NeRF Fundamentals

Discussion #9

Written by:
Justin, Chung Min

1 Alpha Compositing

Suppose we're trying to overlay 3 colored R G B circles with translucent opacity on a white background. The formula for overlaying (compositing) color c_a with opacity (alpha) α_a over an opaque c_b is: $c_{new} = c_a \cdot \alpha_a + c_b \cdot (1 - \alpha_a)$.

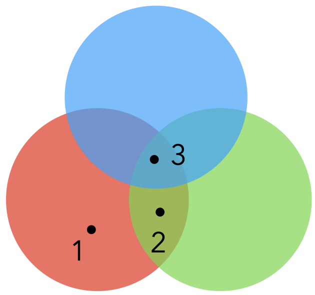

1.1: Enhance! Expand the alpha compositing formula to express the output pixel values at points 1, 2, and 3. Don't forget the background! Also don't worry about simplifying, just expand.

1.2: Simplifying Write your answer for point 3 with 4 terms for r, g, b, and white.

$C_3 = \alpha_b \cdot C_b \quad + \quad \_\_\_\_\_\_\_\_\_\_ \alpha_g \cdot C_g \quad + \quad \_\_\_\_\_\_\_\_\_\_\_\_\_\_\_ \alpha_r \cdot C_r \quad + \quad \_\_\_\_\_\_\_\_\_\_\_\_\_\_\_\_\_\_\_\_ C_w$

1.3: Transmittance! How does the form of this compositing formula look similar to transmittance from a NeRF? $T(n) = \prod_{k=1}^{n-1}(1 - \alpha_k)$.

Flatland NeRF

In this sheet, we consider a simplified NeRF in a flat plane. In this setup, **1)** Each camera has a 1D sensor array which measures color. **2)** All poses and rays exist only in 2 dimensions. **3)** The background is known to be white. **4)** Only output color c and density σ per-point (we ignore view direction).

2 Volume Rendering

Density in NeRF represents where geometry lies in a scene. During NeRF rendering there are three quantities we derive from density to render color from a ray: **alpha** (opacity), **transmittance**, and **rendering weights**. For simplification, the equations below *omit ray sample spacing* δ_k as it is constant.

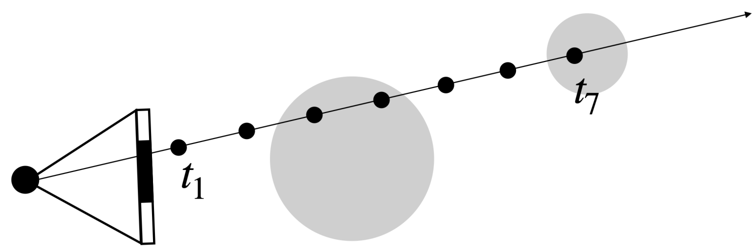

2.1: Rendering a cloud In the plots below, plot the values for the scenario above assuming the circles have low density values $\sigma = \ln 2$ ($\alpha \approx 0.5$).

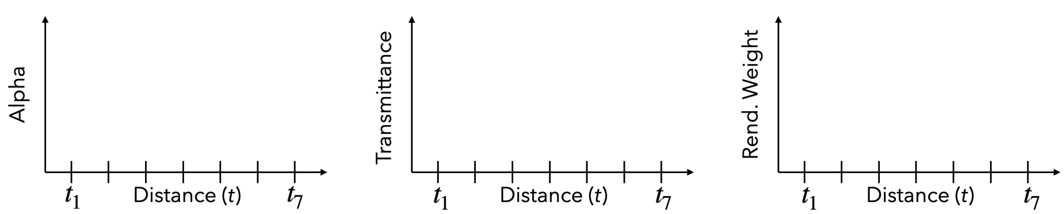

$\alpha(k) = 1 - \exp(-\sigma_k)$ $\quad$ $T(n) = \prod_{k=1}^{n-1} (1 - \alpha_k)$ $\quad$ $w(k) = T(k) \cdot \alpha(k)$

2.2: Rendering a solid In the plots below, plot the values for the scenario above assuming the circles have very high density values $\sigma = 1000$.

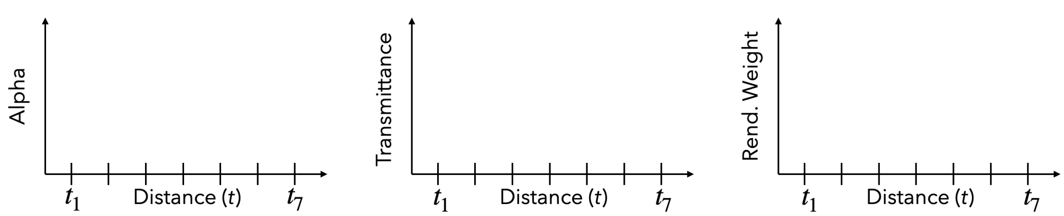

$\alpha(k) = 1 - \exp(-\sigma_k)$ $\quad$ $T(n) = \prod_{k=1}^{n-1} (1 - \alpha_k)$ $\quad$ $w(k) = T(k) \cdot \alpha(k)$

2.3: Analysis Explain which quantity accounts for occlusion in 3D, and how it accomplishes this.

3 Multiview in Flatland

The diagram below shows 2 cameras, each observing a circle in 2D, which projects to a black stripe in their images. On the right are illustrated 4 points in 2D along with their projections into the two cameras.

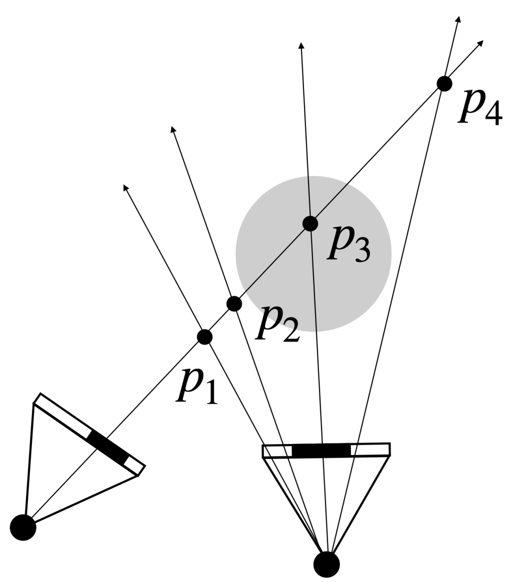

3.1: 4-point NeRF. Based on these 4 points, come up with an assignment of density and color to each point p_i which satisfies the input images.

3.2: Where's the density? Extending the logic from 3.1, indicate in the diagram below *all possible* region over which the density in the scene could lie. (Hint: It's larger than the circle.)

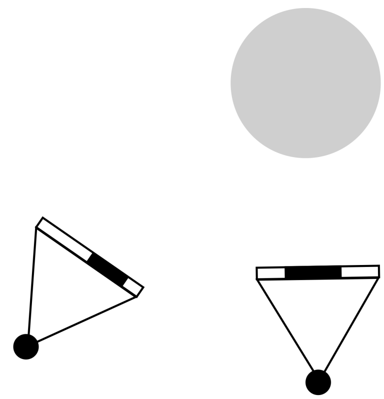

3.3: Adding a new view Based on the existing 2 views, where should I take a new image to give me the most information about where density lives in the scene? Draw the new view and sketch the new region of density.