Volume Rendering Digest (for NeRF)

Andrea Tagliasacchi[1,2] Ben Mildenhall[1]

[1]Google Research [2]Simon Fraser University

Neural Radiance Fields [3] employ simple volume rendering as a way to overcome the challenges of differentiating through ray-triangle intersections by leveraging a probabilistic notion of visibility. This is achieved by assuming the scene is composed by a cloud of light-emitting particles whose density changes in space (in the terminology of physically-based rendering, this would be described as a volume with absorption and emission but no scattering [4, Sec 11.1]. In what follows, for the sake of exposition simplicity, and without loss of generality, we assume the emitted light *does not* change as a function of view-direction. This technical report is a condensed version of previous reports [1, 2], but rewritten in the context of NeRF, and adopting its commonly used notation[1].

Transmittance. Let the density field $\sigma(\mathbf{x})$, with $\mathbf{x}\in\mathbb{R}^3$ indicate the differential likelihood of a ray hitting a particle (i.e. the probability of hitting a particle while travelling an infinitesimal distance). We reparameterize the density along a given ray $\mathbf{r}=(\mathbf{o},\mathbf{d})$ as a scalar function $\sigma(t)$, since any point $\mathbf{x}$ along the ray can be written as $\mathbf{r}(t)=\mathbf{o}+t\mathbf{d}$. Density is closely tied to the transmittance function $\mathcal{T}(t)$, which indicates the probability of a ray traveling over the interval $[0,t)$ without hitting any particles. Then the probability $\mathcal{T}(t+dt)$ of *not* hitting a particle when taking a differential step dt is equal to $\mathcal{T}(t)$, the likelihood of the ray reaching t, times $(1-dt\cdot\sigma(t))$, the probability of not hitting anything during the step:

$$\mathcal{T}(t+dt) = \mathcal{T}(t)\cdot(1-dt\cdot\sigma(t)) \tag{1}$$

$$\frac{\mathcal{T}(t+dt)-\mathcal{T}(t)}{dt} \equiv \mathcal{T}'(t) = -\mathcal{T}(t)\cdot\sigma(t) \tag{2}$$

This is a classical differential equation that can be solved as follows:

$$\mathcal{T}'(t) = -\mathcal{T}(t)\cdot\sigma(t) \tag{3}$$

$$\frac{\mathcal{T}'(t)}{\mathcal{T}(t)} = -\sigma(t) \tag{4}$$

$$\int_a^b \frac{\mathcal{T}'(t)}{\mathcal{T}(t)}\,dt = -\int_a^b \sigma(t)\,dt \tag{5}$$

$$\log\mathcal{T}(t)\big|_a^b = -\int_a^b \sigma(t)\,dt \tag{6}$$

$$\mathcal{T}(a\to b) \equiv \frac{\mathcal{T}(b)}{\mathcal{T}(a)} = \exp\left(-\int_a^b \sigma(t)\,dt\right) \tag{7}$$

where we define $\mathcal{T}(a\to b)$ as the probability that the ray travels from distance a to b along the ray without hitting a particle, which is related to the previous notation by $\mathcal{T}(t)=\mathcal{T}(0\to t)$.

[1]If you are interested in borrowing the LaTeX notation, please refer to: `https://www.overleaf.com/read/fkhpkzxhnyws`

Probabilistic interpretation. Note that we can also interpret the function $1-\mathcal{T}(t)$ (often called *opacity*) as a cumulative distribution function (CDF) indicating the probability that the ray *does* hit a particle sometime before reaching distance t. Then $\mathcal{T}(t)\cdot\sigma(t)$ is the corresponding probability density function (PDF), giving the likelihood that the ray stops precisely at distance t.

Volume rendering. We can now calculate the expected value of the light emitted by the particles in the volume as the ray travels from t=0 to D, composited on top of a background color. Since the probability density for stopping at t is $\mathcal{T}(t)\cdot\sigma(t)$, the expected color is

$$\boldsymbol{C} = \int_0^D \mathcal{T}(t)\cdot\sigma(t)\cdot\mathbf{c}(t)\;dt \;+\; \mathcal{T}(D)\cdot\mathbf{c}_{\text{bg}} \tag{8}$$

where $\mathbf{c}_{\text{bg}}$ is a background color that is composited with the foreground scene according to the residual transmittance $\mathcal{T}(D)$. Without loss of generality, we omit the background term in what follows.

Homogeneous media. We can calculate the color of some homogeneous volumetric media with constant color $\mathbf{c}_a$ and density σ_a over a ray segment $[a, b]$ by integration:

$$\begin{aligned}
\boldsymbol{C}(a\to b) &= \int_a^b \mathcal{T}(a\to t)\cdot\sigma(t)\cdot\mathbf{c}(t)\;dt && (9)\\
&= \sigma_a\cdot\mathbf{c}_a\int_a^b \mathcal{T}(a\to t)\;dt && \text{constant density/radiance} \quad (10)\\
&= \sigma_a\cdot\mathbf{c}_a\int_a^b \exp\left(-\int_a^t \sigma(u)\;du\right)\;dt && \text{substituting (7)} \quad (11)\\
&= \sigma_a\cdot\mathbf{c}_a\int_a^b \exp\left(-\;\sigma_a u\big|_a^t\right)\;dt && \text{constant density (again)} \quad (12)\\
&= \sigma_a\cdot\mathbf{c}_a\int_a^b \exp\left(-\sigma_a(t-a)\right)\;dt && (13)\\
&= \sigma_a\cdot\mathbf{c}_a\cdot\left.\frac{\exp\left(-\sigma_a(t-a)\right)}{-\sigma_a}\right|_a^b && (14)\\
&= \mathbf{c}_a\cdot\left(1-\exp\left(-\sigma_a(b-a)\right)\right) && (15)
\end{aligned}$$

Transmittance is multiplicative. Note that transmittance factorizes as follows:

$$\begin{aligned}
\mathcal{T}(a\to c) = &= \exp\left(-\left[\int_a^b \sigma(t)\;dt + \int_b^c \sigma(t)\;dt\right]\right) && (16)\\
&= \exp\left(-\int_a^b \sigma(t)\;dt\right)\exp\left(-\int_b^c \sigma(t)\;dt\right) && (17)\\
&= \mathcal{T}(a\to b)\cdot\mathcal{T}(b\to c) && (18)
\end{aligned}$$

This also follows from the probabilistic interpretation of $\mathcal{T}$, since the probability that the ray does not hit any particles within $[a, c]$ is the product of the probabilities of the two independent events that it does not hit any particles within $[a, b]$ or within $[b, c]$.

Transmittance for piecewise constant data. Given a set of intervals $\{[t_n, t_{n+1}]\}_{n=1}^N$ with constant density σ_n within the n-th segment, and with t_1=0 and δ_n=$t_{n+1}-t_n$, transmittance is equal to:

$$\mathcal{T}_n = \mathcal{T}(t_n) = \mathcal{T}(0\to t_n) = \exp\left(-\int_0^{t_n}\sigma(t)\;dt\right) = \exp\left(\sum_{k=1}^{n-1} -\sigma_k\delta_k\right) \tag{19}$$

Volume rendering of piecewise constant data. Combining the above, we can evaluate the volume rendering integral through a medium with piecewise constant color and density:

$$\boldsymbol{C}(t_{N+1}) = \sum_{n=1}^{N} \int_{t_n}^{t_{n+1}} \mathcal{T}(t) \cdot \sigma_n \cdot \mathbf{c}_n \, dt \qquad \text{piecewise constant} \qquad (20)$$

$$= \sum_{n=1}^{N} \int_{t_n}^{t_{n+1}} \mathcal{T}(0 \to t_n) \cdot \mathcal{T}(t_n \to t) \cdot \sigma_n \cdot \mathbf{c}_n \, dt \qquad \text{from (18)} \qquad (21)$$

$$= \sum_{n=1}^{N} \mathcal{T}(0 \to t_n) \int_{t_n}^{t_{n+1}} \mathcal{T}(t_n \to t) \cdot \sigma_n \cdot \mathbf{c}_n \, dt \qquad \text{constant} \qquad (22)$$

$$= \sum_{n=1}^{N} \mathcal{T}(0 \to t_n) \cdot (1 - \exp(-\sigma_n (t_{n+1} - t_n))) \cdot \mathbf{c}_n \qquad \text{from (15)} \qquad (23)$$

This leads to the volume rendering equations from NeRF [3, Eq.3]:

$$\boldsymbol{C}(t_{N+1}) = \sum_{n=1}^{N} \mathcal{T}_n \cdot (1 - \exp(-\sigma_n \delta_n)) \cdot \mathbf{c}_n, \quad \text{where} \quad \mathcal{T}_n = \exp\left(\sum_{k=1}^{n-1} -\sigma_k \delta_k\right) \qquad (24)$$

Finally, rather than writing these expressions in terms of volumetric density, we can re-express them in terms of alpha-compositing weights $\alpha_n \equiv 1 - \exp(-\sigma_n \delta_n)$, and by noting that $\prod_i \exp x_i = \exp(\sum_i x_i)$ in (19):

$$\boldsymbol{C}(t_{N+1}) = \sum_{n=1}^{N} \mathcal{T}_n \cdot \alpha_n \cdot \mathbf{c}_n, \quad \text{where} \quad \mathcal{T}_n = \prod_{n=1}^{N-1} (1 - \alpha_n) \qquad (25)$$

Alternate derivation. By making use of the earlier connection between CDF and PDF that $(1-\mathcal{T})' = \mathcal{T}\sigma$, and by assuming constant color $\mathbf{c}_a$ along an interval $[a, b]$:

$$\int_a^b \mathcal{T}(t) \cdot \sigma(t) \cdot \mathbf{c}(t) \, dt = \mathbf{c}_a \int_a^b (1 - \mathcal{T})'(t) \, dt \qquad (26)$$

$$= \mathbf{c}_a \cdot (1 - \mathcal{T}(t))|_a^b \qquad (27)$$

$$= \mathbf{c}_a \cdot (\mathcal{T}(a) - \mathcal{T}(b)) \qquad (28)$$

$$= \mathbf{c}_a \cdot \mathcal{T}(a) \cdot (1 - \mathcal{T}(a \to b)) \qquad (29)$$

Combined with constant per-interval density, this identity yields the same expression for color as (24).

References

[1] Nelson Max and Min Chen. Local and global illumination in the volume rendering integral. Technical report, Lawrence Livermore National Lab (LLNL), Livermore, CA (United States), 2005.

[2] Nelson Max and Min Chen. Local and global illumination in the volume rendering integral. Technical report, Schloss Dagstuhl, Leibniz Center for Informatics (Germany), 2010. `https://drops.dagstuhl.de/opus/volltexte/2010/2709/pdf/18.pdf`.

[3] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.

[4] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016. `https://pbr-book.org/3ed-2018/Volume_Scattering/Volume_Scattering_Processes`.

Acknowledgements

Thanks to Daniel Rebain, Soroosh Yazdani and Rif A. Saurous for a careful proofread.