Efficiency of the Cross-Entropy Method for Markov Chain Problems

Ad Ridder[1] Bruno Tuffin[2]

[1] *Vrije Universiteit, Amsterdam, Netherlands*
aridder@feweb.vu.nl
http://staff.feweb.vu.nl/aridder/

[2] *INRIA, Rennes, France*
bruno.tuffin@irisa.fr

Rare Event Simulation Workshop 2010,
Cambridge
21 June 2010

Introduction

- Suppose that $\{A_n : n = 1, 2, \ldots\}$ is a family of events in a probability space $(\Omega, \mathcal{A}, \mathbb{P})$,
- such that $\mathbb{P}(A_n) \to 0$ as $n \to \infty$.
- Furthermore, suppose that $\mathbb{P}(A_n)$ is difficult to compute analytically or numerically,
- but easy to estimate by simulation.

Research Question

- There might be many simulation algorithms.
- Denote by Y_n the associated unbiased estimator of $\mathbb{P}(A_n)$.

Can we give conditions for strong efficiency (bounded relative error)

$$\limsup_{n\to\infty} \frac{E[Y_n^2]}{(E[Y_n])^2} < \infty,$$

or logarithmic efficiency (asymptotic optimality),

$$\lim_{n\to\infty} \frac{\log \mathbb{E}[Y_n^2]}{\log(E[Y_n])^2} = 1\,?$$

Specific or General?

Many studies in the rare event simulation literature show efficiency of a specific algorithm for a specific problem.

For instance, concerning asymptotic optimality.

- Specific: importance sampling with exponentially twisted distribution for a level crossing probability ([1]).
- More general: Dupuis and co-authors ([2], [3]) developed an importance sampling method based on a control-theoretic approach to large deviations, which is applicable for a large class of problems involving Markov chains and queueing networks.
- More abstract: assume large deviations probabilities, i.e., $\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}(A_n) = -\theta$, and derive conditions under which exponentially twisted importance sampling distribution is asymptotically optimal ([4], [5]).

References

1. Siegmund, D. 1976. *Annals of Statistics* 4, 673-684.
2. Dupuis, P., and Wang, H. 2005. *Annals of Applied Probability* 15, 1-38.
3. Dupuis, P., Sezer, D., and Wang, H. 2007. *Annals of Applied Probability* 17, 1306-1346.
4. Sadowsky, J.S. 1996. *Annals of Applied Probability* 6, 399-422.
5. Dieker, T., and Mandjes, M. 2005. *Advances in Applied Probability* 37, 539-552.

BRE Studies

Examples with strong efficiency.

- Importance sampling with a biasing scheme in a highly reliable Markovian system [1].
- Zero-variance approximation importance sampling in a highly reliable Markovian system [2].
- Combination of conditioning and importance sampling for tail probabilities of geometric sums of heavy tailed rv's [3].
- State-dependent importance sampling (based on zero-variance approximation) for sums of Gaussian rv's [4].

References

1. Shahabuddin, P. 1994. *Management Science* 40, 333-352.
2. L'Ecuyer, P. and Tuffin B. 2009. *Annals of Operations Research*, to appear.
3. Juneja, S. 2007. *Queueing Systems* 57, 115-127.
4. Blanchet, J.H. and Glynn, P.W. 2006. *Proceedings ValueTools 2006*.

More References

vrije Universiteit amsterdam

More or less general studies.

1. Heidelberger, P. 1995. *ACM Transactions on Modeling and Computer Simulation* 5, 43-85.
2. Asmussen, S. and Rubinstein, R. 1995. In *Advances in Queueing Theory, Methods, and Open problems*, 429-462.
3. L'Ecuyer, P., Blanchet, J.H., Tuffin, B. and Glynn, P.W. 2010. *ACM Transactions on Modeling and Computer Simulation* 20, 6:1-6:41.

Cross-Entropy Method

- The cross-entropy method is a heuristic for rare event simulation to find the importance sampling distributions within a parameterized class (book Kroese and Rubinstein 2004).
- A summary on one of the next slides.
- Then proving analytically the efficiency of the resulting estimator is ‘impossible’.
- The usual approach is to estimate the efficiency by empirical (simulation) data.

Contribution

This paper gives sufficient conditions for the cross-entropy method to be efficient for a certain type of rare event problems in Markov chains.

The Rare Event Problem

$\mathbb{P}(A_n)$ is an absorption probability in a finite-state discrete-time Markov chain.

We allow two versions.

A. The rarity parameter n is associated with the problem size which is increasing in n.

B. We assume a constant problem size and we let the rarity parameter to be associated with transition probabilities that are decreasing in n.

For ease of notation, drop rarity parameter n.

Notation

- Markov chain is $\{X(t) : t = 0, 1, \ldots\}$.
- Statespace $\mathcal{X}$; transition prob's $p(x, y)$.
- Markov chain starts off in a reference state 0, $X(0) = 0$.
- A 'good' set $\mathcal{G} \subset \mathcal{X}$ of absorbing states.
- A failure set $\mathcal{F} \subset \mathcal{X}$ of absorbing states.
- No other absorbing states.
- The time to absorption is $T = \inf\{t > 0 : X(t) \in \mathcal{G} \cup \mathcal{F}\}$.
- Absorption probabilities $\gamma(x) = \mathbb{P}(X(T) \in \mathcal{F} | X(0) = x)$.
- Rare event $A = 1\{X(T) \in \mathcal{F}\}$ with probability

$$\mathbb{P}(A) = \gamma(0).$$

Illustration

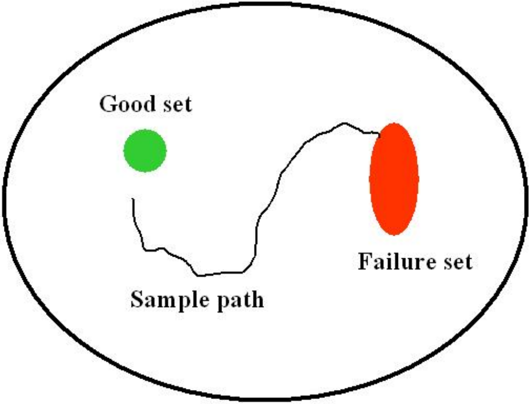

Importance Sampling

Importance sampling simulation implements a change of measure $\mathbb{P}^*$ to obtain unbiased importance sampling estimator

$$Y = 1\{A\} \frac{d\,\mathbb{P}}{d\,\mathbb{P}^*},$$

where $d\,\mathbb{P}/d\,\mathbb{P}^*$ is the likelihood ratio.

Feasibility

Restrict to changes of measure for which $p(x,y) > 0 \Leftrightarrow p^*(x,y) > 0$.

Notation: probability measure $\mathbb{P}$ (or $\mathbb{P}^*$) and associated matrix of transition probabilities P (or P^*) are used for the same purpose whenever convenient.

Zero Variance

- Optimal change of measure $\mathbb{P}^{\text{opt}} = \mathbb{P}(\cdot|A)$ gives $\text{Var}^{\text{opt}}[Y] = 0$.
- $\mathbb{P}^{\text{opt}}$ is feasible,

$$p^{\text{opt}}(x,y) = p(x,y)\frac{\gamma(y)}{\gamma(x)},$$

- not implementable, since it requires knowledge of the unknown absorption probabilities.

Cross-Entropy Minimization

vrije Universiteit amsterdam

Find $\mathbb{P}^*$ by minimizing the Kullback-Leibler distance (or cross-entropy) within the class of feasible changes of measure:

$$\inf_{P^* \in \mathcal{P}} \mathcal{D}(d\mathbb{P}^{\mathrm{opt}}, d\mathbb{P}^*),$$

where the cross-entropy is defined by

$$\mathcal{D}(d\mathbb{P}^{\mathrm{opt}}, d\mathbb{P}^*) = \mathbb{E}^{\mathrm{opt}}\left[\log\left(\frac{d\mathbb{P}^{\mathrm{opt}}}{d\mathbb{P}^*}(\boldsymbol{X})\right)\right] = \mathbb{E}\left[\frac{d\mathbb{P}^{\mathrm{opt}}}{d\mathbb{P}}(\boldsymbol{X})\ \log\left(\frac{d\mathbb{P}^{\mathrm{opt}}}{d\mathbb{P}^*}(\boldsymbol{X})\right)\right].$$

Notation: $\boldsymbol{X}$ is a random sample path of the Markov chain from the reference state 0.

Cross-Entropy Solution

Solution denoted $\mathbb{P}^{\min}$ has (after some algebra)

$$p^{\min}(x,y) = \frac{\mathbb{E}[1\{A\}N(x,y)]}{\mathbb{E}\left[1\{A\}\sum_{z\in\mathcal{X}} N(x,z)\right]},$$

where $N(x,y)$ is the number of times that transition (x,y) occurs.

Equivalence

Lemma

$p^{\min}(x,y) = p^{\text{opt}}(x,y)$ *for all* $x, y \in \mathcal{X}$.

Proof.

(a) Indirect way: $\mathbb{P}^{\text{opt}}$ is a feasible change of measure for the minimization.

(b) Direct way: we can show analytically that the expressions of $p^{\min}(x,y)$ and $p^{\text{opt}}(x,y)$ given above are equal.

ZVA and ZVE

ZVA: an importance sampling estimator based on *approximating* the numerators $\widehat{\gamma(x)}$ of the zero-variance transition probabilities p^{opt}.

ZVE: an importance sampling estimator based on *estimating* the numerators $\widehat{\mathbb{E}}[1\{A\}N(x,y)]$ of the zero-variance transition probabilities $p^{\min}$.

Cross-Entropy Based ZVE

vrije Universiteit amsterdam

Easy:

$$\inf_{P^* \in \mathcal{P}} \mathcal{D}(d\mathbb{P}^{\text{opt}}, d\mathbb{P}^*) \quad \Leftrightarrow \quad \sup_{P^* \in \mathcal{P}} \mathbb{E}[1\{X(T) \in \mathcal{F}\} \log d\,\mathbb{P}^*(\boldsymbol{X})],$$

where, by a change of measure:

$$\mathbb{E}[1\{X(T) \in \mathcal{F}\} \log d\mathbb{P}^*(\boldsymbol{X})] = \mathbb{E}^{(0)}\left[\frac{d\mathbb{P}}{d\mathbb{P}^{(0)}} 1\{X(T) \in \mathcal{F}\} \log d\mathbb{P}^*(\boldsymbol{X})\right].$$

Estimate and iterate:

$$P^{(j+1)} = \arg\max_{P^* \in \mathcal{P}} \frac{1}{k} \sum_{i=1}^{k} \frac{d\mathbb{P}}{d\mathbb{P}^{(j)}}(\boldsymbol{X}^{(i)}) 1\{X^{(i)}(T) \in \mathcal{F}\} \log d\mathbb{P}^*(\boldsymbol{X}^{(i)}).$$

After a few iterations (convergence?): ZVE $\mathbb{P}^{\text{ce}}$.

Asymptotic Optimality

Notation: $\mathbb{P}_n^{\mathrm{ce}}$ and $\mathbb{P}_n^{\mathrm{opt}}$ for explicitly indicating that the change of measure depends also on the rarity parameter.

Theorem

Assume

$$\mathcal{D}(\mathbb{P}_n^{\mathrm{opt}}, \mathbb{P}_n^{\mathrm{ce}}) = o(\log \mathbb{P}(A_n))$$

as $n \to \infty$, then the associated importance sampling estimator is asymptotically efficient.

Proof

$$\mathcal{D}(\mathbb{P}_n^{\text{opt}}, \mathbb{P}_n^{\text{ce}}) = \mathbb{E}^{\text{opt}}[\log d\mathbb{P}_n^{\text{opt}}/d\mathbb{P}_n^{\text{ce}}(\boldsymbol{X})] \geq 0.$$

$$\mathbb{E}^{\text{ce}}\left[Y_n^2\right] = \mathbb{E}^{\text{ce}}\left[\left(\frac{d\mathbb{P}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\,\mathbb{1}\{A_n\}\right)^2\right] = \mathbb{E}^{\text{ce}}\left[\left(\frac{d\mathbb{P}}{d\mathbb{P}_n^{\text{opt}}}(\boldsymbol{X})\,\mathbb{1}\{A_n\}\right)^2\left(\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right)^2\right]$$

$$= \mathbb{P}(A_n)^2\,\mathbb{E}^{\text{ce}}\left[\left(\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right)^2\right] = \mathbb{P}(A_n)^2\,\mathbb{E}^{\text{opt}}\left[\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right].$$

So, we can conclude

$$\frac{\log \mathbb{E}^{\text{ce}}[Y_n^2]}{\log \mathbb{P}(A_n)} = \frac{\log(\mathbb{P}(A_n))^2 + \log \mathbb{E}^{\text{opt}}\left[\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right]}{\log \mathbb{P}(A_n)} = 2 + \frac{\log \mathbb{E}^{\text{opt}}\left[\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right]}{\log \mathbb{P}(A_n)},$$

with

$$\lim_{n\to\infty} \frac{\log \mathbb{E}^{\text{opt}}\left[\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right]}{\log \mathbb{P}(A_n)} = \lim_{n\to\infty} \frac{\log \mathbb{E}^{\text{opt}}\left[\frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right]}{\mathbb{E}^{\text{opt}}\left[\log \frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right]} \frac{\mathbb{E}^{\text{opt}}\left[\log \frac{d\mathbb{P}_n^{\text{opt}}}{d\mathbb{P}_n^{\text{ce}}}(\boldsymbol{X})\right]}{\log \mathbb{P}(A_n)} = 0.$$

A Simple Example: M/M/1

$\{X(t) : t = 0, 1, \ldots\}$ on $\{0, 1, \ldots\}$ is the discrete-time Markov chain by embedding at jump times of the $M/M/1$ queue.

The rare event is hitting state n before returning to the zero state:

$$\gamma(0) = \mathbb{P}((X(t)) \text{ reaches } n \text{ before } 0 | X(1) = 1).$$

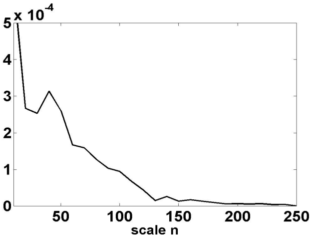

$\mathcal{D}(\mathbb{P}_n^{\text{opt}}, \mathbb{P}_n^{\text{ce}}) / \log \mathbb{P}(A_n).$

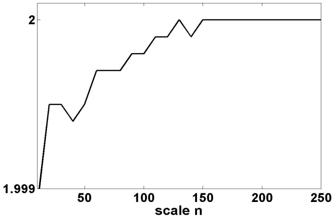

$\log \mathbb{E}^{\text{ce}}[Y_n^2] / \log \mathbb{P}(A_n).$

Bounded Relative Error

From a probabilistic point of view, the cross-entropy method is a randomized algorithm that delivers (unbiased) *estimators* $\widehat{P}_n(x,y)$ of the zero-variance transition probabilities $p_n^{\text{opt}}(x,y)$.

Theorem

Assume that for any $\alpha < 1$ *there is* $K > 0$ *such that*

$$\limsup_{n\to\infty} P\left(\max_{\substack{(x,y)\in\mathcal{X}\times\mathcal{X}\\ p(x,y)>0}} \frac{p_n^{\text{opt}}(x,y)}{\widehat{P}_n(x,y)} \le K\right) \le \alpha,$$

then the associated importance sampling estimator Y_n *is strongly efficient (with probability* α*).*

(Notice that the expectation of Y_n given the estimators $\widehat{P}_n(x,y)$ is a rv.)

Proof

Similar as in the proof of Theorem 1:

$$\mathbb{E}^{\mathrm{ce}}\left[Y_n^2|\widehat{P}_n\right] = \mathbb{P}(A_n)^2\,\mathbb{E}^{\mathrm{opt}}\left[\frac{d\mathbb{P}_n^{\mathrm{opt}}}{d\mathbb{P}_n^{\mathrm{ce}}}(\boldsymbol{X})|\widehat{P}_n\right].$$

Now apply the condition of the theorem, and follow proof of Theorem 2 in [L'Ecuyer, P. and Tuffin B., AOR 2010].

Conclusion and Outlook

- Cross-entropy coincides with zero-variance in rare-event problems of Markov chains.
- Sufficient condition for logarithmic efficiency.
- Further investigations include verification of this condition.
- Probabilistic condition for strong efficiency.
- Further investigations to elaborate the strong efficiency and illustrate.