Gradient Estimation for Random Horizon Experiments

Bernd Heidergott

Vrije Universiteit Amsterdam and *Tinbergen Institute*

joint work with

Felisa Vázquez-Abad

University of Melbourne

14 - 18 July, 2008, Sandbjerg Gods

Efficient Monte Carlo: From Variance Reduction to Combinatorial Optimization.

I A brief overview on gradient estimation: What to do with the derivative of a density?

II Random horizon experiments: our metro example

III Differentiating a Radon-Nikodym derivative

IV The Score function and the weak derivative interpretation

V Conclusion

I.1 A Helicopter View on Gradient Estimation

Let X_θ be random variable with Lebesgue density f_θ. Under suitable differentiability properties it holds that

$$\text{IPA} \qquad \mathbb{E}\left[\frac{d}{d\theta} g(X_\theta)\right] = \frac{d}{d\theta}\mathbb{E}[g(X_\theta)] = \begin{cases} \mathbb{E}\left[g(X_\theta)\frac{d}{d\theta}\ln(f_\theta(X_\theta))\right] & \text{score function} \\ c_\theta(\mathbb{E}[g(X_\theta^+)] - \mathbb{E}[g(X_\theta^-)]) & \text{weak derivative}, \end{cases}$$

for suitable random variables X_θ^+ and X_θ^- and constant c_θ.

I.2 The Relation between the Score Function and a Weak Derivative

A standard version of weak derivative can be constructed as follows. Introduce densities

$$f_\theta^+(x) = \frac{1}{c_\theta}\max\left(\frac{d}{d\theta}f_\theta(x), 0\right), \quad f_\theta^-(x) = \frac{1}{c_\theta}\max\left(-\frac{d}{d\theta}f_\theta(x), 0\right),$$

where

$$c_\theta = \int \max\left(\frac{d}{d\theta}f_\theta(x), 0\right) dx.$$

Then, letting $X_\theta^\pm$ have density $f_\theta^\pm$, yields a weak derivative. Note that

$$\frac{d}{d\theta}f_\theta(x) = c_\theta(f_\theta^+(x) - f_\theta^-(x))$$

and

$$\frac{d}{d\theta}\ln(f_\theta(x)) = c_\theta\left(\frac{f_\theta^+(x)}{f_\theta(x)} - \frac{f_\theta^-(x)}{f_\theta(x)}\right).$$

I.3 A Taxonomy of Gradient Estimators

	IPA	Score Function	Weak Derivatives
Single run	+	+	-
Flexibility	- -	++	++
Variance	+	- -	++
Computational burden	+	+	-
Work normalized variance*	+	- -	-

*The work normalized variance is given by the product of the variance and the expected work per run balancing computational effort and estimator variance (Glynn and Whitt 1992).

Gradient Estimation for Random Horizon Experiments

I A brief overview on gradient estimation: What to do with the derivative of a density?

II **Random horizon experiments: our metro example**

III Differentiating a Radon-Nikodym derivative

IV The Score function and the weak derivative interpretation

V Conclusion

II.1 The Motivating Problem

We consider the following model of a station of the Montreal metro system (no fixed timetable only frequencies are given):

- $Y_\theta(n)$ denotes the nth interarrival time of a train on the θ-line where θ is a scaling parameter
- customers arrive according to a Poisson process to the platform from the outside
- there are bulk arrivals generated by the arrival of other trains (giving connection)

We are interested in minimizing the accumulated waiting time of passengers for the θ-line over a fixed period of time.

References are

[F. Vázquez-Abad and L. Zubieta, *DEDS*, 2005] and [B. Heidergott and F. Vzquez-Abad, *TOMACS*, under review]

We remark that in the above papers $Y_\theta(n)$ is assumed to be normally distributed with mean θ and standard deviation $\theta\sigma$ for some $\sigma > 0$.

II.2 Formalizing a Random Horizon Experiment

Let $\{Y_\theta(n) : n \in \mathbb{N}\}$ be a sequence of i.i.d. random variables with density f_θ.

Define the stopping time

$$\tau_\theta = \min\left\{ n : \sum_{i=1}^{n} Y_\theta(i) > T \right\}, \qquad T > 0.$$

Consider a measurable real-valued functional H_T of the process $\{Y_\theta(i) : i \in \mathbb{N}\}$ such that

$$H_T(Y_\theta(1), Y_\theta(2), \ldots) = \sum_{n \geq 1} h(n; Y_\theta(1), \ldots, Y_\theta(n)) \mathbf{1}_{\{\tau_\theta = n\}} \qquad \text{a.s.}$$

for some measurable mapping $h(n; Y_\theta(1), \ldots, Y_\theta(n))$, for $n \in \mathbb{N}$.

Let $\Theta = [a, b]$, with $0 < a < \theta < b < \infty$, be a neighborhood of θ and set

$$K(x) = \sup_{\theta \in \Theta} \frac{\frac{d}{d\theta} f_\theta(x)}{f_\theta(x)}.$$

I A brief overview on gradient estimation: What to do with the derivative of a density?

II Random horizon experiments: our metro example

III **Differentiating a Radon-Nikodym derivative**

IV The Score function and the weak derivative interpretation

V Conclusion

III.1 The Intermediate Theorem

The key technical assumptions are the following:

Assume that a random variable B exists such that

(i) for all $n \geq 1$ and any vector of possible entries $(y_1, \ldots, y_n)$ it holds that

$$|h(n; y_1, \ldots, y_n)| \leq B \quad \text{a.s.}$$

(ii)

$$\sup_{\tilde{\theta} \in \Theta} \mathbb{E}\left[B \sum_{k=1}^{\tau_{\tilde{\theta}}} K(Y_\theta(k)) \right] < \infty$$

III.2 The Intermediate Theorem

Given that (i) and (ii) hold and under some additional smoothness conditions it holds that

$$\frac{d}{d\theta}\mathbb{E}[H_T(Y_\theta(1), Y_\theta(2), \ldots)]$$

$$= \mathbb{E}\left[H_T(Y_1(1), \ldots, Y_1(\tau_1), Y_\theta(\tau_1 + 1), \ldots) \left(\frac{d}{d\theta} \prod_{i=1}^{\tau_1} f_\theta(Y_1(i)) \right) \left(\prod_{i=1}^{\tau_1} f_1(Y_1(i)) \right)^{-1} \right],$$

where the random variables $\{Y_1(i) : i \in \mathbb{N}\}$ are i.i.d. with density $f_{\theta=1}$ and τ_1 is the corresponding stopping time.

Note that the change of measure applies only to the first τ_1 elements.

I A brief overview on gradient estimation: What to do with the derivative of a density?

II Random horizon experiments: our metro example

III Differentiating a Radon-Nikodym derivative

IV **The Score function and the weak derivative interpretation**

V Conclusion

IV.1 The Score Function Theorem

Under suitable conditions we obtain that

$$\frac{d}{d\theta}\mathbb{E}[H_T(Y_\theta(1), Y_\theta(2), \ldots)]$$

$$= \mathbb{E}\left[H_T(Y_1(1), \ldots, Y_1(\tau_1), Y_\theta(\tau_1 + 1), \ldots) \left(\frac{d}{d\theta} \prod_{i=1}^{\tau_1} f_\theta(Y_1(i)) \right) \left(\prod_{i=1}^{\tau_1} f_1(Y_1(i)) \right)^{-1} \right],$$

$$= \mathbb{E}\left[H_T(Y_1(1), \ldots, Y_1(\tau_1)) \left(\frac{d}{d\theta} \prod_{i=1}^{\tau_1} f_\theta(Y_1(i)) \right) \left(\prod_{i=1}^{\tau_1} f_1(Y_1(i)) \right)^{-1} \right].$$

IV.2 Manipulating the Derivative of the Radon-Nikodym Derivative

By simple algebra

$$\begin{aligned}\frac{d}{d\theta}\prod_{i=1}^{\tau_1} f_\theta(Y_1(i))\left(\prod_{i=1}^{\tau_1} f_1(Y_1(i))\right)^{-1} \\ &= \sum_{i=1}^{\tau_1}\frac{\frac{d}{d\theta}f_\theta(Y_1(i))}{f_\theta(Y_1(i))} \\ &= c_\theta\sum_{i=1}^{\tau_1}\frac{f_\theta^{+}(Y_1(i))}{f_\theta(Y_1(i))} - c_\theta\sum_{i=1}^{\tau_1}\frac{f_\theta^{-}(Y_1(i))}{f_\theta(Y_1(i))}.\end{aligned}$$

Interpretation: The ith occurrence of $Y_\theta(n)$ has density $f_\theta^{\pm}$ whereas all other occurrences have density f_θ.

This effects the stopping time and the new stopping time is denoted by $\tau_\theta^{\pm}(i)$.

IV.3 The MVD Theorem

Under suitable smoothness conditions it holds that

$$\frac{d}{d\theta}\mathbb{E}[H_T(Y_\theta(1), Y_\theta(2), \ldots)]$$

$$c_\theta \mathbb{E}\left[\sum_{i=1}^{\tau_\theta} H_T(Y_\theta(1), \ldots, Y_\theta(i-1), Y_\theta^+(i), Y_\theta(i+1), \ldots) - \sum_{i=1}^{\tau_\theta} H_T(Y_\theta(1), \ldots, Y_\theta(i-1), Y_\theta^-(i), Y_\theta(i+1), \ldots)\right],$$

In words, replace one occurrence of $Y_\theta(n)$ by $Y_\theta^+(n)$ and $Y_\theta^-(n)$, respectively, and terminated the experiment when the stopping criterium is satisfied.

In our metro model, the "+" and "-" phantoms can be coupled in such a way that they can be computed from the nominal path in a simple way.

Conclusions

- We presented gradient estimation for stopped experiments via analyzing the derivative of the Radon-Nikodym derivative.
- We explained the difference between the weak derivative representation and the score function representation of the basic interchange result.
- Our approach illustrates how to conveniently switch from a Score Function estimator to a phantom estimator of weak differentiation type and vice versa.
- We believe that this leads to a variance reduction technique for the Score Function, or, single run implementations of weak derivatives

Our goal is to find gradient estimators with low worknormalized variance.