

Neural Radiosity: Errata

- It turns out that the relative residual in section 4.1 is biased due to the correlation between the Monte Carlo samples in the numerator and the denominator. This bias results in darker results. To tackle the issue:
 - One could use a different set of MC samples to compute RHS in the denominator (normalizer) to avoid the correlation. Since this would double the training time, we are not interested.
 - (provided in the code) Another solution would be to use loss annealing, we start with the biased normalizer and anneal the loss to a regular MSE.
 - (provided in the code) Another solution would be to use LHS normalizer instead of $(LHS+RHS)/2$ normalizer. This variant works best in most cases, but sometimes leads to artifacts.