

Multi-feature Radiance Baking Neural Networks for Instant Volumetric Rendering

SIGGRAPH 2026
Los Angeles 19-23 JUL



Jiaming Liang¹, Hongliang Yuan², Meng Gai¹, Guoping Wang¹, Sheng Li¹
1: Peking University 2: Xiaomi Company

Project Page


Background

Photorealistic volumetric rendering is slow by traditional path tracing.

RPNN family tries to predict radiance directly from the volume, achieving:

- (1) RPNN: ~minutes to converge;
- (2) MRPNN: 40 FPS per sample.

But still too costly for real-time application!




PT: 6 minutes; Ours: **200 FPS**; Prior SOTA: 40 FPS

Core Contribution


1. Theory: Reformulate in-scattering integral in VRTE by Spherical Harmonics, eliminating explicit integration in theory.
2. Architecture : Propose a compact latent-space neural baking that can be decoded by our dual-pattern feature query and lightweight network.
3. Application: Implement a real-time volumetric renderer that achieves 200 FPS and supports spatially varying albedo, heterogeneous phase functions and improved temporal stability.

Method


1. Multi-resolution Encoding: w.r.t. scattering field (with SH operator), light direction, material, etc.
2. Dual-pattern sampling: local sampling (for low-frequency features) and towards-light cone sampling (for high-frequency features).




3. Tiny network: MLP with Residual-Connected blocks + next-level feature fusion, < 30K params.



Pipeline



Experiments



Model credit: Disney Clouds; OpenVDB Samples; CS Motion Design; JangaFX; and data from MRPNN.

