ARTICLE IN PRESS

Neurocomputing xx (xxxx) xxxx–xxxx

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

On learning the visibility for joint importance sampling of low-order scattering

Guo Zhou[a,b,c,*], Dengming Zhu[a,b], Ting Li[d], Zhaoqi Wang[a,b], Yongquan Zhou[e]

[a] Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China
[b] Beijing Key Laboratory of Mobile Computing and Pervasive Device, Beijing 100190, China
[c] University of Chinese Academy of Sciences, Beijing 100049, China
[d] CNCERT/CC, Beijing 100029, China
[e] College of Information Science and Engineering, Guangxi University for Nationalities, Nanning 530006, China

ARTICLE INFO

Keywords:
Light transport simulation
Participating media
Online expectation-maximization
Importance sampling

ABSTRACT

Volumetric path tracing relies on importance sampling to stochastically construct light transport paths from an emitter to the sensor. Existing techniques incrementally sample path vertices or segments with respect to the local scattering property incorporating the geometry and scattering terms. Thus the joint probability density for drawing a path results in a product of the conditional densities each for a local sampling decision. We present a joint path sampling technique that additionally accounts for the spatially varying visibility due to transmittance and occlusion along a double scattering path. The directional density is formulated as a Gaussian mixture model being fitted to single scattered radiance by the online expectation–maximization algorithm. It is first trained with samples oblivious to the visibility, then incrementally consumes an arbitrary number of samples being drawn from the actual scene. The resulting density in turn guides the directional sampling decision for both isotropic and anisotropic scattering. We demonstrate the benefit of our approach by integrating it into the unidirectional path tracing algorithm. The image noise is effectively reduced, even while rendering the heterogeneous participating media in the presence of complex opaque surfaces.

1. Introduction

Robust and efficient light transport simulation in participating media is challenging due to the enclosed particles, which incur a large number of scattering and absorption events. Among many algorithms the computer graphics community has developed, the Monte Carlo integration based volumetric path tracing [1] is widely adopted to render participating media. The total amount of light arriving at the sensor film is described as an integral over the space of possible light transport paths, which is stochastically sampled to achieve an unbiased estimate. A path is successively constructed as a random walk by sampling a probability density function (PDF) to decide a direction or distance. This class of algorithms is able to flexibly handle various scenes, but suffer from much image noise due to the curse of dimensionality.

To improve convergence rate, the importance sampling technique is adopted to steer each decision with respect to the local scattering property. A proposal PDF is introduced to distribute more samples to where the value of the integrand is relatively high, thus variance of the estimator is reduced expectedly. This integrand is a product of properties consisting of the geometry, scattering, transmittance and occlusion terms. It is often intractable to formulate an analytical proposal which is able to match the whole integrand. Existing methods strive for sampling the geometry and scattering terms separately or jointly only. Thereby such local construction of paths is oblivious to the global distribution of radiance. Its efficiency is ensured only in scenes where the transmittance and the occlusion terms do not introduce high variations to the integrand.

We address this visibility issue due to varying optical thickness and spatial occlusion by enhancing the directional decision in the case of low-order scattering. Its sampling PDF is formulated as a Gaussian mixture model (GMM), which subsumes the product of the transmittance and scattering terms over multiple consecutive decisions. Then it can be trained by an arbitrary number of path samples using the online expectation–maximization (EM) algorithm, progressively adapting to spatially varying visibility in the path space. Since the visibility has no closed-form expression, we employ a conjugate prior of inverse Wishart distribution to avoid over-fitting. This directional PDF is initialized by samples being drawn from the joint PDF of the scattering and geometry terms only, then it is incrementally fitted to the actual scene and guides

* Corresponding author at: Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China.
E-mail address: zhouguo@ict.ac.cn (G. Zhou).

http://dx.doi.org/10.1016/j.neucom.2016.09.086
Received 26 December 2015; Received in revised form 4 April 2016; Accepted 3 September 2016
Available online xxxx
0925-2312/

the directional sampling afterwards.

Along with the distance PDFs, our approach enables a joint importance sampling technique which is able to draw a path of single or double scattering entirely. Our main contributions are:

- a directional PDF which progressively adapts to spatially varying visibility in the path space and guides random decisions of future path sample;
- a learning procedure which avoids of over-fitting due to complex visibility and ensures the joint form of both the transmittance and scattering terms;
- a joint sampling technique which improves the volumetric path tracing algorithm by drawing a path entirely thus variance of the estimator is further reduced.

We discuss the related work (Section 2) to position our visibility-aware joint importance sampling technique first. After which the background (Section 3) is given to illustrate the basics about the volumetric path tracing. Then our approach (Section 4) describes how to integrate the visibility with the existing joint importance sampling method via the GMM learning. The results and implementation details (Section 5) are presented to demonstrate the effectiveness. Finally we conclude the paper and discuss future work (Section 6).

2. Related work

The path tracing algorithm [2] was first introduced to sample light transport paths connecting emitters to the sensor in the context of surface rendering. The scenes are assumed to be in a vacuum and the path segments bouncing between surfaces are constant. It was formalized as a solution to the path integral [3], which was then extended to the case of participating media [4]. The path integral is a product of the geometry, visibility (i.e. transmittance and occlusion) and scattering terms (see Section 3), thus it is often impossible to find a closed-form proposal PDF which is ideally proportional to that integrand. Many approximate approaches have been devised to sample the path space with respect to partial integrand or reuse sampled paths for variance reduction.

Path sampling PDF factorization: The natural approach of path construction in participating media alternatively samples the scattering term (e.g. phase function) for the direction of the path segment, and then the transmittance for the traveled distance along this segment [4–6]. A geometry term contributes more variations than the transmittance under some circumstances, thus it is adopted to determine the distance instead for single scattering in isotropic media [7,8]. In this case the proposal PDF is just a product of conditional PDFs exactly proportional to respective term only. For double scattering the geometry and scattering terms over three path segments can be considered jointly [9]. The proposal PDF is chosen to be proportional to their product and then it is factorized into conditional PDFs in canonical coordinate system. And the conditional PDFs corresponding to scattering are numerically tabulated due to lack of closed-form formula. These proposals are combined with the natural approach for path sampling by the multiple importance sampling (MIS) [7]. But the resulting proposal is only a linear combination of individual proposal PDFs rather than a product of all terms in the path integrand. We augment the joint proposal with the visibility term to propose a guided joint path sampling technique for more effective variance reduction.

Caching and reusing of sampled paths: Paths can be bidirectionally constructed by caching path segments starting from emitters, then they are shared among segments from the sensor. The overall light propagating between these two kinds of segments is considered thus more samples can be obtained at once [8]. The Metropolis–Hastings algorithm based path tracing [3,4] strives to implicitly sample path with contribution proportional to the whole integrand. An existing path is treated as a state on the path space, then it is mutated to obtain a new path. Although paths with large contribution appear more frequently, they usually suffer from low acceptance rate resulting from high variations in the integrand. To alleviate this problem, the integrand is transformed to eliminate its significant factors and mutation is applied in a warped path space [10–12]. These approaches suffer from sample correlation and are usually outperformed by the independent samples in practical cases. The GMM learning has been introduced to keep the information of sampled paths [13] for volumetric photon mapping, and an online EM algorithm based on maximum a posteriori (MAP) has been adopted to fit the conditional PDF of surfaces scattering and occlusion in complex scenes [14]. Since the EM algorithm only approaches a local mode, both of them are sensitive to the initial batch of samples. We extend this approach to low-order scattering in participating media. The radiance caching [15] stores the gradient of the reduced radiance, single and multiple scattering to exploit the local smoothness of the radiance field for extrapolation. Similarly, we utilize the gradient of double scattering but to determine the validity radius of the radiance PDF instances in the medium. And the visibility being ignored by the gradient due to its discontinuity in previous work is additionally handled.

3. Background

In this section we review the path integral framework for scenes with participating media, following a presentation similar to [9]. This framework describes how a sensor (e.g. camera) receives light from an emitter (e.g. light bulb) to render a realistic image.

The intensity of a pixel is a high-dimensional integral over the space Ω of light transport paths

$$I = \int_{\Omega} f(\bar{\mathbf{x}})\mathrm{d}\bar{\mathbf{x}} \tag{1}$$

where $\bar{\mathbf{x}} = \mathbf{x}_0, \ldots, \mathbf{x}_M$ is a path with $M \geq 1$ segments and $M + 1$ vertices. The first vertex $\mathbf{x}_0$ is on an emitter while the last vertex $\mathbf{x}_M$ is on a sensor. Others are spatial locations where the path is scattered, either on surfaces or in participating media. The differential measure $\mathrm{d}\bar{\mathbf{x}}$ on Ω represents area integration over surfaces or volume integration over region. The measurement contribution function $f(\bar{\mathbf{x}})$ is a product of geometry throughput $G(\bar{\mathbf{x}})$, scattering throughput $\rho(\bar{\mathbf{x}})$, transmittance throughput $T(\bar{\mathbf{x}})$, and occlusion throughput $O(\bar{\mathbf{x}})$

$$f(\bar{\mathbf{x}}) = G(\bar{\mathbf{x}})\rho(\bar{\mathbf{x}})T(\bar{\mathbf{x}})O(\bar{\mathbf{x}}) \tag{2}$$

where each throughput is a product of corresponding term over all the path segments $\mathbf{x}_m\mathbf{x}_{m+1}$ or vertices $\mathbf{x}_m$. The scattering term is

$$\rho(\mathbf{x}_m) = \begin{cases} L_e(\mathbf{x}_0 \to \mathbf{x}_1), & \text{if } m = 0 \\ W_e(\mathbf{x}_{M-1} \to \mathbf{x}_M), & \text{if } m = M \\ \rho_s(\mathbf{x}_{m-1} \to \mathbf{x}_m \to \mathbf{x}_{m+1}), & \text{if } \mathbf{x}_m \text{ on surface} \\ \rho_p(\mathbf{x}_{m-1} \to \mathbf{x}_m \to \mathbf{x}_{m+1})\sigma_s(\mathbf{x}_m), & \text{if } \mathbf{x}_m \text{ in medium} \end{cases} \tag{3}$$

where L_e is the emitter emission, W_e is the sensor importance, ρ_s is the bidirectional scattering distribution function (BSDF), ρ_p is the phase function and σ_s is the scattering coefficient. The geometry term is

$$G(\mathbf{x}_m, \mathbf{x}_{m+1}) = \frac{D(\mathbf{x}_m \to \mathbf{x}_{m+1})D(\mathbf{x}_{m+1} \to \mathbf{x}_m)}{\|\mathbf{x}_m - \mathbf{x}_{m+1}\|^2} \tag{4}$$

where $D(\mathbf{x} \to \mathbf{y}) = |\mathbf{n}_\mathbf{x} \cdot \mathbf{w}_{\mathbf{xy}}|$ if $\mathbf{x}$ is on surfaces, and $D(\mathbf{x} \to \mathbf{y}) = 1$ if $\mathbf{x}$ is in media. It is responsible for converting the product of solid angle and length measure to the volume measure (i.e. Jacobian determinant). The transmittance term is

$$T(\mathbf{x}, \mathbf{y}) = \exp\left(-\int_0^{\|\mathbf{x}-\mathbf{y}\|} \sigma_t(\mathbf{x} + t\mathbf{w}_{\mathbf{xy}})\mathrm{d}t\right) \tag{5}$$

where σ_t is the sum of the absorption coefficient and scattering coefficient. The occlusion term $O(\mathbf{x}_m, \mathbf{x}_{m+1})$ is zero if segment $\mathbf{x}_m\mathbf{x}_{m+1}$ is intercepted by any surface or one otherwise.

Monte Carlo integration empirically solves Eq. (1) by drawing