Bayesian online regression for adaptive direct illumination sampling: Supplemental document

PETR VÉVODA, Charles University, Prague and Render Legion, a. s.
IVO KONDAPANENI, Charles University, Prague
JAROSLAV KŘIVÁNEK, Charles University, Prague and Render Legion, a. s.

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

1 OUR IMPLEMENTATION OF DONIKIAN ET AL. [2006]

Donikian et al. [2006] divide the image into blocks and process them one by one. For each block, they first fix one shading point for each pixel and then process the block pixels in iterations until convergence. In each iteration, they sample direct illumination at the shading points a fixed number of times (1.5× the light count), and subsequently update light contribution estimates at the block and pixel levels, respectively. The next iteration then uses a sampling distribution which mixes distributions at the block and pixel levels with the uniform distribution. The mixing weights are oblivious to the observed samples and depend solely on the iteration count. They change during the first 10 iterations only and remain fixed after that. This process is repeated until a convergence criterion is met for all pixels in the block; then a new block is started.

To make this method more compatible with ours, we made it progressive by computing all blocks at once. Furthermore, we find a new shading point for every pixel sample. One iteration then corresponds to taking one sample from all image pixels. The rendering time in our tests is set long enough for this method to complete enough iterations to learn (i.e., at least 10). Finally, we set our method in these tests to sample direct illumination at each shading point the same number of times (i.e., 1.5× the light count instead of the default 16×).

2 CLUSTERING

In Fig. 1, we analyze the effect of light clustering on the performance of our method, in particular the effect of ϵ, the fraction of the estimated contribution of the entire cut, used as a threshold for stopping the cut refinement. With higher values the cuts are smaller and faster to compute, the maximum value of 1 would cluster all lights into a single cluster. With lower values the cuts are more accurate, the minimum value of 0 either clusters each light in its own cluster (less than 100 lights) or into a maximum cut of 100 clusters (more than 100 lights).

As expected, the clustering has an important impact on the performance in the City scene, which contains more than 5000 lights (the optimum values yields more than 20× speedup in comparison with the least suitable value). On the other hand, the clustering has much smaller effect in the Hall scene with less than 100 lights (the speedup is only 1.3×). We used $\epsilon = 0.1$ in all our tests, which is optimal in the City scene and close to optimal in the Hall scene.

3 HYPERPARAMETERS

Our default choice of the hyperparameter values yields an uninformed prior distribution over the model parameters, and works robustly across all our tests. In particular, we use $\hat{N}_o = 2, \hat{N}_v = 2, \hat{N} = 1, \hat{N}_\alpha = 1, \beta = 1\mathrm{e}{-6}$. We tried to individually vary each of these values but we did not see any significant change in the resulting image quality (see Fig. 2). Only setting $\hat{N}_v = 1$ or $\hat{N} = 0$ causes sudden increase of image noise since our method with these values essentially degenerates into the maximum likelihood solution.

4 PRIOR ACCURACY

To better understand the importance of the prior of our model and its accuracy, we tested our method with a less precise prior. In particular, we replaced the upper bound $\overline{\cos}\theta_\mathbf{x}$ on the surface cosine in the $\tilde{L}_c(\mathbf{x})$ estimate (Eq. (21)) with a trivial bound of 1. This modification had only a minor effect in most of the scenes except in the City scene, where the trivial bound noticeably increased the image noise (see Fig. 3). This observation is in line with our expectation that the prior is important but our method is not too sensitive to its exact value as it quickly learns the actual lights' contributions.

REFERENCES

Michael Donikian, Bruce Walter, Kavita Bala, Sebastian Fernandez, and Donald P. Greenberg. 2006. Accurate direct illumination using iterative adaptive sampling. *IEEE Transactions on Visualization and Computer Graphics* 12, 3 (2006), 353–363.

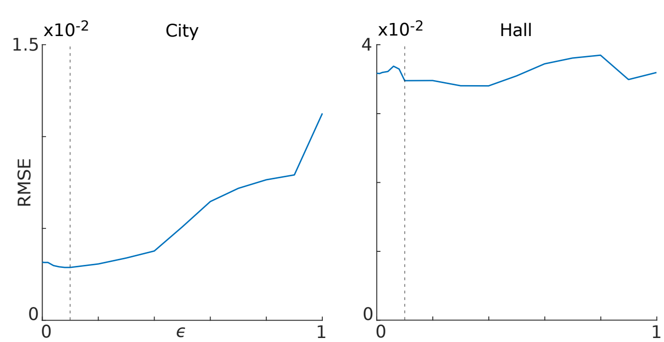

Fig. 1. Plots of RMSE (after 60 s) with respect to the clustering precision ϵ in a direct illumination setting. The dashed line denotes $\epsilon = 0.1$, the value we used in all our tests.

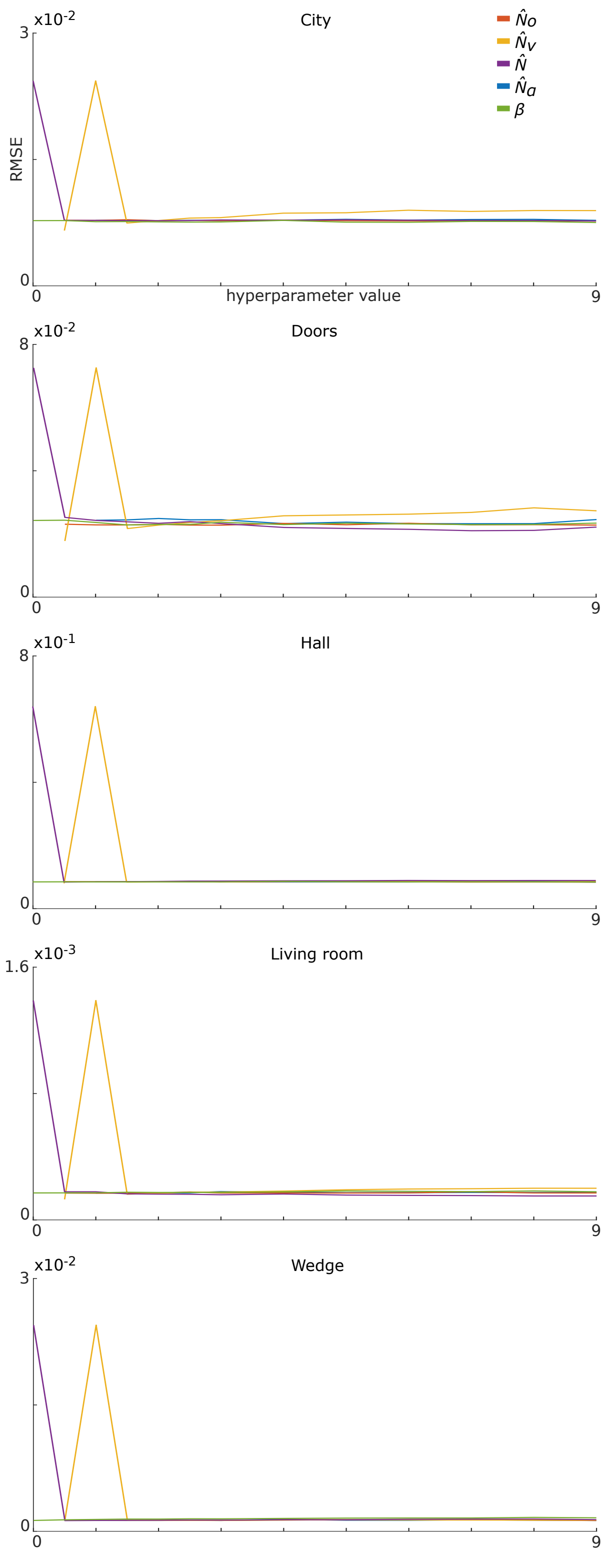

Fig. 2. Plots of RMSE (after 10 s) with respect to different values of the hyperparameters in a direct illumination setting.

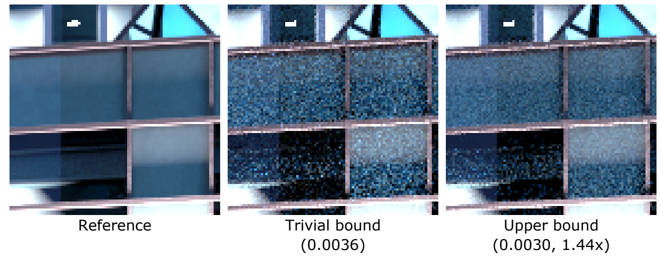

Fig. 3. Equal-time time comparison (60 s) of using a trivial bound on the surface cosine for the model prior against using the upper bound $\overline{\cos}\theta_{\mathbf{x}}$ in a direct illumination setting.