Variance estimation for splitting and killing schemes

Robert J. Webber* and Jonathan Weare*
**Courant Institute of Mathematical Sciences, New York University, New York, USA*

Summary. Splitting and killing is a flexible approach to Monte Carlo sampling that is popular in chemistry and geophysics. While splitting and killing can lead to high-accuracy estimates, it is difficult to gauge the quality of the estimates published in research papers because error bars are frequently not provided. Many researchers are unaware of the existence of variance estimators for splitting and killing schemes. Additionally, the variance estimators that have been proposed are not sufficiently well understood either theoretically or empirically. In response to this critical gap in understanding, we review the variance estimators that have been proposed, establish theoretical properties for these estimators, and present numerical numerical tests evaluating their quality.

Background

The idea of splitting and killing was first envisioned by Ulam and Von Neumann in the late 1940s [1], and it has become increasingly popular over the past decades. Today, splitting and killing schemes include forward flux sampling, which is used in chemistry to compute reaction rates [2], and diffusion Monte Carlo, which is used in geophysics to study extreme weather and climate events [3]. In a splitting and killing scheme, the motivating idea is to clone ("split") particles that move toward a rare event of interest and terminate ("kill") particles that move away from the event of interest, while preserving the overall number of particles. There are many variations on this splitting and killing idea, yet they are all designed to more accurately evaluate rare event statistics.

The algorithm

The splitting and killing approach begins with an initialization step and then iterates over selection and mutation steps:

Initialization. Initialize a stochastic model with a set of states $\left(\xi_0^i\right)_{1\le i\le N}$ and weights $w_0^i = \frac{1}{N}$ for $1 \le i \le N$.

Selection. For each state ξ_t^i, choose a target number of children particles T_t^i so that $T_t^i > 0$ and $\sum_{i=1}^N T_t^i = N$. Use a resampling algorithm [4] to round each number T_t^i to a nonnegative integer so that $\mathrm{E}\left[N_t^i\right] = T_t^i$ and $\sum_{i=1}^N N_t^i = N$. Replace the weighted ensemble $\left(w_t^i, \xi_t^i\right)_{1\le i\le N}$ with a new ensemble $\left(\hat{w}_t^i, \hat{\xi}_t^i\right)_{1\le i\le N}$ where each particle ξ_t^i is represented by N_t^i children particles $\hat{\xi}_t^j$, all with weights w_t^j / T_t^j.

Mutation. Run the stochastic model from each state $\hat{\xi}_t^i$ to a new state ξ_{t+1}^i and set $w_{t+1}^i = \hat{w}_t^i$ for $1 \le i \le N$.

Figure 1: Splitting and killing is used to sample rare, high positions of the x coordinate. White circles indicate that samples are killed, while black circles indicate that samples are preserved and possibly split.

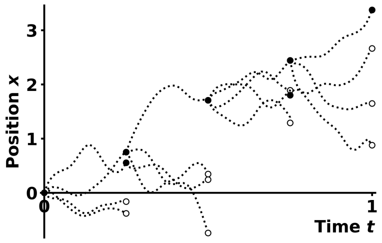

Estimates and error bars

Estimates

For any Markov process X_t, the splitting and killing approach can be used to produce unbiased estimates of expectations $\mathrm{E}\left[f\left(X_T\right)\right]$ using the estimator $\sum_{i=1}^N w_T^i f\left(\xi_T^i\right)$ [4].

Variance estimators

Two variance estimators have been proposed for splitting and killing schemes. The *ancestral history estimator*, proposed in [6], takes the form

$$\operatorname{Var}\left[\sum_{i=1}^N w_T^i f\left(\xi_T^i\right)\right] \approx N \underset{1\le i\le N}{\operatorname{EmpVar}}\left[\sum_{\operatorname{anc}\left(\xi_T^j\right)=i} w_T^j f\left(\xi_T^j\right)\right] \tag{1}$$

Here, $\operatorname{EmpVar}_{1\le i\le N}\left[x^i\right] = \frac{1}{N}\sum_{i=1}^N \left|x^i\right|^2 - \left|\frac{1}{N}\sum_{i=1}^N x^i\right|^2$ denotes the empirical variance of data points $x^1, \ldots, x^N$, while $\operatorname{anc}\left(\xi_T^j\right)$ denotes the index of the original ancestor of ξ_T^j at time $t = 0$. The ancestral history estimator is motivated by assuming quantities $\sum_{\operatorname{anc}\left(\xi_T^j\right)=i} w_T^j f\left(\xi_T^j\right)$ are uncorrelated and applying a sum of squares variance decomposition.

The *importance sampling estimator*, which is used for the *importance sampling* method for Monte Carlo sampling, takes the form

$$\operatorname{Var}\left[\sum_{i=1}^{N} w_T^i f\left(\xi_T^i\right)\right] \approx N \underset{1\leq i\leq N}{\operatorname{EmpVar}}\left[w_T^i f\left(\xi_T^i\right)\right] \tag{2}$$

The importance sampling estimator is motivated by assuming quantities $w_T^i f\left(\xi_T^i\right)$ are uncorrelated and applying a sum of squares variance decomposition.

Theoretical results

The ancestral history estimator has been shown to converge as $N \to \infty$ if the numbers N_t^i are uncorrelated [6]. However, in most applications of splitting and killing, the numbers N_t^i are negatively correlated, so our theoretical analysis shows the ancestral history estimator is biased too high and converges to a number higher than the true variance.

In contrast, the importance sampling estimator is biased too low and converges to a number lower than the true variance if we assume that T_t^i is proportional to $w_t^i\, \mathrm{E}\left[f\left(X_T\right) \middle| X_t = \xi_t^i\right]$. We expect but cannot prove the importance sampling estimator is biased too low in most situations even without this assumption.

Numerical results

In two test problems we apply splitting and killing to an AR(1) process $X_t = \alpha X_{t-1} + Z_t$, where $X_0 = 0$ and $Z_t \sim \mathcal{N}(0, 2)$ are independent Gaussian increments. In the first test we set $\alpha = .9$, while in the second test we set $\alpha = .1$. In both tests, we apply splitting and killing at times $t = 1, \ldots, 9$ with targets T_t^i proportional to $w_t^i \exp\left(\xi_t^i\right)$. We use $N = 1000$ particles, and we examine the relative variance when calculating $\mathrm{E}\left[\exp\left(X_t\right)\right]$ for $T = 0, \ldots, 10$.

In both tests, the ancestral history estimator is biased too high. The importance sampling estimator is biased very low in the first test because it does not account for the long time correlations in the dynamics, but it is more accurate in the second test due to the shorter time correlations.

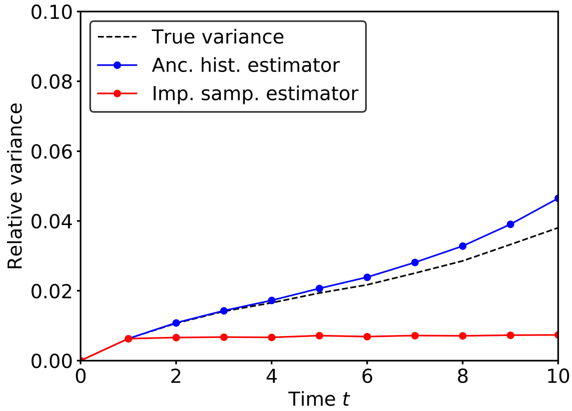

Figure 2: First test. The ancestral history estimator is biased slightly too high while the importance sampling estimator is biased far too low.

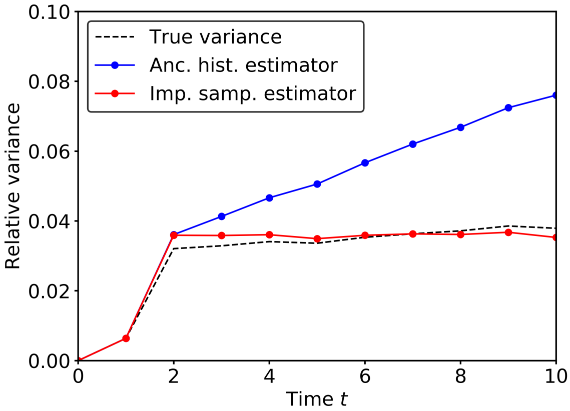

Figure 3: Second test. The ancestral history estimator is biased far too high while the importance sampling estimator is close to the true variance.

Conclusions

From our theoretical and numerical results, we recommend using the ancestral history estimator as an upper bound on the variance and the importance sampling estimator as an approximate lower bound on the variance. We anticipate that these variance estimators will prove useful for uncertainty quantification of splitting and killing estimates in the future.

References

[1] G. E. Forsythe and R. A. Leibler, "Matrix inversion by a Monte Carlo method," *Mathematics of Computation*, vol. 4, no. 31, pp. 127–129, 1950.

[2] R. J. Allen, D. Frenkel, and P. R. ten Wolde, "Forward flux sampling-type schemes for simulating rare events: Efficiency analysis," *The Journal of chemical physics*, vol. 124, no. 19, p. 194111, 2006.

[3] R. J. Webber, D. A. Plotkin, M. E. O'Neill, D. S. Abbot, and J. Weare, "Practical rare event sampling for extreme mesoscale weather," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, no. 5, p. 053109, 2019.

[4] R. J. Webber, "Unifying sequential Monte Carlo with resampling matrices," *arXiv preprint arXiv:1903.12583*, 2019.

[5] F. Cérou, P. Del Moral, F. Le Gland, and P. Lezaud, "Genetic genealogical models in rare event analysis," Research Report RR-5878, INRIA, 2006.

[6] H. P. Chan, T. L. Lai, *et al.*, "A general theory of particle filters in hidden Markov models and some applications," *The Annals of Statistics*, vol. 41, no. 6, pp. 2877–2904, 2013.