

Embodied Vision

Rendering and 3D Scene Reconstruction

Tsung-Wei Ke

Spring 2026



Disclaimer

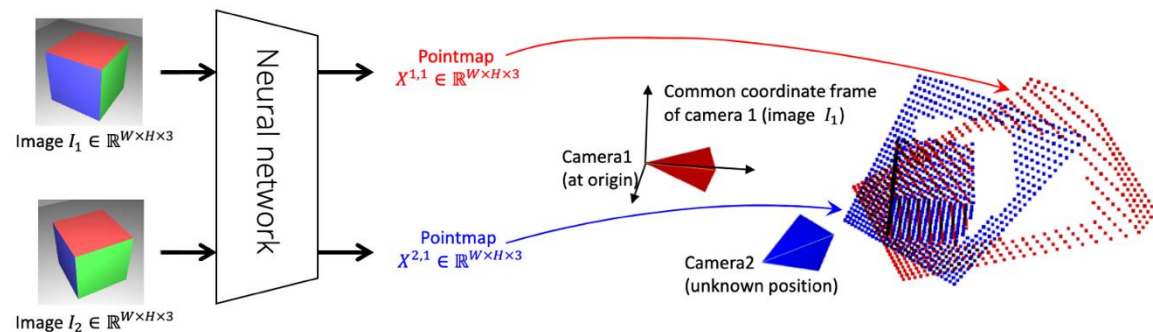
- This lecture borrows contents heavily from
 - [Learning for 3D Vision](#) by Shubam Tulsiani at CMU
 - [Image Synthesis Techniques](#) by Kayvon Fatahalian, Doug James and Matt Pharr at Stanford
 - [Computer Graphics and Imaging](#) by Ren Ng at UC Berkeley
 - [Machine Learning for Inverse Graphics](#) by Vincent Sitzmann at MIT

Recap

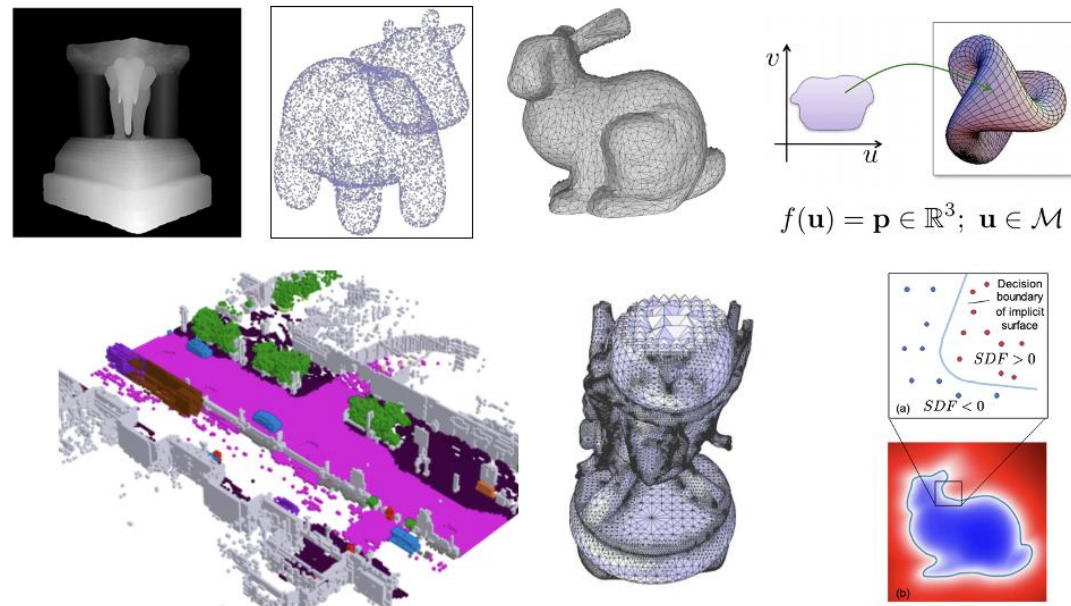
Structure from Motion




DUST3R



3D Representations



Conversion between 3D Representations



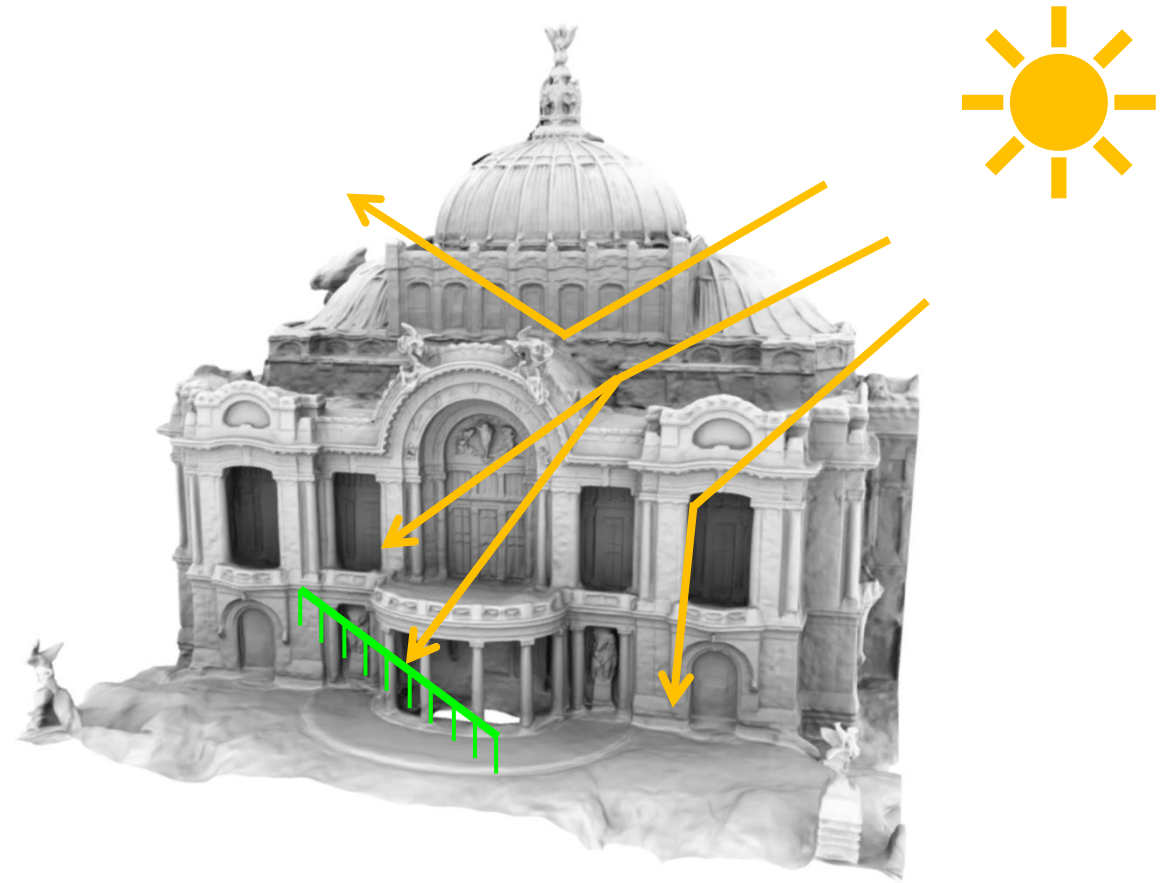
Content

- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting


Content

- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting


We Know How to Model the 3D World, Let's Revisit Image Formation



Recap: 2D Image Formation from the 3D World--A simplified model of photometric image formation



Absorption: Light Gets Absorbed when Traveling In a Medium



Absorption: Light Gets Absorbed when Traveling In a Medium



Low density smoke



High density smoke

Absorption: Light Gets Absorbed when Traveling In a Medium



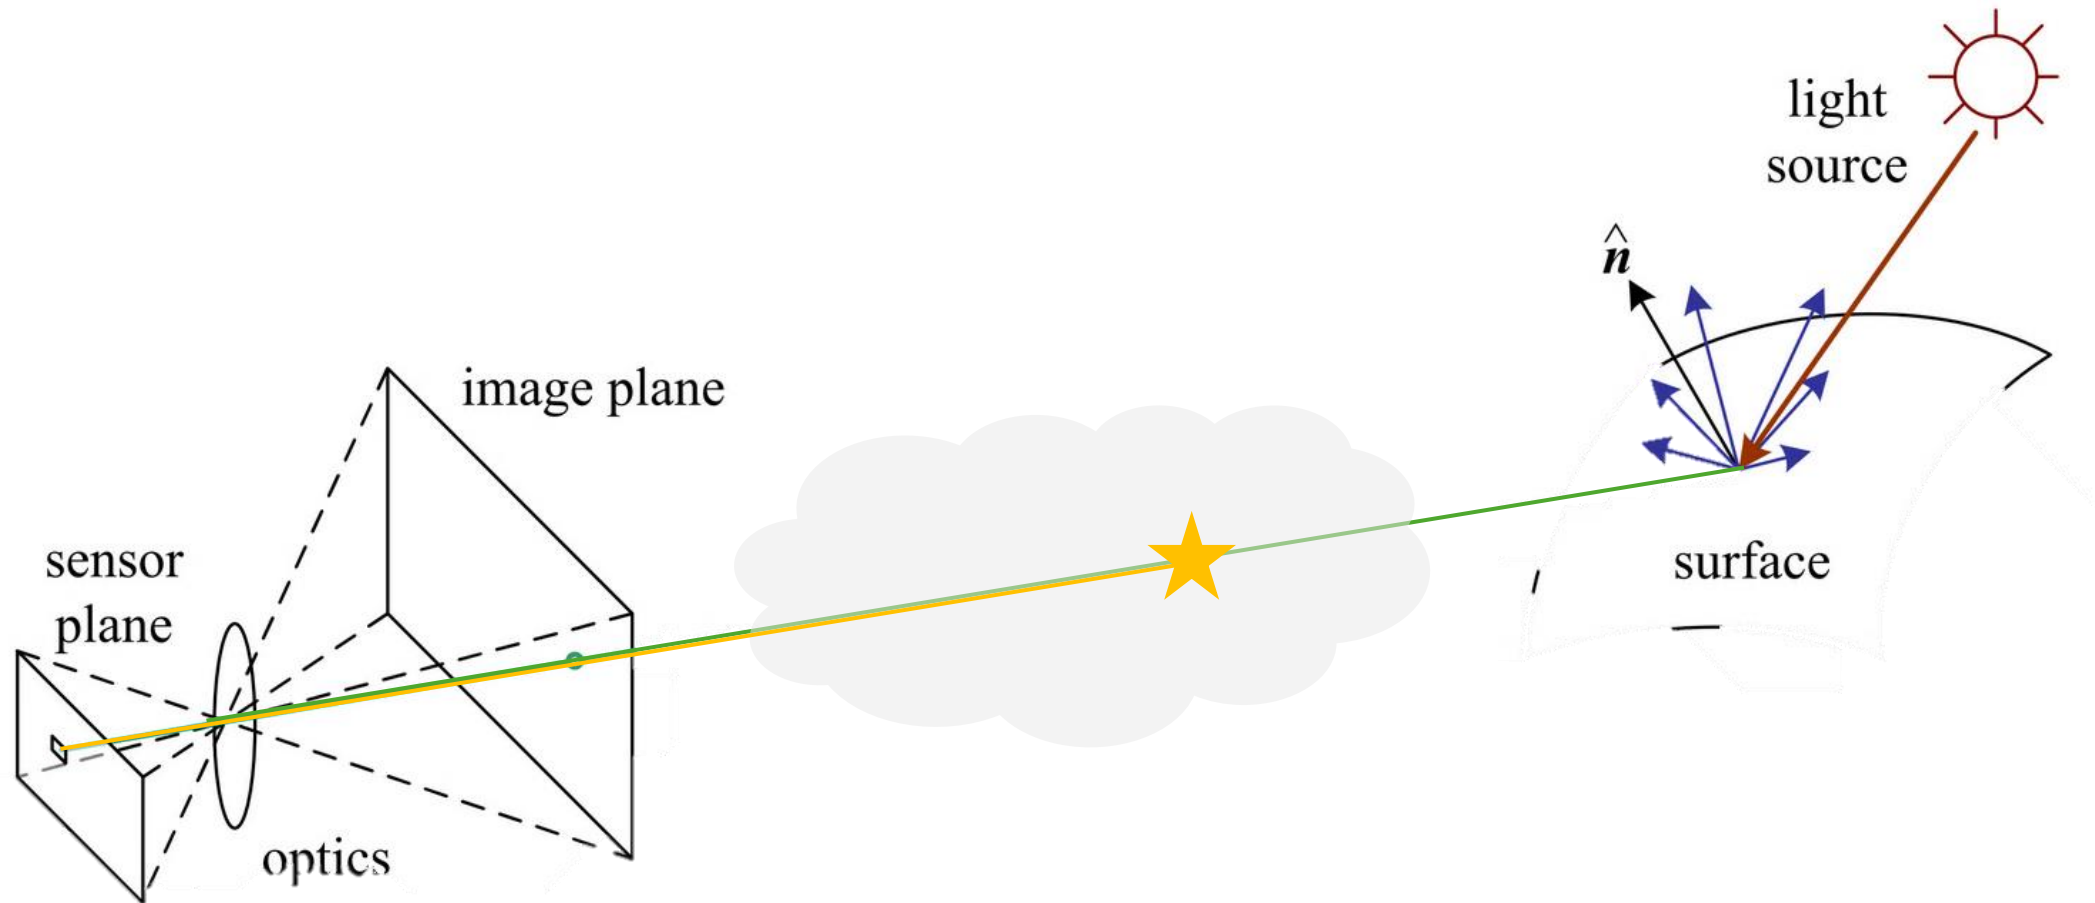
Absorption: Light Gets Absorbed when Traveling In a Medium



Thus, if one is to be five times as distant, make it five times bluer.

—Treatise on Painting, Leonardo Da Vinci, pp 295, circa 1480.


Emission: The Medium Emits Light




Emission: The Medium Emits Light



Scattering: Light Along the Ray Comes from Different Source




Scattering: Light Along the Ray Comes from Different Source



**More Backward
Scattering**

**More Forward
Scattering**

Scattering: Light Along the Ray Comes from Different Source



EPFL RGL Material Database

Surface Reflection Oversimplifies Image Formation



Greg Zaal: <https://polyhaven.com/a/rathaus>


Content

- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting

Let's Derive the Mathematical Model for Image Rendering


How to Measure Light?

Flux – What's the Density of Photons Flowing Through a Sensor?




From London and Upton

Different Types of Light Measurements



**Light Emitted
From A Source**

“Radiant Intensity”



**Light Falling
On A Surface**


“Irradiance”



**Light Traveling
Along A Ray**

“Radiance”

Radiant Intensity: The Power Per Unit Solid Angle Emanating from a Point Source




Radiant intensity:

$$I(\omega) \equiv \frac{d\Phi}{d\omega}$$

Φ : the radiant energy emitted, reflected, transmitted and received per unit time

ω : solid angle

Radiant Intensity: The Power Per Unit Solid Angle Emanating from a Point Source



Radiant intensity:

$$I(\omega) \equiv \frac{d\Phi}{d\omega}$$





Image source
<https://www.jensign.com/LEDIntensity>

Solid Angle

Angle

$$\theta = \frac{l}{r}$$


⇒ circle has 2π radians




Solid angle

$$\Omega = \frac{A}{r^2}$$

⇒ sphere has 4π steradians



Infinitesimal Solid Angle



$$\begin{aligned} dA &= (rd\theta)(r \sin \theta d\phi) \\ &= r^2 \sin \theta d\theta d\phi \end{aligned}$$


Infinitesimal solid angle:

$$d\omega = \frac{dA}{r^2} = \sin \theta d\theta d\phi$$

Total solid angle:

$$\begin{aligned} \Omega &= \oint_{S^2} d\omega = \int_0^\pi \int_0^{2\pi} \sin \theta d\theta d\phi \\ &= \int_{-1}^1 \int_0^{2\pi} d \cos \theta d\phi = 4\pi \end{aligned}$$


Isotropic Point Source



$$\begin{aligned}\Phi &= \int_{S^2} I \, d\omega \\ &= 4\pi I\end{aligned}$$

$$I = \frac{\Phi}{4\pi}$$

Irradiance: The Power Per Unit Area Incident on a Surface



Irradiance

$$E(x) \equiv \frac{d\Phi}{dA}$$

$$A = A' \cos \theta$$

$$E' = E \cos \theta$$

Irradiance Falloff

Assume light is emitting flux Φ in a uniform angular distribution

Compare irradiance at surface of two spheres:



intensity here: E

$$E = \frac{\Phi}{4\pi}$$

intensity here: E/r^2

$$E' = \frac{\Phi}{4\pi r^2} = \frac{E}{r^2}$$

Radiance: The Power Per Unit Area Per Unit Solid Angle Leaving/Received by a Surface



- Radiance

$$L(x, \omega) \equiv \frac{dI(x, \omega)}{dA_{\perp}} = \frac{d^2\Phi(x, \omega)}{d\omega dA \cos \theta}$$

- For geometric optics, the fundamental carrier of light is a ray
- The amount of light traveling along a ray is radiance

Incident Surface Radiance



$$L_i(x, \omega) \equiv \frac{dE(x, \omega)}{d\omega \cos \theta}$$

Exiting Surface Radiance



$$L_o(x, \omega) \equiv \frac{dI(x, \omega)}{dA \cos \theta}$$

Irradiance from the Environment

Computing flux per unit area on surface, due to incoming light from all directions.


$$dE(p, \omega) = L_i(p, \omega) \cos \theta d\omega \quad \leftarrow \text{Contribution to irradiance from light arriving from direction } \omega$$

$$E(p) = \int_{H^2} L_i(p, \omega) \cos \theta d\omega$$



Light meter


Hemisphere: H^2



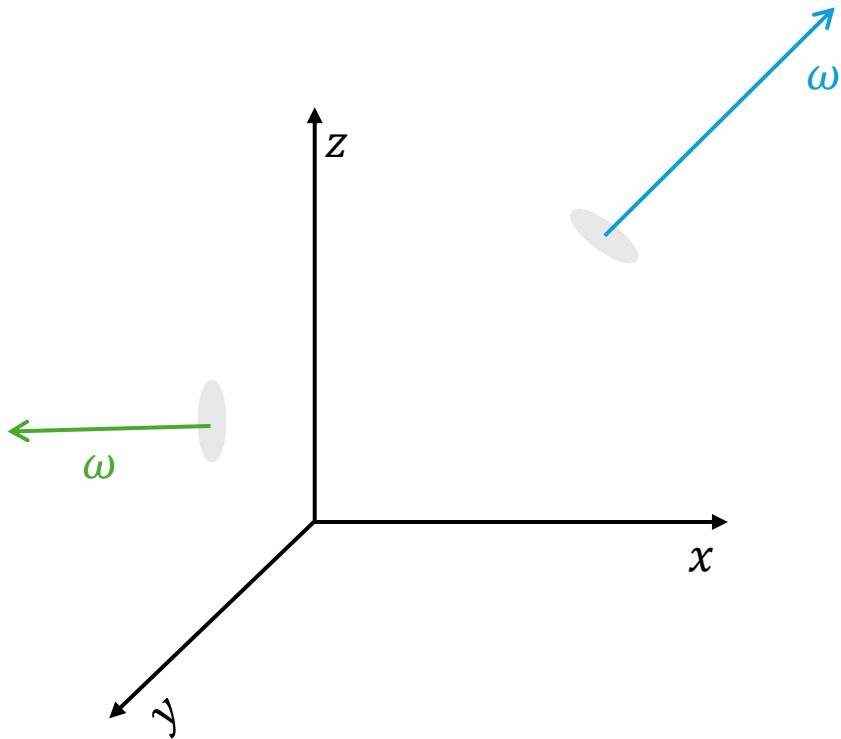
Irradiance from the Environment

$$\begin{aligned} E(\mathbf{p}) &= \int_{H^2} L \cos \theta \, d\omega \\ &= L \int_0^{2\pi} \int_0^{\frac{\pi}{2}} \cos \theta \sin \theta \, d\theta \, d\phi \\ &= L \pi \end{aligned}$$

Note: integral of cosine over hemisphere is only 1/2 the area of the hemisphere.



Light Field




- Definition: The field radiance (luminance) at a point in space in a given direction is the power per unit solid angle per unit area perpendicular to the direction


Content

- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting



Capture the Light Field and Reconstruct the Scene



Volume Rendering from Multi-View Images





Volume Rendering vs. Multi-View Stereo



- Multi-view stereo obtains 3D by correspondence
- Volume rendering obtains 3D by modeling reception of light


Volume Rendering a Ray



How much light is lost/gained along the differential beam due to interactions of light with the medium?

How does $L(x, \omega)$ vary along a ray?

Absorption



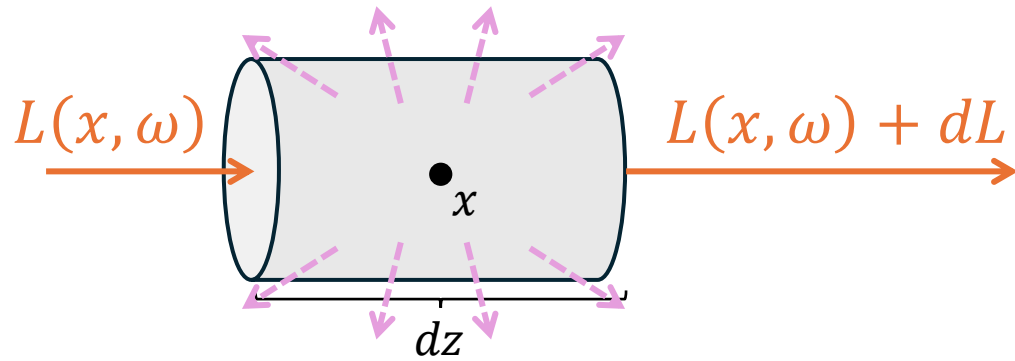
$$dL(x, \omega) = -\sigma_a(x)L(x, \omega)dz$$

Absorption across section σ_a :

- Probability of being absorbed per unit length
- Units: 1 / distance



Emission




$$dL(x, \omega) = \sigma_a(x) L_e(x, \omega) dz$$

Emission across section σ_a :

- Probability of emitting light per unit length
- Units: 1 / distance





Out Scattering




$$dL(x, \omega) = -\sigma_s(x)L(x, \omega)dz$$

Scattering across section σ_s :

- Probability of being scattered per unit length
- Units: 1 / distance



In Scattering



$$dL(x, \omega) = \sigma_s(x)L_s(x, \omega)dz$$

Scattering across section σ_s :


- Probability of gathering scattered light per unit length
- Units: 1 / distance



$$L_s(x, \omega) = \int_{S^2} p(\omega' \rightarrow \omega)L_i(x, \omega')d\omega'$$

Where $p(\omega' \rightarrow \omega)$ denotes the phase function and $L_i(x, \omega)$ also depend on other scattered light

Radiative Transfer Equation



$$dL(x, \omega) = -\sigma_a(x)L(x, \omega)dz - \sigma_s(x)L(x, \omega)dz + \sigma_a(x)L_e(x, \omega)dz + \sigma_s(x)L_s(x, \omega)dz$$


Absorption

Out
scattering

Emission

In
scattering

Radiative Transfer Equation



$$\begin{aligned}dL(x, \omega) &= -\sigma_a(x)L(x, \omega)dz - \sigma_s(x)L(x, \omega)dz + \sigma_a(x)L_e(x, \omega)dz + \sigma_s(x)L_s(x, \omega)dz \\ &= -\sigma_a(x)L(x, \omega)dz + \sigma_a(x)L'(x, \omega)dz \\ &= \sigma(x)(L'(x, \omega) - L(x, \omega))dz\end{aligned}$$

- Model absorption-only case
- Allow emission

$$L'(x, \omega) = L_e(x, \omega) + \frac{\sigma_s}{\sigma_a} (L_s(x, \omega) - L(x, \omega))$$

Absorption-only Volume Rendering

- Volume rendering in a heterogeneous medium:

$$dL(x, \omega) = -\sigma(x)L(x, \omega)dz \Rightarrow L(x_0 + \omega z, \omega) = e^{-\int_{t=0}^z \sigma(x+\omega t)dt} L(x_0, \omega)$$

- Homogeneous medium: $\sigma(x) = \sigma$
- Volume rendering in a homogeneous medium:

$$dL(x, \omega) = -\sigma L(x, \omega)dz \Rightarrow L(x_0 + \omega z, \omega) = e^{-\sigma z} L(x_0, \omega)$$

Transmittance

$$T(\mathbf{x}, \mathbf{y})$$

What fraction of radiance at \mathbf{x} in direction of \mathbf{y} , reaches \mathbf{y} ?
(along a straight line under absorption-only model)


Q: should this be symmetric?

Homogeneous Medium:


$$e^{-\sigma \|\mathbf{x}-\mathbf{y}\|}$$

Heterogeneous Medium:

$$e^{-\int_{t=0}^{\|\mathbf{x}-\mathbf{y}\|} \sigma(\mathbf{x}+\omega\mathbf{t}) dt}$$




Volume Rendering




$$L(x, \omega) = T(x_0, x_z)L(x_z, \omega)$$

Volume Rendering



$$L(x, \omega) = T(x_0, x_z)L(x_z, \omega) + \int_0^z T(x_0, x_t)\sigma_t L_e(x_t, \omega)dt$$


Volume Rendering in a Homogeneous Emitting Medium



$$L(x, \omega) = e^{-\sigma z} L(x_z, \omega) + \int_0^z e^{-\sigma t} \sigma L_e(x_t, \omega) dt$$

Homogeneous emitting medium:
 $C \equiv L_e(x_t, \omega)$
 $T(x, y) = e^{-\sigma \|x-y\|}$


Volume Rendering in a Homogeneous Emitting Medium



$$\begin{aligned}
 L(x, \omega) &= e^{-\sigma z} L(x_z, \omega) + \int_0^z e^{-\sigma t} \sigma L_e(x_t, \omega) dt \\
 &= e^{-\sigma z} L(x_z, \omega) + \int_0^z e^{-\sigma t} \sigma C dt \quad \int \lambda e^{-\lambda x} dx = -e^{-\lambda x} + c \\
 &= e^{-\sigma z} L(x_z, \omega) + (1 - e^{-\sigma z}) C
 \end{aligned}$$

Homogeneous emitting medium:
 $C \equiv L_e(x_t, \omega)$
 $T(x, y) = e^{-\sigma \|x - y\|}$

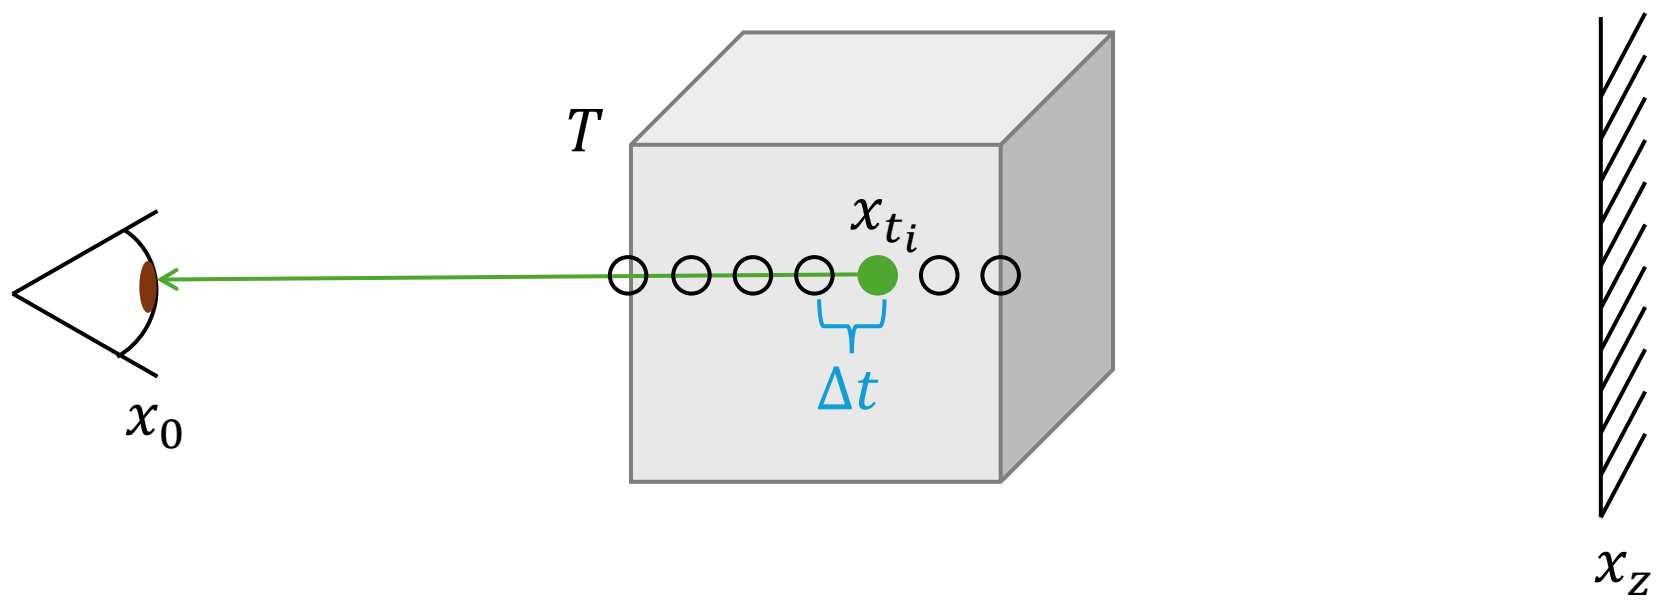
Volume Rendering in a Heterogeneous Medium



$$L(x, \omega) = T(x_0, x_z)L(x_z, \omega) + \int_0^z T(x_0, x_t)\sigma_t L_e(x_t, \omega) dt$$

Very difficult to solve!

Let's Ignore Background Light But Emission from the Medium




$$L(x, \omega) = \int_0^z T(x_0, x_t) \sigma_t L_e(x_t, \omega) dt$$

Approximate with discrete sum

$$\approx \sum_{i=1}^N T(x_0, x_{t_i}) \left(\text{emission}_{t_i}^{t_{i+1}} \right)$$

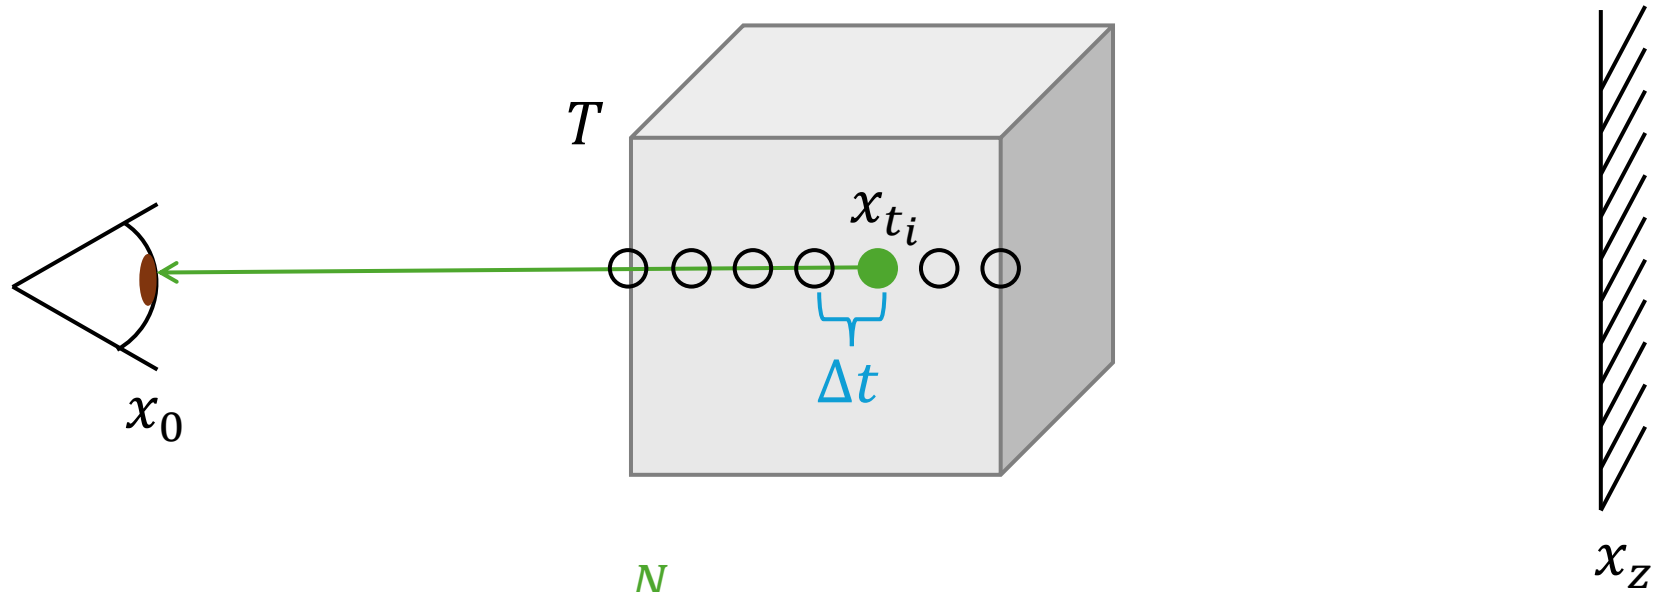
Let's Ignore Background Light But Emission from the Medium



$$L(x, \omega) \approx \sum_{i=1}^N T(x_0, x_{t_i}) \left(\text{radiance}_{t_{i+1}}^{t_i} \right)$$
$$= \sum_{i=1}^N T(x_0, x_{t_i}) (1 - e^{-\sigma_{t_i} \Delta t}) L_e(x_{t_i}, \omega)$$

Assume a homogeneous emitting medium between t_i and t_{i+1}

Let's Ignore Background Light But Emission from the Medium



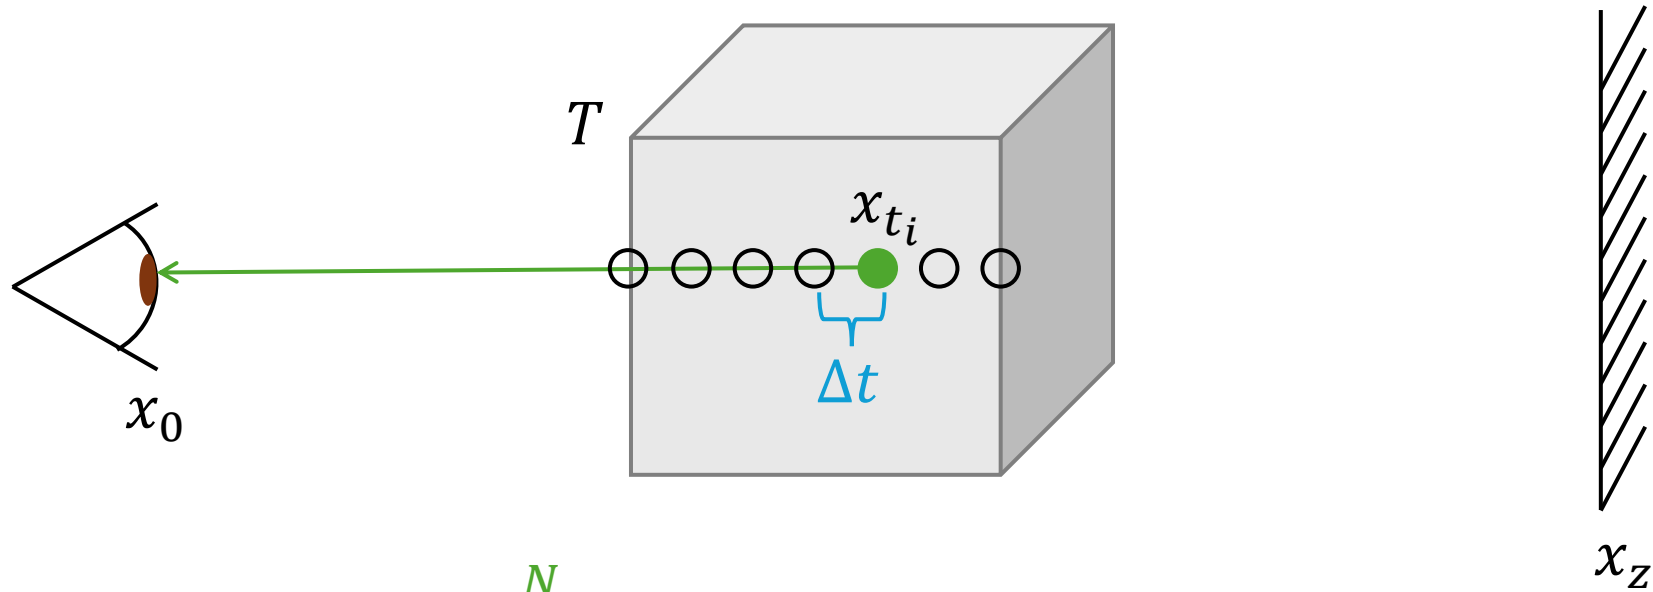
$$L(x, \omega) \approx \sum_{i=1}^N T(x_0, x_{t_i}) \left(\text{radiance}_{t_{i+1}}^{t_i} \right)$$

$$= \sum_{i=1}^N \boxed{T(x_0, x_{t_i})} (1 - e^{-\sigma_{t_i} \Delta t}) L_e(x_{t_i}, \omega)$$

This is unclear

Assume a homogeneous emitting medium between t_i and t_{i+1}


Let's Ignore Background Light But Emission from the Medium



$$L(x, \omega) \approx \sum_{i=1}^N T(x_0, x_{t_i}) (1 - e^{-\sigma_{t_i} \Delta t}) L_e(x_{t_i}, \omega)$$

$$T(x_0, x_{t_i}) = e^{-\int_0^{t_i} \sigma_t dt} = e^{-\int_0^{t_{i-1}} \sigma_t dt} e^{-\int_{t_{i-1}}^{t_i} \sigma_t dt} = T(x_0, x_{t_{i-1}}) e^{-\sigma_{t_{i-1}} \Delta t}$$


Volume Rendering in a Homogeneous Emitting Medium



$$L(x, \omega) \approx \sum_{i=1}^N T(x_0, x_{t_i}) (1 - e^{-\sigma_{t_i} \Delta t}) L_e(x_{t_i}, \omega) + T(x_0, x_{t_{N+1}}) L(x_z, \omega)$$


$$T(x_0, x_{t_i}) = e^{-\int_0^{t_i} \sigma_t dt} = e^{-\int_0^{t_{i-1}} \sigma_t dt} e^{-\int_{t_{i-1}}^{t_i} \sigma_t dt} = T(x_0, x_{t_{i-1}}) e^{-\sigma_{t_{i-1}} \Delta t}$$

Computational Volume Rendering: Ray Marching




1. Draw uniform samples along a ray (N segments, or N+1 points)
2. Compute transmittance between camera and each sample
3. Aggregate contributions across segments to get overall radiance (color)

Computational Volume Rendering: Ray Marching



1. Draw **non-uniform** samples along a ray
2. Compute transmittance between camera and each sample
3. Aggregate contributions across segments to get overall radiance (color)

Computational Volume Rendering: Ray Marching





$$L(x, \omega) \approx \sum_{i=1}^N T(x_0, x_{t_i}) (1 - e^{-\sigma_{t_i} \Delta t}) L_e(x_{t_i}, \omega)$$

$$T(x_0, x_{t_i}) = T(x_0, x_{t_{i-1}}) e^{-\sigma_{t_{i-1}} \Delta t}$$

- We can render any ray through the medium, if we know at each point t_i :
 - Absorption coefficient σ_{t_i}
 - Emitted light $L_e(x_{t_i}, \omega)$


How to Render Volumes from Multi-View Images?



- Given a set of multi-view images, we want to learn at each point:
 - Opacity (Absorption coefficient σ_{t_i})
 - Color (Emitted light $L_e(x_{t_i}, \omega)$)
- Loss: image reconstruction




How to Render Volumes from Multi-View Images?



$$C(x, \omega) \approx \sum_{i=1}^N T(x_0, x_{t_i}) (1 - e^{-\sigma_{t_i} \Delta t}) C(x_{t_i}, \omega), \text{ where } T(x_0, x_{t_i}) = T(x_0, x_{t_{i-1}}) e^{-\sigma_{t_{i-1}} \Delta t}$$

Learning objective: $\min_{\sigma_{t_i}, C(x_{t_i}, \omega)} \|C(x, \omega) - \text{g.t.}\|$

How to Render Volumes from Multi-View Images?

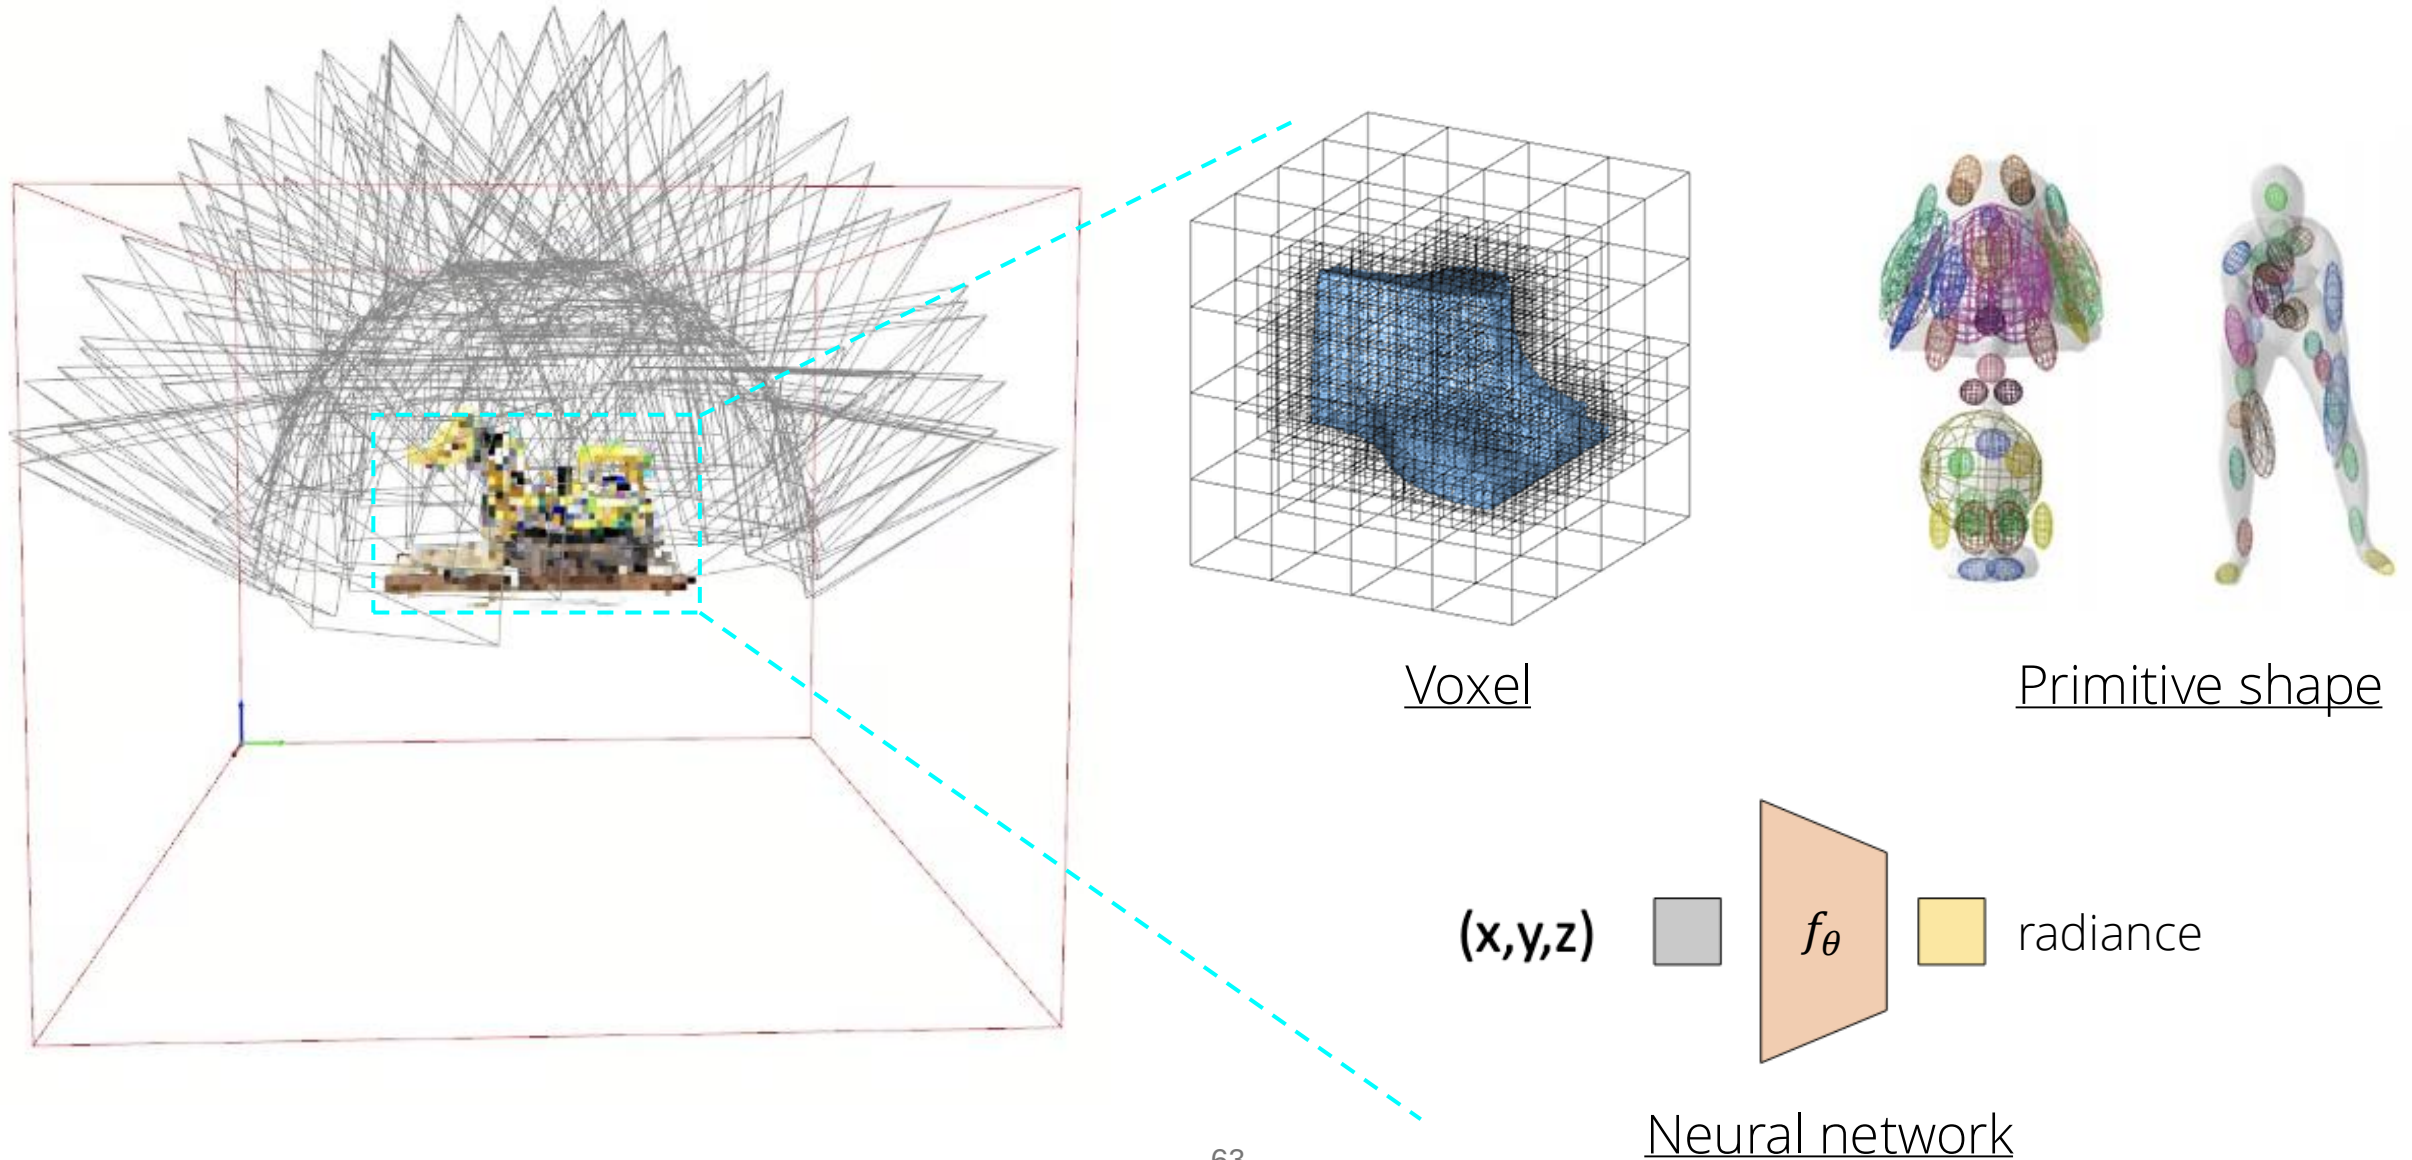


$$C(x, \omega) \approx \sum_{i=1}^N \left(\prod_{j=1}^{i-1} e^{-\sigma_j \Delta t_j} \right) (1 - e^{-\sigma_i \Delta t_i}) C_i$$

T_i α_i

Learning objective: $\min_{\sigma_{t_i}, C_i} \|C(x, \omega) - \text{g.t.}\|$


What is the Representation of Volume?





Content

- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting

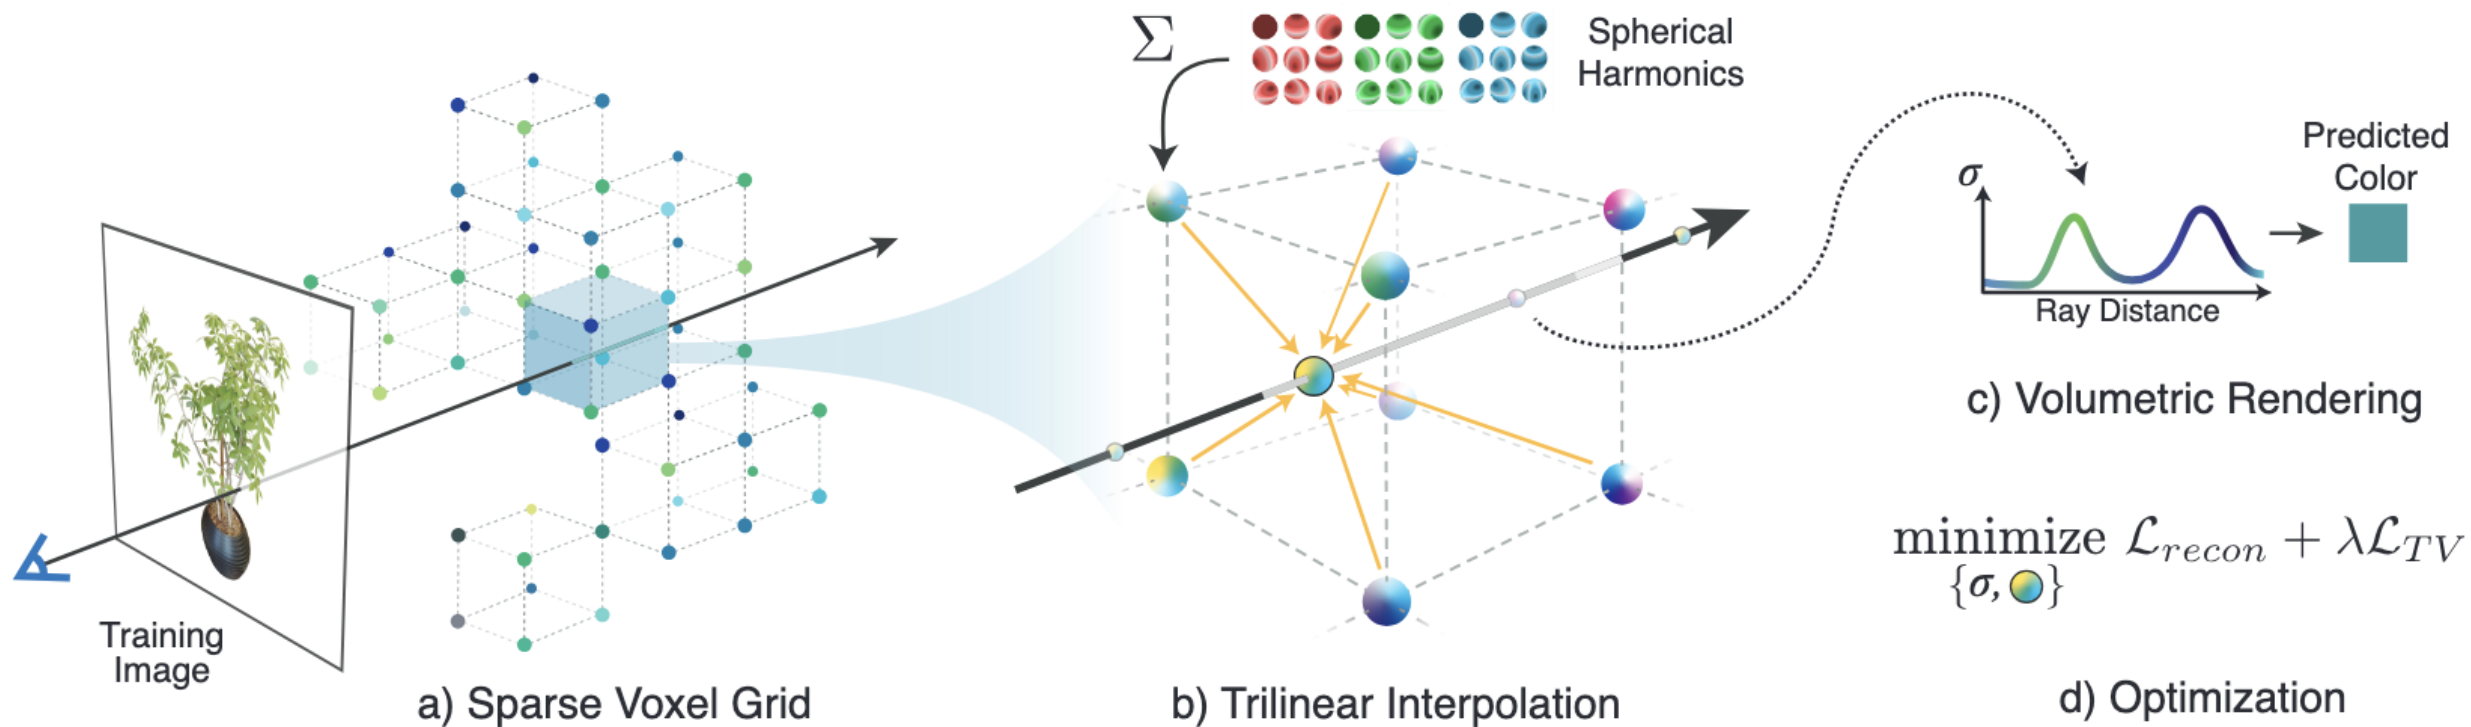
Volume Rendering with Voxel Grid




Storing RGB Color Cannot Handle View-Dependent Effect



Volume Rendering with Voxel Grid



Spherical Harmonics Model View-Dependent Lighting



- For each channel in RGB:
 - Predict coefficients $\mathbf{k} = \{k_l^m\}_{l,m}$
 - Decide the lighting

$$c(\mathbf{d}; \mathbf{k}) = S \left(\sum_{l=0}^{l_{\max}} \sum_{m=-l}^l k_l^m Y_l^m(\mathbf{d}) \right)$$

where S denotes sigmoid function

Spherical Harmonics Model View-Dependent Effect




Figure 5. The first 5 SH bands plotted as unsigned spherical functions by distance from the origin and by colour on a unit sphere. Green (light gray) are positive values and red (dark gray) are negative.

$$y_l^m(\theta, \varphi) = \begin{cases} \sqrt{2}K_l^m \cos(m\varphi)P_l^m(\cos\theta), & m > 0 \\ \sqrt{2}K_l^m \sin(-m\varphi)P_l^{-m}(\cos\theta), & m < 0 \\ K_l^0 P_l^0(\cos\theta), & m = 0 \end{cases}$$

$$K_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{4\pi(l+|m|)!}}$$

P_l^m denotes the associated Legendre polynomials

$$(l-m)P_l^m = x(2l-1)P_{l-1}^m - (l+m-1)P_{l-2}^m$$

$$P_m^m = (-1)^m (2m-1)!! (1-x^2)^{m/2}$$

$$P_{m+1}^m = x(2m+1)P_m^m$$


Spherical Harmonics Model View-Dependent Effect



Content


- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting

NEural Radiance Field (NERF)




- Use a neural network f_{θ} to predict at each point \mathbf{x} and view direction ω :
 - Opacity σ
 - Color \mathbf{c}


NEural Radiance Field (NERF)



- Use a neural network f_{θ} to predict at each point \mathbf{x} and view direction ω :
 - Opacity σ
 - Color \mathbf{c}



Problem 1: The Network Takes Inputs as \mathbf{x}, ω Coordinates Fails to Capture High-Frequency Features



- Use a neural network f_θ to predict at each point \mathbf{x} and view direction ω :
 - Opacity σ
 - Color \mathbf{c}




Ground Truth



No Positional Encoding

Solution: Maps as \mathbf{x}, ω Coordinates to High-Dimensional Feature Space

$$\gamma(v) = [\sin(v), \cos(v), \sin(2v), \cos(2v), \dots]$$



- Use a neural network f_θ to predict at each point \mathbf{x} and view direction ω :
 - Opacity σ
 - Color \mathbf{c}



Ground Truth





No Positional Encoding




Complete Model


Solution: Maps as x, ω Coordinates to High-Dimensional Feature Space



Problem 2: Uniform Sampling Points Along a Ray Ignores Regions with Fine Details



Solution: Hierarchically Sampling Points Along a Ray Based on the Opacity




$$L(\mathbf{x}, \omega) = \sum_i w_i \mathbf{c}_i; \quad w_i = T_i (1 - e^{-\sigma_{t_i} \Delta t})$$

weights indicate how much each segment contributes to the pixel value

- 1) Sample points uniformly
- 2) Sample *more* points per segment depending on weight

NeRF Captures View-Dependent Effect



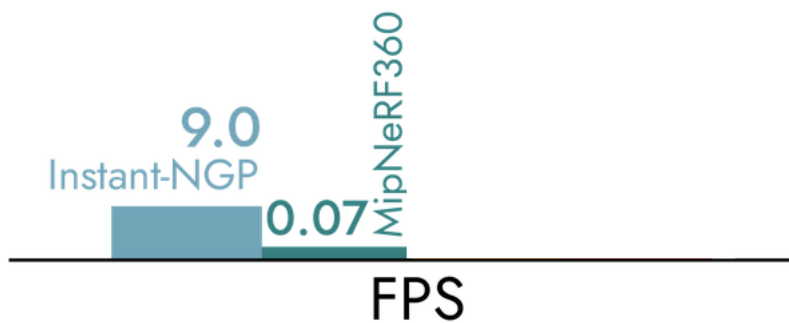
NeRF Decouples View Direction and Position




Content

- Light Sources
- Light Measurement
- A mathematical model for image rendering
- Volume Rendering
 - Voxel
 - Nerf
 - Gaussian Splatting

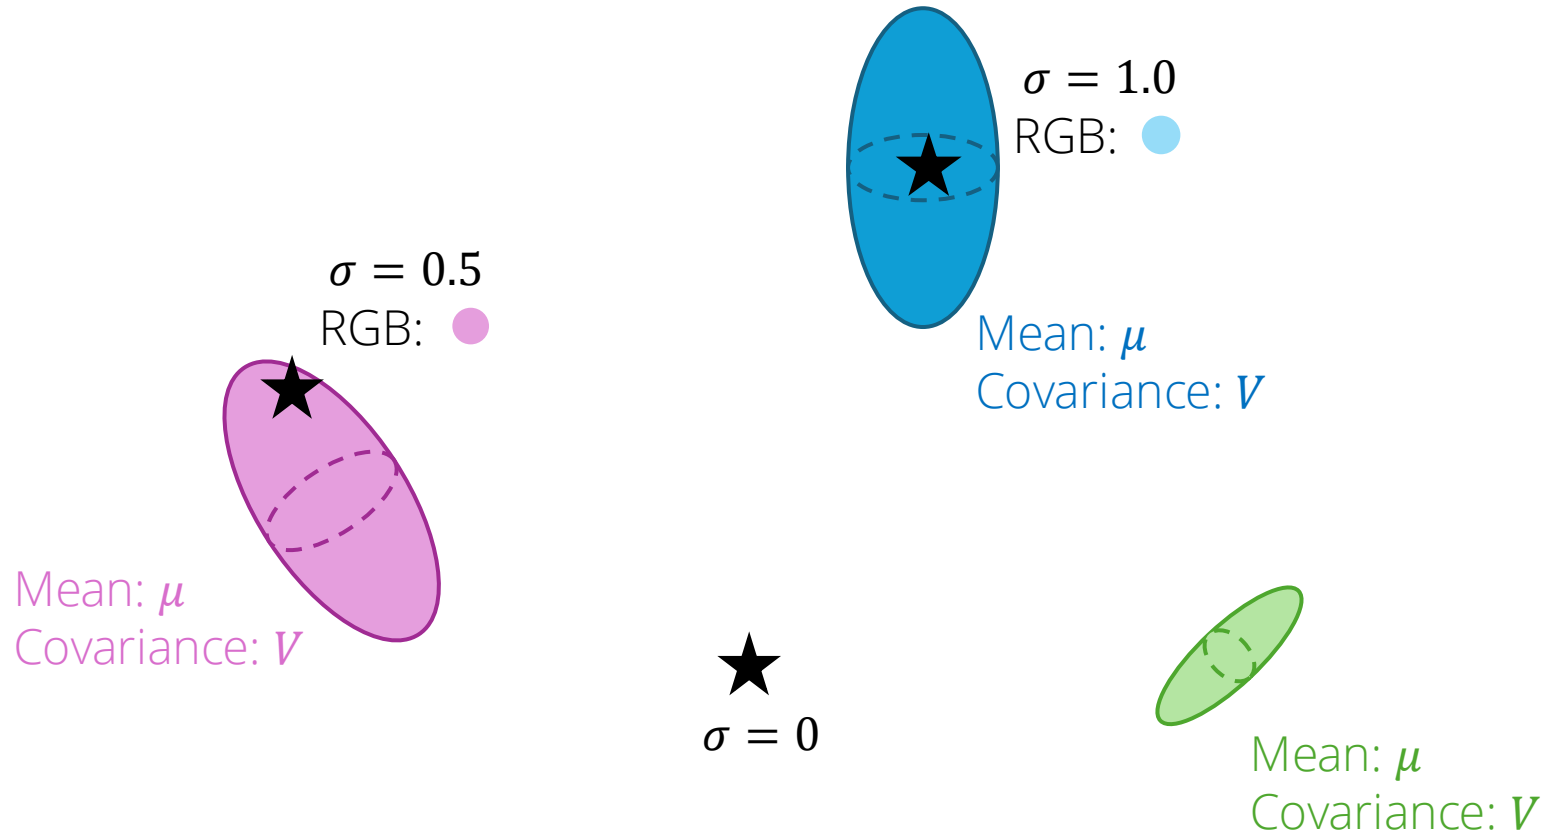
NeRF is Slow...



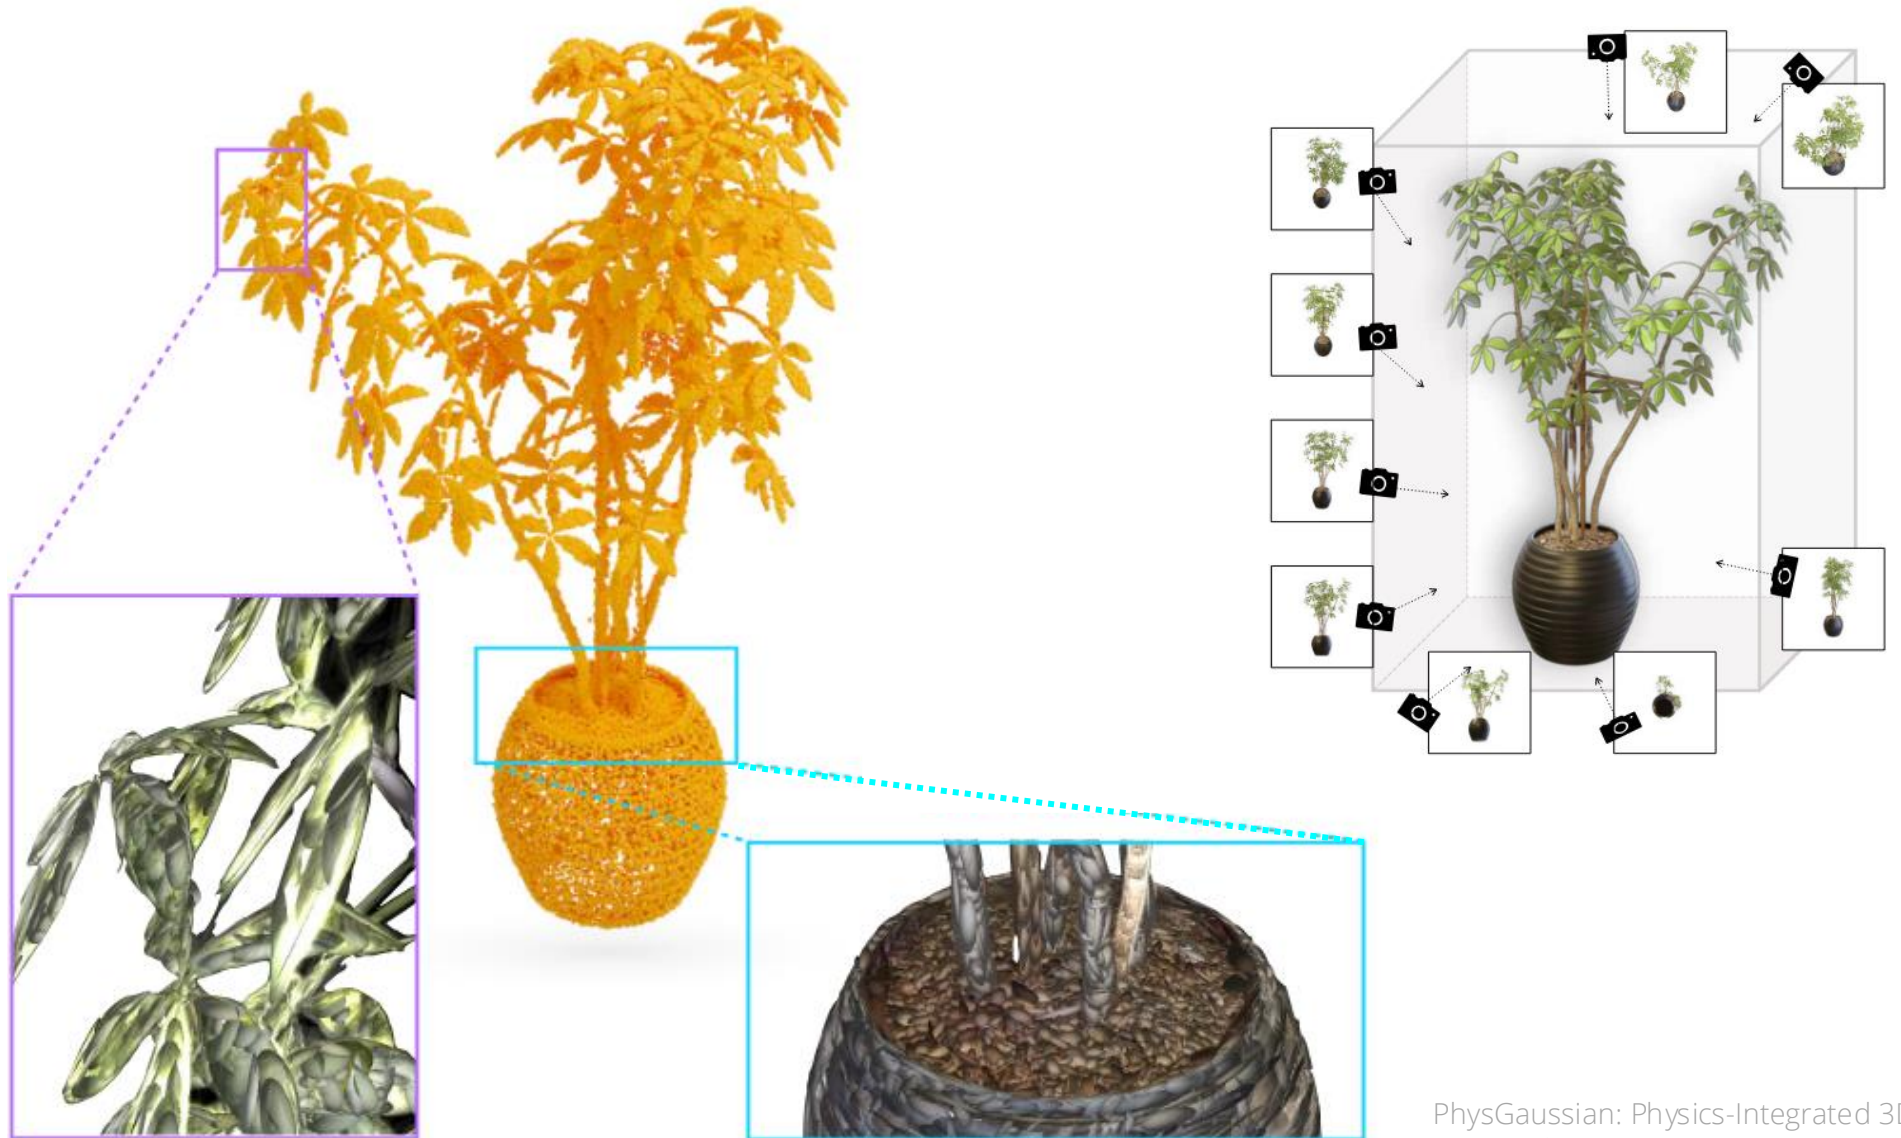
A New Approach for Real-time Rendering




Key Idea: Parameterize Radiance Field Sparsely, Only where Density is Nonzero




Volume Rendering with 3D Gaussian Splatting



How to Render? Same Volume Rendering Integral!

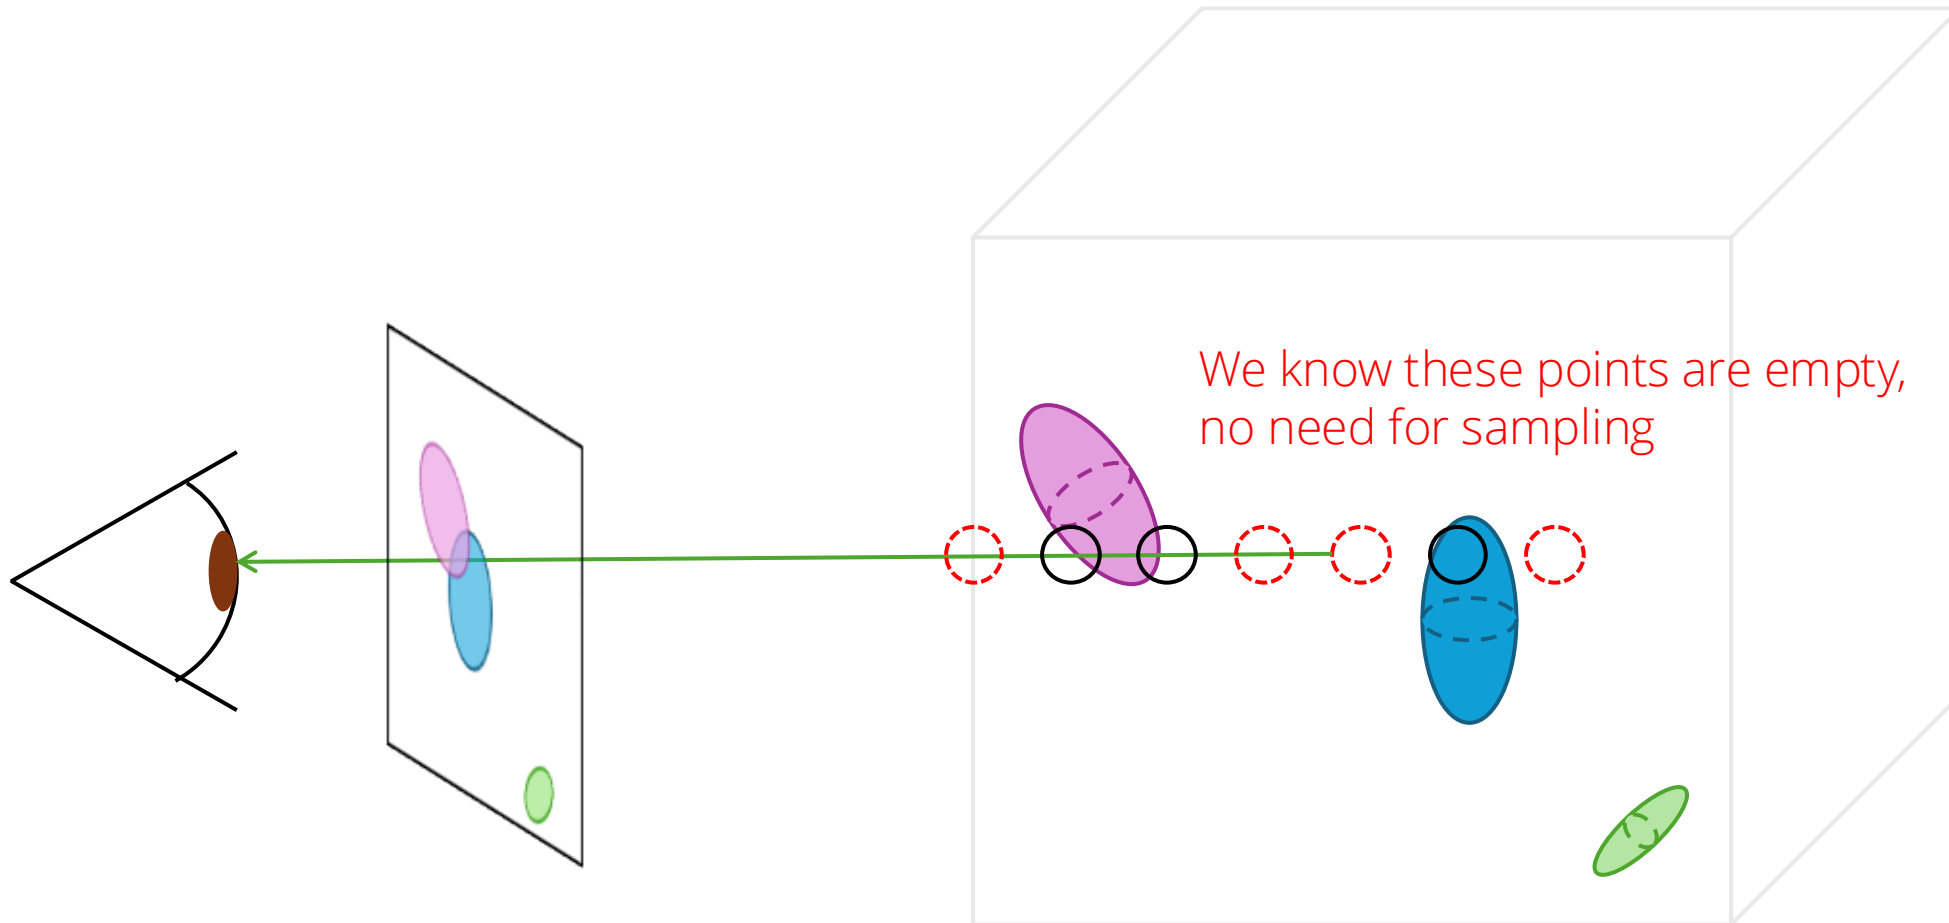


Volume Rendering Integral Samples Many Points from Empty Space, which is Slow





Can We Render from 3D Gaussians Smartly?

Yes! Project 3D Gaussians to 2D




Idea: Solve Image Rendering by Looking Up 2D Gaussians



$$\mathbb{I}_{\mathbf{x}} = \frac{\sum_{k=0}^{N-1} \rho_k(\mathbf{x}) \mathbf{w}_k}{\sum_{k=0}^{N-1} \rho_k(\mathbf{x})}.$$

$\rho_k(\mathbf{x})$: weightings derived from 2D Gaussian at pixel \mathbf{x}
 \mathbf{w}_k : point attributes (e.g. opacity) at pixel \mathbf{x}

Gaussians are Closed Under Integrals



- Integrate along x_2 axis with fixed x_0, x_1

$$\int_{\mathbb{R}} \mathcal{G}_V^3(x - p) dx_2 = \mathcal{G}_{\hat{V}}^2(\hat{x} - \hat{p})$$


3D Gaussian 2D Gaussian

$$x = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} \Leftrightarrow \hat{x} = \begin{bmatrix} x_0 \\ x_1 \end{bmatrix}$$

$$V = \begin{bmatrix} a & b & c \\ b & d & e \\ c & e & f \end{bmatrix} \Leftrightarrow \hat{V} = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$$

- Integrate along x_2 axis with fixed x_0, x_1

Gaussians are Closed Under Integrals



- Integrate along x_2 axis with fixed x_0, x_1

$$\int_{\mathbb{R}} \mathcal{G}_V^3(x - p) dx_2 = \mathcal{G}_{\hat{V}}^2(\hat{x} - \hat{p})$$


3D Gaussian 2D Gaussian

$$x = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} \Leftrightarrow \hat{x} = \begin{bmatrix} x_0 \\ x_1 \end{bmatrix}$$


$$V = \begin{bmatrix} a & b & c \\ b & d & e \\ c & e & f \end{bmatrix} \Leftrightarrow \hat{V} = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$$

- Different rays have different axes
- How to define a space that describe each ray with x_0, x_1, x_2 ?

Camera Space vs. Ray Space





- In the camera 3D space, a point is described by (X, Y, Z) position with the camera coordinate frame
- Each point along a ray can be described as u_0, u_1, u_2
- When $u_2 = 1$, u_0, u_1 denotes the intersecting point on the camera plane



- In the ray space space, a ray is described by (x_0, x_1, x_2)
- x_0, x_1 denotes the pixel location
- x_2 describes the extent of the ray

Distorting the Camera 3D Space into Ray Space


Different rays have different levels of distortion



- In the camera 3D space, a point is described by (X, Y, Z) position with the camera coordinate frame
- Each point along a ray can be described as u_0, u_1, u_2
- When $u_2 = 1$, u_0, u_1 denotes the intersecting point on the camera plane

- In the ray space space, a ray is described by (x_0, x_1, x_2)
- x_0, x_1 denotes the pixel location
- x_2 describes the extent of the ray

We Can Compute 2D Gaussians on the Image to Obtain Integrals of 3D Gaussians in the Ray Space



- Integrate along x_2 axis with fixed x_0, x_1


$$\int_{\mathbb{R}} \mathcal{G}_V^3(x - p) dx_2 = \mathcal{G}_{\hat{V}}^2(\hat{x} - \hat{p})$$

3D Gaussian 2D Gaussian

- Integrate along x_2 axis in the Ray space: we have the same x_0, x_1 along the ray!
- Integrals of 3D Gaussians along the x_2 axis in the Ray Space are Equivalent to 2D Gaussians at x_0, x_1

However, 3D Gaussians are in the World Frame

- What is missing: describe 3D Gaussians from different 3D space




Gaussians are Closed Under Affine Transforms

$$\mathcal{G}_V(x - p) = \frac{1}{2\pi|V|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-p)^T \underbrace{V^{-1}}_{\substack{\text{3D covariance} \\ \downarrow}} (x-p)}$$

- With affine mapping $\phi = Mx + p$ of coordinates (such as cam2world matrix)


$$\mathcal{G}_V(\phi^{-1}(u) - p) = \frac{1}{|M^{-1}|} \mathcal{G}_{MVM^T}(u - \phi(p))$$



Step 1: Transform 3D Gaussians from the World Coordinates to the Camera Coordinates


- Cam-to-World transformation is affine mapping $\phi = W\mathbf{x} + \mathbf{p}$:

$$\mathcal{G}_{V_k''}(\phi^{-1}(\mathbf{u}) - t_k) = \frac{1}{|W^{-1}|} \mathcal{G}_{V_k'}(\mathbf{u} - \mathbf{u}_k)$$

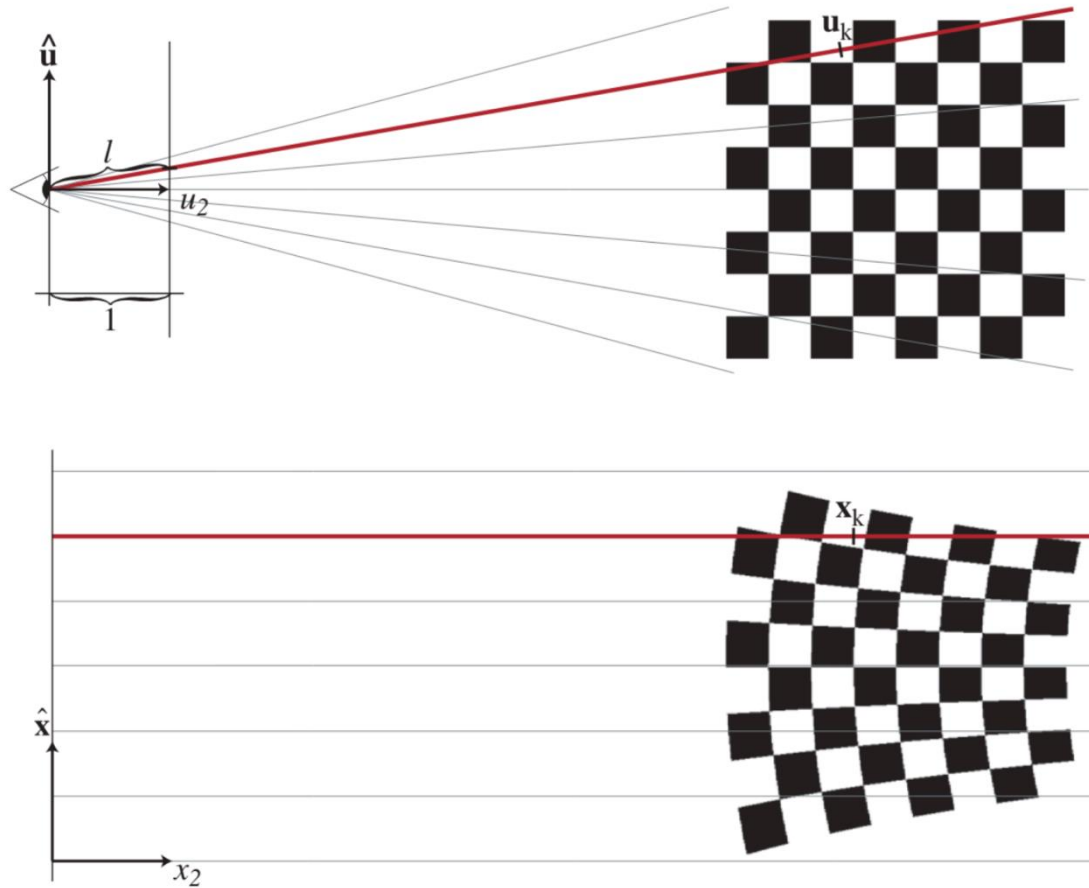


However, 3D Gaussians are in the World Frame

$$\begin{aligned} & \mathcal{G}_{V_k''}(\phi^{-1}(\mathbf{u}) - t_k) \\ &= \frac{1}{|W^{-1}|} \mathcal{G}_{V_k'}(\mathbf{u} - \mathbf{u}_k) \end{aligned}$$



Step 2: Convert 3D Camera Space to Ray Space



- Ray space $\mathbf{x} = \begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix}$: Given a center of projection and a projection plane, x_0, x_1 specifies the intersecting points of a ray on the plane; x_2 specifies the Euclidean distance from the center of projection to a point on the viewing ray
- Camera space $\mathbf{u} = \begin{pmatrix} u_0 \\ u_1 \\ u_2 \end{pmatrix}$: 3D points along the same ray in the camera coordinate frame
- The mapping of ray space and camera space is:

$$\begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix} = m(\mathbf{u}) = \begin{pmatrix} u_0/u_2 \\ u_1/u_2 \\ \|\mathbf{u}\| \end{pmatrix}$$

Approximate the Mapping Function with Jacobian $J_{\mathbf{u}_k}$

- Projection is not affine mapping

$$\begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix} = m(\mathbf{u}) = \begin{pmatrix} u_0/u_2 \\ u_1/u_2 \\ \|\mathbf{u}\| \end{pmatrix}$$

- We can approximate with 1st order Taylor Expansion at mean \mathbf{u}_k and \mathbf{x}_k of two spaces as:

$$m_{\mathbf{u}_k}(\mathbf{u}) = \mathbf{x}_k + J_{\mathbf{u}_k} \cdot (\mathbf{u} - \mathbf{u}_k), \quad \text{where } J_{\mathbf{u}_k} = \frac{\partial m}{\partial \mathbf{u}}(\mathbf{u}_k)$$

\downarrow
 $m(\mathbf{u}_k)$

Approximate the Mapping Function with Jacobian $J_{\mathbf{u}_k}$

- We can approximate with 1st order Taylor Expansion at mean \mathbf{u}_k and \mathbf{x}_k of two spaces as:

$$m_{\mathbf{u}_k}(\mathbf{u}) = \mathbf{x}_k + J_{\mathbf{u}_k} \cdot (\mathbf{u} - \mathbf{u}_k), \quad \text{where } J_{\mathbf{u}_k} = \frac{\partial m}{\partial \mathbf{u}}(\mathbf{u}_k)$$

- 3D Gaussians in the Ray space:

$$\frac{1}{|W^{-1}| |J^{-1}|} \mathcal{G}_{V_k}(\mathbf{x} - \mathbf{x}_k)$$
$$V_k = J V_k' J^\top = J W V_k'' W^\top J^\top$$

V_k'' : covariance matrix in the 3D world space

V_k' : covariance matrix in the 3D camera space

V_k : covariance matrix in the 3D ray space

W : world to camera transformation matrix

J : camera space to ray space Jacobian matrix

Putting Everything Together

- We can approximate with 1st order Taylor Expansion at mean \mathbf{u}_k and \mathbf{x}_k of two spaces as:

$$m_{\mathbf{u}_k}(\mathbf{u}) = \mathbf{x}_k + J_{\mathbf{u}_k} \cdot (\mathbf{u} - \mathbf{u}_k), \quad \text{where } J_{\mathbf{u}_k} = \frac{\partial m}{\partial \mathbf{u}}(\mathbf{u}_k)$$

- 3D Gaussians in the Ray space:

$$\frac{1}{|W^{-1}| |J^{-1}|} \mathcal{G}_{V_k}(\mathbf{x} - \mathbf{x}_k)$$
$$V_k = J V_k' J^\top = J W V_k'' W^\top J^\top$$

- Integrate along rays:

$$q_k(\hat{\mathbf{x}}) = \int_{\mathbb{R}} \frac{1}{|W^{-1}| |J^{-1}|} \mathcal{G}_{V_k}(\hat{\mathbf{x}} - \hat{\mathbf{x}}_k, x_2 - x_{k2}) dx_2 = \frac{1}{|W^{-1}| |J^{-1}|} \mathcal{G}_{\hat{V}_k}(\hat{\mathbf{x}} - \hat{\mathbf{x}}_k)$$


Putting Everything Together

For each 3D Gaussian k :


- Compute \mathbf{x}_k and \mathbf{u}_k
- Compute Jacobian $J_{\mathbf{u}_k} = \frac{\partial m}{\partial \mathbf{u}}(\mathbf{u}_k)$
- Compute new covariance matrix $V_k = J W V_k'' W^T J^T$
- Obtain 2D Gaussian $\frac{1}{|W^{-1}| |J^{-1}|} \mathcal{G}_{\hat{V}_k}(\hat{\mathbf{x}} - \hat{\mathbf{x}}_k)$

For $\mathbf{x}_0, \mathbf{x}_1$ in all pixel locations:

- Find the set of 3D Gaussian intersected with the ray
- Sort the 3D Gaussians by depth
- For each 3D Gaussian k :
 - Retrieve the corresponding 2D Gaussian
 - Compute weightings $\rho_k(\mathbf{x}_0, \mathbf{x}_1)$ based on the 2D Gaussian
 - Compute the opacity $\rho_k(\mathbf{x}_0, \mathbf{x}_1) \sigma_k$



However, 3D Gaussians are in the World Frame



How to Obtain Jacobian?

- Jacobian $J_{\mathbf{u}_k}$ is given by the derivative of m at the point \mathbf{u}_k :

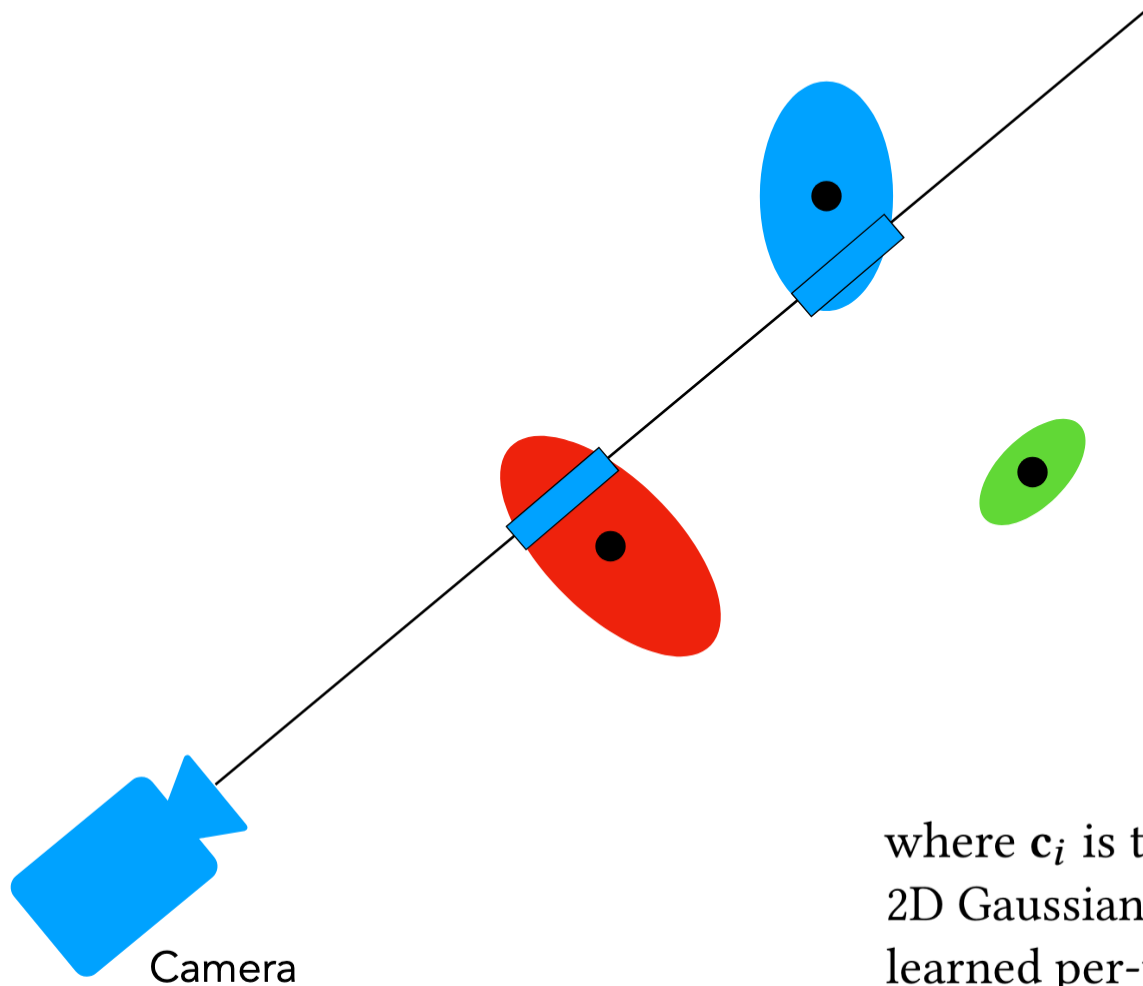
$$J_{\mathbf{u}_k} = \frac{\partial m}{\partial \mathbf{u}}(\mathbf{u}_k) = \begin{bmatrix} \frac{1}{u_{k,2}} & 0 & -\frac{u_{k,0}}{u_{k,2}^2} \\ 0 & \frac{1}{u_{k,2}} & -\frac{u_{k,1}}{u_{k,2}^2} \\ \frac{u_{k,0}}{\|\mathbf{u}_k\|} & \frac{u_{k,1}}{\|\mathbf{u}_k\|} & \frac{u_{k,2}}{\|\mathbf{u}_k\|} \end{bmatrix}$$

How to Obtain Jacobian?

- Jacobian $J_{\mathbf{u}_k}$ is given by the derivative of m at the point \mathbf{u}_k with focal length f_x, f_y :

$$J_{\mathbf{u}_k} = \frac{\partial m}{\partial \mathbf{u}}(\mathbf{u}_k) = \begin{bmatrix} \frac{1}{u_{k,2}} f_x & 0 & -\frac{u_{k,0}}{u_{k,2}^2} f_x \\ 0 & \frac{1}{u_{k,2}} f_y & -\frac{u_{k,1}}{u_{k,2}^2} f_y \\ \frac{u_{k,0}}{\|\mathbf{u}_k\|} & \frac{u_{k,1}}{\|\mathbf{u}_k\|} & \frac{u_{k,2}}{\|\mathbf{u}_k\|} \end{bmatrix}$$

Can compute volume rendering integral without ever sampling a single 3D point in space!



For each ray, ignore the non-overlapping 3D Gaussians, and sort the rest by ascending depth. Render with the known formulation:

$$C = \sum_{i \in \mathcal{N}} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (3)$$

where c_i is the color of each point and α_i is given by evaluating a 2D Gaussian with covariance Σ [Yifan et al. 2019] multiplied with a learned per-point opacity.

Problem 1: Local Minima


Initialize the mean of each 3D Gaussian with a random position



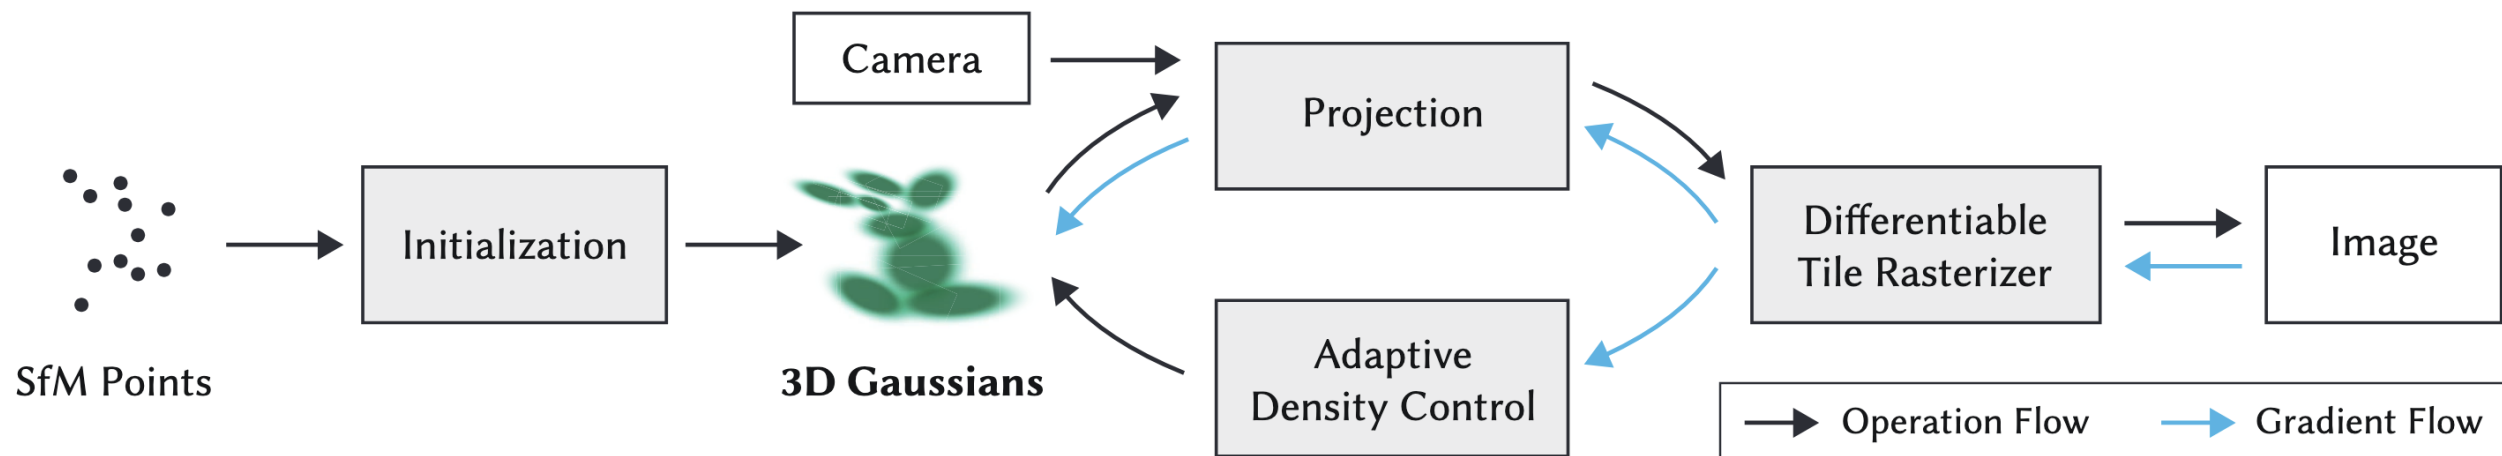
Initialize the mean of each 3D Gaussian with a 3D position reconstructed by Structure from Motion



Problem 2: Unevenly Distributed 3D Gaussians



The Overall Framework of 3D Gaussian Splatting



Render Scenes with 3D Gaussian Splatting



Render Scenes with 3D Gaussian Splatting



3D Gaussians are Explicit, Facilitating Tracking




3D Gaussians are Explicit, Facilitating Tracking




3D Gaussians are Explicit, Easy to Compose



3D Gaussians are Explicit, Facilitating Physical Modeling



Learn More About Rendering



What We Will Cover Next Week

- Modeling in 4D