

Lecture 19:


Volumes and Points

Computer Graphics: Rendering, Geometry, and Image Manipulation
Stanford CS248A, Winter 2024


Today's subject

- **Rendering from geometry representations that are not meshes**
 - **Volume rendering**
 - **Point rendering**
- **And their implications to modern progress in scene capture**

Volumetric effects



Absorption in a volume



$$\mathbf{p} = (x, y, z)$$


$$\boldsymbol{\omega} = (\phi, \theta)$$

$$dL(p, \omega) = -\sigma_a(p) L(p, \omega) ds$$

$$\frac{dL(p, \omega)}{ds} = -\sigma_a(p) L(p, \omega)$$

- $L(p, \omega)$ radiance along a ray from p in direction ω
- Absorption cross section at point in space: $\sigma_a(p)$
 - Probability of being absorbed per unit length
 - Units: 1/distance

Absorption in a volume



$$\mathbf{p} = (x, y, z)$$

$$\boldsymbol{\omega} = (\phi, \theta)$$

$$\frac{dL(p, \omega)}{L(p, \omega)} = -\sigma_a(p) ds$$


$$L(\mathbf{p} + s\boldsymbol{\omega}, \omega) = e^{-\int_0^s \sigma_a(\mathbf{p} + s'\boldsymbol{\omega}) ds'} L(\mathbf{p}, \omega) = T(s) L(\mathbf{p}, \omega)$$

Transmittance: $T(s) = e^{-\int_0^s \sigma_a(\mathbf{p} + s'\boldsymbol{\omega}, \omega) ds'}$


Absorption: lower density



Absorption: higher density



Out scattering



$$dL(p, \omega) = -\sigma_s(p) L(p, \omega) ds$$

- **Scattering cross section at point in space: σ_s**
 - **Probability of being scattered per unit length**
 - **Units: 1/distance**

Absorption and out scattering diminish radiance

Total cross section:

$$\sigma_t = \sigma_a + \sigma_s$$

$$dL(p, \omega) = -\sigma_t(p) L(p, \omega) ds$$

$$L(p + s\omega, \omega) = T(s) L(p, \omega)$$

Where total transmittance is:

$$T(s) = e^{-\int_0^s \sigma_t(p + s'\omega) ds'} = e^{-\tau(s)}$$

$$\tau(s) = \int_0^s \sigma_t(p + s'\omega) ds'$$

“Optical distance” (from absorption and scattering)

Ray marching to compute transmittance

Step through volume in small steps

Given "camera ray" from point \mathbf{o} in direction ω ...

$$\mathbf{r}(t) = \mathbf{o} + t\omega$$


And volume with density

$$\sigma(\mathbf{p})$$

Estimate optical thickness as:


$$\tau(s) \approx \frac{s}{N} \sum_i^N \sigma_t(\mathbf{p}_i)$$

$$\mathbf{p}_i = \mathbf{o} + \frac{i + 0.5}{N} \omega$$



In scattering

- Light going in other directions also scatters into the direction ω
- In scattering increases radiance along ω



$$S(p, \omega) = \sigma_s(p) \int_{S^2} p(\omega' \rightarrow \omega) L(p, \omega') d\omega'$$

Phase function: $p(\omega' \rightarrow \omega)$


Energy conservation: $\int_{S^2} p(\omega' \rightarrow \omega) d\omega' = 1$

Direct illumination in a volume

- Can treat like direct illumination on a surface
- e.g., sample light sources (or from phase function's distribution)

$$S_d(p', \omega) = \sigma_s(p') \int_{S^2} p(\omega' \rightarrow \omega) L_d(p', \omega') d\omega'$$

- But computing direct lighting can be expensive
- Why?
 - Hint: requires more than a shadow ray




Single scattering




Multiple scattering (not discussed today)

- Appearance of volume is not just due to a single scattering event (light directly from light source scattering in direction of eye)
- Light scatters many times in the volume before existing in the direction of the eye
- Advanced rendering topic: Monte Carlo estimate of multiple scattering events (“volume rendering equation”)




Single scattering



Multiple scattering

Let's ignore lighting for a moment

Consider representing a scene as a volume



Volume rendered CT scan

Volume density and "color" at all points in space.

$$\sigma(p)$$

$$c(p, \omega) = c(x, y, z, \phi, \theta)$$

The reflectance off surface at point p in direction ω



Volume rendered scene
(Mildenhall et al.)

Rendering volumes

Given “camera ray” from point \mathbf{o} in direction \mathbf{w} ...

$$\mathbf{r}(t) = \mathbf{o} + t\mathbf{w}$$


And volume with density and directional radiance.

$$\sigma(\mathbf{p})$$

← Volume density and color at all points in space.

$$c(\mathbf{p}, \omega)$$

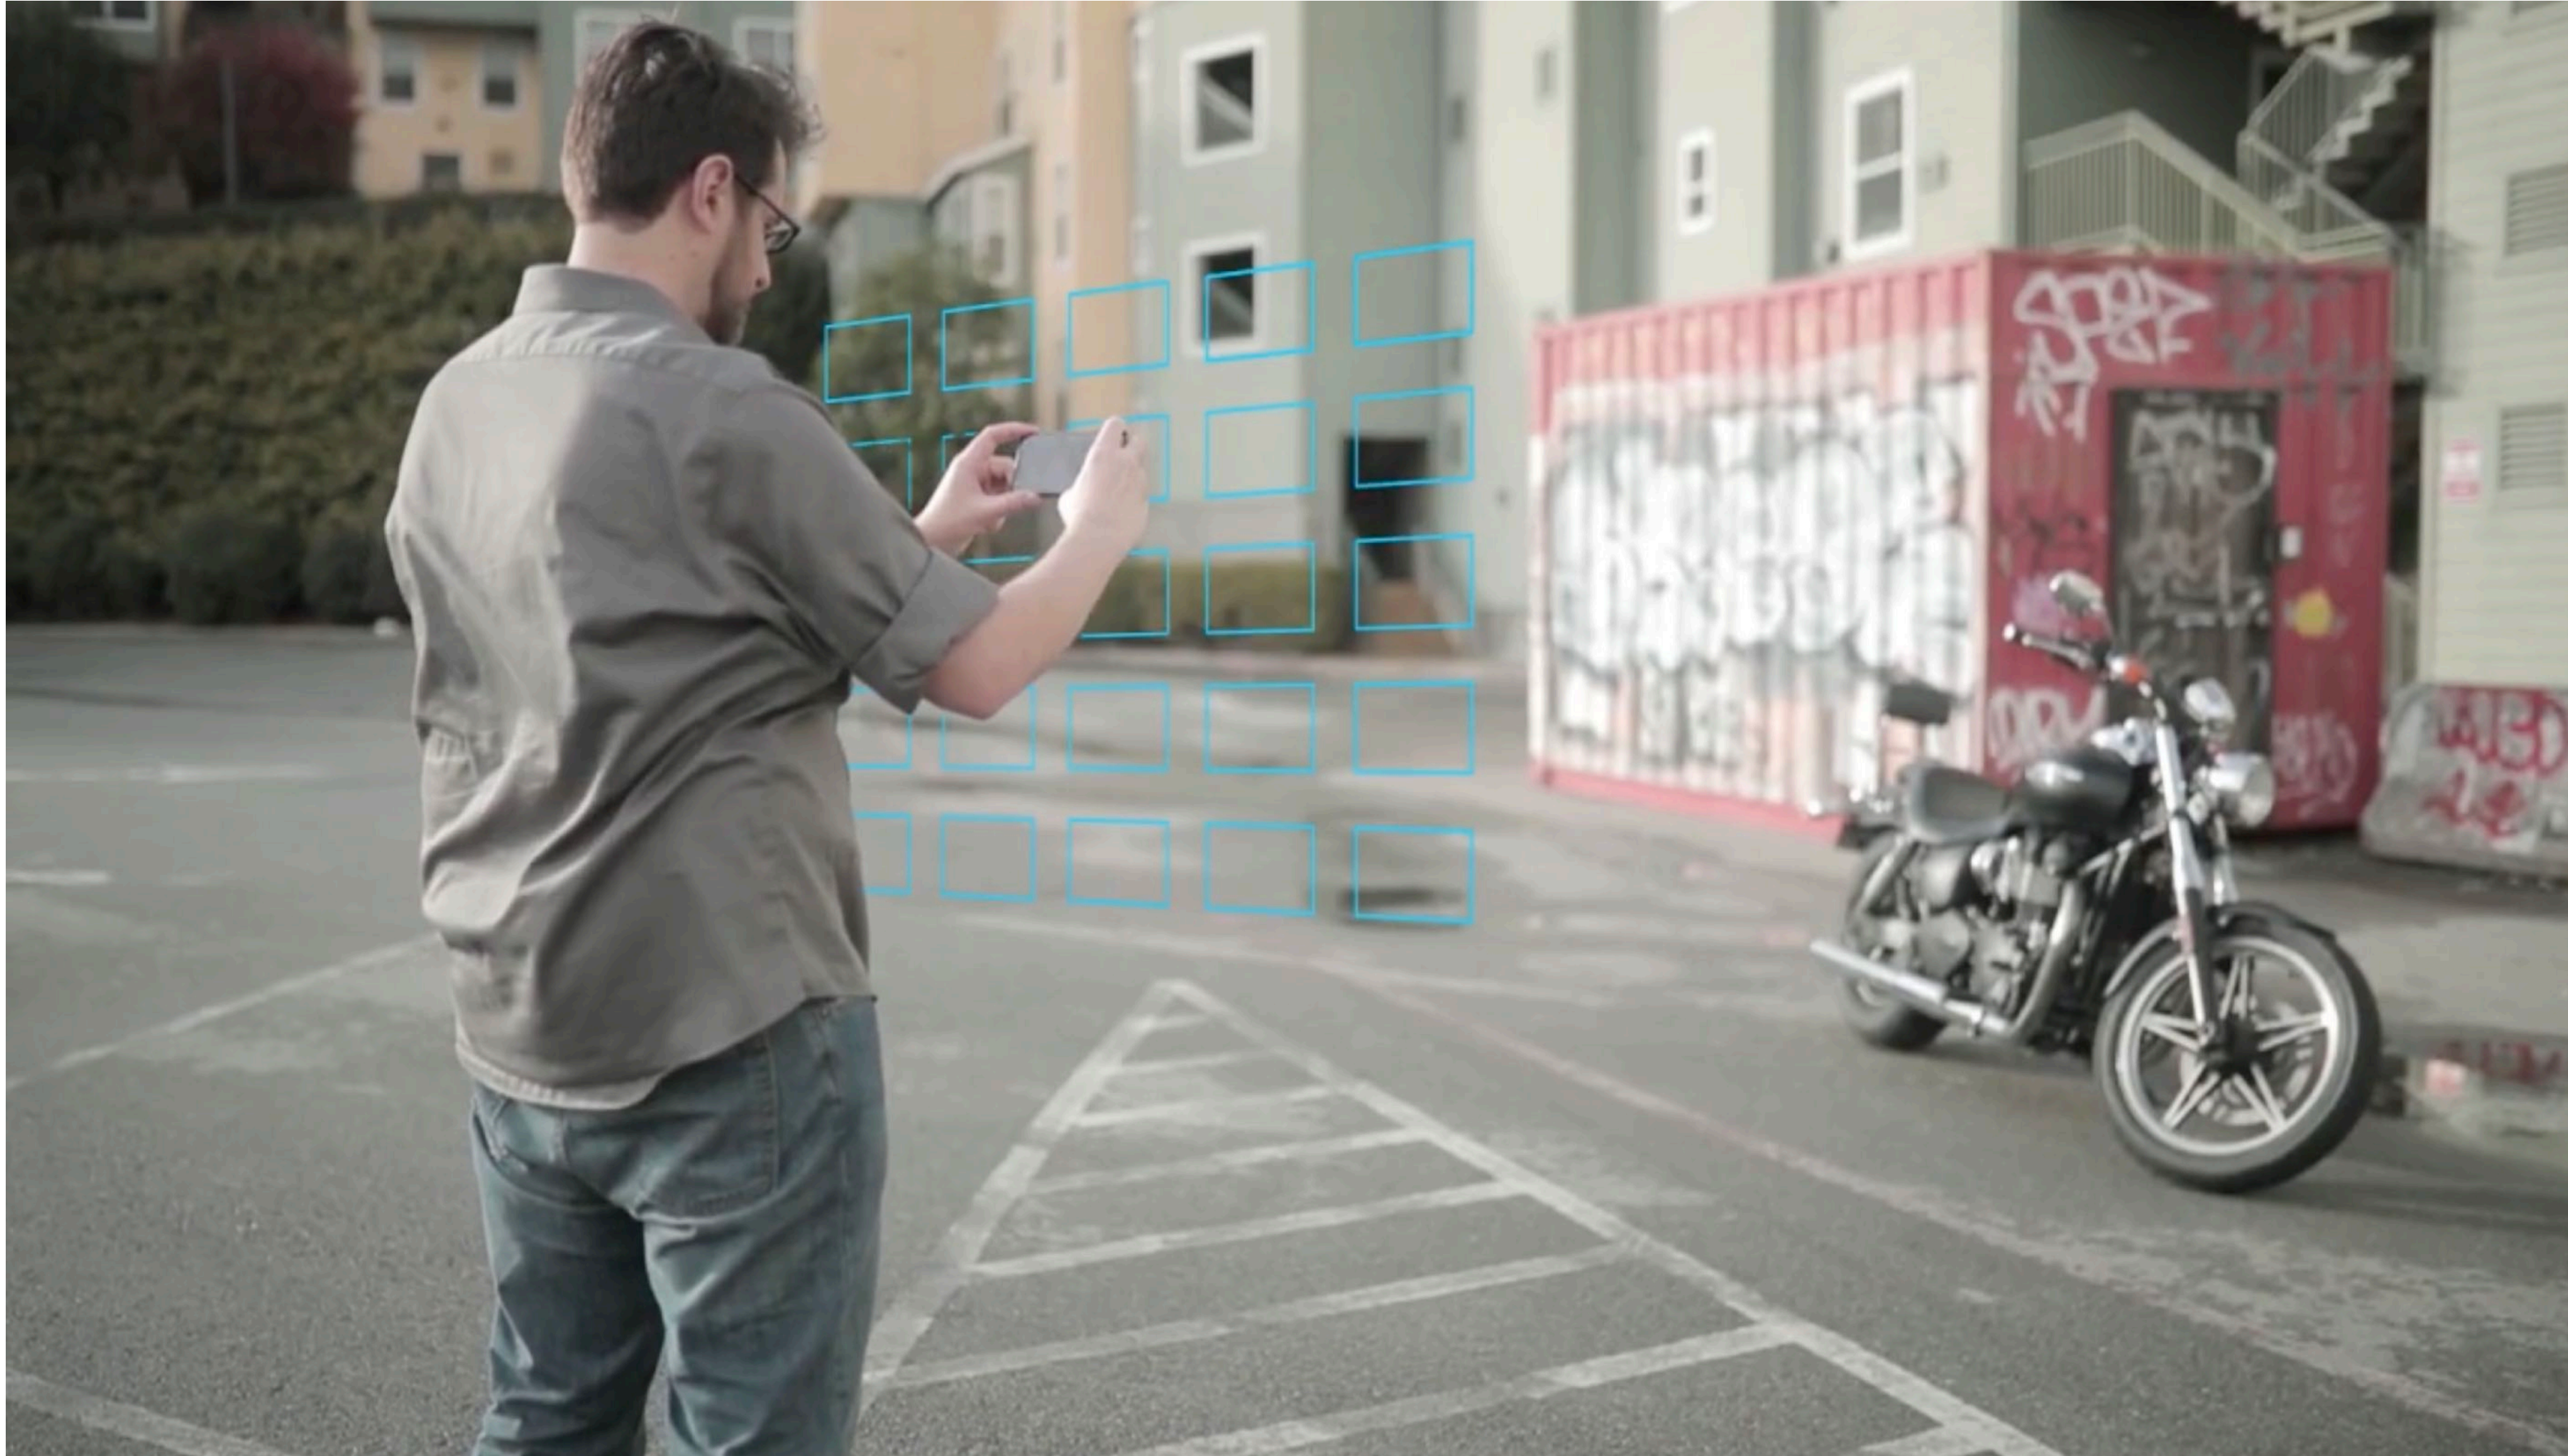
Step through the volume to compute radiance along the ray.



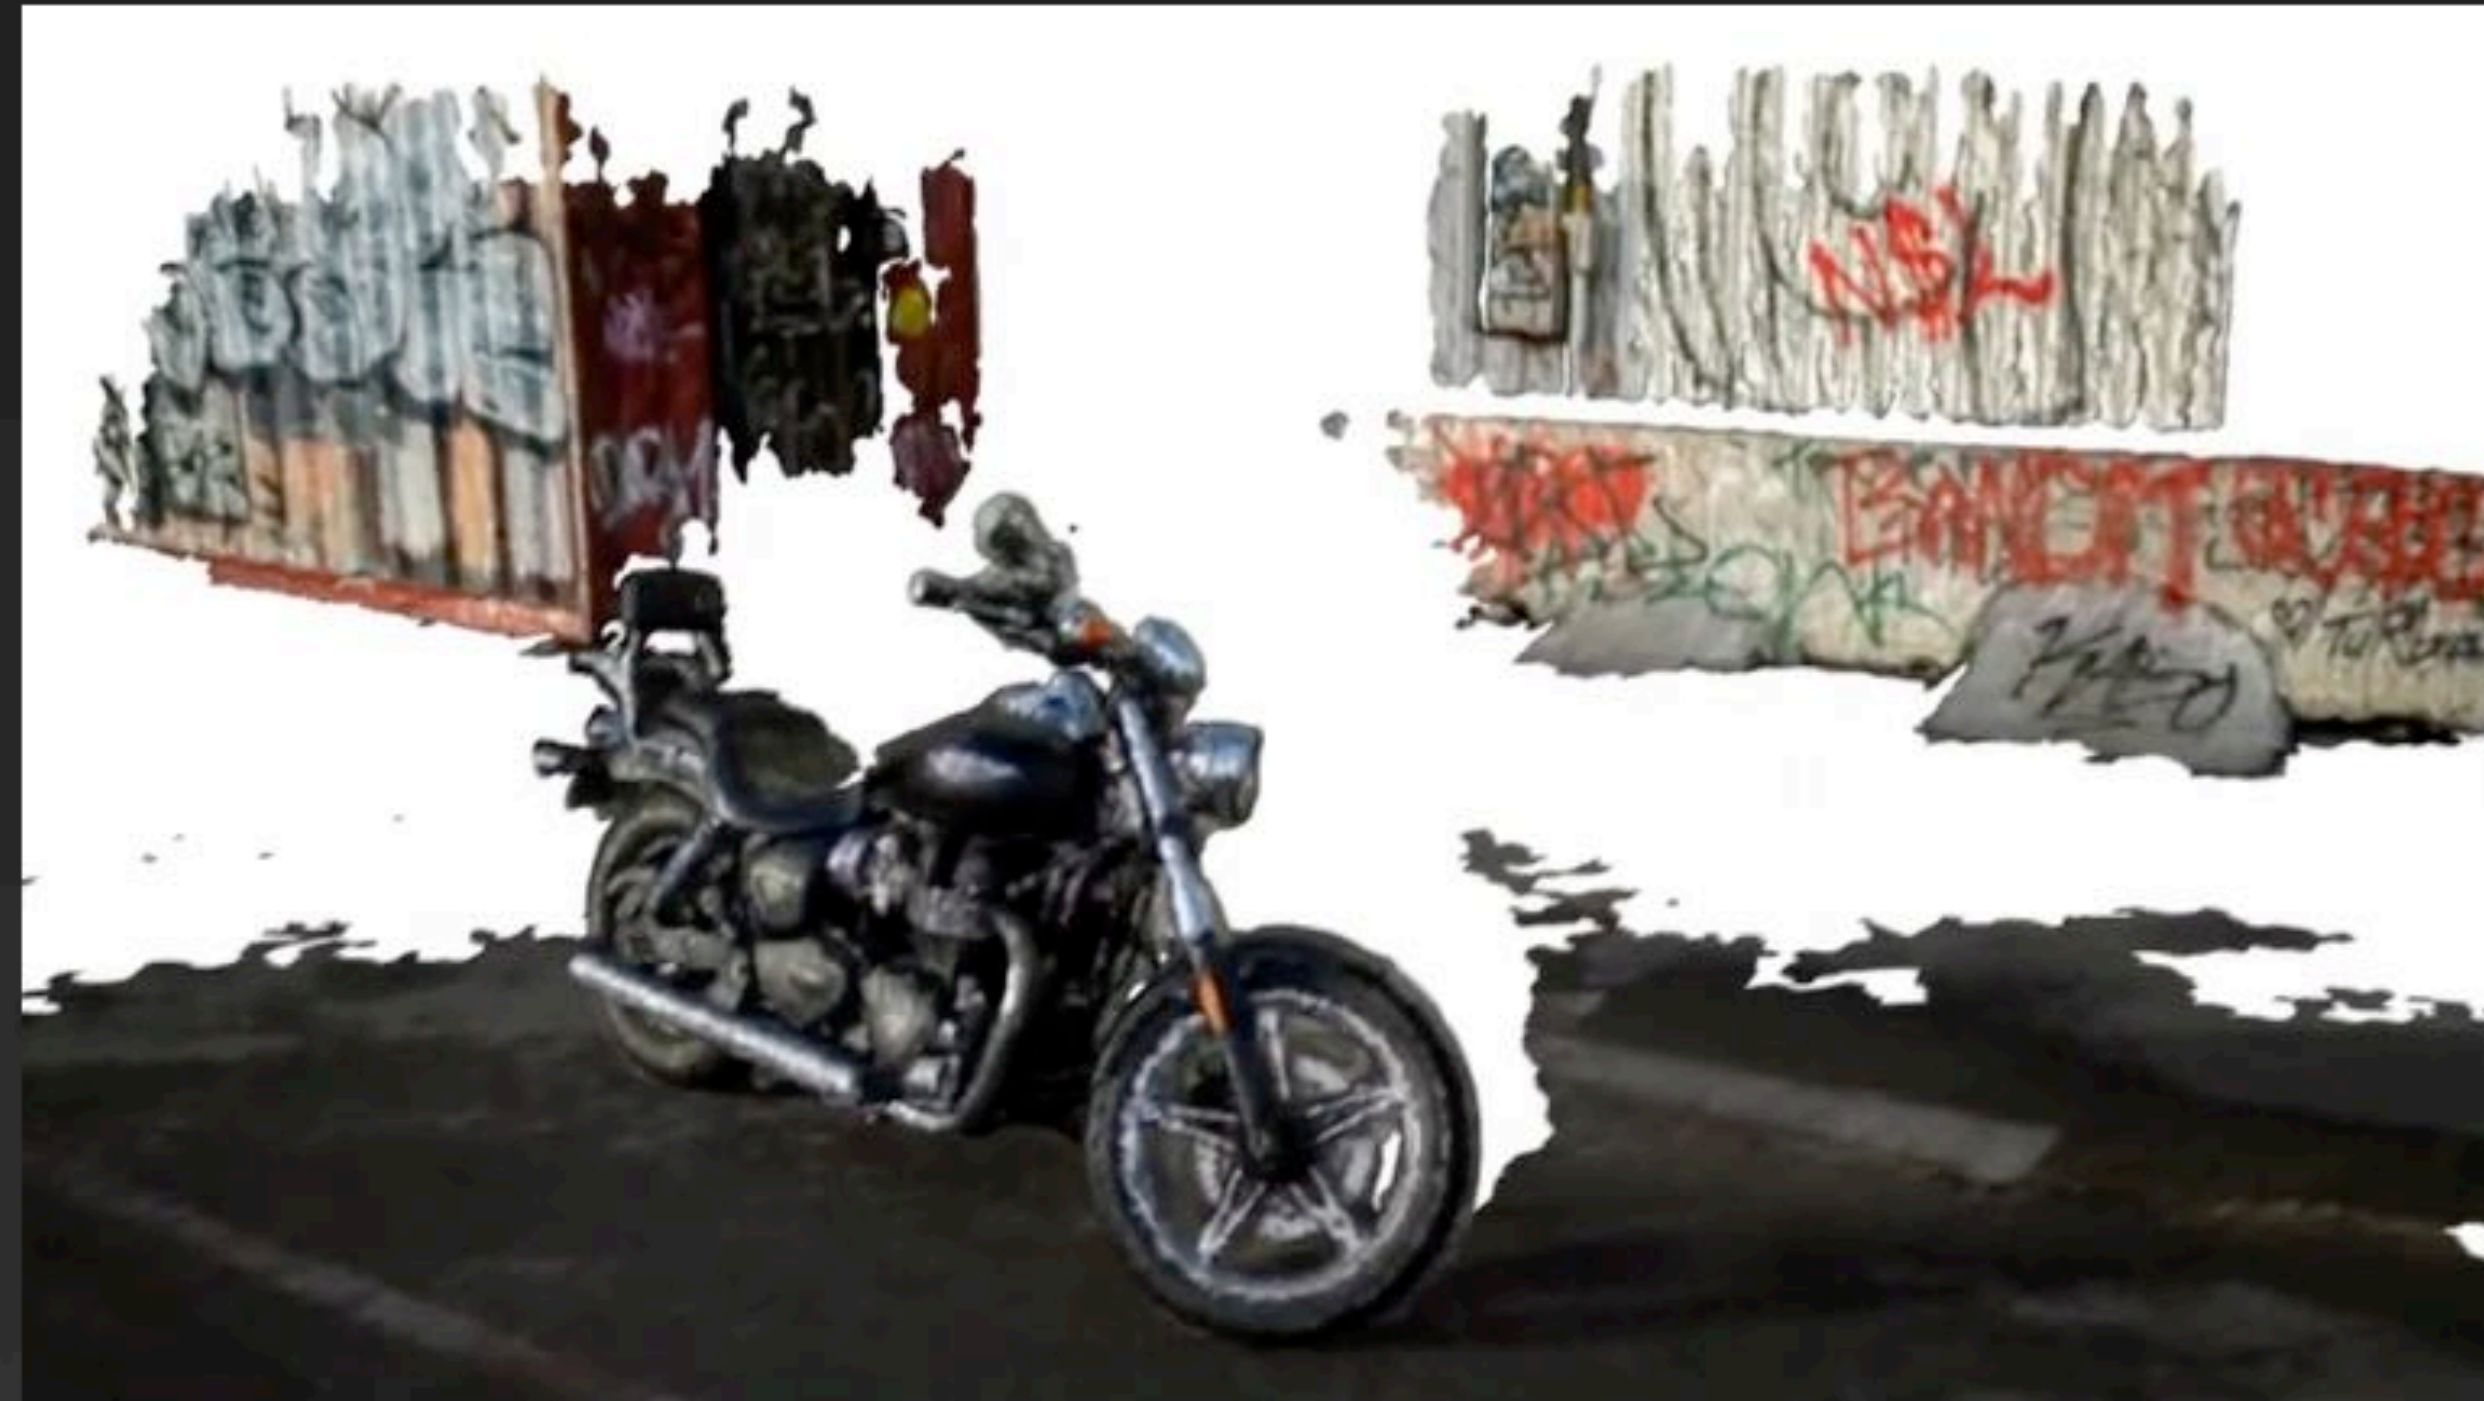
$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \quad \text{where } T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right)$$

An interesting task

- Given a collection of photographs (from known camera viewpoints)
- Compute a 3D reconstruction of the scene (surface locations + color at each point on surface)



Estimating mesh geometry is tricky



Reconstructed Mesh


Re-interest in using volume rendering (circa 2018)

- Let's just drop this triangle-based representation entirely, it's much simpler (and more versatile when it's unclear what the geometry is anyway) to emit a single volumetric representation



Regular 3D grid representation of a volume

- Dense 3D grid
 - $V[i,j,k] = \text{rgba}$
- Note, this representation treats surface as diffuse, since: $c(p, \omega) = c(p)$
- Would need $\sigma[i,j,k]$ and $c[i,j,k,\text{phi},\text{theta}]$ to represent directional distribution of color

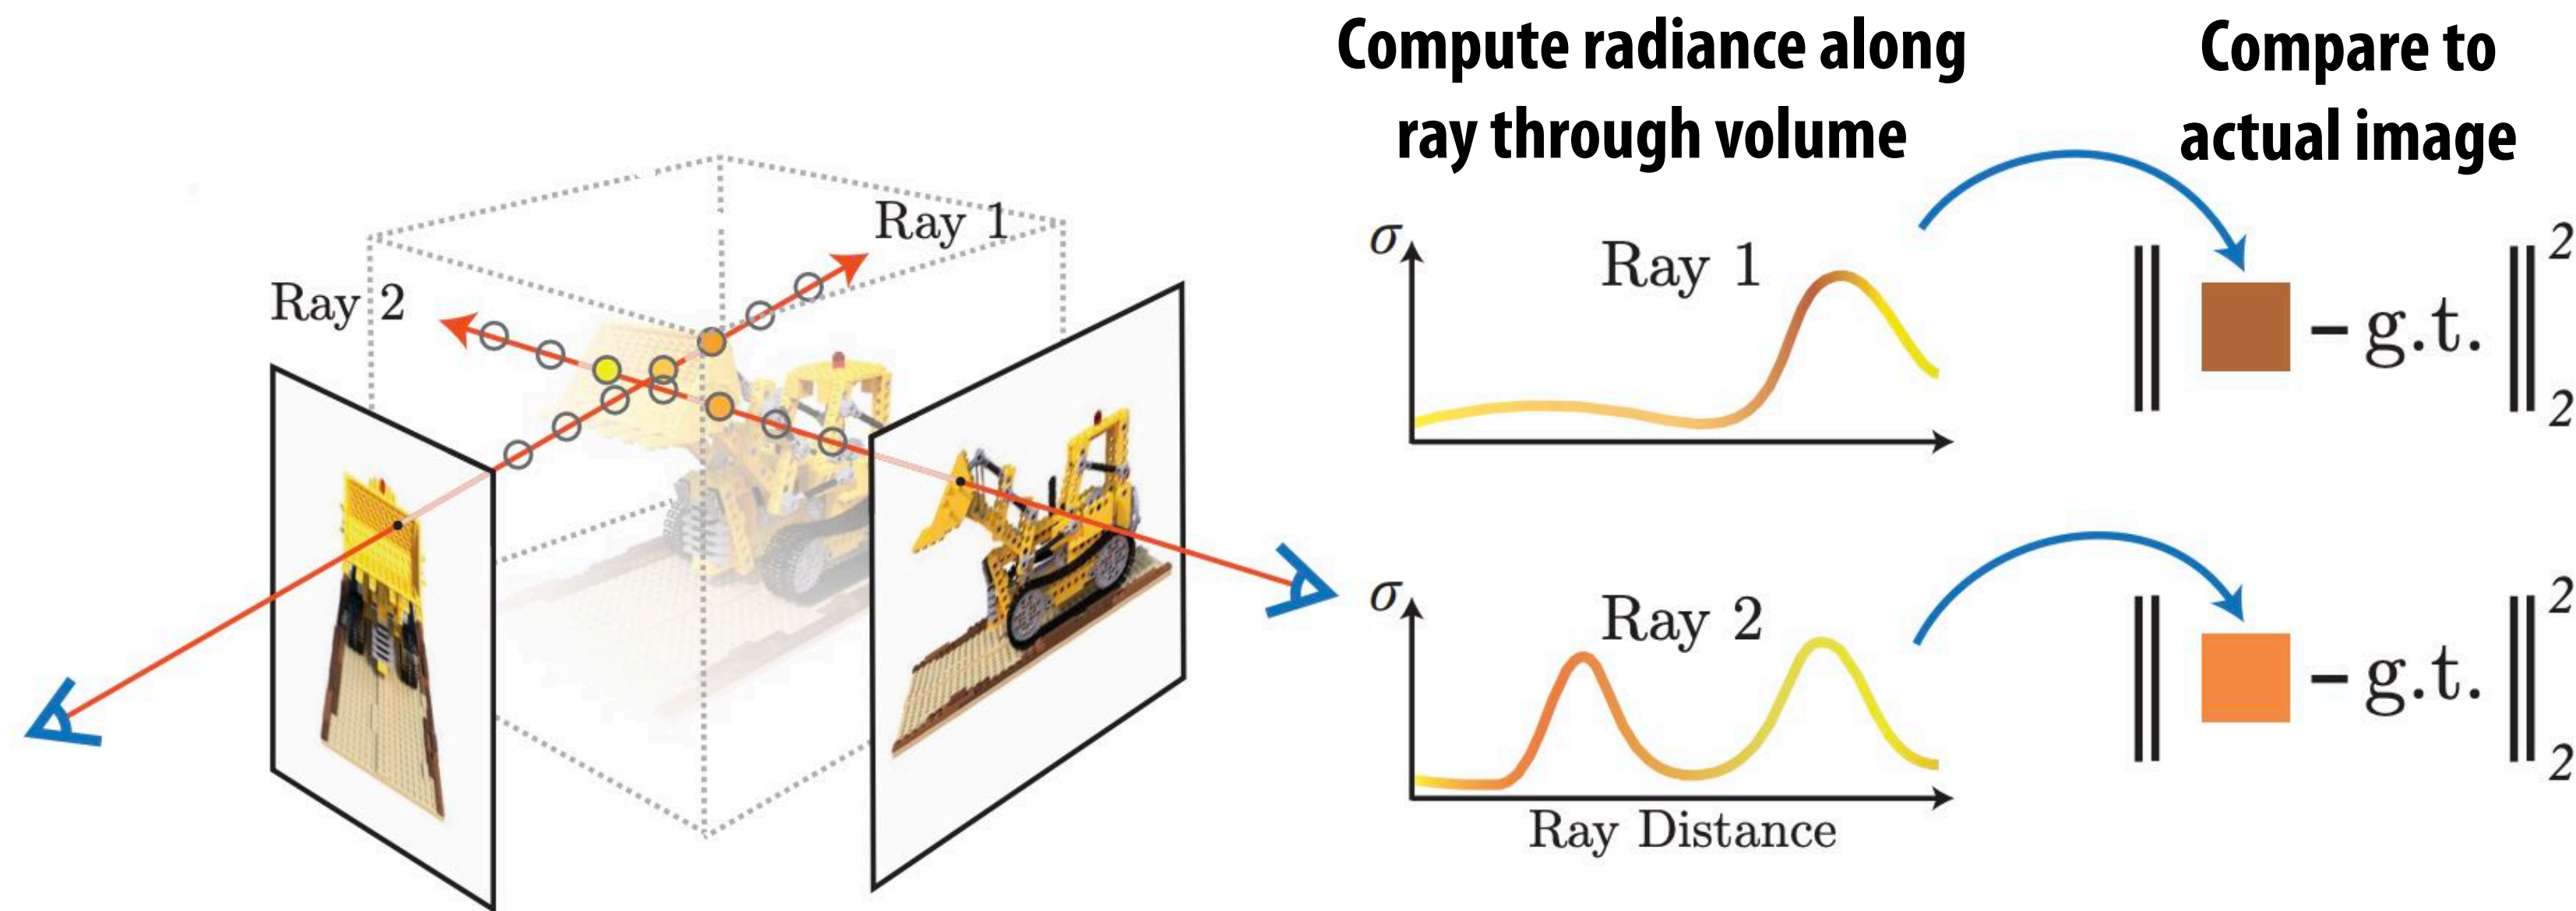



Optimizing volumes

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \text{ where } T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right)$$

Idea: optimize volume values (opacity and color) so that $C(\mathbf{r})$ matches that of photos.

For many rays.... trace through volume... see if the result matches the photo... use error to update volume opacity/color values



Regular 3D grid representation of a volume

Consider storage requirements:

1024^3 cells

Ignore directional dependency: rgb 4 bytes/cell

~ 4GB

Now consider directional dependency of color
on (ϕ, θ) ... much worse storage cost




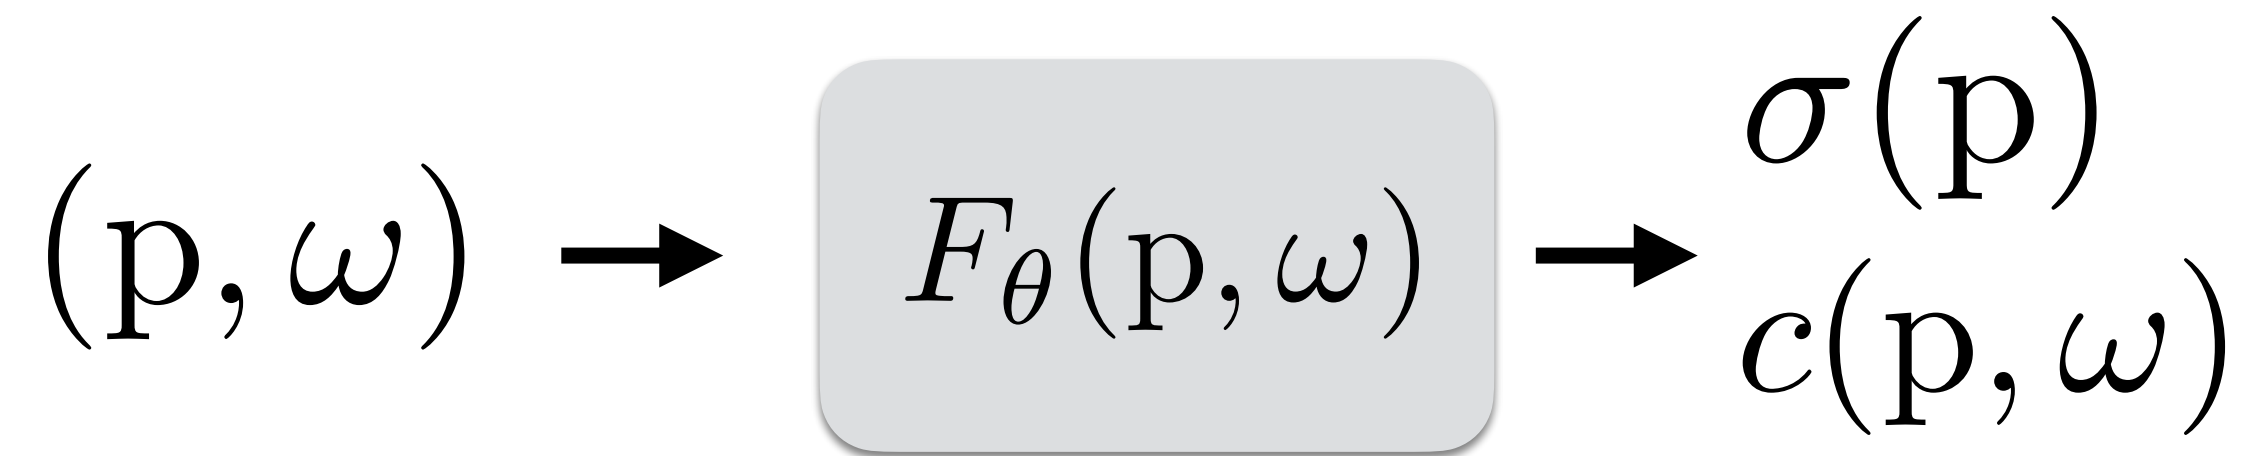
Typical challenge:
limited resolution





Credit: Voxel Ville NFT (voxelville.io)

Learning (compressed) representations

Why not just learn an approximation to the continuous function that matches observations from different viewpoints?



Learning neural radiance fields (NeRF)




Key idea: differentiable volume renderer to compute $dC/d(\text{color})d(\text{opacity})$

Great visual results!



What just happened?

- **Continuous coordinate-based representation vs regular grid: DNN “learns” how to use its weights to produce high-resolution output where needed... given input data**
- **Compact representation: trades-off space for expensive rendering**
 - **Good: a few MBs = effectively very high-resolution dense grid**
 - **Bad: must evaluate DNN every step during ray marching**
 - **And the DNN is a “big” MLP (8-layer x 256)** ← MLP must do real work to associate weights with 5D locations
 - **Bad: must step densely (because we don't know where the surface is)**
- **Compact representation: optimization can learn to interpolate views despite complexity of volume density and radiance function**
 - **Only prior is the separation into positional σ and directional rgb**
 - **Training time: hours to a day to learn a good NeRF**



Improving rendering performance

- **Main idea: move to a different point in the compression-compute trade-off space**
- **Main ideas:**
 - **Avoid stepping densely through empty space (costly to evaluate the DNN to find density = 0)**
 - **Shrink the size of the DNN**
 - **Avoid evaluating the DNN when you can**


Recall quad-tree / octree

Quad-tree: nodes have 4 children (partitions 2D space)

Octree: nodes have 8 children (partitions 3D space)

Like uniform grid: easy to build (don't have to choose partition planes)



Has greater ability to adapt to location of scene geometry than uniform grid.



Simple two-level sparse quad tree



Quad-tree: nodes have 4 children (partitions 2D space)

Octree: nodes have 8 children (partitions 3D space)




Let's just run optimization for a bit...

- Optimization will push some opacity values to 0
- DNN tells us where the empty space is!
- Then convert dense opacity grid to an octree representation that's more efficient to render from...
- With the octree structure ***fixed***, we can continue to optimize color/density at leaves




Use the initial MLP to densely sample volume
(Find the empty space that's used to build the octree)



Note:
This implementation uses 2-level octree

What just happened?

- We performed initial training... a la original NeRF
- Once we get a sense of where the empty space is, we add a traditional acceleration structure to replace the “big” DNN. Can use a little DNN at leaves.
- That structure speeds up rendering (a lot), and also speeds up “fine tuning” training, since the initial “big” DNN need not be trained to convergence




Credit: Yu 2021

- Cost? Octree structure now 100's of MBs instead of a few MBs for MLP

Another idea: spherical harmonics

- Useful basis for representing directional information
- Analogy: cosine basis on the sphere



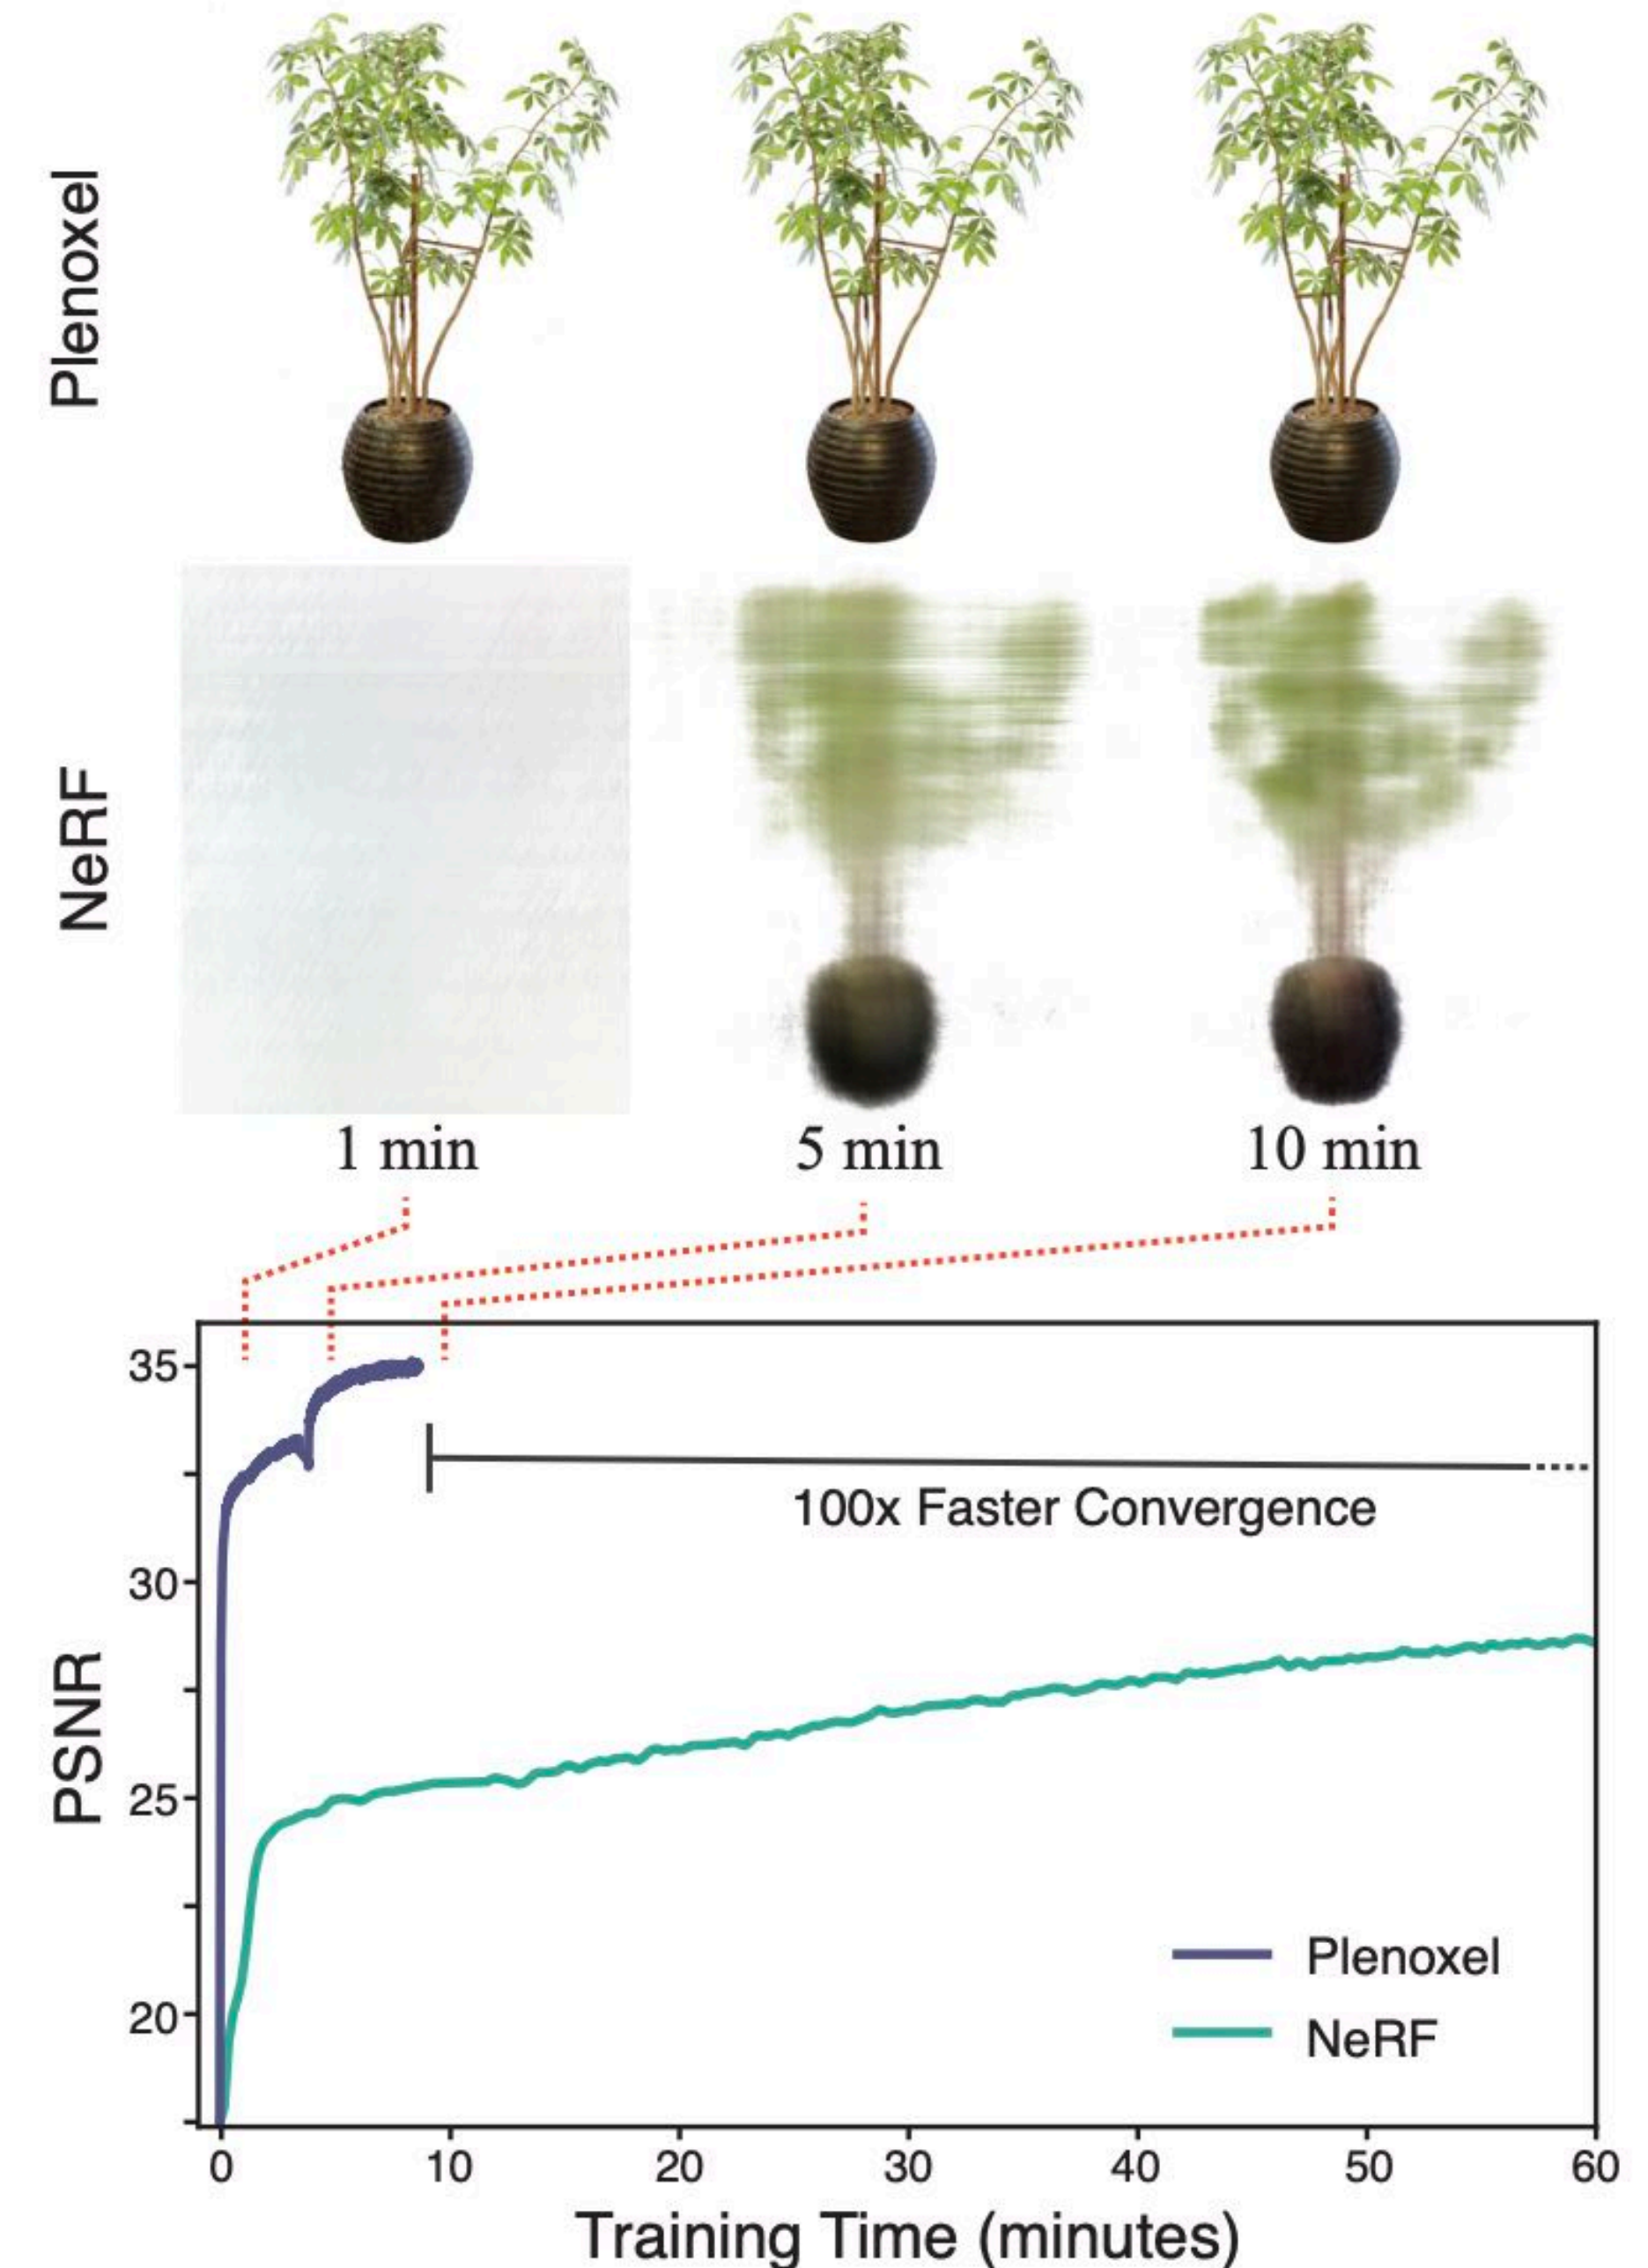
- Represent $c(p, \omega)$ compactly by projecting into basis of SH.

$$Y_l^m(\omega)$$

Finally...back to where we began

Plenoxels [CVPR 22]


- Start with a dense 3D grid of SH coefficients, optimize those coefficients at low resolution
- Now move to a sparse higher resolution representation (octree)
- Directly optimize for opacities and SH coefficients using differentiable volume rendering
- **No neural networks. Just optimizing the octree representation of baked SH lighting**
- **Takeaway: conventional computer graphics representations are efficient representations to learn/optimize on**



Summarizing it all: the “template”

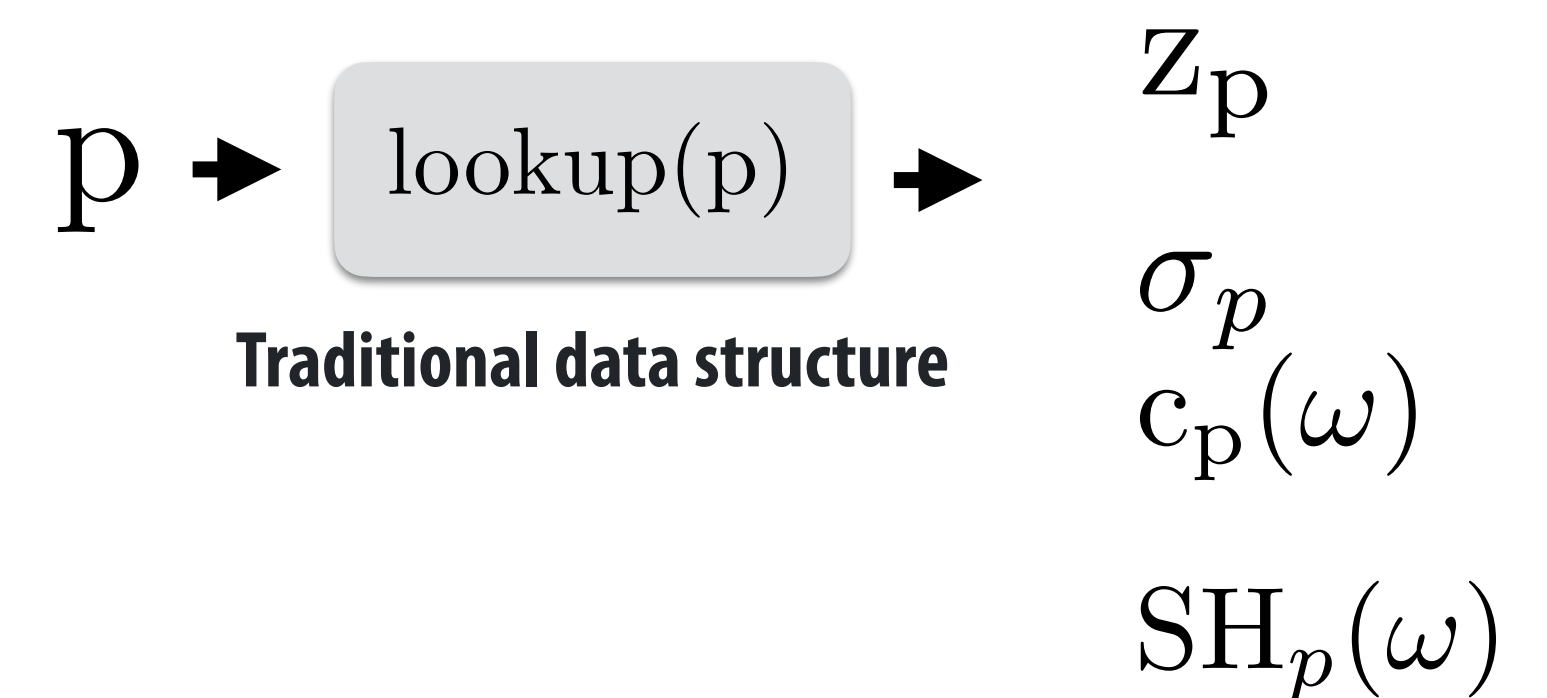
- **Train a DNN to gain understanding of 3D occupancy (where the surface is)**

- **Little to no geometric priors (so need position encoding tricks, etc)**



- **Then move to a traditional sparse encoding of occupancy (sparse volumetric structure)**

- **Now the “topology” of the irregular data structure is fixed**
- **Representation of surface/appearance/etc is stored at the nodes of this structure (spherical harmonics, neural code, etc.)**
- **Most of the heavy lifting is now performed by the data structure**




- **Continue optimization on the fixed, sparse representation**

- **Leverages differential volume rendering on sparse structure**
- **What we’re now learning is how to represent/compress the local details**

More recent innovations

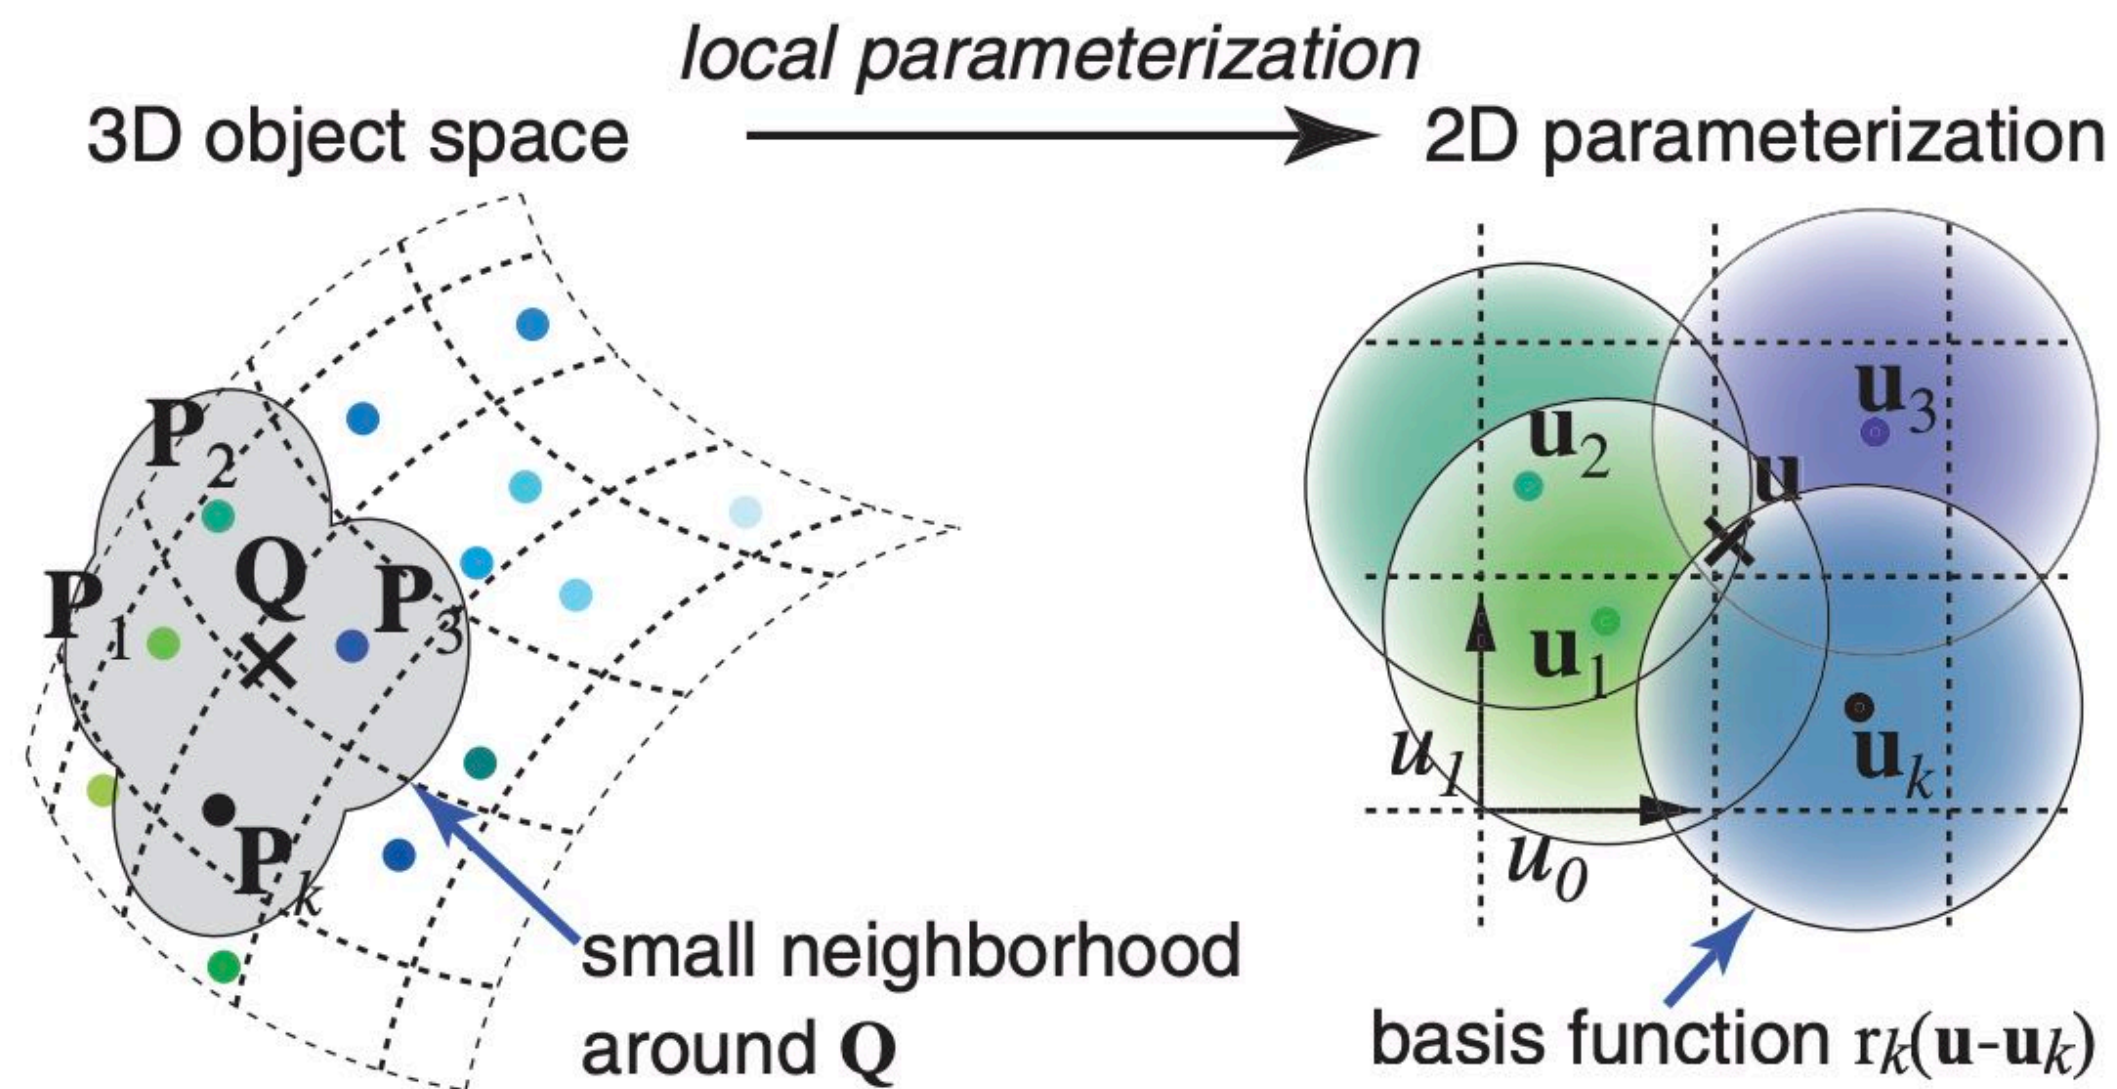
- **Best practices since 2022 replace this two-step process with a single optimization pass using a data structure called a “hash grid”**
- **See “instant neural graphics primitives” (NGP)**

Rendering point clouds

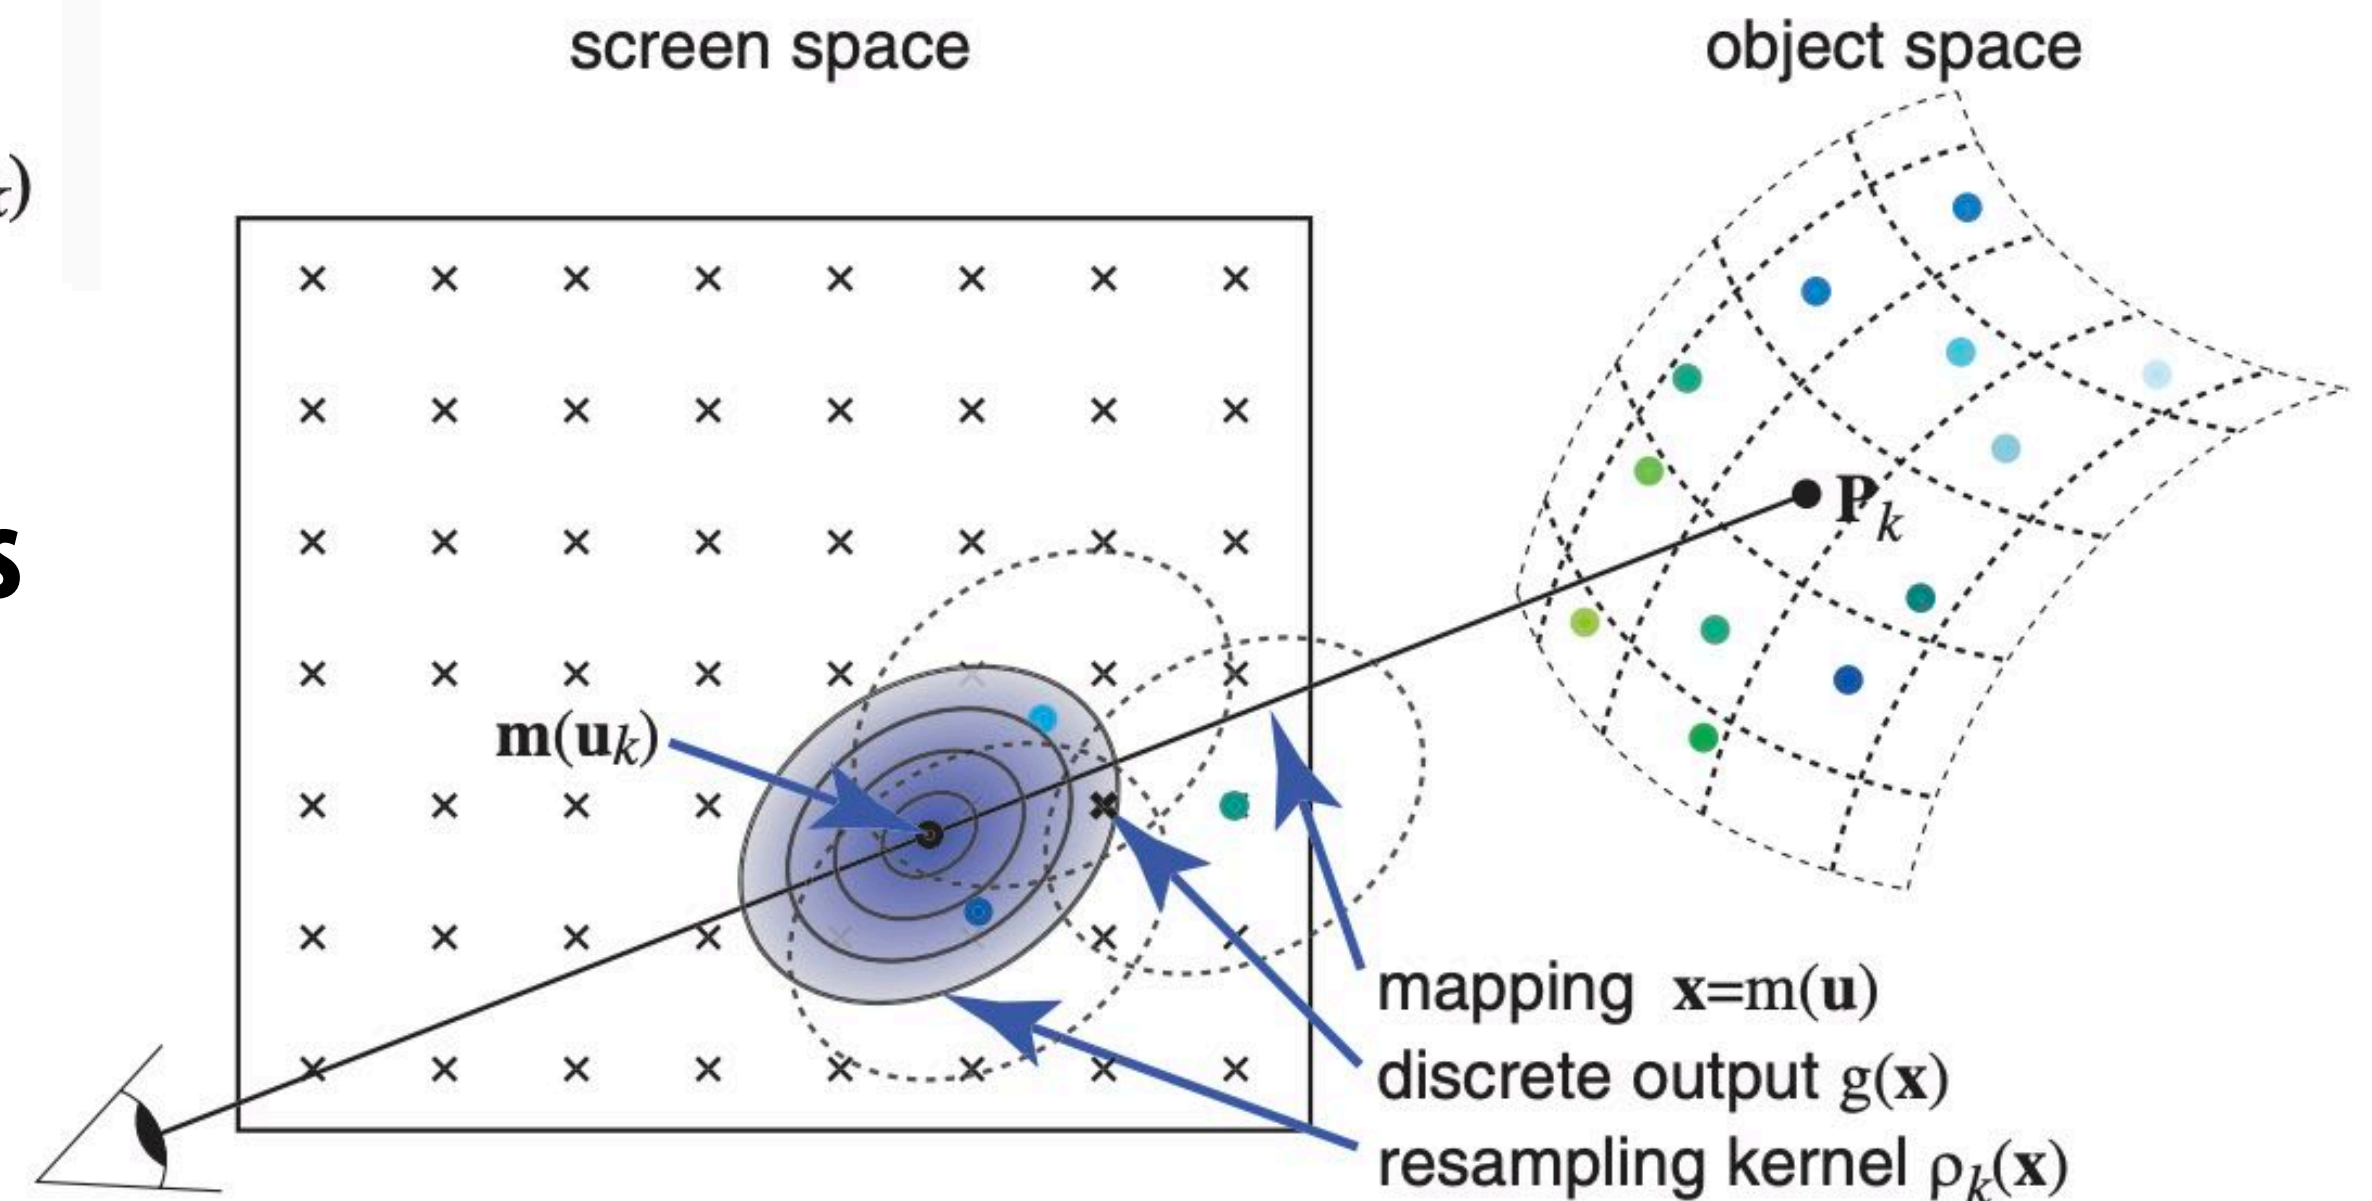


Anti-aliasing point clouds

- Treat surface as a collection of “Gaussian blobs” (convolve points with Gaussian filter)



- 3D Gaussians turn into oriented 2D gaussians when projected onto the 2D screen
- Can render the blobs back to front (requires alpha compositing)




“Point splatting”



Optimization to produce Gaussians, not voxels

- Earlier in lecture: optimization produces color and opacity at each voxel
- Now: same idea, but optimization chooses color, position, and radius of the Gaussians
 - Now: also need to decide on the number of Gaussians (a bit trickier)



Key idea: differentiable Gaussian splatting rendering to compute $dC/d(\text{color})d(\text{radius})d(\text{location})$

See “3D Gaussian Splatting for Real-Time Radiance Field Rendering” [Kerbl 2023]

Summary

- **Volumes (voxels) and Gaussian points as two alternative representations of geometry and materials**
- **Significant interest in modern times due to ease of writing differentiable renderers for these representations**
- **Heavy use in reconstructing scenes from (potentially sparse) set of photos**
 - **Surprising effectiveness of large-scale optimization**
- **Some of these solutions employ interesting combinations of neural structures (learned DNN weights) and “traditional” graphics primitives**
 - **Takeaway for graphics students in 2024: need to be a master of both!**

Summary

- **Thanks to Matt Pharr, Pat Hanrahan for materials in these slides**