MIS Compensation: Optimizing Sampling Techniques in Multiple Importance Sampling - supplemental material

ONDŘEJ KARLÍK, Chaos Czech a. s.
MARTIN ŠIK, Chaos Czech a. s.
PETR VÉVODA, Charles University, Prague and Chaos Czech a. s.
TOMÁŠ SKŘIVAN, IST Austria
JAROSLAV KŘIVÁNEK, Charles University, Prague and Chaos Czech a. s.

This supplemental document provides several additional results to complete the main article.

1 APPLICATION I: IMAGE-BASED LIGHTING

1.1 Closeness of the MIS-compensated and optimal pdfs

We test the increase of estimator's second moment and of variance when using the *MIS-compensated* pdf instead of the *optimal* pdf in practice by devising the following experiment. Given a discrete integration space $\mathcal{X}$ and two sampling techniques, we search for an integrand f and a fixed pdf p_2 that yield the highest possible increase of the estimator's second moment when using our *MIS-compensated* pdf $\tilde{p}_1$ instead of the *optimal* pdf p_1^*. For $c_1 = 0.5$ the highest found increase of estimator's second moment was 1.026 corresponding to 1.080 increase of variance. For values of $c_1 \in [0.1, 0.9]$ we found that the variance never increased more than 1.6 times. This suggests that in practice the *MIS-compensated* pdf $\tilde{p}_1$ is often close to the *optimal* pdf p_1^*. Fig. 1 demonstrates the worst case scenario obtained in our experiment for $c_1 = 0.5$.

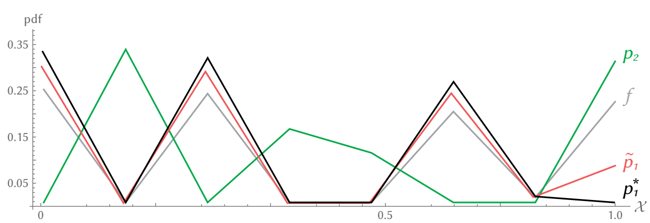

Fig. 1. Comparison of the *MIS-compensated* pdf $\tilde{p}_1$ (red) and the *optimal* pdf p_1^* (black) in the experimentally found worst case scenario (for $c_1 = 0.5$), where using $\tilde{p}_1$ leads to the highest increase of estimator's second moment over using p_1^*.

1.2 Sensitivity of the MIS-compensated pdf to integral approximation

Since the *MIS-compensated* pdf requires knowledge of the target integral L_{dir} we test how its performance changes when an approximation of varying precision is used instead of L_{dir}. Fig. 2 shows how variance of the MIS estimator using the *MIS-compensated* pdf changes with approximation error in the same flatland setup as used in Sec. 6.3 in the paper. For approximation error within 40% of L_{dir} the variance increase is almost negligible and only once the error is at least 200% of the precise value, the *MIS-compensated* pdf becomes worse than the original unmodified HDR map pdf.

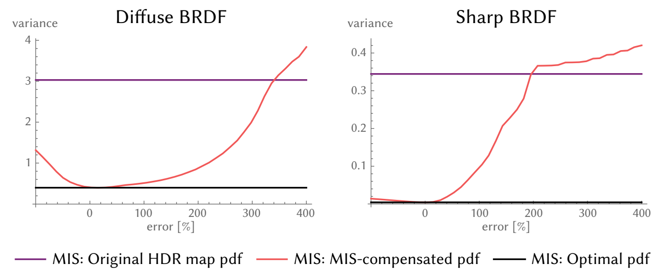

Fig. 2. Plot of variance of the MIS estimator using the *MIS-compensated* pdf with respect to the relative error of an approximation used in the pdf formula Eq. (12) instead of the precise integral L_{dir}. Variance of using the original HDR map pdf and the *optimal* pdf are shown for reference.

1.3 Render time comparison

Here we show all render time data for the IBL application we used in the paper. We rendered each scene shown in the paper using MIS with the original HDR map pdf (Baseline), our *practical normal-independent* pdf (Ours-ni), our *practical normal-dependent* pdf (Ours-nd), and using the resampled importance sampling method (RIS). All methods rendered the same number of samples per pixel (160), Table 1 lists the resulting render times and overhead (relative difference of render time with respect to Baseline). From the table we can draw a conclusion that our *practical normal-independent* solution does not incur any overhead nor it is faster, since averages match almost perfectly, and the maximum of ±3.28% in difference of performance can be contributed to external factors (such as other programs using the CPU). We can also see that the *practical normal-dependent* solution incurs approximately 10% overhead while the RIS method is about 60% slower.

1.4 MIS heuristics comparison

To draw our conclusions on MIS heuristics, we rendered all our real production scenes using four heuristics: the balance heuristic, the power heuristic with power factor $\beta = 2$, the cutoff heuristic with cutoff threshold $q_{max} = 0.1$, and the maximum heuristic. The same number of samples per pixel (160) were used for all of them. Our method was not used in this comparison since it is optimized only for the balance heuristic. Table 2 shows the measured normalized mean squared errors (NMSE). We can see that, apart from the maximum heuristic which performs the worst, other heuristics yield similar performance. The balance heuristic comes slightly on top overall.

Table 1. Rendering time comparison for the IBL application

	Baseline	Ours-ni		Ours-nd		RIS	
Scene name	Time [ms]	Time [ms]	Overhead [%]	Time [ms]	Overhead [%]	Time [ms]	Overhead [%]
Synthetic (low contrast)	3905	3958	1.36	4672	19.64	7169	83.59
Synthetic (high contrast)	4016	4022	0.15	4391	9.34	7304	81.87
Synthetic (mid contrast)	3996	3865	-3.28	4317	8.03	6649	66.39
Car	17864	18024	0.90	20450	14.48	28625	60.24
Pills	77872	80151	2.93	78474	0.77	99215	27.41
Room	46193	45876	-0.69	50519	9.37	75500	63.44
Average overhead			**0.23**		**10.27**		**63.82**

Table 2. MIS heuristics comparison for the IBL application

	NMSE (scaled by 10^3)				NMSE compared to balance		
Scene	balance	power	cutoff	max	power	cutoff	max
Synthetic (low contrast)	0.297	0.293	0.289	0.453	98.76%	97.25%	152.50%
Synthetic (high contrast)	0.304	0.311	0.304	0.346	102.49%	100.22%	114.04%
Synthetic (mid contrast)	0.479	0.508	0.489	0.605	105.96%	101.96%	126.13%
Car	0.187	0.200	0.187	0.423	107.19%	100.10%	226.96%
Pills	1.028	1.053	1.026	1.225	102.51%	99.87%	119.25%
Rooom	5.883	5.998	5.927	6.396	101.97%	100.76%	108.73%
Average					**103.14%**	**100.03%**	**141.27%**

Table 3. Rendering time and average path length comparison for the path guiding application

	Baseline		Ours-ni		
Scene name	Avg. path length	Time [m]	Avg. path length	Time [m]	Overhead [%]
Kitchen	7.15	2.22	6.62	2.03	-8.5
Pool	4.90	2.53	4.60	2.35	-7.1
Average overhead					**-7.8**

2 APPLICATION II: PATH GUIDING

2.1 Render time comparison

Here we show all render time data for the path guiding application we used in the paper. We rendered both scenes shown in the paper using the Müller et al.'s [2017] method in its original version (Baseline) and with our *practical normal-independent* MIS compensation applied (Ours-ni). Both versions rendered the same number of samples per pixel (140 in the *Kitchen* scene and 260 in the *Pool* scene). Table 3 lists the resulting render times, overhead (relative difference of render time with respect to Baseline) and an average length of traced paths. From the table we can draw a conclusion that our *practical normal-independent* solution does not incur any overhead. In fact, by guiding paths more efficiently towards a light source it decreases the average path length. As a result, the rending time is reduced by 7.8% on average.

3 ADDITIONAL RESULTS

We provide additional rendering results for scenes in the paper accessible via the html file: `image_comparisons/comparison.html`

REFERENCES

Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical path guiding for efficient light-transport simulation. *Comput. Graph. Forum (EGSR '17)* 36, 4 (June 2017).