carlo approximation}}$$

$$\lim_{t\to\infty}\theta(t) = \theta^* \quad \text{with probability one.}$$

Numerical Example

Target density π is a mixture of Gaussians

$$\pi = \sum_{i=m}^{M} w_m p_m \quad \text{where} \sum_{m=1}^{M} w_m = 1, \quad p_m := \mathcal{N}(\mu_m, \sigma_m^2)$$

Eg: $M = 2 \qquad \mu_1 = -1, \mu_2 = 1 \qquad \sigma_1 = \sigma_2 = 0.4 \qquad w_1 = w_2 = 0.5$

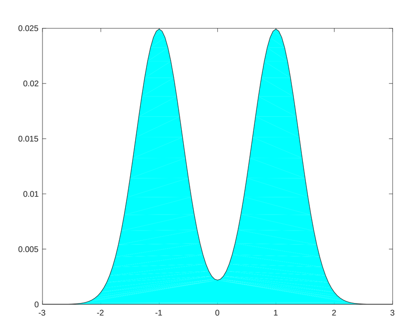

Parameterization: A linear parameterization of the form

$$h_\theta(x) = \sum_{i=1}^{d}\theta_i\psi_i(x)$$

$$\{\psi_i(x) : x \in \mathbb{R},\ 1 \leq i \leq d\} = \{x^k q_m(x) : 0 \leq k \leq d_0,\ m = 1, 2, d = 2d_0\}$$

$$q_m \sim \mathcal{N}(\mu_m, \lambda_m^2) \quad \text{such that } \lambda_m^2 \geq \sigma_m^2.$$

Asymptotic variance reduction

Asymptotic variance reduction with d: Basis 1 with $\lambda_m^2 = \sigma_m^2$ V Basis 2 with $\lambda_m^2 = 2\sigma_m^2$.

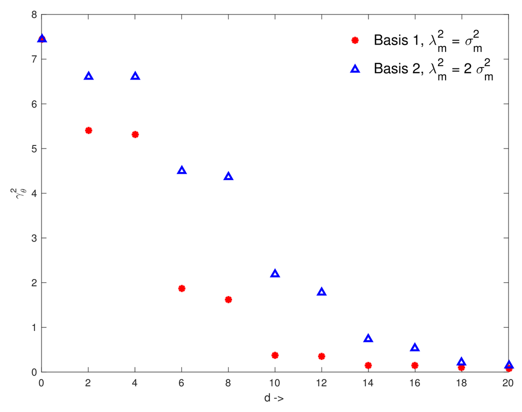

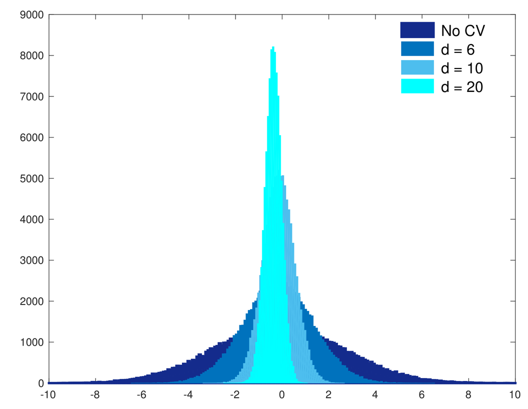

Application to Metropolis-Hastings algorithm

$\mathcal{D}h_\theta$ remains a control variate for the MH algorithm with target density π. Although not “optimal”, reduction in asymptotic variance achieved is similar to Langevin case.

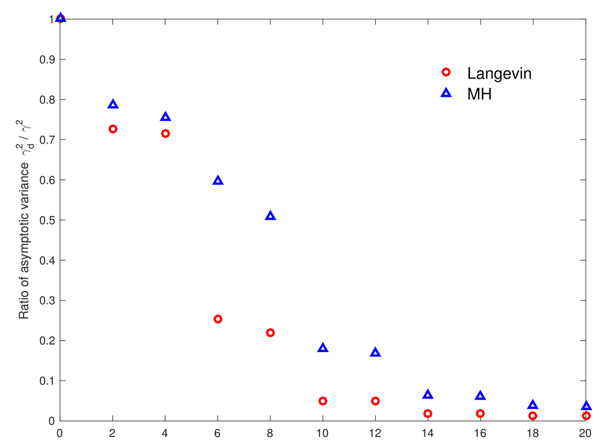

Variance v Asymptotic variance

In prior research [1, 2] the control variate is chosen by minimizing the ordinary variance rather than the asymptotic variance.

Well-motivated only if the samples are i.i.d. The unbiased estimator is based on a function c_ϑ:

$$c_\vartheta = c + \mathcal{D}h_\vartheta\,, \qquad \text{where} \qquad h_\vartheta = \sum_{i=1}^{d}\vartheta_i\psi_i = \vartheta^T\psi$$

Optimal parameter ϑ^* is obtained by minimizing the variance:

$$\vartheta^* := \arg\min_\vartheta \sigma_\vartheta^2 = \arg\min_\vartheta \|c_\vartheta - \eta\|_\pi^2 = -\overline{M}^{-1}\overline{b}$$

$$\overline{M} := \langle \mathcal{D}\psi\,,\, \mathcal{D}\psi\rangle_\pi \quad , \quad \overline{b} := \langle \mathcal{D}\psi\,,\, \tilde{c}\rangle_\pi$$

But, variance minimization $\neq$ asymptotic variance minimization.

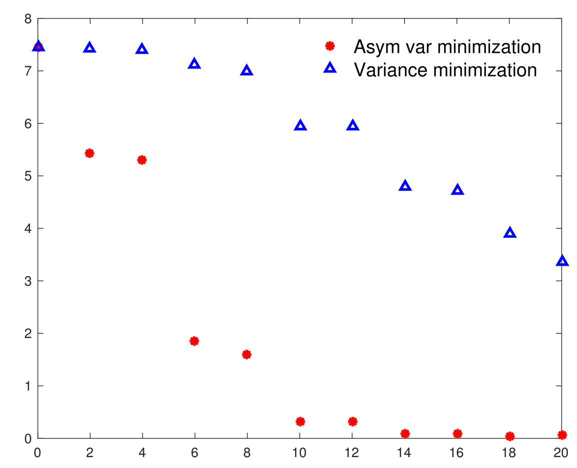

Control variates c_θ^* and c_ϑ^* imply that **“oscillations are good”.**

$$\int_{-\infty}^{0} c_{\theta^*}(x)p(x)\, dx \approx \int_0^\infty c_{\theta^*}(x)p(x)\, dx \approx \eta = 0$$

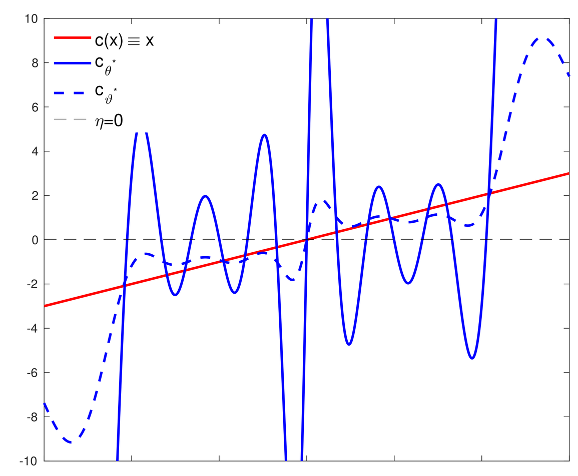

Conclusion

Optimization of asymptotic variance is remarkably simple for the Langevin diffusion.

In parallel research, these techniques are applied to gain approximation for the feedback particle filter.

Ongoing research includes:

- Basis selection for multidimensional models.
- Investigating a basis independent approach using reproducing kernels.

References

[1] Antonietta Mira, Reza Solgi, and Daniele Imparato.
Zero variance markov chain monte carlo for bayesian estimators.
Statistics and Computing, 23(5):653–662, September 2013.

[2] Chris J. Oates, Mark Girolami, and Nicolas Chopin.
Control functionals for monte carlo integration.
Journal of the Royal Statistical Society: Series B (Statistical Methodology), 79(3):695–718, 2017.

Acknowledgements

UF UNIVERSITY of FLORIDA — The Foundation for The Gator Nation

COGNITION & CONTROL IN COMPLEX SYSTEMS