

Sampling Clear Sky Models using Truncated Gaussian Mixtures Supplemental Material

N. Vitsas¹ and K. Vardis¹  and G. Papaioannou¹ 

¹Department of Informatics, Athens University of Economics and Business, Greece

Many modern rendering frameworks rely on Quasi Monte Carlo integration in order to distribute a given number of samples over the sampling domain more evenly. These simulations use low-discrepancy sequences for random number generation or stratified sampling. In Figures 1 and 2, we demonstrate how samples are dispersed on the hemispherical sky domain for a Halton sequence and stratified rectangular grid sampling. We also present the corresponding mapped points using a SAT-based approach and show the resulting illumination on a simple environment, where potential differences can be more easily spotted.

Certain applications require that the stratification of the initial point distribution on the unit square is maintained after the inverse

CDF mapping. The standard and most efficient way to sample from GMMs and by extension, tGMMs, does not preserve the strata, since one random variable is also used for switching between Gaussian lobes. However, this has no meaningful negative implication for sampling tasks similar to ours, as demonstrated in the above experiments but also in the extreme case of a regular grid of samples on the unit square, which is shown in Figure 3.

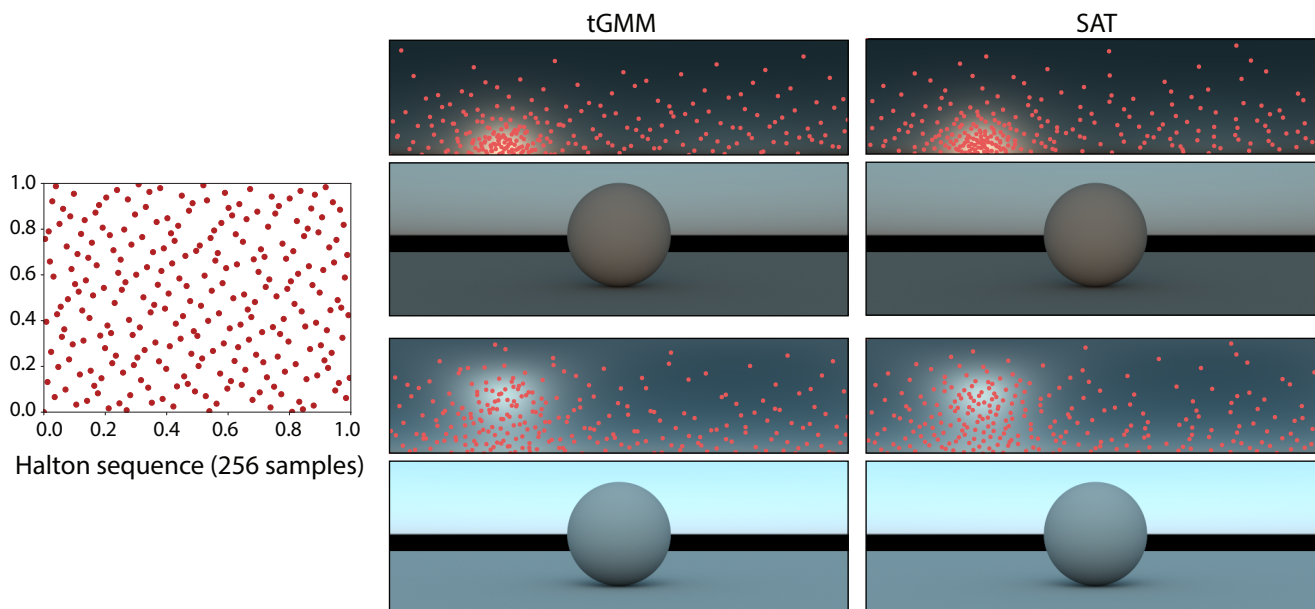


Figure 1: The mapping of a Halton sequence of 256 samples from the unit square to the sky dome for two different clear sky instances. We compare the results from our sampling approach with typical SAT-based environment sampling.

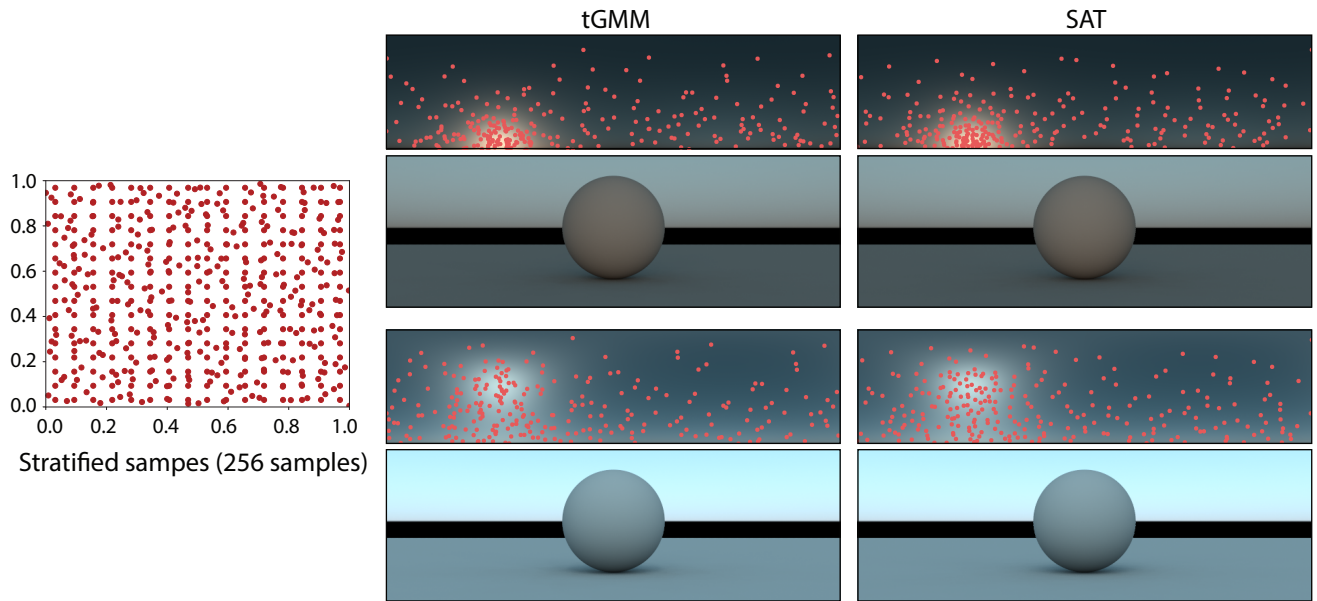


Figure 2: The mapping of a set of 256 stratified samples from the unit square to the sky dome for two different clear sky instances. We compare the results from our sampling approach with typical SAT-based environment sampling.

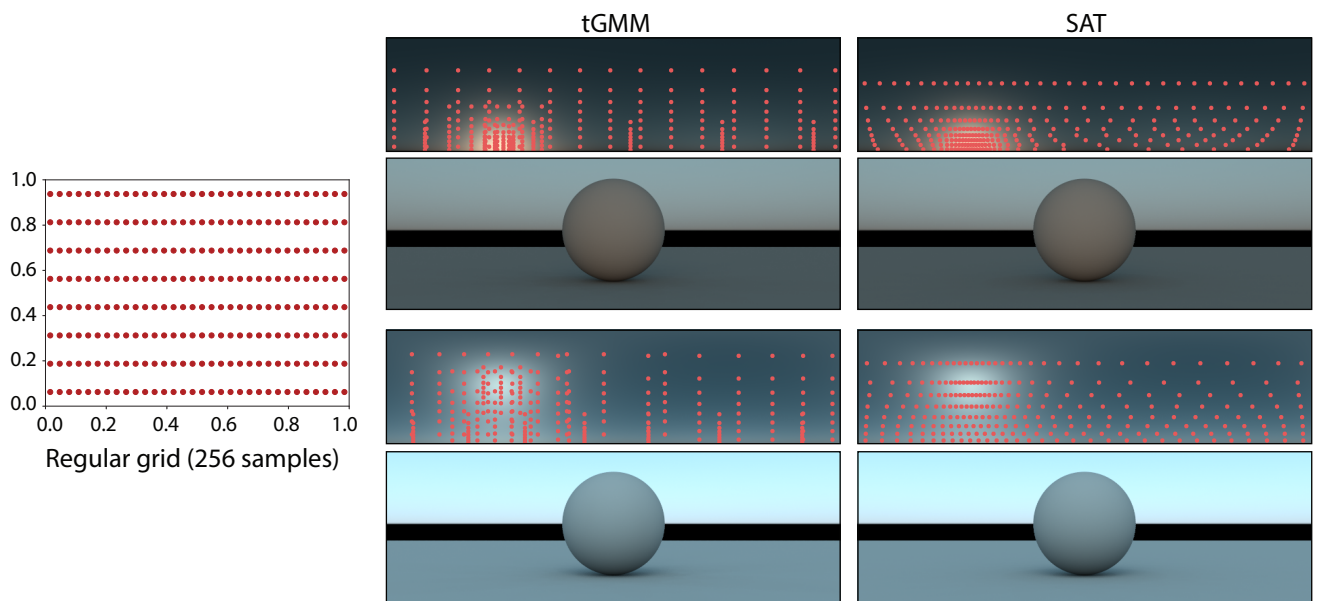


Figure 3: The mapping of a regular grid of samples from the unit square to the sky dome for two different clear sky instances. We compare the results from our sampling approach with typical SAT-based environment sampling. Sample density matches the proportions of the output domain ($32 \times 8 = 256$)