Sample Matching for Joint Extinction Gradient Estimation in Differentiable Volume Rendering: Supplementary Material

This supplementary document gives an RTE-centric tour of our theory in the main paper and an unbiasedness proof for estimators in the path space.

1 An RTE-Centric View of Sample Matching

The main paper splits paths into prefix and suffix to expose the segment structure for derivatives. With the language of the radiative transfer equation (RTE), the segment can be directly exposed and handled. The contribution of path measure will be turned into a product containing throughput and in-scattering radiance.

1.1 Volumetric Light Transport along a Single Ray

In volume rendering, light transport is modeled as the interaction between light and microscopic particles distributed throughout space. Our work focuses on scattering and absorption between light and microscopic particles. For heterogeneous media, these interactions are usually characterized by spatially varying albedo $\alpha(\boldsymbol{x})$ and extinction coefficient $\sigma_t(\boldsymbol{x})$, where $\boldsymbol{x}$ denotes a spatial position. A camera ray with origin $\boldsymbol{x}_0$ and direction $\boldsymbol{\omega}$ is parameterized as $\boldsymbol{x}(t) = \boldsymbol{x}_0 + t\,\boldsymbol{\omega}$, where $t \geq 0$ represents the distance along the ray.

The extinction coefficient induces attenuation along the ray, quantified by the transmittance

$$T(t) := \exp\left(-\int_0^t \sigma_t(s)\,\mathrm{d}s\right). \tag{37}$$

The transmittance $T(t)$ represents the fraction of radiance originating at $\boldsymbol{x}(t)$ that reaches $\boldsymbol{x}_0$ without being absorbed or scattered.

Using the integral form of the radiative transfer equation (IRTE), also known as the formal solution of the RTE [Chandrasekhar 1960], and considering only scattering and absorption, the radiance accumulated along a ray can be written as

$$\begin{aligned} L(\boldsymbol{x}_0, -\boldsymbol{\omega}) = &\int_0^{t_{\max}} T(t)\,\sigma_t(t)\,\alpha(t)\,L_\mathrm{s}(\boldsymbol{x}(t), -\boldsymbol{\omega})\,\mathrm{d}t \\ &+ T(t_{\max})\,L_\mathrm{o}(\boldsymbol{x}(t_{\max}), -\boldsymbol{\omega}), \end{aligned} \tag{38}$$

where $t_{\max}$ denotes the distance to the nearest boundary intersection along the ray. Here, $L(\boldsymbol{x}_0, -\boldsymbol{\omega})$ represents the outgoing radiance at $\boldsymbol{x}_0$ in direction $\boldsymbol{\omega}$, and $L_\mathrm{o}(\boldsymbol{x}(t_{\max}), -\boldsymbol{\omega})$ denotes the outgoing radiance at the boundary.

The in-scattered radiance $L_\mathrm{s}(\boldsymbol{x}, \boldsymbol{\omega}_s)$ accounts for light arriving at $\boldsymbol{x}$ from all directions $\boldsymbol{\omega}'$ and scattered toward $\boldsymbol{\omega}_s$:

$$L_\mathrm{s}(\boldsymbol{x}, \boldsymbol{\omega}_s) = \int_{S^2} L(\boldsymbol{x}, \boldsymbol{\omega}')\,\rho(\boldsymbol{\omega}' \to \boldsymbol{\omega}_s; \boldsymbol{x})\,\mathrm{d}\boldsymbol{\omega}', \tag{39}$$

where $\rho(\boldsymbol{\omega}' \to \boldsymbol{\omega}_s; \boldsymbol{x})$ denotes the phase function describing scattering from direction $\boldsymbol{\omega}'$ to $\boldsymbol{\omega}_s$ at $\boldsymbol{x}$. Note that L_s is the integral for L and accounts for multiple bounces of light transport paths which need recursive computation. This generates a light path $\bar{x} = \{\boldsymbol{x}_0, \ldots, \boldsymbol{x}_K\}$ from the camera into the scene, where each $\boldsymbol{x}_j$ is a scattering event position, also serving as a path vertex.

Sampling interaction distances. Efficient Monte Carlo evaluation of volumetric light transport requires careful sampling of the distance t at which scattering events occur along a ray. An efficient sampling strategy draws t from the PDF

$$p(t) \propto T(t)\,\sigma_t(t), \tag{40}$$

commonly referred to as *free-flight sampling*. In homogeneous media, this distribution admits an analytic inverse-CDF and can be sampled exactly. In heterogeneous volumes, however, direct sampling becomes non-trivial due to the spatially varying extinction coefficient. Delta tracking [Butcher and Messel 1958; Woodcock et al. 1965] addresses this challenge by introducing a majorant extinction coefficient and simulating a sequence of null-collisions, which allows free-flight distances to be sampled without explicitly inverting the heterogeneous transmittance.

1.2 Differential Radiative Transfer Equation

Prior works including those by Zhang et al. [2019] and Nimier-David et al. [2022] derived the derivative of the above formulations with respect to a parameter. In this work, we handle the differentiation of the RTE with respect to extinction coefficients $\sigma_t(\boldsymbol{x})$ parameterized using a finite-dimensional vector $\boldsymbol{\theta}$ (e.g., voxel coefficients). We denote by θ_k the k-th element of $\boldsymbol{\theta}$.

Applying the product rule of differentiation to Equation 38, the derivative $\frac{\partial L}{\partial \theta_k}$ decomposes into three components as follows. The first one—which we refer to as the **scattering component**—corresponds to differentiating the product directly related to the extinction coefficient:

$$\left(\frac{\partial L}{\partial \theta_k}\right)_{scat} = \int_0^{t_{\max}} T(t)\,L_\mathrm{s}(t)\,\alpha(t)\,\frac{\partial \sigma_t(t)}{\partial \theta_k}\,\mathrm{d}t. \tag{41}$$

The second component results from differentiating the transmittance. We refer to this as the **transmittance component**:

$$\begin{aligned} \left(\frac{\partial L}{\partial \theta_k}\right)_{trans} = &\int_0^{t_{\max}} T(t)\,\sigma_t(t)\alpha(t)\,L_\mathrm{s}(t)\left(-\int_0^t \frac{\partial \sigma_t(s)}{\partial \theta_k}\,\mathrm{d}s\right)\mathrm{d}t \\ &+ T(t_{\max})\left(-\int_0^{t_{\max}} \frac{\partial \sigma_t(s)}{\partial \theta_k}\,\mathrm{d}s\right)L_\mathrm{o}(t_{\max}). \end{aligned} \tag{42}$$

Note that it generates a nested integral that requires sampling an additional position s along the segment for Monte Carlo estimation.

The third component arises from the changes in the in-scattering radiance:

$$\left(\frac{\partial L}{\partial \theta_k}\right)_{radin} = \int_0^{t_{\max}} T(t)\,\sigma_t(t)\alpha(t)\,\frac{\partial L_\mathrm{s}(t)}{\partial \theta_k}\,\mathrm{d}t. \tag{43}$$

There remains a radiance at the boundary $L_\mathrm{o}(t_{\max})$ to be considered. We assume that once a ray reaches the boundary of the medium, it exits the bounding box and does not re-enter. Besides, we assume boundary illumination (e.g., environment map lighting or emitters) is known. Under this assumption, $L_\mathrm{o}(t_{\max})$ is independent of the optimized extinction parameters $\boldsymbol{\theta}$.

As L_s is defined as an integral over the radiance L (Equation 39), differentiating it with respect to θ_k induces the same product-rule decomposition on subsequent ray segments, yielding terms structurally identical to Equation 41–Equation 43. Due to this recursive structure, we can focus our derivation on the scattering and transmittance components of gradients (Equation 41 and Equation 42) for a single segment without loss of generality.

1.3 Free-Flight and DRT Estimators

In practice, gradient estimation with respect to a scene parameter θ_k is carried out along the light path $\bar{x} = \{x_0, \ldots, x_K\}$ sampled during forward volume rendering described in Section 1.1. Under this formulation, both the radiance estimator and its corresponding gradient estimators rely on the same sampling distribution Equation 40.

This PDF naturally handles the in-scattering radiance term (Equation 43), yielding the estimator:

$$\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle_{radin} = \alpha(t)\, \frac{\partial L_s(t)}{\partial \theta_k}. \tag{44}$$

To simplify the following derivation, we introduce the notation function $h(t)$, which encapsulates both volumetric in-scattering and boundary radiance:

$$h(t) := \begin{cases} \alpha(t)\, L_s(t), & t \le t_{\max}, \\ L_o(t_{\max}), & t > t_{\max}. \end{cases} \tag{45}$$

Using the above free-flight sampling PDF, the ray undergoes no interaction along the segment and exits the medium with probability $T(t_{\max})$. Under this sampling distribution, the estimator of Equation 42 becomes:

$$\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle_{trans} = -h(t) \int_0^{\min(t, t_{\max})} \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s. \tag{46}$$

Both estimators (Equation 44 and Equation 46) are unbiased. The scattering estimator under the same free-flight PDF, however, has two issues.

First, as observed by Nimier-David et al. [2022], the scattering gradient estimator is biased:

$$\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle_{scat} = \frac{\frac{\partial \sigma_t(t)}{\partial \theta_k}\, L_s(t)\, \alpha(t)}{\sigma_t(t)}, \tag{47}$$

since the sampling probability vanishes in regions where $\sigma_t(t) = 0$ but the integrand may still be nonzero. Moreover, when $\sigma_t(t) \neq 0$, the division by $\sigma_t(t)$ leads to a rapid increase in variance in regions of low extinction, resulting in extremely noisy gradient estimates in practice.

The second issue is that the scattering and transmittance components evaluate $\frac{\partial \sigma_t}{\partial \theta_k}$ at different positions on the same segment: the scattering term at the segment endpoint t, the transmittance term at an interior point $s \in [0, t]$. Both terms probe the same extinction field along the same ray, but at uncorrelated points, so any statistical coupling between them is left unexploited.

Nimier-David et al. [2022] address the first issue by decoupling the sampling strategy for the scattering component (Equation 41). Instead of sampling from $T(t)\sigma_t(t)$, they use Ratio Tracking [Novák et al. 2014] to generate the transmittance distribution $T(t)$, and Equation 41 is evaluated using samples drawn from

$$p_{\mathrm{DRT}}(t) \propto T(t), \tag{48}$$

which removes the problematic division by $\sigma_t(t)$. Under this choice, the resulting Monte Carlo estimator of Equation 41 becomes

$$\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle_{scat} = \frac{\partial \sigma_t(t)}{\partial \theta_k}\, L_s(t)\, \alpha(t), \tag{49}$$

which is unbiased and significantly reduces variance for the scattering component on its own.

The second issue, however, has never been recognized or addressed: the scattering and transmittance components evaluate $\frac{\partial \sigma_t}{\partial \theta_k}$ at different positions most of the time, leaving negative covariance unexploited. Forcing them in the same evaluation position is nontrivial, since the marginal PDF of transmittance gradient samples is not available in closed form and the two components are defined over different integration domains.

1.4 Joint Reformulation as a Nested Integral

We aim to derive a single **joint estimator** for the extinction gradient $\frac{\partial L}{\partial \theta_k}$ within each ray segment (x_{j-1}, x_j) to utilize the potential negative correlation between the scattering and transmittance components. Unlike Nimier-David et al. [2022], which decouples the estimations of Equation 41 and Equation 42 using different importance sampling methods, our formulation merges them into a single integral as follows.

At this stage, Equation 41 remains a single integral, whereas Equation 42 involves a nested integral over the domain $s \le t$. To align their integration domains, we leverage the fundamental identity derived from the definition of transmittance (Equation 37):

$$T(s) - T(t_{\max}) = \int_s^{t_{\max}} \sigma_t(t) T(t)\, \mathrm{d}t. \tag{50}$$

After rewriting Equation 41 as the integral over s:

$$\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle_{scat} = \int_0^{t_{\max}} T(s)\, L_s(s)\, \alpha(s) \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s, \tag{51}$$

we can substitute $T(s)$ above with Equation 50. Then Equation 51 can be expanded into a nested integral combined with a boundary term:

$$\begin{aligned} &\int_0^{t_{\max}} T(s)\alpha(s)\, L_s(s) \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s \\ &\quad = \int_0^{t_{\max}} \frac{\partial \sigma_t(s)}{\partial \theta_k} \left[\int_s^{t_{\max}} \sigma_t(t) T(t) \alpha(s)\, L_s(s)\, \mathrm{d}t \right] \mathrm{d}s \\ &\quad \underbrace{+\, T(t_{\max}) \int_0^{t_{\max}} \alpha(s)\, L_s(s) \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s}_{\text{boundary term}}. \end{aligned} \tag{52}$$

Equation 52 and Equation 42 now share similar forms, which allows us to merge them into a single joint formulation. However, the nested integral parts still have different integral limits. Applying Fubini's theorem to Equation 52 over the domain $\{(s, t) \mid 0 \le s \le t\}$

yields:

$$\int_0^{t_{\max}} \frac{\partial \sigma_t(s)}{\partial \theta_k} \int_s^{t_{\max}} \sigma_t(t) T(t) \alpha(s)\, L_\mathrm{s}(s)\, \mathrm{d}t\, \mathrm{d}s = \int_0^{t_{\max}} \int_0^{t} \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \sigma_t(t) T(t) \alpha(s)\, L_\mathrm{s}(s)\, \mathrm{d}s\, \mathrm{d}t. \quad (53)$$

Finally, we can merge Equation 52, Equation 53 and Equation 42, leading to the joint formulation:

$$\frac{\partial L}{\partial \theta_k} = \int_0^{t_{\max}} \int_0^{t} \sigma_t(t)\, T(t)\, \left[\alpha(s)\, L_\mathrm{s}(s) - \alpha(t)\, L_\mathrm{s}(t)\right] \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s\, \mathrm{d}t + T(t_{\max}) \left(\int_0^{t_{\max}} \left[\alpha(s)\, L_\mathrm{s}(s) - L_\mathrm{o}(t_{\max})\right] \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s \right). \quad (54)$$

Equation 54 contains a nested integral over the triangular domain $\{(s,t) \mid 0 \le s \le t\}$. In principle, this structure permits joint sampling of the pair (s,t) using an arbitrary joint PDF $p(s,t)$, yielding:

$$\frac{\partial L}{\partial \theta_k} = \mathbb{E}\left[\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle \right]_{(s,t)\sim p(s,t)}. \quad (55)$$

However, the light path for our method is pre-sampled. Under the path-construction setting described in Section 1.1, the sampling PDF of t is fixed and its marginal PDF is given by Equation 40. According to the definition of transmittance, when sampling the distance t with this PDF, the boundary event $t > t_{\max}$ occurs with probability $T(t_{\max})$, corresponding to the case where no scattering event occurs before the ray exits the segment boundary. Similar to Equation 46, and using the notation defined in Equation 45, the estimator of the extinction gradient can be written as:

$$\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle = \left[h(s) - h(t)\right] \int_0^{\min(t, t_{\max})} \frac{\partial \sigma_t(s)}{\partial \theta_k}\, \mathrm{d}s. \quad (56)$$

Only the PDF of s remains to be specified, which we denote as the conditional PDF $p(s \mid t)$, defined on the interval $[0,t]$. For simplicity of implementation, we choose a uniform distribution over the segment interior:

$$p(s \mid t) = \frac{1}{t}, \qquad 0 \le s \le \min(t, t_{\max}). \quad (57)$$

It yields the estimator:

$$\boxed{\left\langle \frac{\partial L}{\partial \theta_k} \right\rangle = \min(t, t_{\max}) \cdot \left[h(s) - h(t)\right]\, \frac{\partial \sigma_t(s)}{\partial \theta_k}.} \quad (58)$$

Discussion. Equation 58 defines our final gradient estimator for the extinction coefficient within a single ray segment. Unlike prior work that decouples the extinction gradient into separate transmittance and scattering components, our formulation naturally yields a single joint estimator which aligns with physical intuition. While Nimier-David et al. [2022] handles bias by changing the PDF of the scattering component, our estimator is unbiased by construction: the sampling distribution has full support over the integration domain $[0, t_{\max}]$, ensuring that all regions where $\frac{\partial \sigma_t(s)}{\partial \theta_k} \neq 0$ can be sampled with non-zero probability.

1.5 Connection to the Path-Space Formulation

The estimator above is segment-level; Section 4 surfaces the same structure path-by-path. The triangular domain $\{(s,t)\}$ corresponds to the line segment $\overline{\boldsymbol{x}_{j-1}\,\boldsymbol{x}_j}$ on which the per-segment integrand $\mathcal{E}_j$ in Equation 25 is integrated over $(\boldsymbol{x}_j, \boldsymbol{y})$, and the two cases of h in Equation 45 correspond respectively to f_j^{ins} (Equation 24) on interior segments and to f^{out} (Equation 29) on the host segment terminating at a light source.

2 Unbiasedness of Our Estimators

Internal estimators (Equation 32, Equation 35). Both take the form

$$\frac{\|\boldsymbol{x}_K - \boldsymbol{x}_{K-1}\|}{p^{\mathrm{int}}}\, \partial_{\boldsymbol{\theta}} \sigma_t(\boldsymbol{y}) \left(\sigma_t(\boldsymbol{x}_K)\, \hat{\rho}(\boldsymbol{x}_0^{\mathrm{suf}} \to \boldsymbol{y} \to \boldsymbol{x}_{K-1}) \cdot s_K(\bar{\mathbf{x}})\, f_{\setminus K}(\bar{\mathbf{y}} + \bar{\mathbf{x}}^{\mathrm{suf}}) - f(\bar{\mathbf{x}} + \bar{\mathbf{x}}^{\mathrm{suf}}) \right). \quad (59)$$

Expanding the measurement contribution and the path-tracing PDF,

$$f(\bar{\mathbf{x}}) \;=\; W_\mathrm{e}\, L_\mathrm{e} \cdot \prod_{j=1}^{K-1} \sigma_t(\boldsymbol{x}_j)\, \hat{\rho}(\boldsymbol{x}_{j+1} \to \boldsymbol{x}_j \to \boldsymbol{x}_{j-1}) \cdot \prod_{j=1}^{K} s_j(\bar{\mathbf{x}}), \quad (60)$$

$$p^{\mathrm{int}}(\bar{\mathbf{x}}) \;\propto\; \prod_{j=1}^{K} \sigma_t(\boldsymbol{x}_j)\, T(\boldsymbol{x}_{j-1} \leftrightarrow \boldsymbol{x}_j), \quad (61)$$

(the latter times bounded angular factors and the strictly positive probe-resampling PDF), every $\sigma_t(\boldsymbol{x}_j)$ that appears in the denominator also appears in the numerator: the $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{K-1}$ factors come from f in the second term and from $f_{\setminus K}$ in the first; the missing $\sigma_t(\boldsymbol{x}_K)$ is supplied by the explicit multiplier in front of the first term and by the K-th vertex factor inside f for the second. The bracket therefore vanishes whenever p^{int} does, and after cancellation the estimator reduces to a bounded combination of albedos. It is unbiased.

Light estimator (Equation 36). Substituting f_K^{ins} (Equation 24) and f^{out} (Equation 29) and pulling the common $s_K(\bar{\mathbf{x}})$, $T(\boldsymbol{x}_K \leftrightarrow \boldsymbol{x}_K^{\perp})$, and $\sigma_t(\boldsymbol{x}_K)$ factors outside the bracket, Equation 36 reads

$$\left(\frac{\hat{\rho}(\boldsymbol{x}_0^{\mathrm{suf}} \to \boldsymbol{x}_K \to \boldsymbol{x}_{K-1})\, f_{\setminus K}(\bar{\mathbf{x}} + \bar{\mathbf{x}}^{\mathrm{suf}})}{p_{\mathrm{pre}}^{\mathrm{light}}\, p_{\mathrm{suf}}^{\mathrm{light}}} - \frac{L_\mathrm{e}(\boldsymbol{x}_K^{\perp} \to \boldsymbol{x}_K)\, f_{\setminus K}(\bar{\mathbf{x}})}{p_{\mathrm{pre}}^{\mathrm{light}}} \right) \cdot\; s_K(\bar{\mathbf{x}})\, T(\boldsymbol{x}_K \leftrightarrow \boldsymbol{x}_K^{\perp})\, \sigma_t(\boldsymbol{x}_K). \quad (62)$$

$f_{\setminus K}(\bar{\mathbf{x}})$ and $f_{\setminus K}(\bar{\mathbf{x}} + \bar{\mathbf{x}}^{\mathrm{suf}})$ are f (Equation 60) with the K-th vertex/segment factors removed, so they carry $\prod_{j=1}^{K-1} \sigma_t(\boldsymbol{x}_j)\, T(\boldsymbol{x}_{j-1} \leftrightarrow \boldsymbol{x}_j)$ along the prefix and an additional $\prod \sigma_t\, T$ chain along the suffix in the first term, matching the corresponding $\sigma_t\, T$ products in $p_{\mathrm{pre}}^{\mathrm{light}}$ and $p_{\mathrm{suf}}^{\mathrm{light}}$. The trailing $s_K(\bar{\mathbf{x}})\, T(\boldsymbol{x}_K \leftrightarrow \boldsymbol{x}_K^{\perp})\, \sigma_t(\boldsymbol{x}_K)$ is exactly the host-segment contribution that the notation split off, so the numerator's σ_t structure cancels the PDF completely. The integrand therefore vanishes when the PDF is 0, and the light estimator is unbiased.

References

John C Butcher and Harry Messel. 1958. Electron number distribution in electron-photon showers. *Physical Review* 112, 6 (1958), 2096.

Subrahmanyan Chandrasekhar. 1960. *Radiative Transfer.* Dover Publications, New York. Originally published by Oxford University Press, 1950.

Merlin Nimier-David, Thomas Müller, Alexander Keller, and Wenzel Jakob. 2022. Unbiased Inverse Volume Rendering with Differential Trackers. *ACM Trans. Graph.* 41, 4, Article 44 (July 2022), 20 pages. doi:10.1145/3528223.3530073

Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph.* 33, 6 (2014), 179–1.

E Woodcock, T Murphy, P Hemmings, and S Longworth. 1965. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In *Proc. Conf. Applications of Computing Methods to Reactor Problems*, Vol. 557. Argonne National Laboratory.

Cheng Zhang, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao. 2019. A differential theory of radiative transfer. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–16.