Sampling of Anisotropic Spatial Gaussians for Path Guiding

S. Lelyakin, V. Schüßler and C. Dachsbacher

Karlsruhe Institute of Technology, Germany

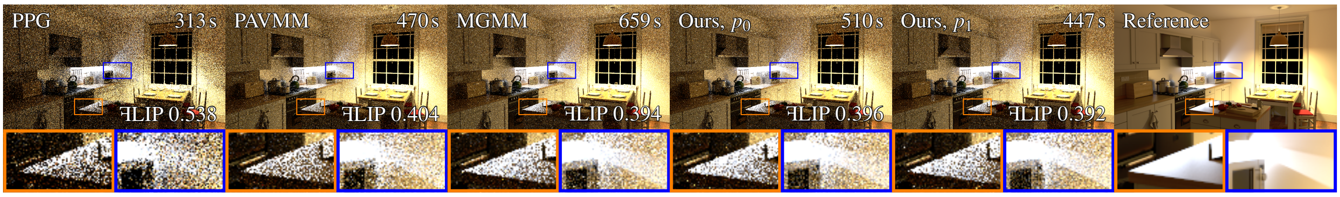

Figure 1: *Anisotropic spatial Gaussians, sampled either using our method or MGMM [SHJD22], improve guiding quality compared to isotropic (PAVMM [RHL20]) and directional models without reprojection (PPG [MGN17]). In comparison to MGMM, our sampling method is faster at approximately equal quality. All renders use 128 samples per pixel.*

Abstract

Directional models in path guiding struggle with representing parallax effects or anisotropic features. Our model instead describes the spatial distribution of a target vertex using a 3D Gaussian mixture model. While this dispenses with the need for reprojection and allows to represent anisotropic features easily, its directional probability density is not readily available, since it involves a marginal integral. In this work, we derive an expression for the PDF of our model in solid angle measure that is practical to evaluate. We demonstrate how our model can improve guiding accuracy in various scenes.

CCS Concepts

• ***Computing methodologies*** → ***Rendering; Ray tracing;***

1. Introduction

A key challenge in path guiding [VHH*19] is to accurately represent the directional distribution of the next path segment, conditional on the last path vertex position. One common approach is to place directional mixture models in cells of a spatial data structure. Its weakness in handling parallax effects within cells has been addressed by blurring [VHH*19, Ch. 10] or reprojection [RHL20]. While the former inevitably reduces guiding precision, the latter is based on isotropic mixture components. Consequently, anisotropic features require many components for a faithful representation.

Instead of a directional model, we use a Gaussian mixture model (GMM) to represent the 3D spatial distribution of a targeted path vertex. To importance sample a direction, we first sample a target vertex, and then normalize the vector from current to target vertex. Our model inherently handles parallax effects and anisotropic features, including projection effects like foreshortening. Unfortunately, its directional probability density function (PDF) is not readily available, due to the normalization step. We solve a marginal integral to derive the PDF p_0 and show how to optionally eliminate the error function in the resulting expression by restricting sampling to a hemisphere. This results in a variant of our sampling method with a simpler PDF p_1.

Related work. Schüßler et al. [SHJD22] also sample spatial 3D Gaussians. However, their sampling only approximates the 3D density by projecting onto a 2D plane in order to marginalize before sampling. Rath et al. [RYS23] sample focal points, which are represented as a spatial density by an octree. They traverse it to evaluate the marginal integral in the directional PDF. Our model could similarly be used to represent focal points, provided with an efficient fitting method. Huang et al. [HIT*24] introduce an anisotropic directional distribution for path guiding, but rely on neural networks and gradient descent for reprojection and fitting, respectively.

2. Sampling Method and Directional PDF

We fit our GMM similar to previous work [RHL20]. For sampling, we reweigh components based on the cosine of the angle between the surface normal and the direction to the lobe mean. In the following, we describe sampling and PDF evaluation of a single component $\mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma})$ with mean $\boldsymbol{\mu}_0 \in \mathbb{R}^3$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{3\times 3}$.

Sampling. To sample a direction $\boldsymbol{\omega}$ for the next ray from vertex $\boldsymbol{x}_0$, we first sample a point $\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma})$ and then use $\boldsymbol{\omega} = \frac{\boldsymbol{x}-\boldsymbol{x}_0}{\|\boldsymbol{x}-\boldsymbol{x}_0\|}$. Optionally, we flip $\boldsymbol{\omega}$ if $\boldsymbol{\omega} \cdot (\boldsymbol{\mu}_0 - \boldsymbol{x}_0) < 0$ to get a simpler PDF.

DOI: 10.2312/egp.20251017

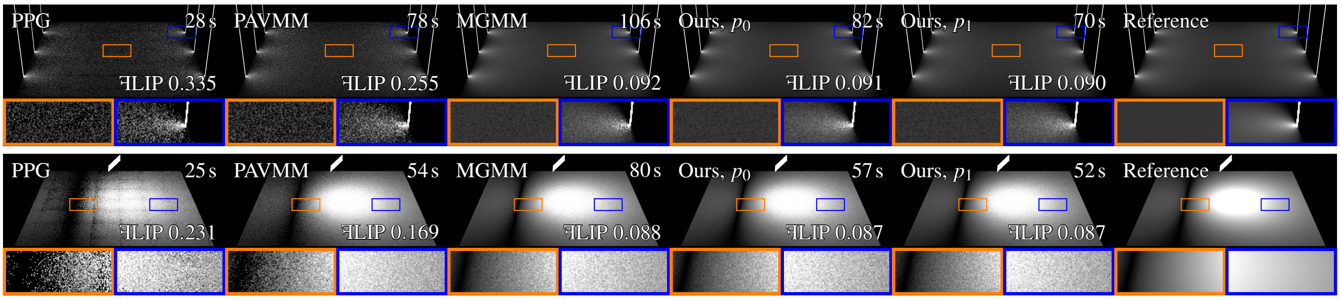

Figure 2: *Scene 1 (top) highlights the potential of anisotropic guiding distributions (MGMM and ours). The blue inset shows an example where the MGMM approximation is slightly suboptimal close to the light source. For scene 2 (bottom), we use one global GMM only (except for PPG). This demonstrates that our model is able to adapt to different vertex positions and to describe foreshortening of a light source well.*

PDF evaluation. We obtain the directional PDF in solid angle measure by integrating the spatial PDF over the entire ray corresponding to the direction $\boldsymbol{\omega}$:

$$p_0(\boldsymbol{\omega}) = \int_0^\infty r^2 \mathcal{N}(r\boldsymbol{\omega} \mid \boldsymbol{\mu}_0 - \boldsymbol{x}_0, \boldsymbol{\Sigma})\, \mathrm{d}r. \tag{1}$$

Introducing some abbreviations, this resolves to

$$\begin{aligned} p_0(\boldsymbol{\omega}) &= (2\pi)^{-\frac{3}{2}} |\boldsymbol{\Sigma}|^{-\frac{1}{2}} \int_0^\infty r^2 \mathrm{e}^{-\frac{a}{2}r^2 + br - \frac{c}{2}}\, \mathrm{d}r \\ &= \frac{1}{2} (2\pi)^{-\frac{3}{2}} |\boldsymbol{\Sigma}|^{-\frac{1}{2}} a^{-\frac{5}{2}} \mathrm{e}^{-\frac{c}{2}} \\ &\quad \left(\sqrt{2\pi} \left(a + b^2\right) \left(\operatorname{erf}\left(\frac{b}{\sqrt{2a}}\right) + 1 \right) \mathrm{e}^{\frac{b^2}{2a}} + 2b\sqrt{a} \right), \end{aligned} \tag{2}$$

where $a = \boldsymbol{\omega}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\omega}$, $b = \boldsymbol{\mu}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\omega}$, $c = \boldsymbol{\mu}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$, and $\boldsymbol{\mu} = \boldsymbol{\mu}_0 - \boldsymbol{x}_0$.

For the variant where we flip the sampled direction towards a predetermined hemisphere, the lower integration bound extends to $-\infty$:

$$\begin{aligned} p_1(\boldsymbol{\omega}) &= \int_{-\infty}^\infty r^2 N(r\boldsymbol{\omega} \mid \boldsymbol{\mu}_0 - \boldsymbol{x}_0, \boldsymbol{\Sigma})\, \mathrm{d}r \\ &= (2\pi)^{-1} |\boldsymbol{\Sigma}|^{-\frac{1}{2}} a^{-\frac{5}{2}} \left(a + b^2\right) \mathrm{e}^{\frac{b^2}{2a} - \frac{c}{2}}. \end{aligned} \tag{3}$$

To improve numerical stability, we shift the subtractive cancellation in the exponent term into a cross product by using the identity

$$\boldsymbol{\omega}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\omega}\, \boldsymbol{\mu}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - \left(\boldsymbol{\omega}^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}\right)^2 = \frac{1}{|\boldsymbol{\Sigma}|} (\boldsymbol{\omega} \times \boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Sigma} (\boldsymbol{\omega} \times \boldsymbol{\mu}). \tag{4}$$

We thus obtain an alternate form of the exponent:

$$\frac{b^2}{2a} - \frac{c}{2} = -\frac{1}{2a|\boldsymbol{\Sigma}|} (\boldsymbol{\omega} \times \boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Sigma} (\boldsymbol{\omega} \times \boldsymbol{\mu}). \tag{5}$$

3. Results

We compare our approach with equal number of samples per pixel (spp) against the guiding methods by Müller et al. [MGN17] (PPG) and Ruppert et al. [RHL20] (PAVMM). We also compare our sampling methods, using either p_0 or p_1, against the approximate sampling method by Schüßler et al. [SHJD22] (MGMM), where both our methods and MGMM are used to sample our proposed 3D GMM model.

In fig. 2 (32 spp), we demonstrate the advantage of our model in representing anisotropic features. We also show how anisotropy can occur with mostly isotropic geometry due to foreshortening. In a more complex scene (fig. 1), improvements are more subtle, but still measurable.

The three methods using our GMM (MGMM and our two) differ only marginally. This is to be expected: Flipped directions only occur if a component contains the sample origin. As this is effectively never the case in the given scenes, the difference between p_0 and p_1 is close to zero. As for MGMM, the approximation it offers is also good in most cases, although a slight deterioration can be observed near the light source in the blue inset of fig. 2, scene 1. Furthermore, both our methods are consistently faster than MGMM.

4. Conclusion

We presented a method for sampling anisotropic spatial Gaussians that is both exact and more computationally efficient than previous work. Making full use of the projective nature of our model may be interesting for future work, e.g. to guide towards focal points.

References

[HIT*24] Huang, Jiawei, Iizuka, Akito, Tanaka, Hajime, et al. "Online Neural Path Guiding with Normalized Anisotropic Spherical Gaussians". *ACM TOG* 43.3 (Apr. 2024). DOI: 10.1145/3649310 1.

[MGN17] Müller, Thomas, Gross, Markus, and Novák, Jan. "Practical Path Guiding for Efficient Light-Transport Simulation". *CGF (Proc. EGSR)* 36.4 (June 2017), 91–100. DOI: 10/gbnvrs 1, 2.

[RHL20] Ruppert, Lukas, Herholz, Sebastian, and Lensch, Hendrik P. A. "Robust Fitting of Parallax-Aware Mixtures for Path Guiding". *ACM TOG (Proc. SIGGRAPH)* 39.4 (July 2020). DOI: 10/gg8xc6 1, 2.

[RYS23] Rath, Alexander, Yazici, Ömercan, and Slusallek, Philipp. "Focal Path Guiding for Light Transport Simulation". *ACM SIGGRAPH Conference Papers*. July 2023. DOI: 10/gsk4jp 1.

[SHJD22] Schüssler, Vincent, Hanika, Johannes, Jung, Alisa, and Dachsbacher, Carsten. "Path Guiding with Vertex Triplet Distributions". *CGF (Proc. EGSR)* 41.4 (2022). DOI: 10/g8rsdj 1, 2.

[VHH*19] Vorba, Jiří, Hanika, Johannes, Herholz, Sebastian, et al. "Path Guiding in Production". *ACM SIGGRAPH Courses*. 2019, 18:1–18:77. DOI: 10.1145/3305366.3328091 1.