# Stanford EE367 / CS448I: Computational Imaging Project Proposal

Daniel Song (djsong), Alex Liu (alexlpc)

February 18, 2026

## 1 Motivation

Recent efforts in novel-view synthesis sparked the use of radiance fields as a method of representing high-fidelity scenes. Approaches such as Neural Radiance Fields (NeRF [4]) introduced using neural networks to represent the underlying scene geometry implicitly. More recently, research has shifted towards using explicit representations such as Plenoxels [6] or 3D Gaussian Splatting [2] to improve the training time and allow real-time inference. However, despite the use of spherical harmonics to represent view-dependent colors, current methods suffer from poor representations of refractive objects, like glass and water. In this project, we aim to create an approach for allowing these explicit volumetric scenes to also model physically accurate light propagation in *refractive* objects.

## 2 Related Work

Modeling view-dependent appearance in radiance fields has been an active area of research since the introduction of Neural Radiance Fields (NeRF) [4]. While NeRF demonstrated impressive performance in synthesizing photorealistic novel views, its formulation assumes straight-line light transport and struggles to correctly model non-Lambertian effects such as reflection and refraction. This limitation becomes especially apparent in scenes containing transparent or refractive objects, where light paths deviate due to changes in the index of refraction.

Several works have attempted to extend NeRF to better handle transparent phenomena. NeRFrac [7] explicitly models refractive surfaces by estimating geometry and applying refraction-aware ray bending during rendering. Similarly, REF2-NeRF [3] incorporates both reflection and refraction into the radiance field formulation, improving reconstruction quality for glass-like objects. Yoon et al. [5] combine NeRF with visual hull priors to better constrain transparent object geometry. While these approaches improve rendering fidelity, they rely on neural implicit representations, which are computationally expensive to train and often require carefully curated capture setups or strong geometric priors.

More recent work has explored explicit radiance field representations for efficiency. Plenoxels [6] replace neural networks with sparse voxel grids storing density and spherical harmonics coefficients, significantly accelerating optimization and enabling real-time rendering. 3D Gaussian Splatting [2] further improves performance by representing scenes as anisotropic Gaussian primitives optimized through rasterization-based splatting. However, these explicit methods largely inherit the straight-ray assumption from NeRF and therefore struggle with refractive transport. TransparentGS [1] introduces transparent Gaussian primitives and deferred refraction strategies to better reconstruct transparent objects, but its formulation focuses on inverse rendering and probe-based lighting rather than differentiable physical light transport through volumetric media.

In contrast to prior work, we propose a differentiable rendering framework built on an explicit voxel radiance field that models refraction as a *continuous volumetric transport phenomenon*. Instead of refracting rays at discrete surfaces, we model light propagation through spatially varying indices of refraction using the Eikonal equation. This allows ray trajectories to bend smoothly within the volume while remaining fully differentiable. By combining explicit voxel grids (for efficiency) with physics-based refractive ray integration (for accuracy), our approach bridges the gap between fast explicit radiance fields and physically grounded light transport for transparent media.

# 3 Project Overview and Goals

To keep the project within scope to complete in a couple of weeks, we first aim to create a simplified differentiable rendering pipeline similar to Plenoxels [6], where we have a uniform voxel field of densities and colors. We will be using standard iterative ray tracing with trilinear interpolation to calculate the interpolated densities and colors along each camera ray. To implement a more accurate physical propagation of light through transparent or translucent surfaces, we will also store the index of refraction at each point. However, since the scene is effectively represented as a continuous field of values, it is hard to concretely define a surface at which the light would be refracted. Instead of using a non-differentiable Snell's Law approach, we will need to use a formulation that works in continuous, differentiable spaces. One method is to use the Eikonal equation, which could be used to model the propagation of light in a medium with a spatially varying index of refraction. The ray trajectory $\mathbf{x}(s)$ parameterized by the arclength could be derived from the Eikonal equation,

$$\frac{d}{ds}\left(n\frac{d\mathbf{x}}{ds}\right) = \nabla n,$$

where $n$ is the index of refraction. Here, $\nabla n$ could be estimated using finite differences. Using this, we can use Euler integration to discretize the update step for the ray position and direction.

The current setup assumes that the entire scene will have varying refractive indices. However, this is usually not true for most natural scenes. Because of this, we will need to experiment with different scene initialization methods and include a regularization term that discourages large variations in refractive indices across the scene, as these variations may hinder convergence during the learning process.

If this process proves effective and time permits, we may expand this approach to also account for both the reflection and transmission of light through the medium. One way to avoid non-differentiable ray splitting would be to estimate a radiance cache through spherical harmonics. We can weigh this term by the magnitude of the density gradient, along with the Fresnel coefficient, to encourage this term to be considered only on "surfaces".

Milestone 1 (week 8): Create a differentiable rendering pipeline for a voxel field of densities and colors (similar to Plenoxels)

Milestone 2 (week 9): Create a simple dataset to test the implementation on scenes with refractive objects (e.g., glass, water, etc.)

Milestone 3 (week 10): Implement ray refraction traversal and experiment with regularization term

Milestone 4: If time permits, experiment with also implementing ray reflections along with the transmission according to the Fresnel equation

# References


[1] Letian Huang, Dongwei Ye, Jialin Dan, Chengzhi Tao, Huiwen Liu, Kun Zhou, Bo Ren, Yuanqi Li, Yanwen Guo, and Jie Guo. Transparentgs: Fast inverse rendering of transparent objects with gaussians. *ACM Trans. Graph.*, 44(4), July 2025.

[2] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023.

[3] Wooseok Kim, Taiki Fukiage, and Takeshi Oishi. Ref2-nerf: Reflection and refraction aware neural radiance field. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7196–7203, 2024.

[4] 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. *CoRR*, abs/2003.08934, 2020.

[5] Heechan Yoon and Seungkyu Lee. Neural Radiance Fields for Transparent Object Using Visual Hull . In *2024 IEEE International Conference on Big Data and Smart Computing (BigComp)*, pages 375–376, Los Alamitos, CA, USA, February 2024. IEEE Computer Society.

[6] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *CoRR*, abs/2112.05131, 2021.

[7] Yifan Zhan, Shohei Nobuhara, Ko Nishino, and Yinqiang Zheng. Nerfrac: Neural radiance fields through refractive surface. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18356–18366, 2023.