Combining Analytic Direct Illumination and Stochastic Shadows

Supplemental Material: Convergence Analysis

Abstract

In this document, we propose a 1D experiment that shows that our ratio estimator has generally lower variance than the full-stochastic and control-variate estimators. This document is referenced in Section 3.4 of the main paper in paragraph *The ratio estimator has low variance.*

Numerical experiment in 1D In order to investigate the evolution of the estimators through increasing number of samples we apply them to a 1D problem where the goal of this experiment is to estimate the integral of the product of two functions w and V over $[0, 1]$:

$$S = \int_0^1 w(x)\, V(x)\, dx, \tag{1}$$

where $w = R\,L/p$ plays the role of the sampling weight of the direct illumination and V is the binary visibility function. The classic *full-stochastic* estimator, the *control-variate* estimator, and our *ratio* estimator are obtained by sampling uniform values $x \in [0, 1]$:

$$\begin{aligned} S_N^{\text{sto}} &= \frac{1}{N} \sum_{n=1}^{N} w(x_n)\, V(x_n), \\ S_N^{\text{CV}} &= \int_0^1 w(x)\, dx - \frac{1}{N} \sum_{n=1}^{N} w(x_n)\ (1 - V(x_n))\,, \\ S_N^{\text{ratio}} &= \int_0^1 w(x)\, dx \times \frac{\sum_{n=1}^{N} w(x_n)\, V(x_n)}{\sum_{n=1}^{N} w(x_n)}, \end{aligned} \tag{2}$$

Note that the variance properties of these estimators are strictly equivalent to their respective direct-illumination estimators since we only change the integration domain. The advantage of doing this is that it makes investigations and visualization easier. In Figure 1 we compare the convergence of these estimators in different scenarios.

Discussion We provide a qualitative summary of the results shown in Figures 1 in Table 1. We interpret these results for different configurations:

(a,b): if V is always (or almost always) 1, the variances of both the control-variate and the ratio estimators are low. This explains why shadowless regions are noise-free with these formulations while a classic full-stochastic evaluation produces noise.

(c,d): if V is always (or almost always) 0, the variances of both a full-stochastic evaluation and our ratio estimator are zero. This explains why shadowed regions are noise-free with these estimators while the control-variate estimator exhibits variance in these regions.

(e,f): as the variance of V increases, the estimators tend to have similar variances until they become equivalent as the variance of w becomes negligible in comparison to the one of V. This explains why penumbra regions are similarly noisy with all of the estimators.

Figure 1	w	V	S_N^{sto}	S_N^{CV}	S_N^{ratio}
(a)	high variance	1	✗	✓	✓
(b)	high variance	near 1	✗	✓	✓
(c)	high variance	0	✓	✗	✓
(d)	high variance	near 0	✓	✗	✓
(e)	high variance	high variance	✗	✗	✗
(f)	low variance	high variance	✗	✗	✗

Table 1: Qualitative summary of the experiments.

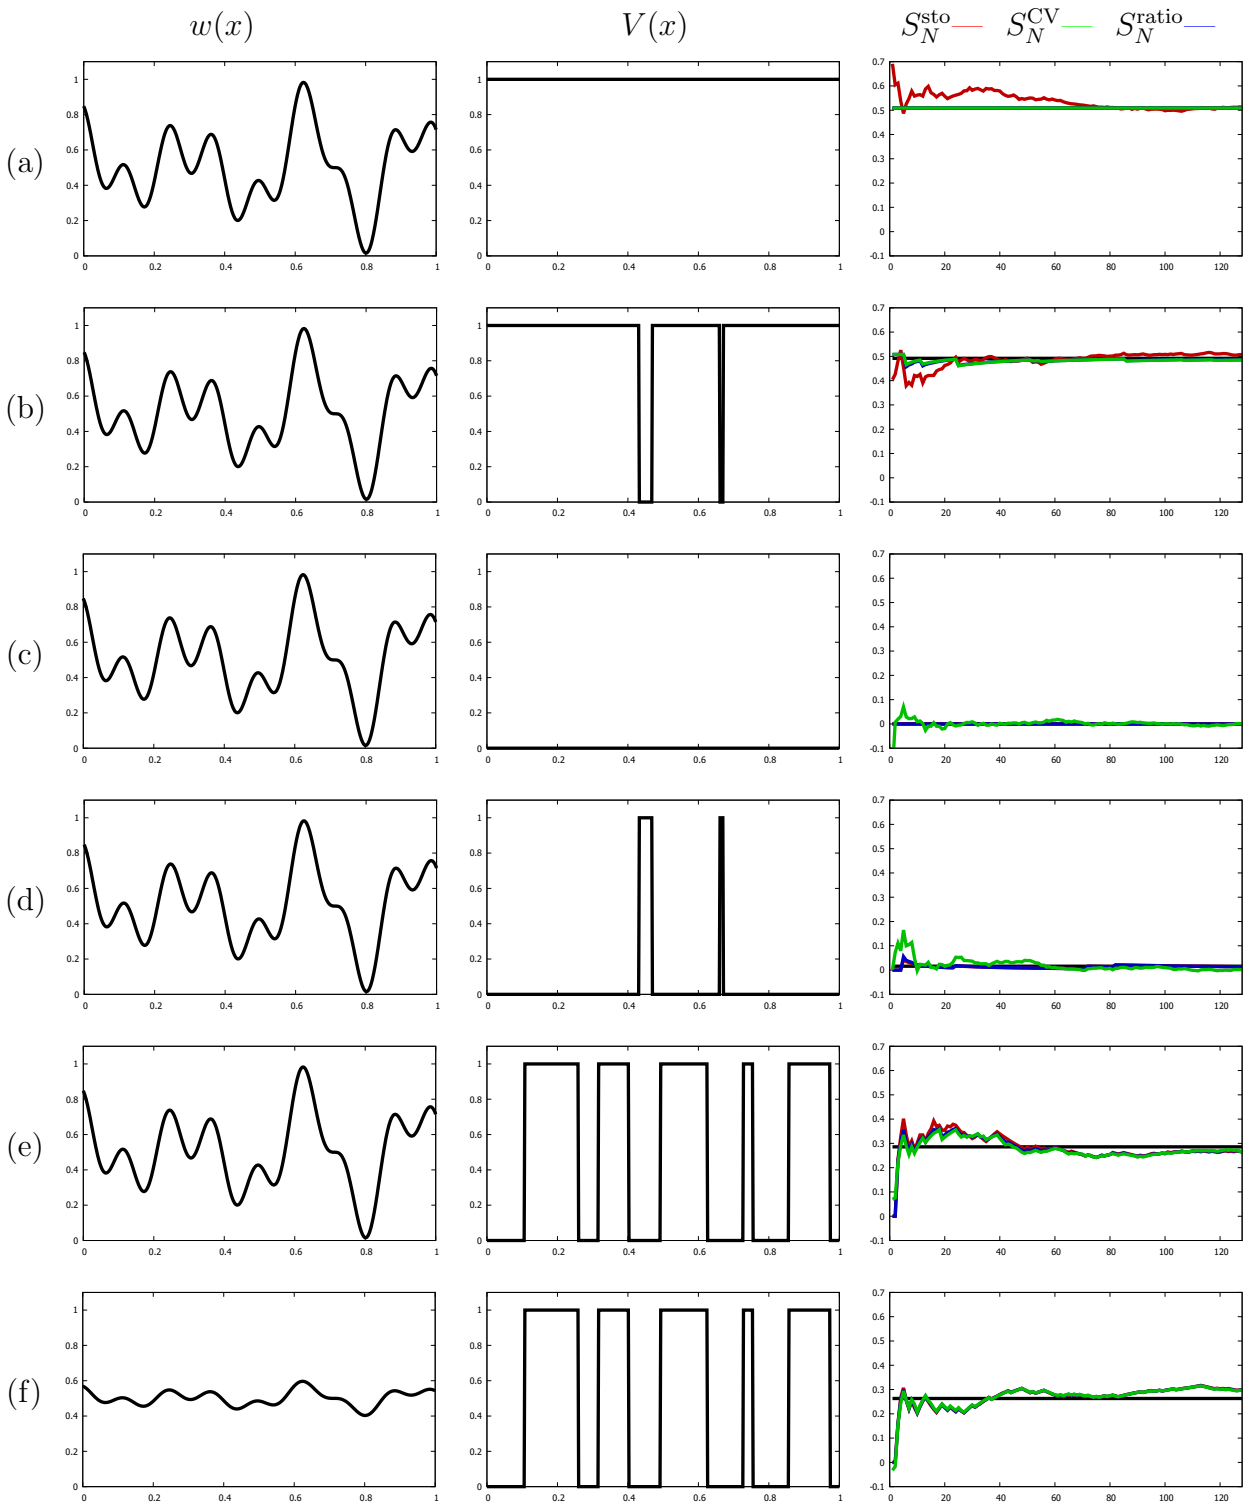

Figure 1: Convergence of the estimators with increasing number of samples.