Neural Radiance Fields pt 3

made with nerfstudio

CS180/280A: Intro to Computer Vision and Computational Photography

Angjoo Kanazawa and Alexei Efros

UC Berkeley Fall 2023

Lots of content from Noah Snavely and Ben Mildenhall, Pratul Srinivasan, and Matt Tancik from [ECCV 2022 Tutorial on Neural Volumetric Rendering for Computer Vision]

Class Choice Award Project 3

Winner:

Chloe Zhong

https://inst.eecs.berkeley.edu/~cs180/fa23/upload/files/proj3/czhongx4

Tied 2nd and 3rd: Irene Geng and Jai Singh

https://inst.eecs.berkeley.edu/~cs180/fa23/upload/files/proj3/irenegeng2/

https://inst.eecs.berkeley.edu/~cs180/fa23/upload/files/proj3/jai.s

Where we are

Now we need to render an image from this 3D representation in a differentiable manner

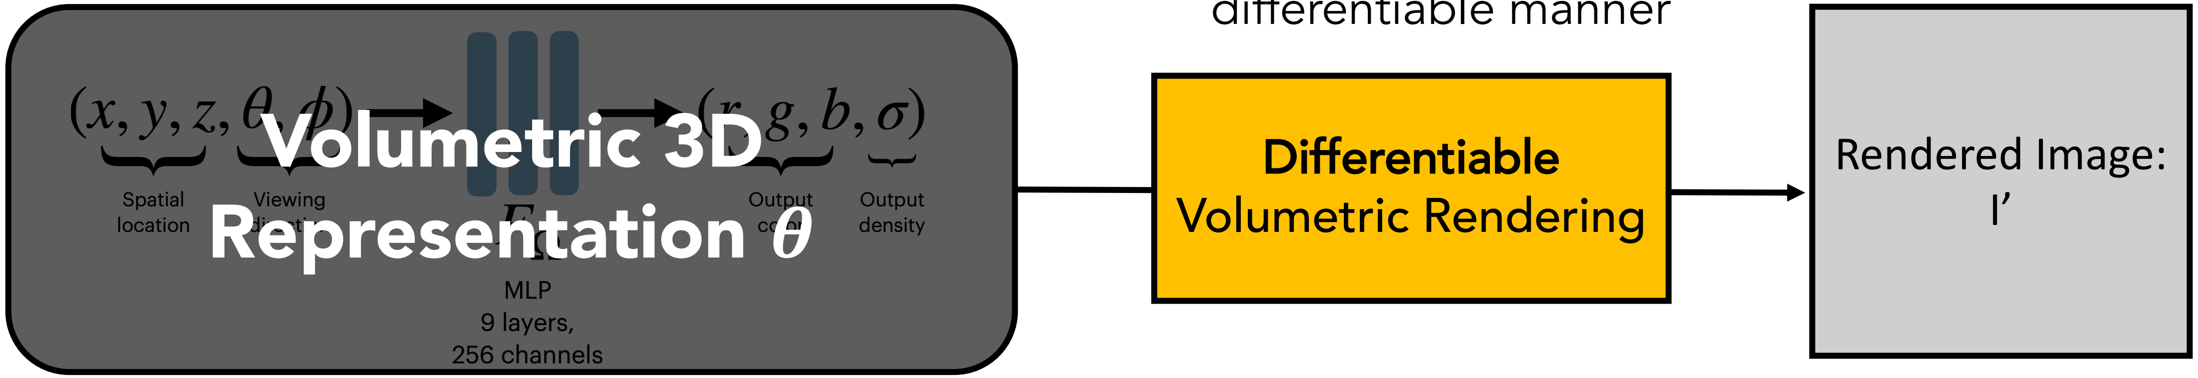

“Training” Objective (aka Analysis-by-Synthesis):

$$\min_{\theta} \left\| \text{Rendered Image: } I' - \text{Observed Image: } I \right\|_2$$

A Precursor: Multi-plane Images

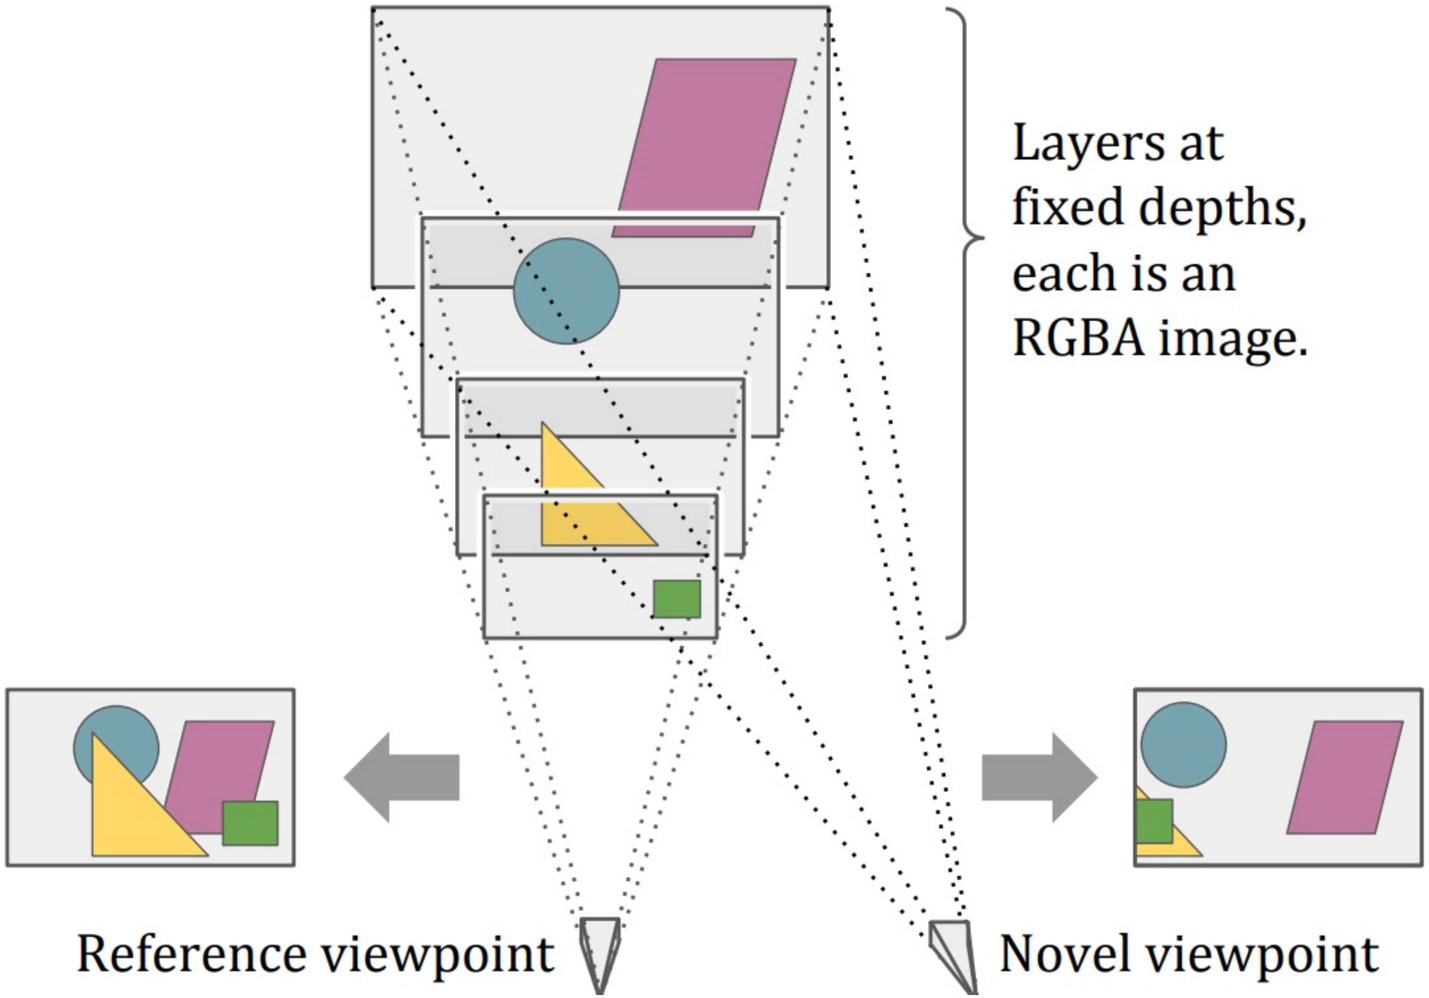

Zou et al. Stereo Magnification, SIGGRAPH 2018

Also called front-to-back compositing or "over" operation

Alpha Blending

for two image case, A and B,
both partially transparent:

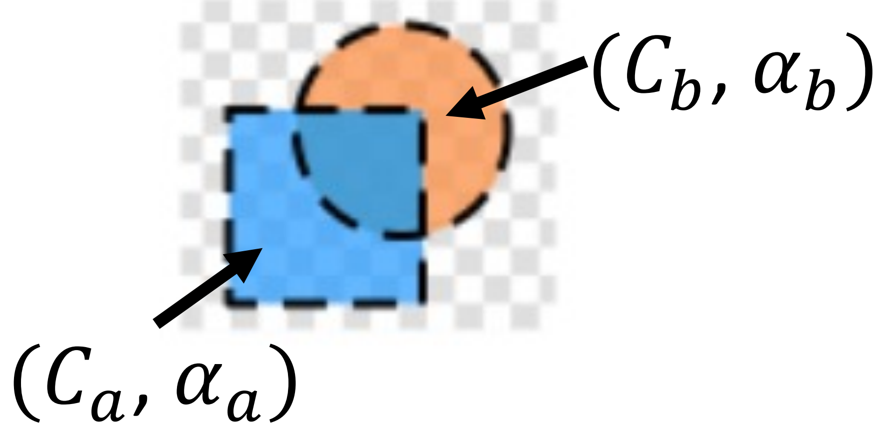

$$I = C_a\,\alpha_a + C_b\,\alpha_b \underbrace{(1-\alpha_a)}_{\text{How much light is the previous layer letting through?}}$$

General D layer case:

$$I = \sum_{i=1}^{D} C_i \alpha_i \prod_{j=1}^{i-1} (1-\alpha_j)$$

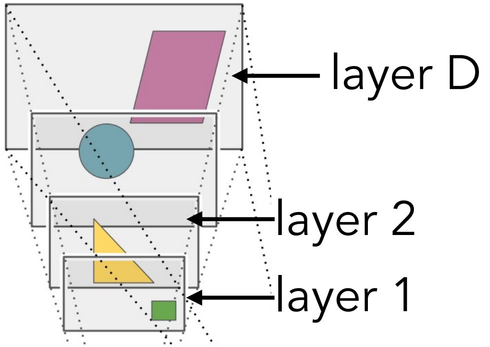

Volumetric formulation for NeRF

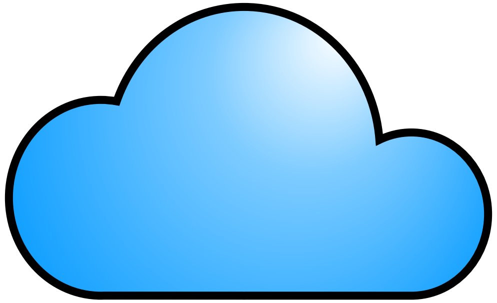

Scene is a cloud of tiny colored particles

Max and Chen 2010, *Local and Global Illumination in the Volume Rendering Integral*

Volumetric formulation for NeRF

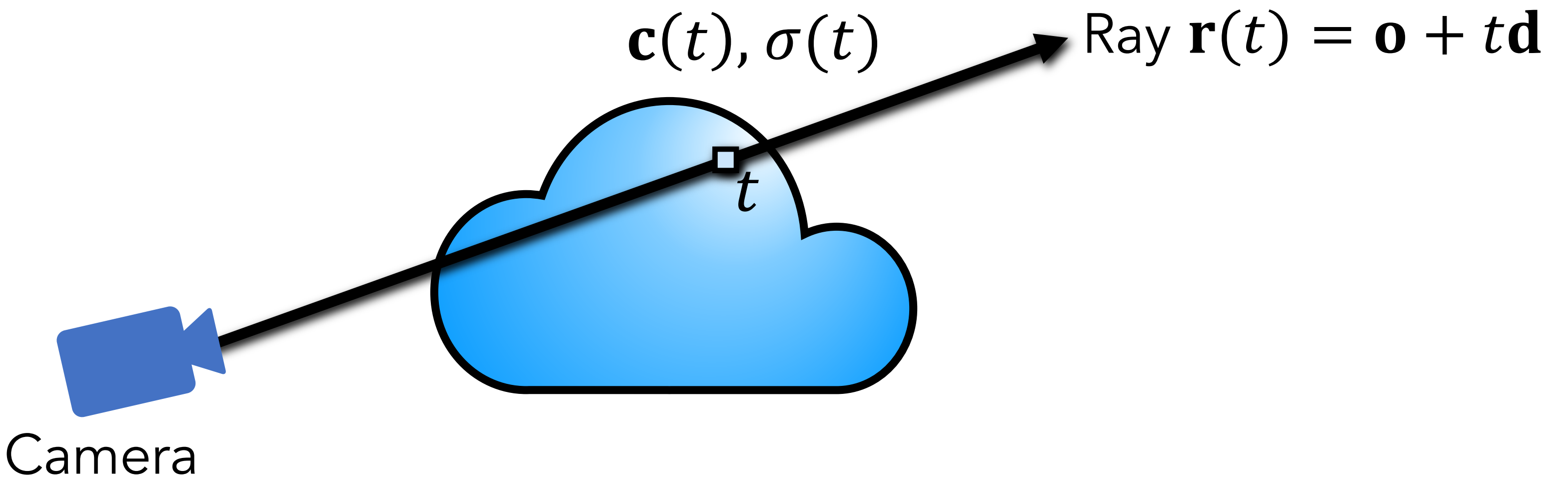

at a point on the ray $\mathbf{r}(t)$, we can query color $\boldsymbol{c}(t)$ and density $\sigma(t)$

How to integrate all the info along the ray to get a color per ray?

Idea: Expected Color

- Pose probabilistically.
- Each point on the ray has a probability to be the first "hit" : $P[first\ hit\ at\ t]$
- Color per ray = Expected value of color with this probability of first "hit"

for a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$:

$$\boldsymbol{c}(\boldsymbol{r}) = \int_{t_0}^{t_1} P[first\ hit\ at\ t]\boldsymbol{c}(t)dt$$

$$\approx \sum_{t=0}^{T} P[first\ hit\ at\ t]\boldsymbol{c}(t)$$

$$\approx \sum_{t=0}^{T} w_t \boldsymbol{c}(t)$$

$$= \sum_{i=1}^{n} T_i \alpha_i \mathbf{c}_i$$

where $T_i = \prod_{j=1}^{i-1}(1-\alpha_j)$ $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$

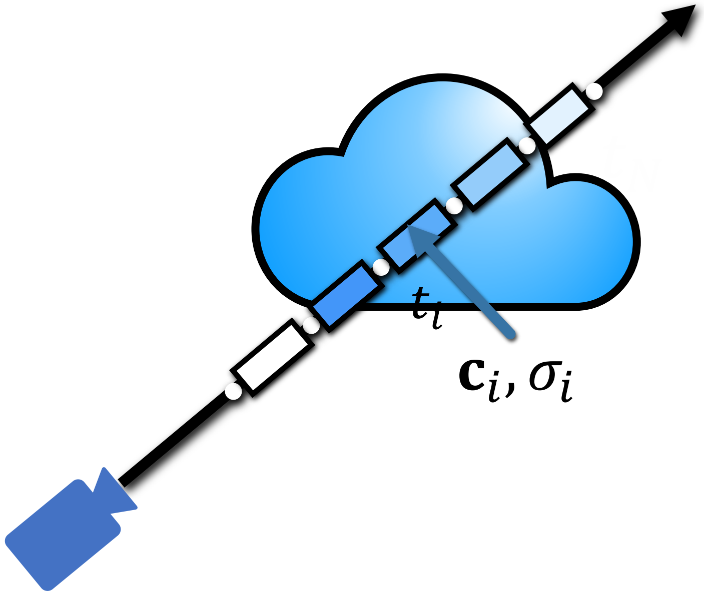

Differentiable Volumetric Rendering Formula

for a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$:

differentiable w.r.t. $\mathbf{c}, \sigma$

$$\mathbf{c} \approx \boxed{\sum_{i=1}^{n} w_i \mathbf{c}_i}$$

colors

weights

How much light is blocked earlier along ray:

$$T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$$

How much light is contributed by ray segment *i*:

$$\alpha_i = 1 - \exp(-\sigma_i \delta_i)$$

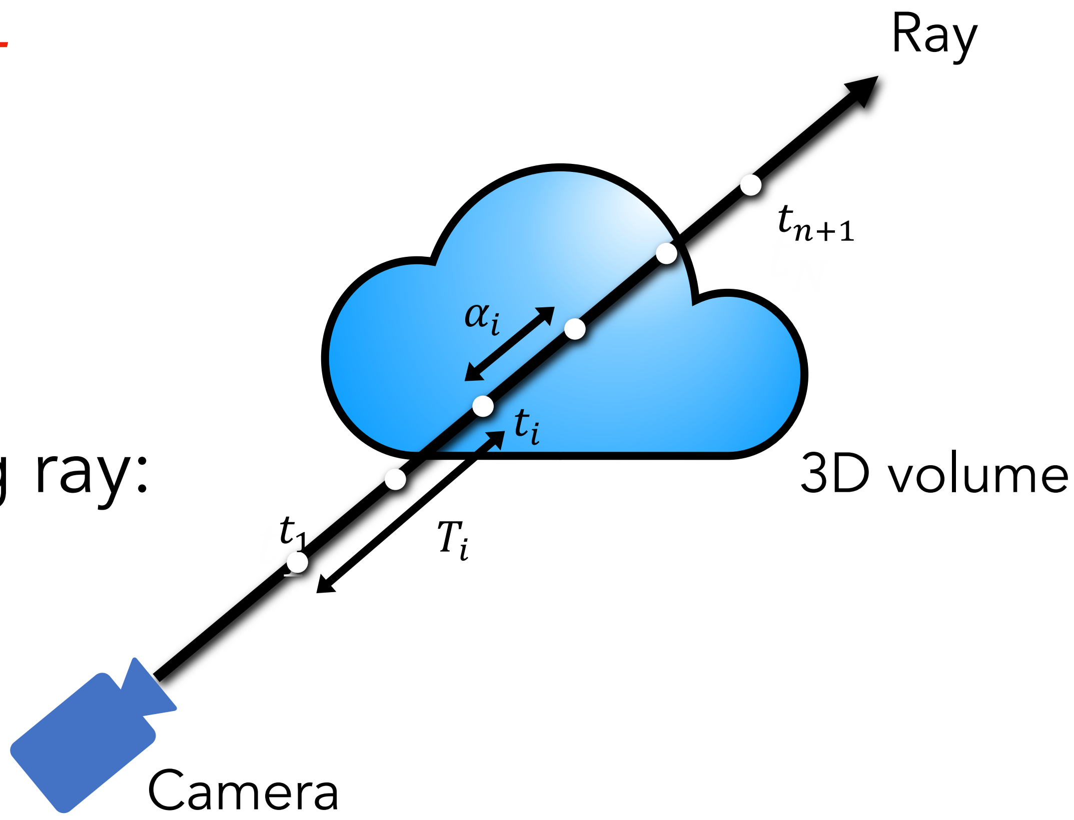

Visual intuition: rendering weights is specific to a ray

$$C \approx \sum_{i=1}^{N} \boxed{T_i \alpha_i} c_i$$

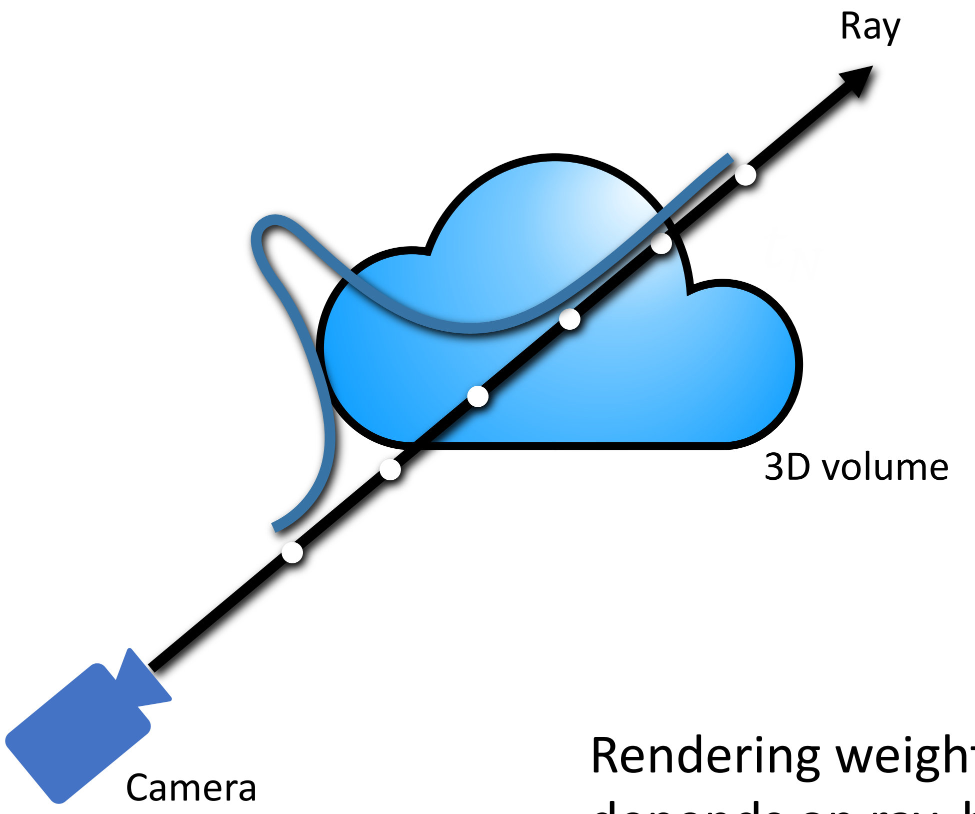

Rendering weights are not a 3D function — depends on ray, because of tranmisttance!

10

Visual intuition: rendering weights is specific to a ray

$$C \approx \sum_{i=1}^{N} \boxed{T_i \alpha_i} c_i$$

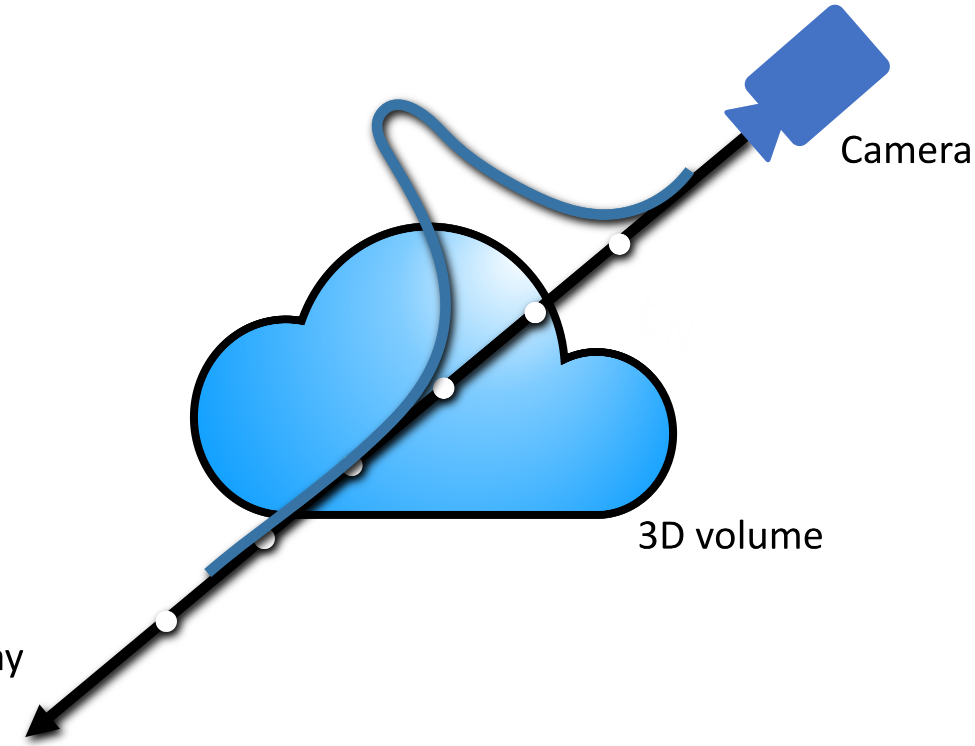

Rendering weights are not a 3D function — depends on ray, because of tranmisttance!

11

What's the point

- Remember, for each pixel or a ray we render a color with this formula based on the Volumetric 3D Representation
- We use this to supervise the 3D Representation (sigma, RGB volume)

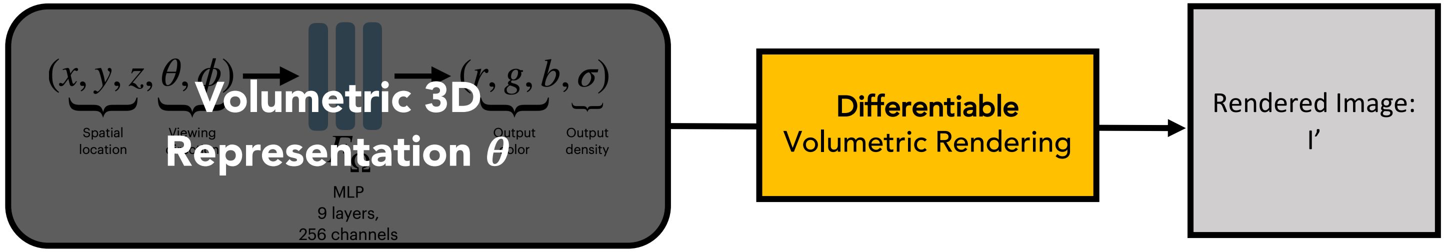

"Training" Objective (aka Analysis-by-Synthesis):

$$\min_{\theta} \left\| \boxed{\text{Rendered Image: } I'} - \boxed{\text{Observed Image: } I} \right\|_2$$

Let's derive this:

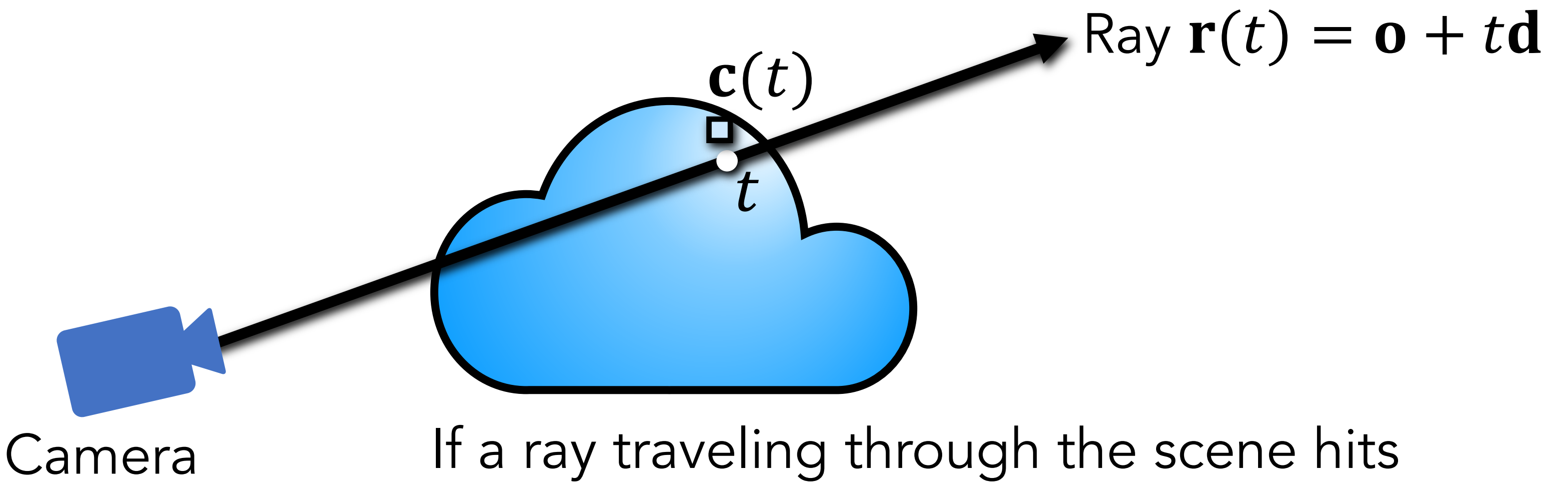

If a ray traveling through the scene hits a particle at distance t along the ray, we return its color $\mathbf{c}(t)$

What does it mean for a ray to "hit" the volume?

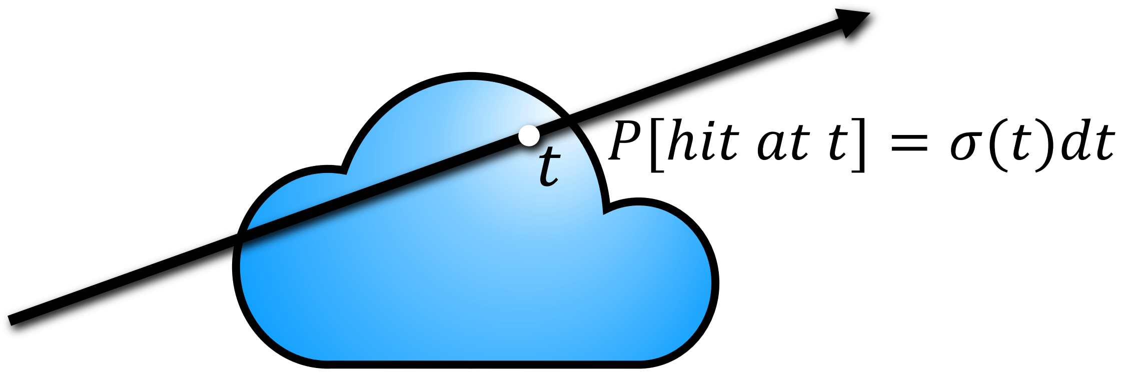

This notion is *probabilistic:* chance that ray hits a particle in a small interval around t is $\sigma(t)dt$.
σ is called the "volume density"

Is it the **first** hit?

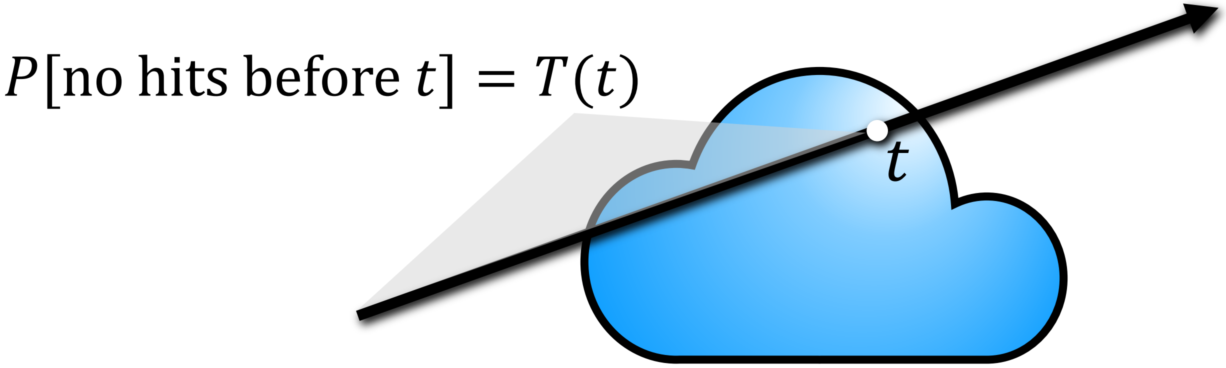

To determine if t is the *first* hit along the ray, need to know $T(t)$: **the probability that the ray makes it through the volume up to** t.

$T(t)$ is called "transmittance"

Define First Fit

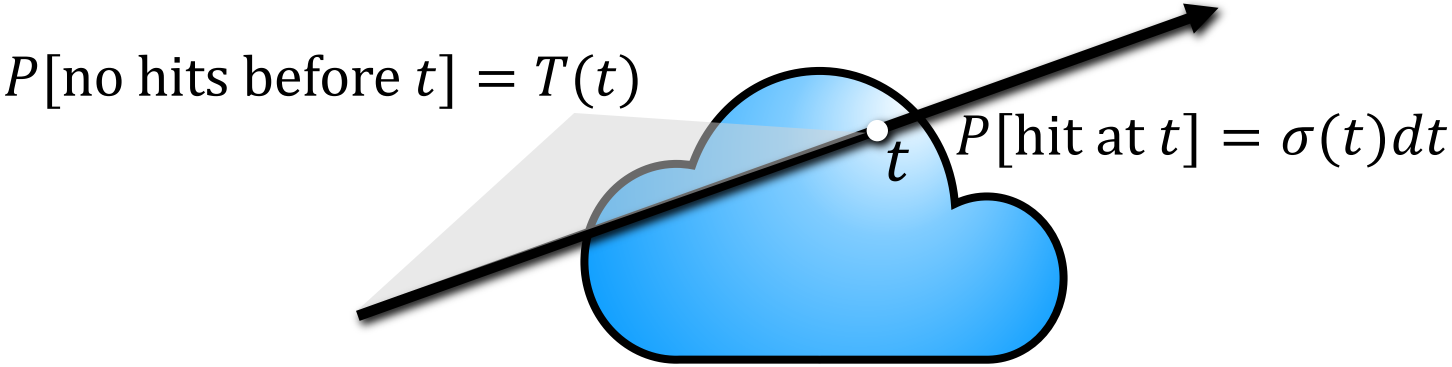

The product of these probabilities tells us how much you see the particles at t:
$P[\text{first hit at } t] = P[\text{no hit before } t] \times P[hit\ at\ t] = T(t)\sigma(t)dt$

Also called Ray Termination
Let's write T as a function of σ ! How?

Calculating T given σ

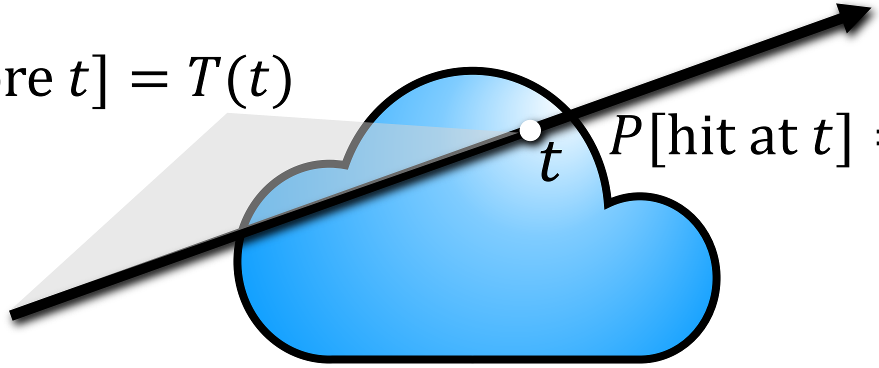

We got: $P[\text{first hit at } t] = P[\text{no hit before } t] \times P[hit\ at\ t] = T(t)\sigma(t)dt$

Now use a slightly different equation to relate σ and T:

$$\underbrace{P[\text{no hit before } t + dt]}_{T(t+dt)} = \underbrace{P[\text{no hit before } t]}_{T(t)} \times \underbrace{P[\text{no hit at } t]}_{(1-\sigma(t)dt)}$$

Now we can solve for T as a function of σ

Solve for T as a function of σ

$$\underbrace{P[\text{no hit before } t+dt]}_{T(t+dt)} = \underbrace{P[\text{no hit before } t]}_{T(t)} \times \underbrace{P[\text{no hit at } t]}_{(1-\sigma(t)dt)}$$

Solve the differential equation $T(t+dt) = T(t)(1-\sigma(t)dt)$

Expanded Righthand side

Taylor expansion $\Rightarrow \cancel{T(t)} + T'(t)dt = \cancel{T(t)} - T(t)\sigma(t)dt$

Rearrange$\Rightarrow \dfrac{T'(t)}{T(t)}dt = -\sigma(t)dt$

Integrate$\Rightarrow \log T(t) = -\int_{t_0}^{t}\sigma(s)ds$

Exponentiate$\Rightarrow T(t) = \exp\left(-\int_{t_0}^{t}\sigma(s)ds\right)$

Integral of:

$$\int \frac{f'(x)}{f(x)}dx = \log f(x)$$

Finally, we can write the ray termination PDF

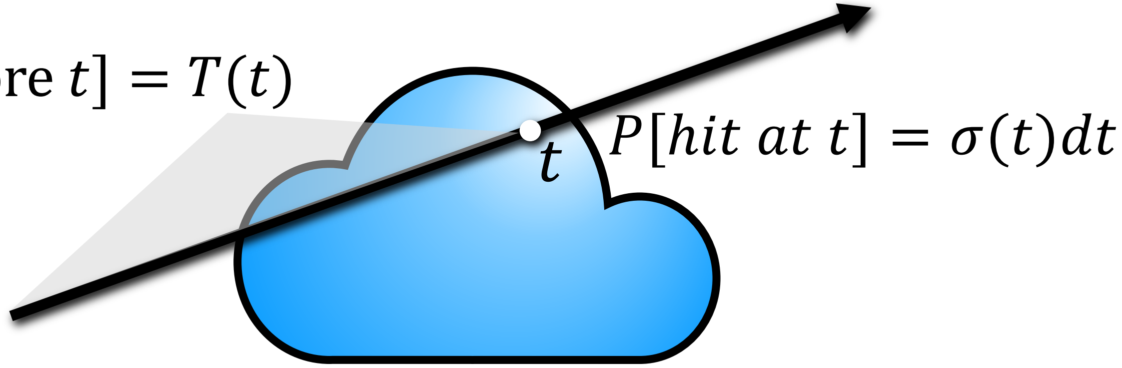

Finally, we can write the probability that a ray terminates at t as a function of only sigma

$$P[\text{first hit at } t] = P[\text{no hit before } t] \times P[\text{hit at } t]$$

$$= T(t)\sigma(t)dt$$

$$= \exp\left(-\int_{t_0}^{t} \sigma(s)ds\right)\sigma(t)dt$$

Finally, Expected Color along the ray

Then, the expected color returned by the ray will be

$$\boldsymbol{c}(\boldsymbol{r}) = \int_{t_0}^{t_1} P[first\ hit\ at\ t]\boldsymbol{c}(t)dt$$

$$= \int_{t_0}^{t_1} T(t)\sigma(t)\mathbf{c}(t)dt$$

$$= \int_{t_0}^{t_1} \exp\left(-\int_{t_0}^{t} \sigma(s)ds\right)\sigma(t)\mathbf{c}(t)dt$$

Note the nested integral!

Approximating the nested integral

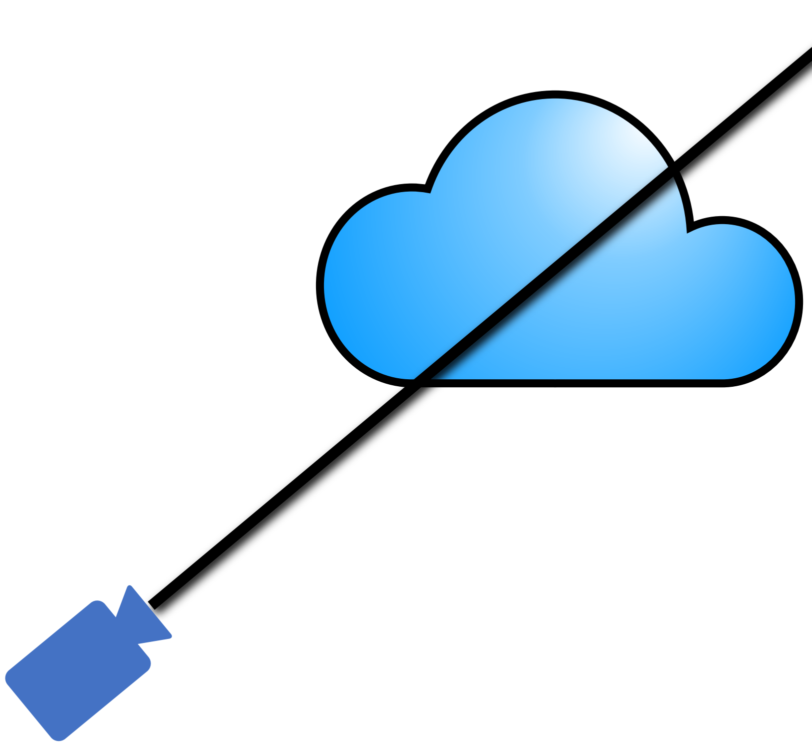

We use quadrature to approximate the nested integral,

Approximating the nested integral

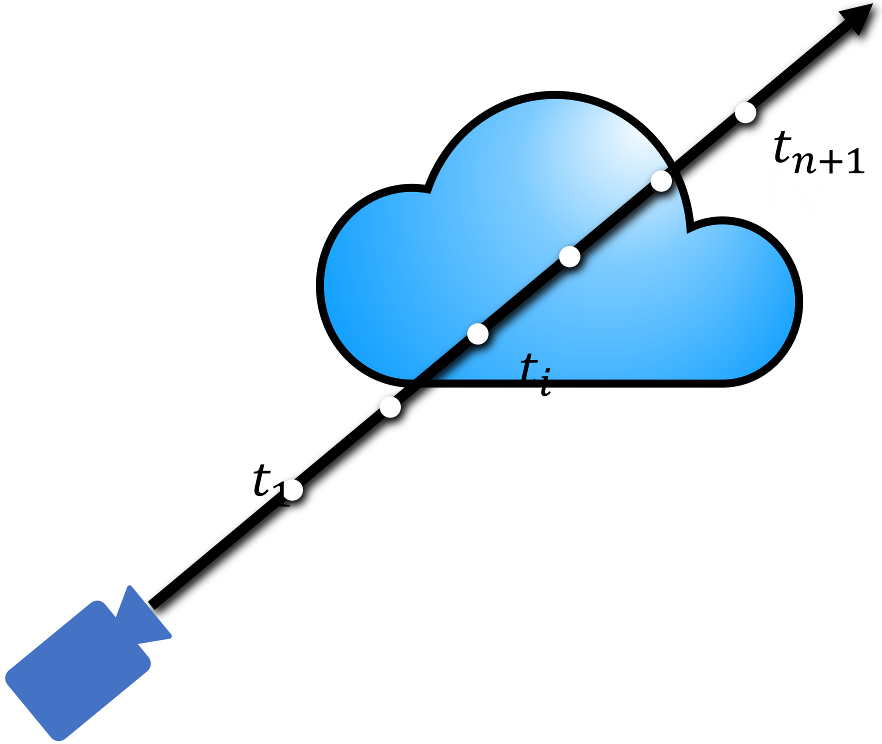

We use quadrature to approximate the nested integral, splitting the ray up into n segments with endpoints $\{t_1, t_2, \ldots, t_{n+1}\}$

Slide Credit: Ben Mildenhall

Approximating the nested integral

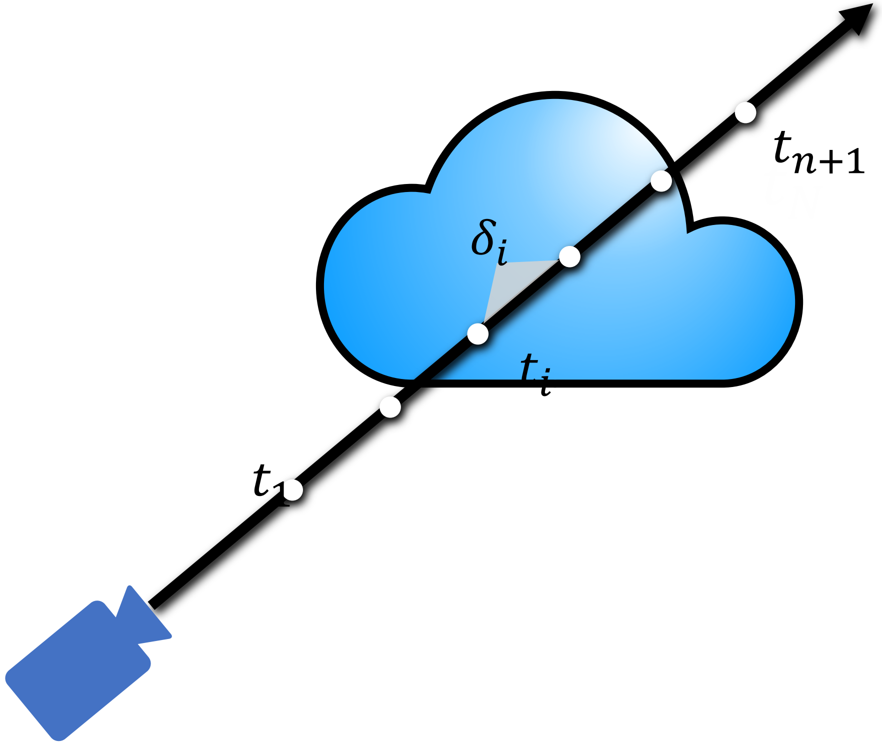

We use quadrature to approximate the nested integral, splitting the ray up into n segments with endpoints $\{t_1, t_2, \ldots, t_{n+1}\}$
with lengths $\delta_i = t_{i+1} - t_i$

Slide Credit: Ben Mildenhall

Approximating the nested integral

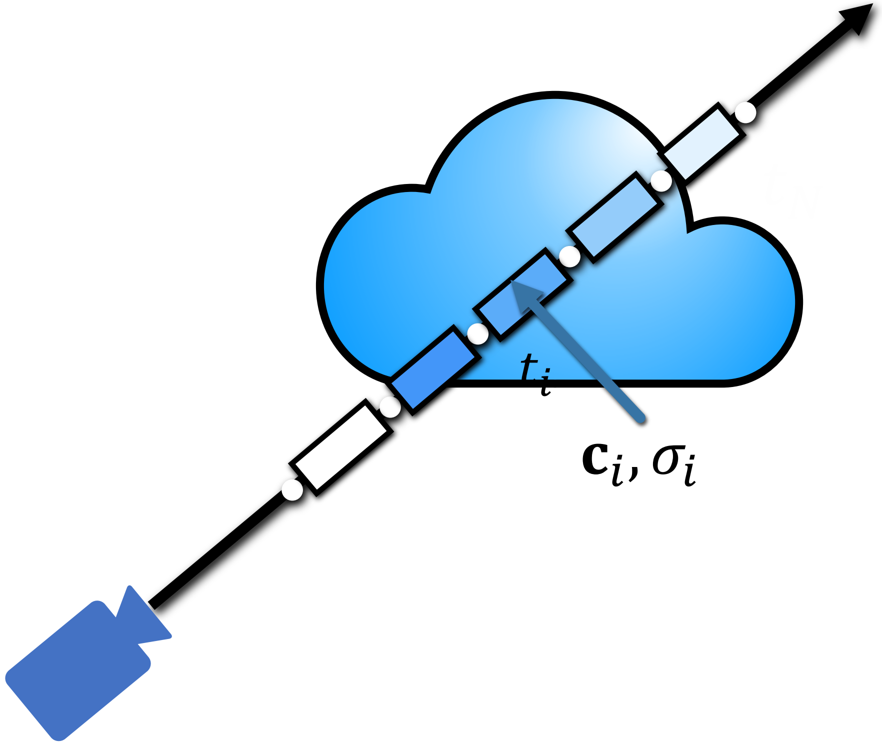

We assume volume density and color are roughly constant within each interval

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

Expected color: $\int T(t)\sigma(t)\mathbf{c}(t)dt \approx$

This allows us to break the outer integral

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

Expected color: $\int T(t)\sigma(t)\mathbf{c}(t)dt \approx \sum_{i=1}^{n}\int_{t_i}^{t_{i+1}} T(t)\sigma_i\mathbf{c}_i dt$

This allows us to break the outer integral
into a sum of analytically tractable integrals

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

$$\int T(t)\sigma(t)\mathbf{c}(t)dt \approx \sum_{i=1}^{n} \int_{t_i}^{t_{i+1}} \boxed{T(t)}\sigma_i \mathbf{c}_i dt$$

Caveat: piecewise constant density and color
do not imply constant transmittance!

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

$$\int T(t)\sigma(t)\mathbf{c}(t)dt \approx \sum_{i=1}^{n} \int_{t_i}^{t_{i+1}} \boxed{T(t)}\sigma_i \mathbf{c}_i dt$$

Caveat: piecewise constant density and color **do not** imply constant transmittance!

Important to account for how early part of a segment blocks later part when σ_i is high

Evaluating T for piecewise constant density

For $t \in [t_i, t_{i+1}]$, $T(t) = \exp\left(-\int_{t_1}^{t_i} \sigma_i ds\right) \exp\left(-\int_{t_i}^{t} \sigma_i ds\right)$

We need to evaluate at continuous t values
that can lie *partway through* an interval

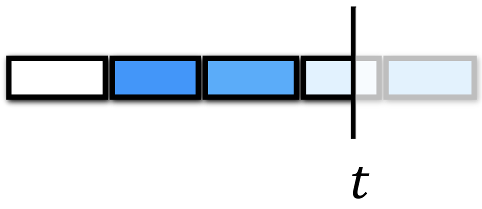

Slide Credit: Ben Mildenhall

Evaluating T for piecewise constant density

For $t \in [t_i, t_{i+1}]$, $T(t) = \exp\left(-\int_{t_1}^{t_i} \sigma_i ds\right) \exp\left(-\int_{t_i}^{t} \sigma_i ds\right)$

$$\exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right) = T_i$$

"How much light is blocked by all previous segments?"

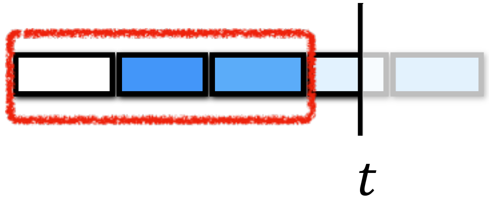

Slide Credit: Ben Mildenhall

Evaluating T for piecewise constant density

For $t \in [t_i, t_{i+1}]$, $T(t) = \exp\left(-\int_{t_1}^{t_i} \sigma_i ds\right) \exp\left(-\int_{t_i}^{t} \sigma_i ds\right)$

"How much light is blocked partway through the current segment?"

$$\exp(-\sigma_i(t - t_i))$$

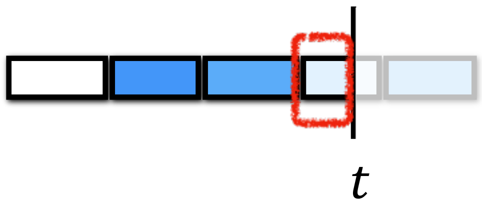

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

$$\int T(t)\sigma(t)\mathbf{c}(t)dt \approx \sum_{i=1}^{n} \int_{t_i}^{t_{i+1}} T(t)\sigma_i \mathbf{c}_i dt$$

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

$$\int T(t)\sigma(t)\mathbf{c}(t)dt \approx \sum_{i=1}^{n} \int_{t_i}^{t_{i+1}} T(t)\sigma_i \mathbf{c}_i dt$$

Substitute $$= \sum_{i=1}^{n} T_i \sigma_i \mathbf{c}_i \int_{t_i}^{t_{i+1}} \exp(-\sigma_i (t - t_i)) dt$$

Slide Credit: Ben Mildenhall

Deriving quadrature estimate

$$\int T(t)\sigma(t)\mathbf{c}(t)dt \approx \sum_{i=1}^{n} \int_{t_i}^{t_{i+1}} T(t)\sigma_i \mathbf{c}_i dt$$

$$= \sum_{i=1}^{n} T_i \sigma_i \mathbf{c}_i \int_{t_i}^{t_{i+1}} \exp(-\sigma_i(t - t_i))dt$$

Integral of Exponential:

$$\int \exp(-ax)\, dx = -\frac{1}{a}\exp(-ax)$$

Integrate $= \sum_{i=1}^{n} T_i \sigma_i \mathbf{c}_i \dfrac{\exp(-\sigma_i(t_{i+1} - t_i)) - 1}{-\sigma_i}$

$$\int_{t_i}^{t_{i+1}} \exp(-\sigma(t - t_i))\, dt = -\frac{1}{\sigma}\exp(-\sigma(t - t_i)) \,\Big|_{t_i}^{t_{i+1}}$$

Cancel σ_i $= \sum_{i=1}^{n} T_i \mathbf{c}_i (1 - \exp(-\sigma_i \delta_i))$

$$\frac{\exp(-\sigma_i(t_{i+1}-t_i)) - \exp(-\sigma_i(t_i - t_i))}{-\sigma_i} = \frac{\exp(-\sigma_i(t_{i+1}-t_i)) - 1}{-\sigma_i}$$

Expected Color $= \sum_{i=1}^{n} T_i \mathbf{c}_i (1 - \exp(-\sigma_i \delta_i))$

Putting it all together

$$\boxed{\text{Expected Color} = \sum_{i=1}^{n} T_i \mathbf{c}_i (1 - \exp(-\sigma_i \delta_i))}$$

where $$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right)$$

Slide Credit: Ben Mildenhall

Connection to alpha compositing

$$\boxed{\text{Expected Color} = \sum_{i=1}^{n} T_i \mathbf{c}_i \underbrace{(1 - \exp(-\sigma_i \delta_i))}_{\text{segment opacity } \alpha_i}}$$

$$\boxed{\text{Expected Color} = \sum_{i=1}^{n} T_i \mathbf{c}_i \alpha_i}$$

where $$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right)$$

$$= \prod_{j=1}^{i-1} (1 - \alpha_j)$$

$$\prod_i \exp(x_i) = \exp(\sum_i x_i)$$
$$\alpha_i = 1 - \exp(\sigma_i \delta_i)$$
$$1 - \alpha_i = -\exp(\sigma_i \delta_i)$$

Summary

for a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$:

differentiable w.r.t. $\mathbf{c}, \sigma$

$$\mathbf{c} \approx \boxed{\sum_{i=1}^{n} \underbrace{w_i}_{\text{weights}} \underbrace{\mathbf{c}_i}_{\text{colors}} = \sum_{i=1}^{n} T_i \alpha_i \mathbf{c}_i}$$

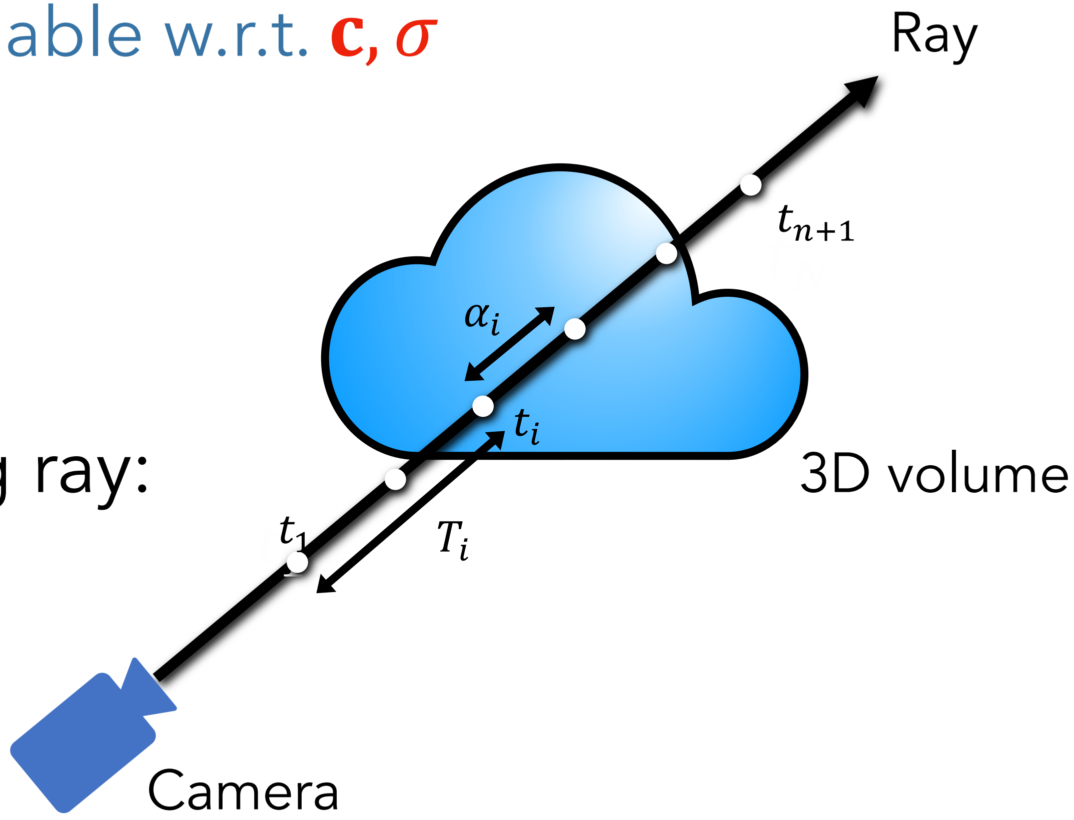

How much light is blocked earlier along ray:

$$T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$$

How much light is contributed by ray segment *i*:

$$\alpha_i = 1 - \exp(-\sigma_i \delta_i)$$