

Radiance Caching for Participating Media

Wojciech Jarosz

Craig Donner

Matthias Zwicker

Henrik Wann Jensen

University of California, San Diego

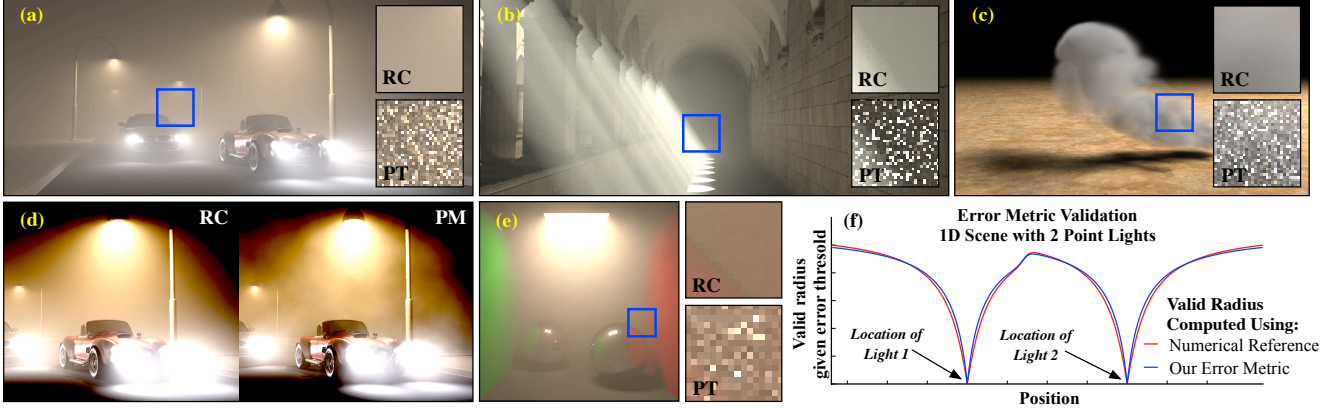


Figure 1: Our method renders the cars scene (a) in 100 minutes using 150K cache points. In a contrast-enhanced, equal-time comparison to photon mapping (d), artifacts remain in the homogeneous medium even after tracing 8M photons due to the large scene extent and many light sources. With volumetric radiance caching, the Sponza scene (b) renders in 72 minutes using only 13K cache points. Sponza is notoriously difficult for photon mapping: even using projection maps to direct photons towards the windows, tracing photons for an equal amount of time results in less than 1K usable photons. A still frame from an animation of heterogeneous smoke (c) renders in 11 minutes using 12K cache points and produces flicker-free animation. Using spherical harmonics our method renders the anisotropic media in the Cornell box (e) in 14 minutes. Images were rendered with a horizontal resolution of 1024 (a, b, c) and 512 (e) with up to 16 samples per pixel. The zoomed insets (a, b, c, e) provide equal-time quality comparisons for radiance caching (RC) and path tracing (PT). A validation of our error metric as compared to a numerically computed optimal radius is shown in (f) for a 1D scene with two point lights of differing strength.

We present an extension of irradiance caching [Ward et al. 1988] to participating media. Our method is based on Monte Carlo ray tracing, and it exploits the smooth nature of scattering in participating media by *caching* radiance accurately at a sparse set of locations and using *interpolation* to evaluate the radiance at arbitrary points in the medium. We compute gradients of the radiative transport equation to estimate the local variation of the scattered radiance and to improve the accuracy of interpolation. Our approach handles both heterogeneous media and anisotropic phase functions, and it is several orders of magnitude faster than path tracing. Furthermore, it is view driven and well suited in large scenes where methods such as photon mapping [Jensen and Christensen 1998] become costly (see Figure 1 a, b and d).

Radiative Transport In participating media, radiative transport is described by the volume rendering equation, which for a non-emissive medium is:

$$L(\mathbf{e}, \vec{\omega}) = \int_0^s T_r(\mathbf{e} \leftrightarrow \mathbf{x}) \sigma_s(\mathbf{x}) L_i(\mathbf{x}, \vec{\omega}) d\mathbf{x} + T_r(\mathbf{e} \leftrightarrow \mathbf{x}') L(\mathbf{x}', \vec{\omega}), \quad (1)$$

where T_r is the transmittance between two points, and $\mathbf{x}' = \mathbf{e} + s\vec{\omega}$. We evaluate the in-scattered radiance L_i as the sum of two components $L_i = L_s + L_m$ (single scattering and multiple scattering).

Single Scattering For single scattering we express the in-scattered radiance at a point \mathbf{x} in direction $\vec{\omega}$ as:

$$L_s(\mathbf{x}, \vec{\omega}) = \int_A p(\vec{\omega}, \mathbf{x}' \rightarrow \mathbf{x}) L_r(\mathbf{x}' \rightarrow \mathbf{x}) V(\mathbf{x}' \leftrightarrow \mathbf{x}) H(\mathbf{x}' \rightarrow \mathbf{x}) d\mathbf{x}', \quad (2)$$

where p is the phase function, L_r is the reduced radiance, V is the visibility term, and H is a geometry term. Assuming constant visibility, we compute the gradients by analytically differentiating w.r.t. \mathbf{x} :

$$\nabla L_s(\mathbf{x}, \vec{\omega}) = \int_A (\nabla p) L_r V H + p (\nabla L_r) V H + p L_r V (\nabla H) d\mathbf{x}'. \quad (3)$$

In anisotropic media, this formulation fully accounts for the change in the phase function. In heterogeneous media, the ∇L_r term accounts for the change in the optical thickness due to varying scattering properties. The gradients are computed in conjunction with the radiance during Monte Carlo integration, and our derivations place no restriction on the sampling pdf. We handle single scattering from point lights, cosine lights, and area lights, as well as scattering from reflective surfaces in the scene due to indirect illumination. Single scattering gradients of new light sources are easily incorporated into our system.

Multiple Scattering For multiple scattering, we compute the radiance, $L_m(\mathbf{x}, \vec{\omega})$, as an integral of random-walk paths starting at \mathbf{x} . We express the gradient as the change of radiance arriving along the paths as \mathbf{x} and the paths are translated *as a whole*. This approach takes into account how radiance changes along the *full* path, and, furthermore allows for importance sampling the phase function at all but the first path vertex.

Caching Our method computes and caches radiance lazily during the rendering process. When evaluating radiance at some location \mathbf{x} , we first check if interpolation is possible by querying nearby cache points in an octree (as in [Ward et al. 1988]). If cache points are found, we perform a smoothly-weighted exponential extrapolation suitable for participating media. When no points are found, the radiance and gradient are computed and cached for future extrapolation. Each cache point has a valid radius within which extrapolation is allowed. This valid radius is computed using a novel, perceptually-based error metric which adapts to the local smoothness of the radiance field. Even in the presence of saddle points when the gradient vanishes, our metric robustly approximates the largest radius that maintains a specific relative L_2 error (Figure 1f). The amount of extrapolation is controlled using an error tolerance parameter as well as minimum and maximum radius constraints in both world and screen space.

In our system, we further separate our cache based on the expected frequency content of the radiance values. Specifically, we maintain distinct cache points for single scattering from lights, single scattering from surfaces, and multiple scattering. This allows our system to sample low-frequency components more sparsely. For isotropic media, we store the fluence and gradients as RGB values. In anisotropic media, we project our radiance and gradient computations onto spherical harmonics (though any spherical basis could be used) and store the full spherical representations with a user-specified number of coefficients.

References

- JENSEN, H. W., AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of SIGGRAPH*.
- WARD, G. J., RUBINSTEIN, F. M., AND CLEAR, R. D. 1988. A ray tracing solution for diffuse interreflection. In *Proceedings of SIGGRAPH*.