5D Radiance Fields From Local Regressions

Rodolphe Valicon De Soete, Colas Schretter, Peter Schelkens, Nikos Deligiannis
Department of Electronics and Informatics, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium
imec, Kapeldreef 75, B-3001 Leuven, Belgium

***Abstract*—We introduce mixtures of spatio-directional (5D) local regressions as a novel radiance field representation based on classical semi-parametric statistical models. The spatial domain is covered with 3D multivariate Gaussian kernels and the directional domain is encoded by 2D Kent distributions. The field of radiometric emissions is defined by mixing local linear regressors. Volume rendering with parallel numerical integrations enables the reconstruction of a sharp and continuous radiance field from few mixture components without noticeable visual artifacts. Additionally, directional distributions are attached to each kernel, enabling the representation of arbitrarily complex anisotropic emissive behaviors. The local support of each reconstruction primitive enables fast view synthesis using hardware rasterization on GPU and distances from arbitrary viewpoints may be computed to define physical bounds from the model, enabling precise collision detection. We present preliminary experiments on a synthetic model and analyze the potential of the technique in terms of accuracy with a ray marching numerical integrator.**

***Index Terms*—Radiance Field, Volume rendering, Image representation, Semi-parametric mixture models, Linear regression.**

I. Introduction

In 1991, the plenoptic function [1] conceptualized the idea of a complete mathematical description of the continuous visual cues that may be perceived, yet without proposing a practical method for measuring and storing the radiance field.

Recent results on representing and reconstructing the 5D spatio-directional radiance field from scattered 2D views have shown groundbreaking abilities at capturing complex realistic scenes. In 2020, Neural Radiance Field (NeRF) [2] proposed a deep neural network-based solution to store fair approximations of the plenoptic function. Unfortunately, the non-local nature of the method makes it suffer from long training times and poor inference performance, preventing real-time view synthesis on mobile devices with large models.

Local primitive splatting techniques [3] are the current state of the art for rendering radiance fields. In 2023, 3D Gaussian Splatting [4] revamped the traditional point-based rendering techniques [5] and in particular, that work leveraged the Elliptical-Weighted-Average (EWA) volume splatting algorithm by Zwicker [6]. This technique offers shorter training time and better inference performance compared to NeRF.

While splatting-based techniques are efficient, they suffer from severe visual artifacts caused by discontinuities between local primitives. Therefore, these models are fairly large. Hierarchical approximation may provide adaptive levels of details depending on the view point [7] and localized streaming.

This work was funded by the UNIMODE project (Nr. G0B3521N) of the Research Foundation - Flanders (FWO) and the ERC Consolidator Grant IONIAN (Nr. 101171240), funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the EU nor the granting authority can be held responsible.

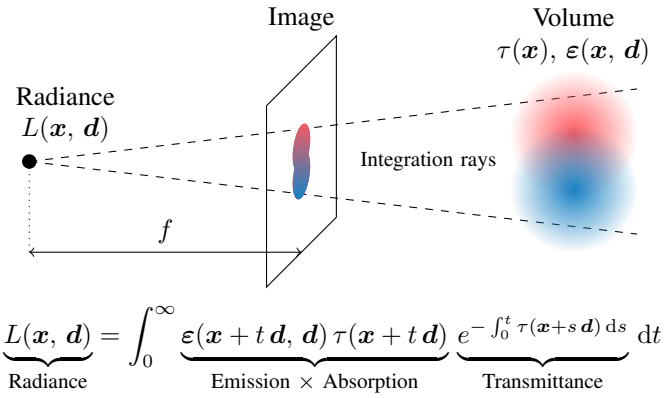

Fig. 1. Image formation from a spatio-directional (5D) radiance field.

In 2025, the original authors of 3DGS identified the need for accurate radiance integration of the radiance along the viewing rays [8]. The results demonstrated better quality and removed the popping artifacts, but they concluded that, despite their improvements, they still need a large number of small kernels.

In this work, we introduce a novel radiance field representation, based on well-known classical statistical models. Our model is a generalization from our prior work on representing images with local linear regressors [9]. We cover the volumetric domain with 5D spatio-directional marginal kernels. We associate to these kernels, local linear regressors that reconstruct the volume's radiance from any spatial point and direction. We implemented a numerical integrator, on GPU, that enables the rendition of sharp views from sparse models.

II. Radiance Field Model

Our radiance field representation is based on the classical volumetric light transport equation within a non-homogeneous medium, assuming no scattering, as illustrated in Fig. 1.

The absorption function $\tau : \mathbb{R}^3 \to \mathbb{R}_+$ gives the density of the volume at any point $\boldsymbol{x} \in \mathbb{R}^3$. The denser the volume is, the more light is absorbed while traversing it. The emission function $\boldsymbol{\varepsilon} : \mathbb{R}^3 \times \mathcal{S}^2 \to \mathcal{C}$ gives the light intensity emitted at any point $\boldsymbol{x}$ and in any unit direction vector $\boldsymbol{d} \in \mathbb{R}^3$.

The radiance field $L : \mathbb{R}^3 \times \mathcal{S}^2 \to \mathcal{C}$ caused by the volume can be evaluated at any point $\boldsymbol{x}$ and in any direction $\boldsymbol{d}$, by casting a ray inside the volume, and integrating the absorption and emission along the ray's path with the equation in Fig. 1.

Spatial (3D) support	**Directional (2D) support**
Position	Orientation
$\boldsymbol{\mu} = (x \quad y \quad z)^\top$ (2)	$\boldsymbol{\omega} = (\theta \quad \phi)^\top$ (15)
Gaussian distribution	Kent distribution
$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_x^2 & \cdot & \cdot \\ \sigma_{xy} & \sigma_y^2 & \cdot \\ \sigma_{xz} & \sigma_{yz} & \sigma_z^2 \end{pmatrix}$ (3)	$\Theta = \begin{Bmatrix} \kappa \\ \beta \\ \psi \end{Bmatrix}$ (16)

Linear regressors

Spatial absorption	Spatial gradients on Y
$\bar{\tau}$ (9)	$\boldsymbol{\delta}^{\boldsymbol{x}} = (\delta_x \quad \delta_y \quad \delta_z)^\top$ (12)
Mean intensity and color	Directional gradients on Y
$\bar{\boldsymbol{\varepsilon}} = \{Y,\, C_b,\, C_r\}$ (10)	$\boldsymbol{\delta}^{\boldsymbol{d}} = (\delta_\theta \quad \delta_\phi)^\top$ (22)

Fig. 2. Notations for spatio-directional mixture model's components.

In the following sections, we define the spatial absorption and spatio-directional emission functions with mixture models whose parameters are summarized in Fig. 2.

A. Spatial absorption function

We approximate the volumetric absorption function with

$$\tau(\boldsymbol{x}) = \sum_{k=1}^{K} \bar{\tau}_k \, p\left(k \mid \boldsymbol{x}\right). \tag{1}$$

This mixture model comprises a set of K peak-normalized 3D multivariate Gaussian kernels, defined by their means

$$\boldsymbol{\mu}_k = (x \quad y \quad z)^\top, \tag{2}$$

and variance-covariance matrices

$$\boldsymbol{\Sigma}_k = \begin{pmatrix} \sigma_x^2 & \cdot & \cdot \\ \sigma_{xy} & \sigma_y^2 & \cdot \\ \sigma_{xz} & \sigma_{yz} & \sigma_z^2 \end{pmatrix}. \tag{3}$$

As the matrices $\boldsymbol{\Sigma}$ are positive definite and symmetric, they can be fully encoded by the six scalars of their lower triangle or the non-zero elements of a Cholesky decomposition.

The marginal density (3D support) of Gaussian kernels is

$$p(\boldsymbol{x} \mid k) = \exp\left(-\frac{1}{2} \left\|\boldsymbol{x} - \boldsymbol{\mu}_k\right\|^2_{\boldsymbol{\Sigma}_k}\right), \tag{4}$$

where the squared Mahalanobis distance between $\boldsymbol{x}$ and $\boldsymbol{\mu}$ is

$$\left\|\boldsymbol{x} - \boldsymbol{\mu}_k\right\|^2_{\boldsymbol{\Sigma}_k} = \left(\boldsymbol{x} - \boldsymbol{\mu}_k\right)^\top \boldsymbol{\Sigma}_k^{-1} \left(\boldsymbol{x} - \boldsymbol{\mu}_k\right). \tag{5}$$

The gating function $p(k \mid \boldsymbol{x})$ is the partition of unity for the k^{th} kernel, representing its spatial relative importance compared to the rest of the mixture at the point $\boldsymbol{x}$:

$$p(k \mid \boldsymbol{x}) = \frac{p(\boldsymbol{x} \mid k)}{\sum_{i=1}^{K} p(\boldsymbol{x} \mid i)}. \tag{6}$$

By expanding the expression of $p(k \mid \boldsymbol{x})$ using (4), we obtain

$$p(k \mid \boldsymbol{x}) = \frac{\exp\left(-\frac{1}{2} \left\|\boldsymbol{x} - \boldsymbol{\mu}_k\right\|^2_{\boldsymbol{\Sigma}_k}\right)}{\sum_{i=1}^{K} \exp\left(-\frac{1}{2} \left\|\boldsymbol{x} - \boldsymbol{\mu}_i\right\|^2_{\boldsymbol{\Sigma}_i}\right)} \tag{7}$$

$$= \operatorname{softmin}\left[\frac{1}{2} \left\|\boldsymbol{x} - \boldsymbol{\mu}\right\|^2_{\boldsymbol{\Sigma}}\right]_k. \tag{8}$$

To reconstruct the continuous absorption function, we attach to each kernel, a constant spatial regressor with mean value

$$\bar{\tau}_k > 0\,. \tag{9}$$

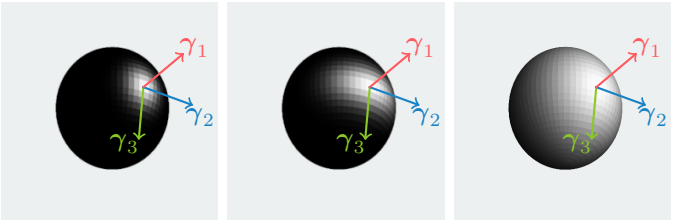

Fig. 3. Directional supports for various parameters of the Kent distribution.

B. Spatio-directional emission function

Similarly, to reconstruct the continuous emission function, we attach to each kernel, a spatial regressor with mean color

$$\bar{\boldsymbol{\varepsilon}}_k = \{Y,\, C_b,\, C_r\}\,. \tag{10}$$

Moreover, we model the spatial variation of the emission intensity using a first-order linear regressor

$$Y_k(\boldsymbol{x}) = Y_k + \boldsymbol{\delta}_k^{\boldsymbol{x}} \cdot \left(\boldsymbol{x} - \boldsymbol{\mu}_k\right), \tag{11}$$

where the vector

$$\boldsymbol{\delta}_k^{\boldsymbol{x}} = (\delta_x \quad \delta_y \quad \delta_z)^\top \in \mathbb{R}^3 \tag{12}$$

contains the three spatial emission gradient parameters. We only use gradients for the luminance component Y because human perception is less sensitive to chrominance variations.

The model presented so far is not capable of representing the directional variations of the emission function. This can be solved by introducing, for each regressor, an additional modulation factor $p(\boldsymbol{d} \mid k)$ that depends on $\boldsymbol{d}$ such that

$$\boldsymbol{\varepsilon}(\boldsymbol{x},\, \boldsymbol{d}) = \sum_{k=1}^{K} \boldsymbol{\varepsilon}_k\left(\boldsymbol{x},\, \boldsymbol{d}\right) p\left(k \mid \boldsymbol{x},\, \boldsymbol{d}\right). \tag{13}$$

with the joint spatio-radiometric density

$$p(\boldsymbol{x},\, \boldsymbol{d} \mid k) = p(\boldsymbol{x} \mid k)\, p(\boldsymbol{d} \mid k)\,. \tag{14}$$

With this addition, each kernel has a distribution of directions conditioning its view-dependent importance.

The direction-dependence $p(\boldsymbol{d} \mid k)$ is modulated by a peak-normalized Kent distribution [10], which is analogous to a wrapped Gaussian distribution on the sphere oriented at

$$\boldsymbol{\omega}_k = (\theta \quad \phi)^\top, \tag{15}$$

whose shape is defined with a set of three parameters

$$\Theta_k = \{\kappa,\, \beta,\, \psi\}\,, \tag{16}$$

where κ controls the distribution's concentration around the mean direction, β controls the ellipticity, and ψ tunes the orientation around the mean direction when $0 < \beta < \kappa$. Fig. 3 shows some visualizations of this statistical distribution.

The marginal density (2D support) of Kent distributions is

$$p(\boldsymbol{d} \mid k) = \exp\left(-\frac{1}{2}\left\|\boldsymbol{d}\cdot\boldsymbol{\gamma}_1\right\|^2_{\Theta_k}\right), \tag{17}$$

with the directional squared Mahalanobis distance

$$\left\|\boldsymbol{d}\cdot\boldsymbol{\gamma}_1\right\|^2_{\Theta_k} = \kappa_k\left(1 - \boldsymbol{d}\cdot\boldsymbol{\gamma}_1\right) + \beta_k\left[(\boldsymbol{d}\cdot\boldsymbol{\gamma}_3)^2 - (\boldsymbol{d}\cdot\boldsymbol{\gamma}_2)^2\right], \tag{18}$$

where a unit mean direction vector is constructed from $\boldsymbol{\omega}_k$:

$$\boldsymbol{\gamma}_1 = (\sin\theta_k\cos\phi_k \quad \sin\theta_k\sin\phi_k \quad \cos\theta_k)^\top \in \mathbb{R}^3, \tag{19}$$

and two additional basis vectors $\boldsymbol{\gamma}_2,\, \boldsymbol{\gamma}_3 \in \mathbb{R}^3$ are constructed orthogonally to $\boldsymbol{\gamma}_1$, using the orientation parameter ψ_k, for defining the major and minor axes.

We combine (18) with the spatial squared Mahalanobis distance (5) for writing the spatio-directional partition of unity

$$p(k \mid \boldsymbol{x},\, \boldsymbol{d}) = \operatorname{softmin}\left[\frac{1}{2}\left(\left\|\boldsymbol{x} - \boldsymbol{\mu}\right\|^2_{\boldsymbol{\Sigma}} + \left\|\boldsymbol{d}\cdot\boldsymbol{\gamma}_1\right\|^2_{\Theta}\right)\right]_k, \tag{20}$$

and using (11), we add a directionally-dependent geodesic regressor as a function of the spherical log-map approximation on the tangent plane defined by $\boldsymbol{\gamma}_2$ and $\boldsymbol{\gamma}_3$ for avoiding explicit computation of the great circle angle $\alpha = \arccos(\boldsymbol{d}\cdot\boldsymbol{\gamma}_1)$:

$$\boldsymbol{\varepsilon}_k(\boldsymbol{x},\, \boldsymbol{d}) = \boldsymbol{\varepsilon}_k(\boldsymbol{x}) + \boldsymbol{\delta}^{\boldsymbol{d}}_k \cdot (\boldsymbol{d}\cdot\boldsymbol{\gamma}_2 \quad \boldsymbol{d}\cdot\boldsymbol{\gamma}_3)^\top \tag{21}$$

where

$$\boldsymbol{\delta}^{\boldsymbol{d}}_k = (\delta_\theta \quad \delta_\phi)^\top \in \mathbb{R}^2 \tag{22}$$

is a vector containing the two directional gradient parameters.

III. View Synthesis

A view from the model is obtained by evaluating the received radiance at the camera's position $\boldsymbol{x}$ for each pixel direction $\boldsymbol{d}$. To numerically compute the light transport integral in Fig. 1, we can use the well-known Nelson Max's quadrature [11]. This approximation takes N consecutive samples along the ray and computes the radiance as

$$L(\boldsymbol{x},\, \boldsymbol{d}) \approx \sum_{n=1}^{N} T_n\left[1 - \exp(-\tau(\boldsymbol{x}_n)\,\Delta)\right]\boldsymbol{\varepsilon}(\boldsymbol{x}_n,\, \boldsymbol{d})\,, \tag{23}$$

where

$$T_n = \exp\left(\sum_{m=1}^{n-1} -\tau(\boldsymbol{x}_m)\,\Delta\right) \tag{24}$$

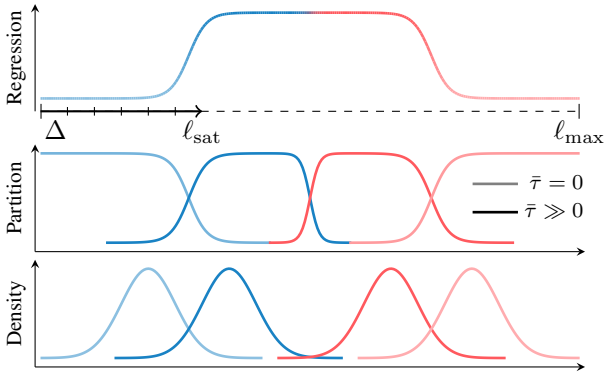

Fig. 4. Integration along a ray within a volumetric radiance field.

is the transmittance at the n^{th} integration step, due to Beer-Lambert's law. The $\boldsymbol{x}_n$ are the sampling positions and Δ is the step size between successive samples, such that

$$\boldsymbol{x}_n = n\,\Delta\,\boldsymbol{d} \quad \text{with} \quad \Delta = \frac{\ell_{\max}}{N}\,, \tag{25}$$

where $\boldsymbol{d}$ is the ray direction vector from the viewpoint $\boldsymbol{x}$, that is computed from the pixel's coordinates $(u \quad v)^\top$.

The reconstructed volume's functions τ and $\boldsymbol{\varepsilon}$ are sampled by evaluating the mixture of local regression along the integration rays. This fixed step size integrator corresponds to a simple Digital Differential Analyzer (DDA) and the accuracy improves with smaller step sizes as shown in Fig. 6.

To improve numerical accuracy, we account for the skipped regions in between successive samples along the ray by convolving the model with an additive smoothing factor and compensate accordingly the spatial absorption factors in order to preserve the constant integrated spatial density. The optimal smoothing variance is set by the least square fit of the uniform sample's spacing as shown in the plots in Fig. 7.

Special care should be taken when sampling the partitions of unity $p(\,k \,|\, \boldsymbol{x}\,)$ and $p(\,k \,|\, \boldsymbol{x},\, \boldsymbol{d}\,)$, as their naive evaluation can be numerically ill-conditioned in regions that are sparsely covered by support kernels. To avoid numerical instabilities, we should compute them using the online safe softmax algorithm [12], which can be done in a single pass as specified in Algorithm 1.

Note that the integral might saturate before the ray reaches maximum distance $\ell_{\max}$, wasting computational resources for every subsequent integration step that won't contribute to the incident radiance. In practice, we stop the numerical integration early at a saturation distance ℓ_{sat}, once the transmittance falls below $T_{\min} = \frac{1}{255}$. As a byproduct, this stopping criterion generates an accurate depth map from the camera's viewpoint.

We implemented the system on a Nvidia RTX 4070 laptop GPU and the ray integrator required 20.1ms for rendering each of the anti-aliased 512×512 pixels images in Fig. 5, with a very large number of 512 integration steps. A pre-render pass may collect lists per pixel of the most important kernels (TopK) such that the integration will be independent of the model size. Only a few kernels modulate the radiance along any given ray, and integration is invariant with the model size.

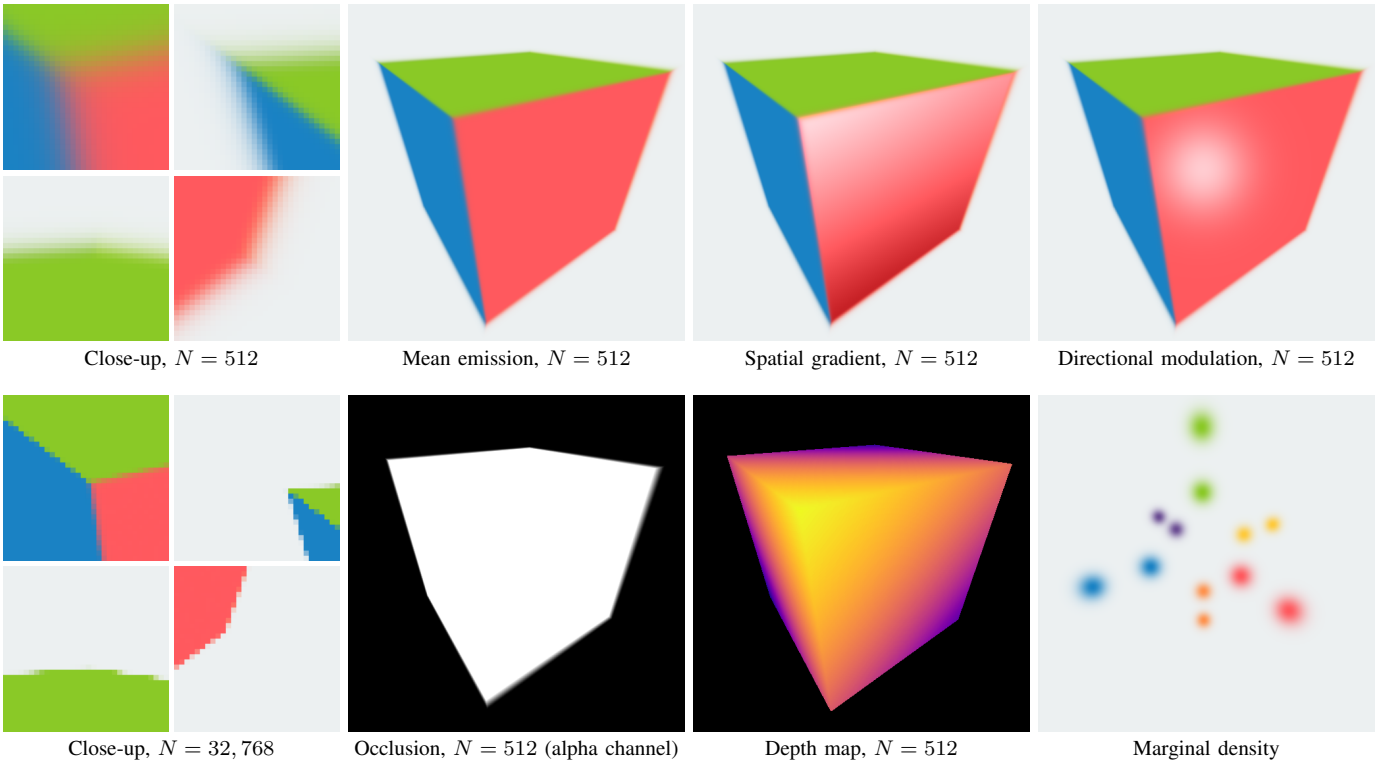

Fig. 5. Rendering of a synthetic model of a cube defined by six pairs of spatial kernels, with spatial gradients and directional emission modulation.

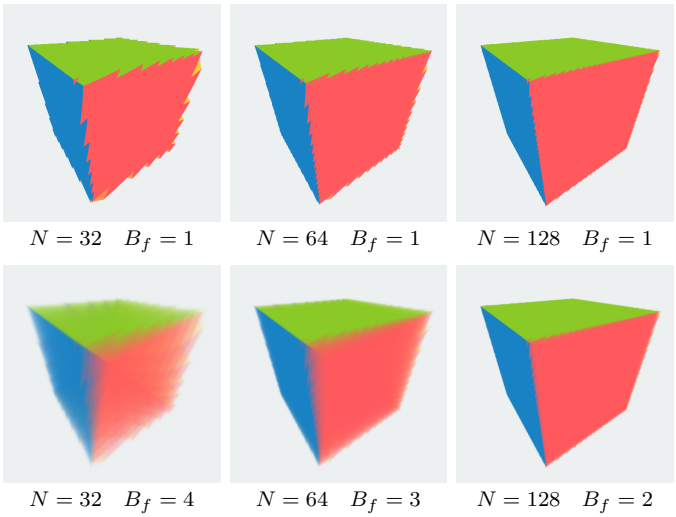

Fig. 6. Progressive improvement with integration steps and model smoothing.

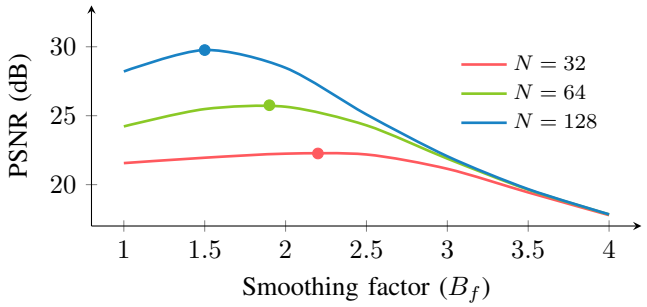

Fig. 7. Quality of numerical integration as a function of model smoothing.

IV. Results

In our experiments, we evaluated the numerical aspects of the integration inside a volume defined with our radiance field representation. We experimented on a numerically challenging, sparse model representing a standard cube with colored faces, as shown in Fig. 5 and Fig. 6.

In Fig. 5, we show the effect of an emission gradient with respect to the scene's vertical axis. Only the component of the red face is modified. For the directional modulation, we superimposed an extra component with white emission, on the inner red face's kernel, with the same spatial shape but with a directional support in the direction of the camera. It gives a highlight effect, as the relative importance of this extra kernel decreases as the rays directions deviate from the kernel's mean. In our model, six pairs of kernels encode the partition between empty and solid regions with emissive colors.

Fig. 6 shows the progressive accuracy with increasing number of integration steps. Adaptive smoothing acts as filter for avoiding aliasing along the sampling path. The optimal smoothing factor depends on the step size as shown in Fig. 7. The reference was obtained on a sharp model (with a smoothing factor $B_f = 1$) with up to 32,768 integration steps. The smooth blend between mixture components performs anti-aliasing out-of-the-box when no singular kernel is present.

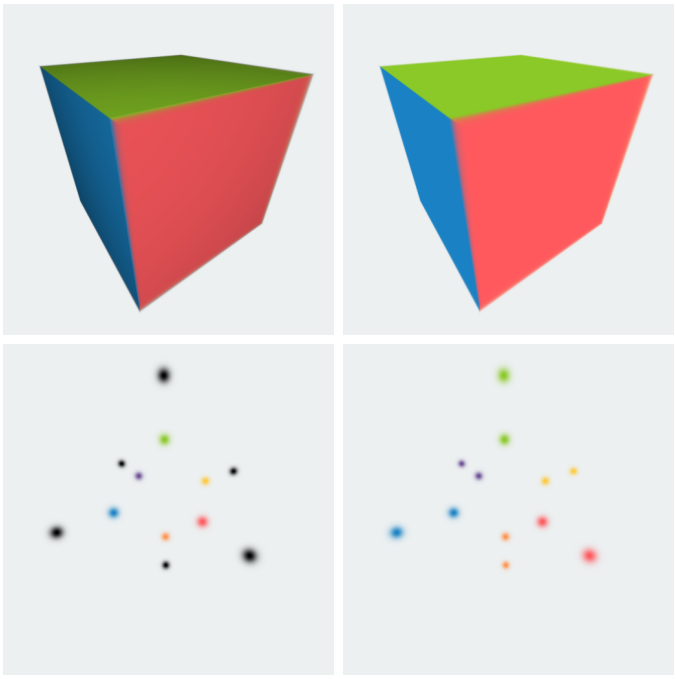

Fig. 8. Effect of the interplay between absorption and emission coefficients.

Fig. 8 demonstrates the darkening effect that depends on the ray's angle with the surface: a slanted incident ray travels a longer distance in this low emission and absorbing layer, than a ray orthogonal to the surface. This generates a visual gradient as if the cube was illuminated from the camera's viewpoint. The outer components have the same mean emission $\boldsymbol{\varepsilon}$ than the inner components. The reason behind this, is that we want the volume to emit the color of the faces as soon as the density increases. When they have zero mean emission, the mix between the face's components constructs a layer with low emission and medium absorption in front of the fully emissive and absorbing layer represented by the inner component. This acts as a semi-transparent dark filter in front of the face, diminishing the perceived radiance.

V. Conclusion

This work introduced a physically-based statistical model based on mixtures of local linear regressions for representing piecewise smooth radiances by representing both the emissive (solid) and the transparent (empty) regions of the volumetric domain. A partition of unity may therefore be computed for recovering a well-defined radiance from arbitrary viewpoints and ray directions within the domain.

In further works, we want to estimate the model parameters from scattered views, in order to perform novel view synthesis of real and synthetic scenes. As the model is based on Gaussian kernels, an expectation-maximisation (EM) method could be used to efficiently estimate the parameters. An adaptive kernel allocation technique could also be used to achieve a trade-off between quality and model complexity. We also want to tackle model compression for progressive streaming.

Algorithm 1 Radiance integration along a ray $(\boldsymbol{x}, \boldsymbol{d})$

$L \leftarrow (0,0,0)$ ▷ Initialize the output radiance
$\eta \leftarrow 0$ ▷ Start without any occlusion
$\Delta \leftarrow \ell_{\text{max}}/N$ ▷ Step size between samples
for $n \leftarrow 1$ **to** N **do** ▷ Numerical integration loop
  $\boldsymbol{x}_n \leftarrow \boldsymbol{x} + n\,\Delta\,\boldsymbol{d}$ ▷ Set sample's position along the ray
  $\tau \leftarrow 0$ ▷ Initialize absorption at $\boldsymbol{x}_n$
  $\boldsymbol{\varepsilon} \leftarrow (0,0,0)$ ▷ Initialize emission at $\boldsymbol{x}_n$ in direction $\boldsymbol{d}$
  $M_0^{\text{3D}} \leftarrow -\infty$ ▷ Initialize max. squared spatial distance
  $S_0^{\text{3D}} \leftarrow 0$ ▷ Initialize spatial normalizer
  $M_0^{\text{5D}} \leftarrow -\infty$ ▷ Initialize max. squared spatio-directional distance
  $S_0^{\text{5D}} \leftarrow 0$ ▷ Initialize spatio-directional normalizer
  for each component $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\omega}_k, \Theta_k, \bar{\tau}_k, \bar{\boldsymbol{\varepsilon}}_k, \boldsymbol{\delta}_k^{\boldsymbol{x}}, \boldsymbol{\delta}_k^{\boldsymbol{d}}\}$ **do**
    $D_k^{\text{3D}} \leftarrow -\frac{1}{2}\|\boldsymbol{x}_n - \boldsymbol{\mu}_k\|^2_{\boldsymbol{\Sigma}_k}$
    $M_k^{\text{3D}} \leftarrow \max(M_{k-1}^{\text{3D}},\, D_k^{\text{3D}})$
    $S_k^{\text{3D}} \leftarrow S_{k-1}\,\exp(M_{k-1}^{\text{3D}} - M_k^{\text{3D}}) + \exp(D_k^{\text{3D}} - M_k^{\text{3D}})$
    $\tau \leftarrow \tau\,\exp(M_{k-1}^{\text{3D}} - M_k^{\text{3D}}) + \bar{\tau}_k\,\exp(D_k^{\text{3D}} - M_k^{\text{3D}})$
    $D_k^{\text{5D}} \leftarrow D_k^{\text{3D}} - \frac{1}{2}\|\boldsymbol{d} \cdot \boldsymbol{\omega}_k\|^2_{\Theta_k}$
    $M_k^{\text{5D}} \leftarrow \max(M_{k-1}^{\text{5D}},\, D_k^{\text{5D}})$
    $S_k^{\text{5D}} \leftarrow S_{k-1}^{\text{5D}}\,\exp(M_{k-1}^{\text{5D}} - M_k^{\text{5D}}) + \exp(D_k^{\text{5D}} - M_k^{\text{5D}})$
    $\boldsymbol{\varepsilon} \leftarrow \boldsymbol{\varepsilon}\,\exp(M_{k-1}^{\text{5D}} - M_k^{\text{5D}}) + \boldsymbol{\varepsilon}_k(\boldsymbol{x}_n, \boldsymbol{d})\,\exp(D_k^{\text{5D}} - M_k^{\text{5D}})$
  end for
  $\tau \leftarrow \tau / S_K^{\text{3D}}$ ▷ Normalize the mixed absorption
  $\boldsymbol{\varepsilon} \leftarrow \boldsymbol{\varepsilon} / S_K^{\text{5D}}$ ▷ Normalize the mixed emission
  $\eta \leftarrow \eta + \tau\,\Delta$ ▷ Accumulate occlusions
  $T \leftarrow \exp(-\eta)$ ▷ Beer-Lambert's transmittance
  $\alpha \leftarrow 1 - \exp(-\tau\,\Delta)$ ▷ Normalized attenuation
  $L \leftarrow L + \boldsymbol{\varepsilon}\,\alpha\,T$ ▷ Accumulate radiance
end for
return L

References

[1] E. H. Adelson and J. R. Bergen, "The plenoptic function and the elements of early vision," in *Computational Models of Visual Processing*, M. Landy and J. A. Movshon, Eds. The MIT Press, 1991.

[2] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: representing scenes as neural radiance fields for view synthesis," *Commun. ACM*, vol. 65, no. 1, p. 99–106, 2021.

[3] L. Westover, "Footprint evaluation for volume rendering," *SIGGRAPH Comput. Graph.*, vol. 24, no. 4, p. 367–376, 1990.

[4] B. Kerbl, G. Kopanas, T. Leimkuehler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, Jul. 2023.

[5] S. Rusinkiewicz and M. Levoy, "QSplat: a multiresolution point rendering system for large meshes," in *Proc. of SIGGRAPH '00*. USA: ACM Press/Addison-Wesley Publishing Co., 2000, p. 343–352.

[6] M. Zwicker, H. Pfister, J. Van Baar, and M. Gross, "EWA volume splatting," in *Proc. of VIS '01*. IEEE, 2001, pp. 29–538.

[7] B. Kerbl, A. Meuleman, G. Kopanas, M. Wimmer, A. Lanvin, and G. Drettakis, "A hierarchical 3D Gaussian representation for real-time rendering of very large datasets," *ACM Trans. Graph.*, vol. 43, no. 4, Jul. 2024. [Online]. Available: https://doi.org/10.1145/3658160

[8] A. Celarek, G. Kopanas, G. Drettakis, M. Wimmer, and B. Kerbl, "Does 3D Gaussian splatting need accurate volumetric rendering?" *Computer Graphics Forum*, vol. 44, no. 2, p. e70032, 2025.

[9] C. Schretter, J. Sun, and P. Schelkens, "Image reconstruction with smoothed mixtures of regressions," in *25th IEEE International Conference on Image Processing (ICIP)*, 2018, pp. 400–404.

[10] J. T. Kent, "The Fisher-Bingham distribution on the sphere," *Journal of the Royal Statistical Society. Series B*, vol. 44, no. 1, pp. 71–80, 1982.

[11] N. Max, "Optical models for direct volume rendering," *IEEE Trans. on Visualization and Computer Graphics*, vol. 1, no. 2, pp. 99–108, 1995.

[12] M. Milakov and N. Gimelshein, "Online normalizer calculation for softmax," 2018. [Online]. Available: http://arxiv.org/abs/1805.02867