Pseudo-marginal Metropolis Light Transport

Authors: Joel Kronander, Thomas B. Schon, and Jonas Unger. 2015
Presenter: Seo Hansol

Review - Scene Aware Audio for 360° Videos

- Synthesize ambisonic audio for 360° Video
- Separate direct sound, ERIR, and LRIR: Combine later
- For ERIR,
  - Build geometry from 360° video
  - Measure room IR response & optimize material absorbance
  - Synthesize ERIR with geometric acoustic model
- LRIR is isotropic; reuse room IR response
- For low frequency, use frequency modulation.

Motivations: Overview

Motivations: Overview

- Naïve Monte Carlo rendering: hard to find eye-to-light path
- Metropolis light transport (MLT): light path sampling algorithm
- MLT uses Markov chain Monte Carlo (MCMC), using the image contribution of a path as probability distribution for estimation.
  - Metropolis-Hastings (MH) algorithm is used.

Motivations: Overview

- But the image contribution function is unknown, and hard to evaluate.
  - Especially in complex scenes, such as a scene with participating media.
- Solution : Estimate!
  - [Pauly et al. 2000] proposed an approximation, but it is biased.
- Need to build an unbiased estimator for MH acceptance probability.

Metropolis
Light Transport

Background: Key Ideas

- Metropolis light transport (MLT)
- Path integral formulation
- Markov chain Monte Carlo (MCMC)
- Metropolis-Hastings (MH) algorithm

Metropolis Light Transport

- Only few paths reach a light in Naïve Monte Carlo path tracing

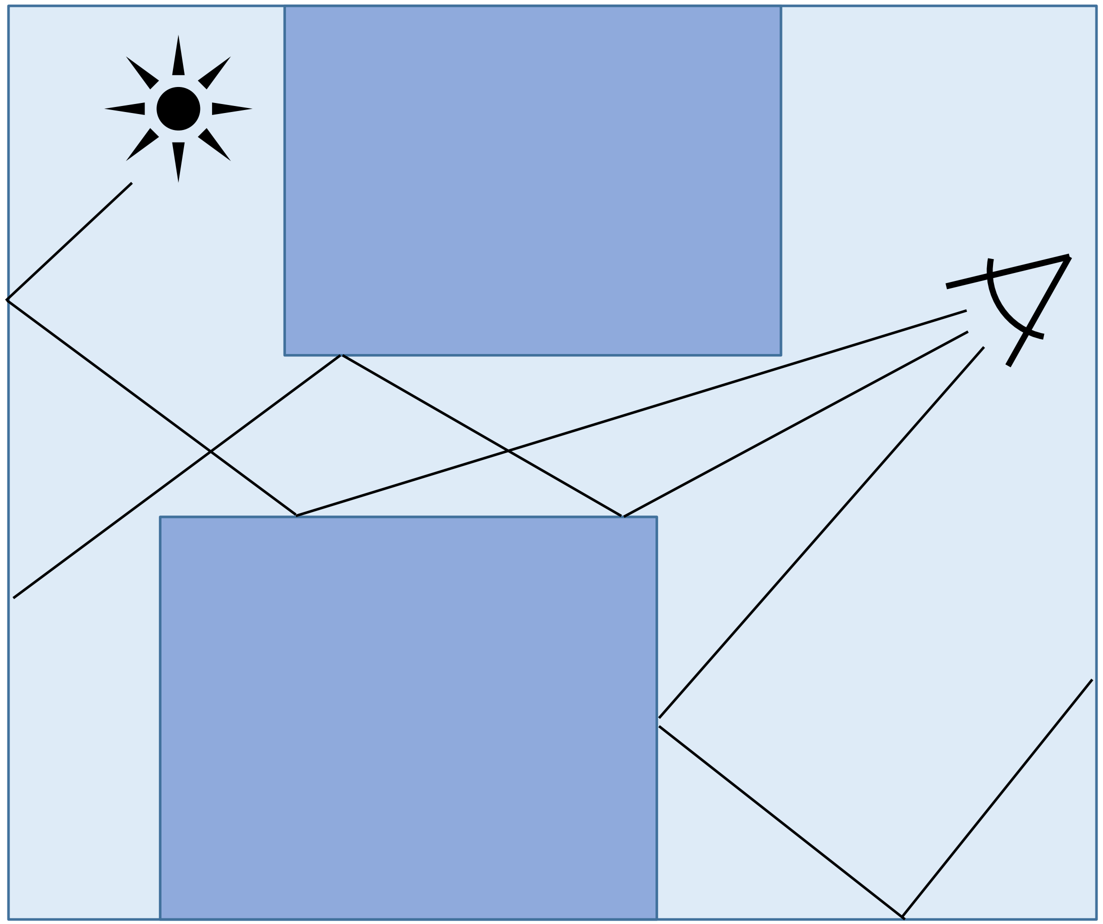

Metropolis Light Transport

- Only few paths reach a light in Naïve Monte Carlo path tracing
- MLT focuses on successful eye-to-light ray

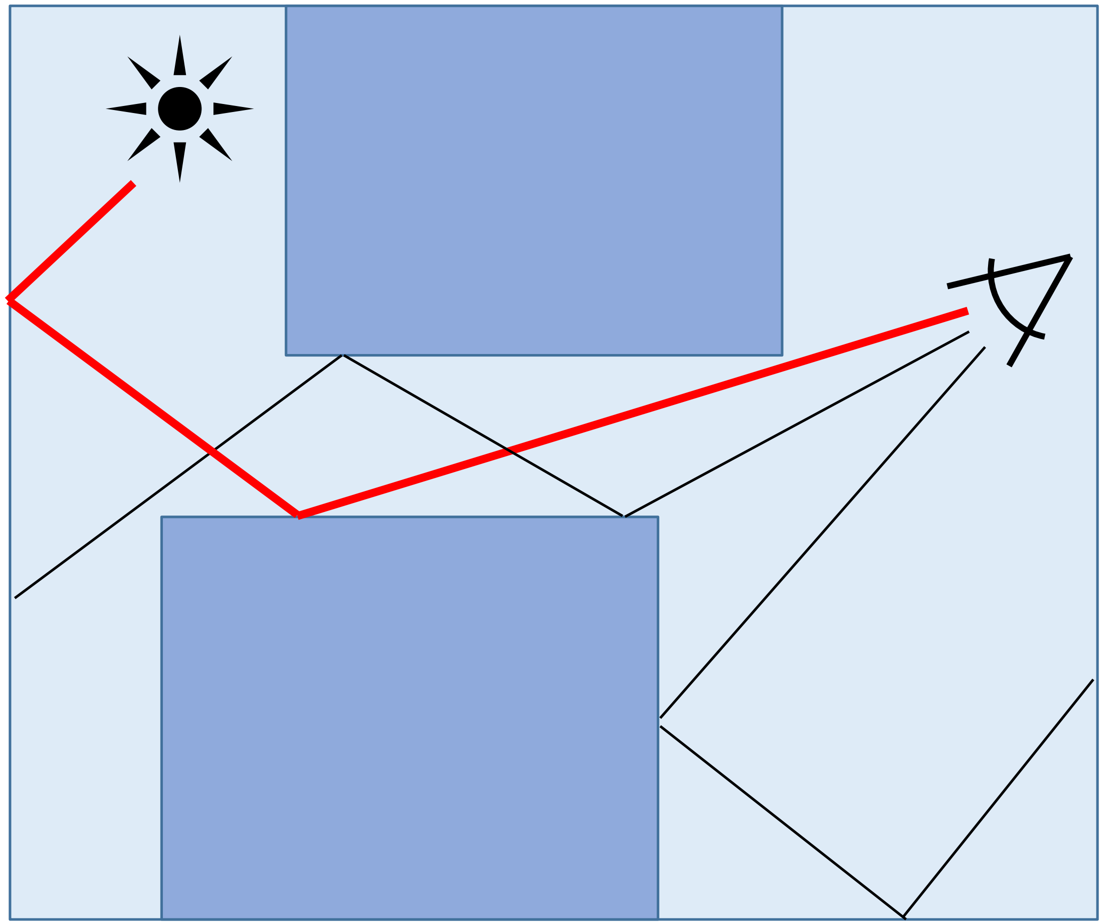

Metropolis Light Transport

- Only few paths reach a light in Naïve Monte Carlo path tracing
- MLT focuses on successful eye-to-light ray
- Make small perturbations to build similar paths

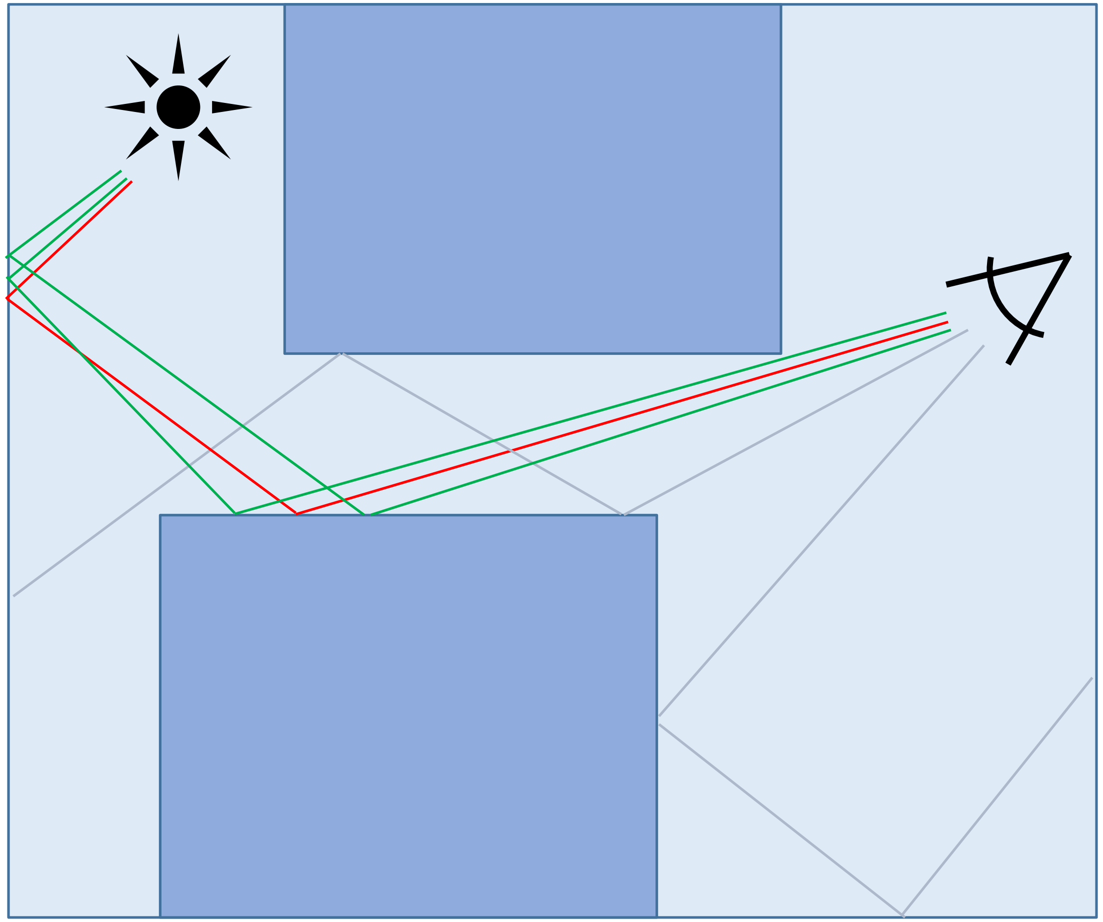

Metropolis Light Transport

- Only few paths reach a light in Naïve Monte Carlo path tracing
- MLT focuses on successful eye-to-light ray
- Make small perturbations to build similar paths
- Make & accept new path probabilistically

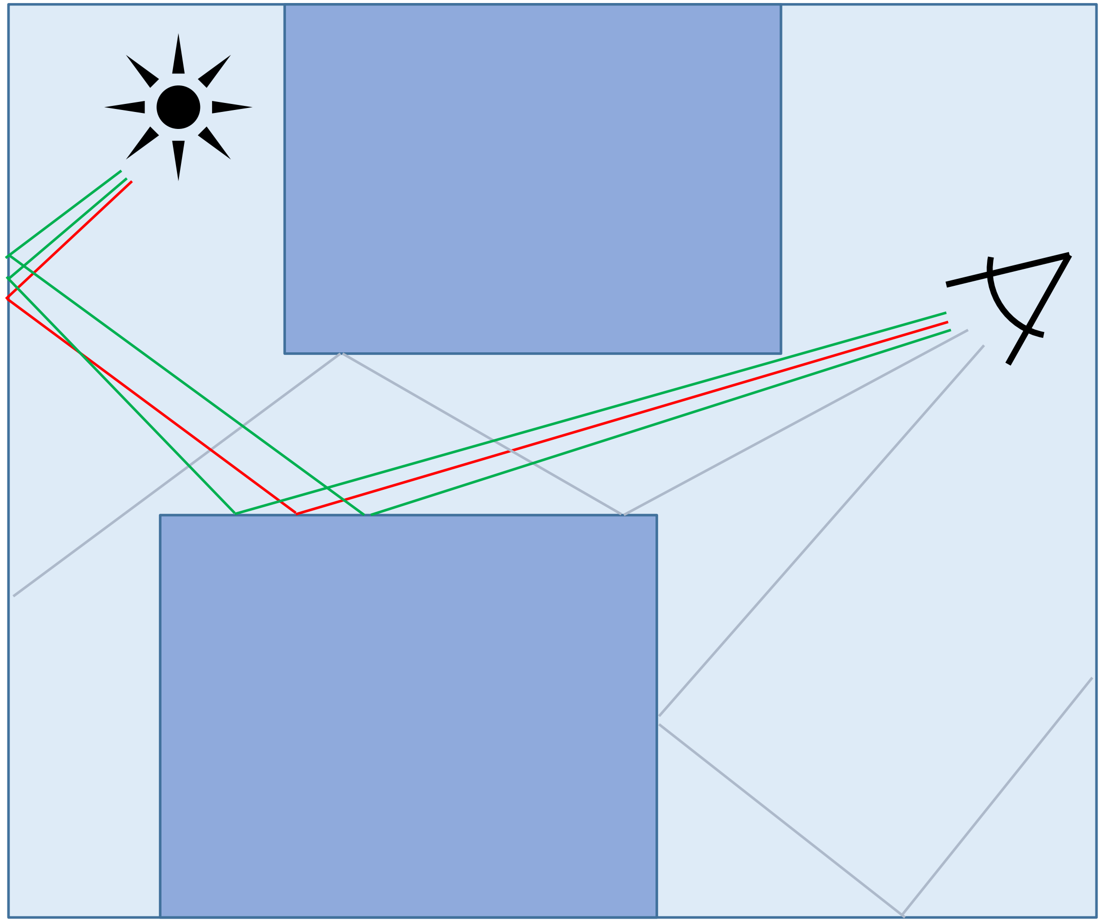

Metropolis
Light Transport

More Mathematical

Path Integration Formulation

- MLT integrates over all valid light path.
- Path integration formulation

$$I_j = \int_{\Omega} f_j(\bar{x})\, W_j\, d\mu(\bar{x})$$

  - I_j is j-th pixel luminance, and W_j is pixel sensitivity.
  - $d\mu(\bar{x}) = \prod_{i=0}^{k} d\mu(x_i)$
  - $d\mu(x_i) = \begin{cases} dA(x_i) & if\ x_i\ is\ on\ surface \\ dV(x_i) & if\ x_i\ is\ in\ media \end{cases}$

- $f_j(\bar{x})$ is the measurement contribution function.

Path Integration Formulation

- $f_j(\bar{x})$ is the measurement contribution function.
  - In other words, luminance contribution of path $\bar{x}$ to pixel j.
  - Defined as

$$f_j(\bar{x}) = L_e \left[\prod_{i=0}^{k-1} G(x_i, x_{i+1}) T(x_i, x_{i+1})\right] \left[\prod_{i=1}^{k-1} \rho(x_i)\right]$$

- G is geometrical factor between vertices, ρ is scattering factor.
- T is the transmittance function along a path edge.

$$T(x_i, x_{i+1}) = \exp\left(-\int_0^d \sigma(x + t\omega_{i,i+1}) dt\right)$$

Path Integration Formulation

- $f_j(\bar{x})$ is the measurement contribution function.
  - In other words, luminance contribution of path $\bar{x}$ to pixel j.
  - Defined as

$$f_j(\bar{x}) = L_e \left[\prod_{i=0}^{k-1} G(x_i, x_{i+1}) T(x_i, x_{i+1}) \right] \left[\prod_{i=1}^{k-1} \rho(x_i) \right]$$

- Hard to evaluate; let's estimate

Path Integration: Estimation

- Estimate $I_j = \int_\Omega f(\bar{x})\, W_j\, d\mu(\bar{x})$.
- Let's sample paths following a probability density function

$$\pi(\bar{x}) = \frac{L(\bar{x})}{\int_\Omega L(\bar{x}) d\mu(\bar{x})}$$

  - $L(\bar{x})$ is a scalar probability function, in other word, contribution function.
- Then estimation of the path integral would be

$$\widehat{I_j} = \frac{Z}{N} \sum_{i=1}^{N} \frac{W_j f(\bar{x}^i)}{L(x)}$$

  - Z is normalization factor; can be computed from traditional path tracing.

Path Integration: Estimation

- But we need to know $\pi(\bar{x})$ in advance to sample paths.
  - Need to evaluate $L(\bar{x})$.
  - $L(\bar{x})$ is intractable for a complex scene with participating media.
- Let’s try using Markov chain Monte Carlo method.

Markov Chain Monte Carlo

- Estimate $I_j = \int_{\Omega} f(\bar{x})\, W_j \, d\mu(\bar{x})$.
- Use Markov chain Monte Carlo to sample light paths.
- Procedure of sampling a new path is only dependent on current path ("Markov chain").
- And probabilistic ("Monte Carlo").
- More specifically, let's use Metropolis-Hastings algorithm.

Metropolis-Hastings Algorithm

- Starting from an eye-to-light path,
- 1. Sample new path $\bar{x}^{i+1}$ with perturbations from current path $\bar{x}^{i}$, with probability distribution $q(\bar{x}^{i+1}|\bar{x}^{i})$.
- 2. Accept with probability

$$r(\bar{x}^{i},\bar{x}^{i+1}) = min\left\{1,\frac{L(\bar{x}^{i+1})}{L(\bar{x}^{i})}\frac{q(\bar{x}^{i}|\bar{x}^{i+1})}{q(\bar{x}^{i+1}|\bar{x}^{i})}\right\}$$

  - L is the contribution function (again)
  - Discard if not accepted, keep if accepted
  - Acceptance probability should be exact for exact (unbiased) result
- As we iterate 1 & 2, path distribution converges to $\pi(\bar{x})$.

Metropolis-Hastings Algorithm

- 2. Accept with probability $r = min\left\{1, \frac{L(\bar{x}^{i+1})}{L(\bar{x}^{i})} \frac{q(\bar{x}^{i}|\bar{x}^{i+1})}{q(\bar{x}^{i+1}|\bar{x}^{i})}\right\}$
- q is given by algorithm, but still we cannot evaluate L.
  - Acceptance probability r should be exact for exact (unbiased) result!
- L is intractable for a complex scene with participating media.
- We should build an estimator for L.

Previous Works on MLT

- Primary sample space Metropolis light transport (PSSMLT) by Kelemen et al. [2002]
  - Sample on primary sample space, not path space.
  - Depends on path tracing algorithm; limited usability.
- Delta tracking by Raab et al. [2008]
  - Unbiased path tracing algorithm for PSSMLT
  - Same limitations
- Ray marching by Pauly et al. [2000]
  - Evaluate acceptance probability L by moving ray tip step-by-step.
  - But ray marching is biased; analogous to quadrature method

Background: Summary

- MLT sample paths from path space, and integrates.
  - $I_j = \int_{\Omega} f(\bar{x})\, W_j \, d\mu(\bar{x})$
- Use MCMC, specifically MH algorithm.
- MH algorithm relies on the acceptance probability.
  - $r = min\left\{1, \frac{L(\bar{x}^{i+1})}{L(\bar{x}^{i})} \frac{q(\bar{x}^{i}|\bar{x}^{i+1})}{q(\bar{x}^{i+1}|\bar{x}^{i})}\right\}$
- We need unbiased estimator of MH acceptance probability.

Pseudo-marginal
Metropolis-Hastings Algorithm

Pseudo-marginal MH: Overview

- MH algorithm (pseudo-marginal version)
  - Starting from an eye-to-light path,
  - 1. Sample new path $\bar{x}^{i+1}$, and $u_{i+1} \sim g(u|x)$
  - 1-1. Build $\hat{L}_{u_{i+1}}(\bar{x}^{i+1})$.
  - 2. Accept by probability $r = min\left\{1, \frac{\hat{L}_{u_{i+1}}(\bar{x}^{i+1})}{\hat{L}_{u_i}(\bar{x}^{i})} \frac{q(\bar{x}^{i}|\bar{x}^{i+1})}{q(\bar{x}^{i+1}|\bar{x}^{i})}\right\}$
  - Iterate 1 & 2.
- This algorithm is unbiased if estimator $\hat{L}_{u_i}(\bar{x}^i)$ is unbiased!

Pseudo-marginal MH: Why Does It Work?

- Simple explanation
- Consider auxiliary variable $u \sim g(u|x)$, and an estimator $\hat{L}_u(\bar{x})$
- Let $p(\bar{x}, u) = \frac{L(\bar{x})}{Z} g(u|\bar{x})$, and $\phi(\bar{x}, u) = \frac{\hat{L}_u(\bar{x})}{Z} g(u|\bar{x})$.
- Then $\int \phi(\bar{x}, u) du = \frac{1}{Z} \int \hat{L}_u(\bar{x}) g(u|\bar{x}) du = \mathrm{E}_{g(u|\bar{x})}\left[\hat{L}_u(\bar{x})\right]$.
- If $\hat{L}_u(\bar{x})$ is unbiased, $\int \phi(\bar{x}, u) du = \mathrm{E}_{g(u|\bar{x})}\left[\hat{L}_u(\bar{x})\right] = L(\bar{x})$.
- In other words, we can use $\hat{L}_u(\bar{x})$ in place of $L(\bar{x})$, while the MH algorithm is kept unbiased.

Pseudo-marginal MH: Why Does It Work?

- Let's look at acceptance probability.
- The sample space is now $(path\ space) \otimes (domain\ of\ u)$.
- Therefore acceptance probability should be

$$r(\bar{x}^i, \bar{x}^{i+1}) = min\left\{1, \frac{\hat{L}_{u_{i+1}}(\bar{x}^{i+1})g(\bar{u}^{i+1}|\bar{x}^{i+1})}{\hat{L}_{u_i}(\bar{x}^i)g(\bar{u}^i|\bar{x}^i)} \frac{q(\bar{x}^i|\bar{x}^{i+1})g(\bar{u}^i|\bar{x}^i)}{q(\bar{x}^{i+1}|\bar{x}^i)g(\bar{u}^{i+1}|\bar{x}^{i+1})}\right\}$$

  - Sampling u is independent from previous states.
- The g terms are reduced and what remains is

$$r(\bar{x}^i, \bar{x}^{i+1}) = min\left\{1, \frac{\hat{L}_{u_{i+1}}(\bar{x}^{i+1})}{\hat{L}_{u_i}(\bar{x}^i)} \frac{q(\bar{x}^i|\bar{x}^{i+1})}{q(\bar{x}^{i+1}|\bar{x}^i)}\right\}!$$

Pseudo-marginal MH: Overview (again)

- MH algorithm (pseudo-marginal version)
  - Starting from an eye-to-light path,
  - 1. Sample new path $\bar{x}^{i+1}$, and $u_{i+1} \sim g(u|x)$
  - 1-1. Build $\hat{L}_{u_{i+1}}(\bar{x}^{i+1})$.
  - 2. Accept by probability $r = min\left\{1, \frac{\hat{L}_{u_{i+1}}(\bar{x}^{i+1})}{\hat{L}_{u_i}(\bar{x}^{i})} \frac{q(\bar{x}^{i}|\bar{x}^{i+1})}{q(\bar{x}^{i+1}|\bar{x}^{i})}\right\}$
  - 3. Iterate 1 & 2.
- This algorithm is unbiased if estimate $\hat{L}_{u_i}(\bar{x}^i)$ is unbiased!
- However, what $\hat{L}_{u_i}(\bar{x}^i)$ should we use?

Unbiased Estimator

- Any unbiased estimator can be $\hat{L}_{u_i}(\bar{x}^i)$.
- However, it is reasonable to make it similar to $f_j(\bar{x})$.
  - $f_j(\bar{x})$ was the measurement contribution function.

$$f_j(\bar{x}) = L_e\left[\prod_{i=0}^{k-1} G(x_i, x_{i+1}) T(x_i, x_{i+1})\right]\left[\prod_{i=1}^{k-1} \rho(x_i)\right]$$

- Let's try using below function

$$\hat{L}_u(\bar{x}) = L_e\left[\prod_{i=0}^{k-1} G(x_i, x_{i+1}) \hat{T}_i(x_i, x_{i+1})\right]\left[\prod_{i=1}^{k-1} \rho(x_i)\right]$$

Unbiased Estimator

- Let's try using below function

$$\hat{L}_u(\bar{x}) = L_e \left[\prod_{i=0}^{k-1} G(x_i, x_{i+1}) \widehat{T_i}(x_i, x_{i+1})\right] \left[\prod_{i=1}^{k-1} \rho(x_i)\right]$$

- The only difference is transmittance term T, which is now an estimator.
  - Other terms (G: geometric term, ρ: scattering term) are much cheaper to evaluate, in scenes with participating media.
- $\widehat{T_i}$ is an unbiased estimator from ratio tracking [Novak et al. 2014], so $\hat{L}_u(\bar{x})$ is unbiased.
- This function is independent to u, so there is no need to actually sample u at all.

Ratio Tracking

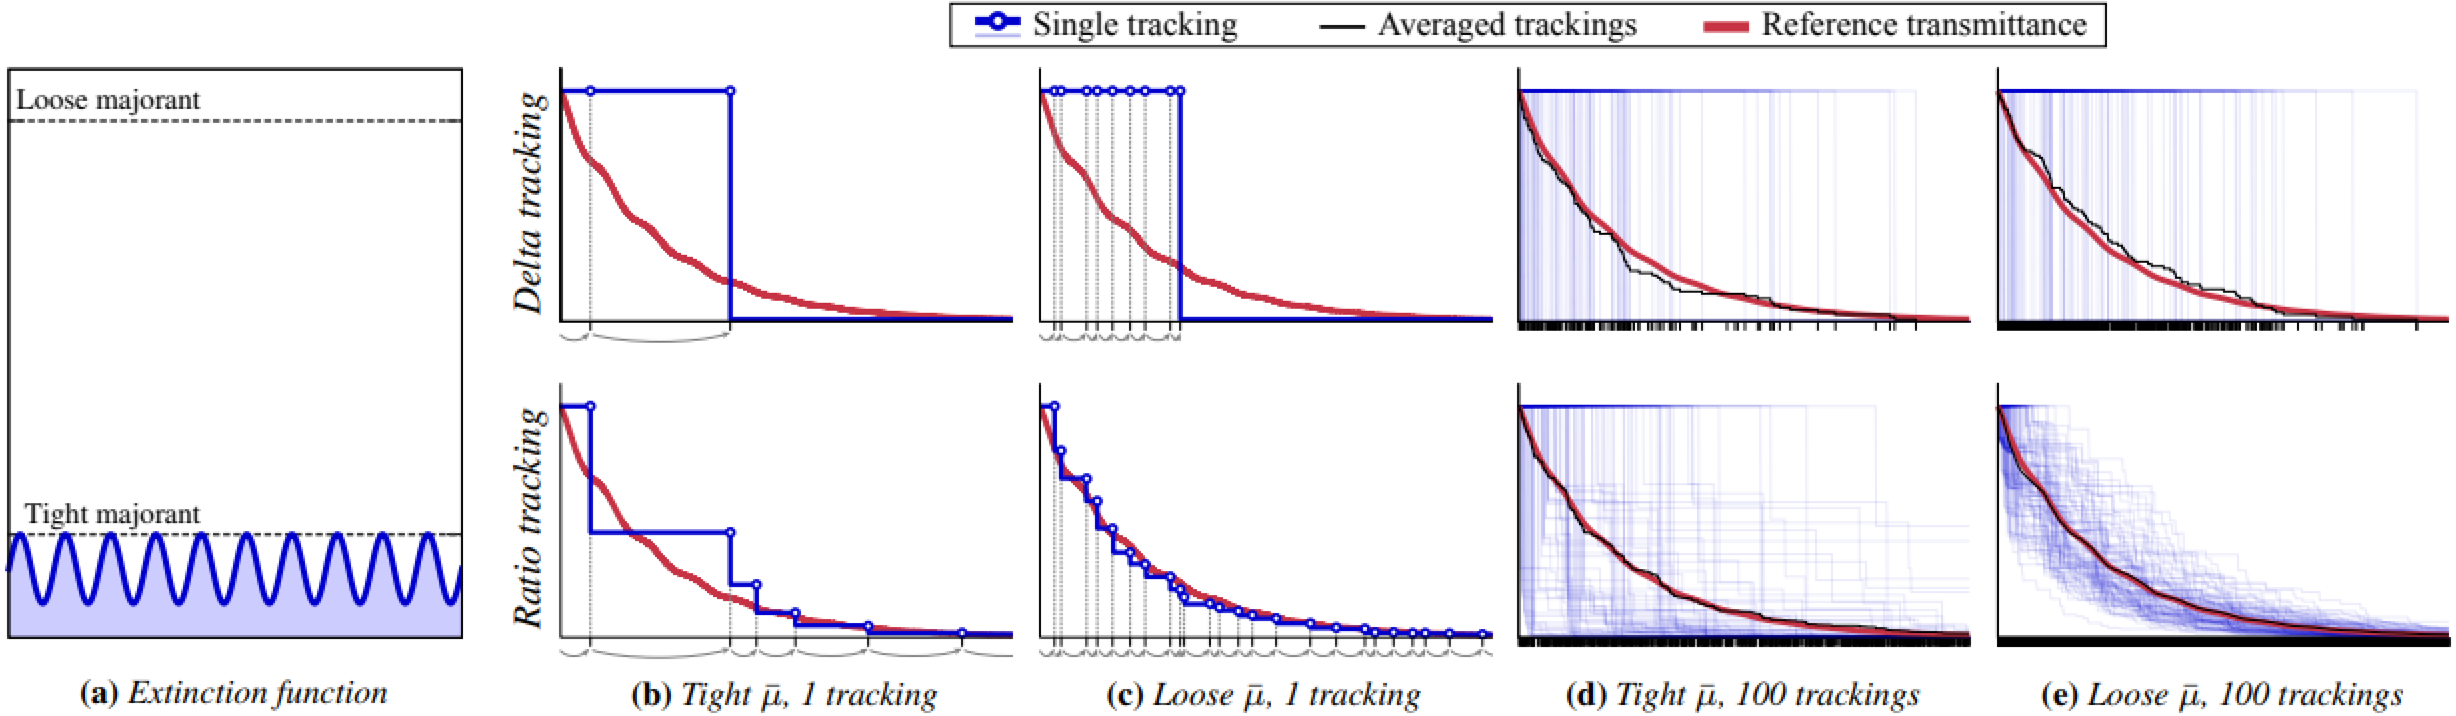

Figure 4: *A comparison of delta and ratio tracking-based estimators in a periodic heterogeneous medium bound by a loose and tight majorant extinction coefficient* (**a**). *In* (**b**, **c**), *we show single instances (blue) of delta (top) and ratio (bottom) trackings for the two different majorants. In* (**d**, **e**), *we show approximations of the transmittance function (black) obtained by averaging 100 instances (thin blue) of the corresponding tracker. Red curves represent the ground truth transmittance function. The ticks on horizontal axes mark all sampled (tentative) collision points.*

Excerpt from [Residual Ratio Tracking for Estimating Attenuation in Participating Media. Novak et al. 2014]

Ray Marching vs Ratio Tracking

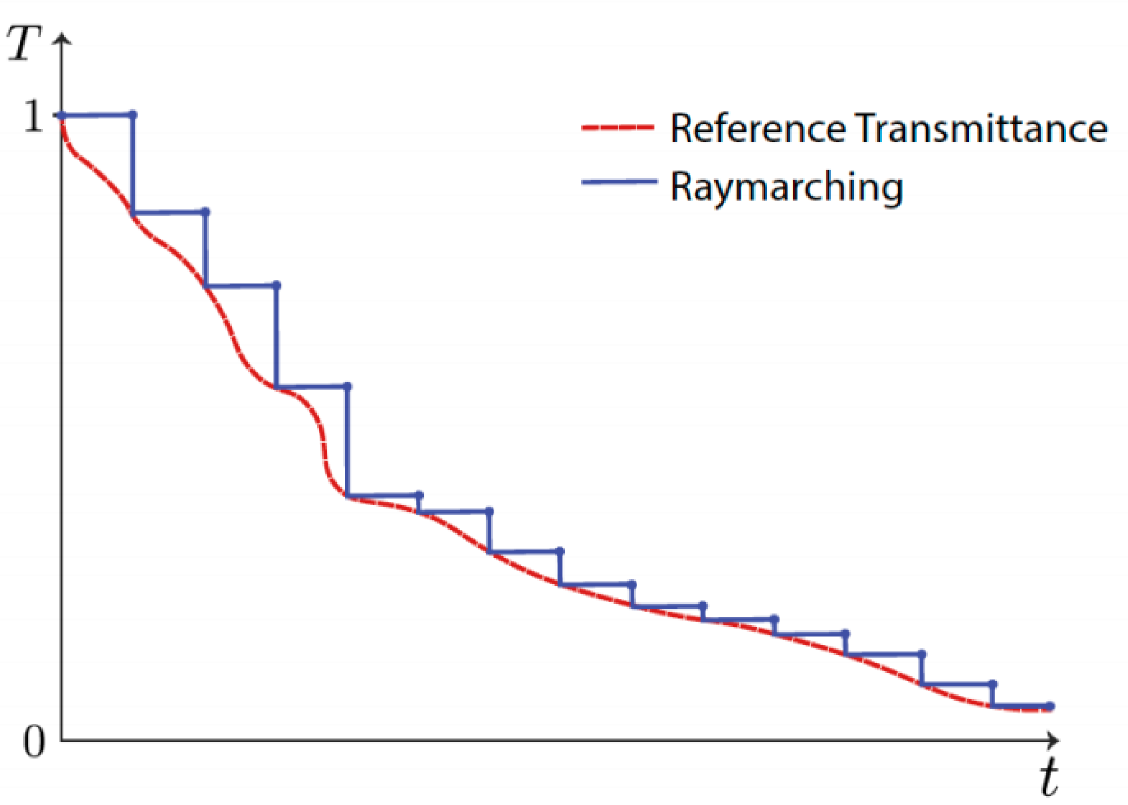

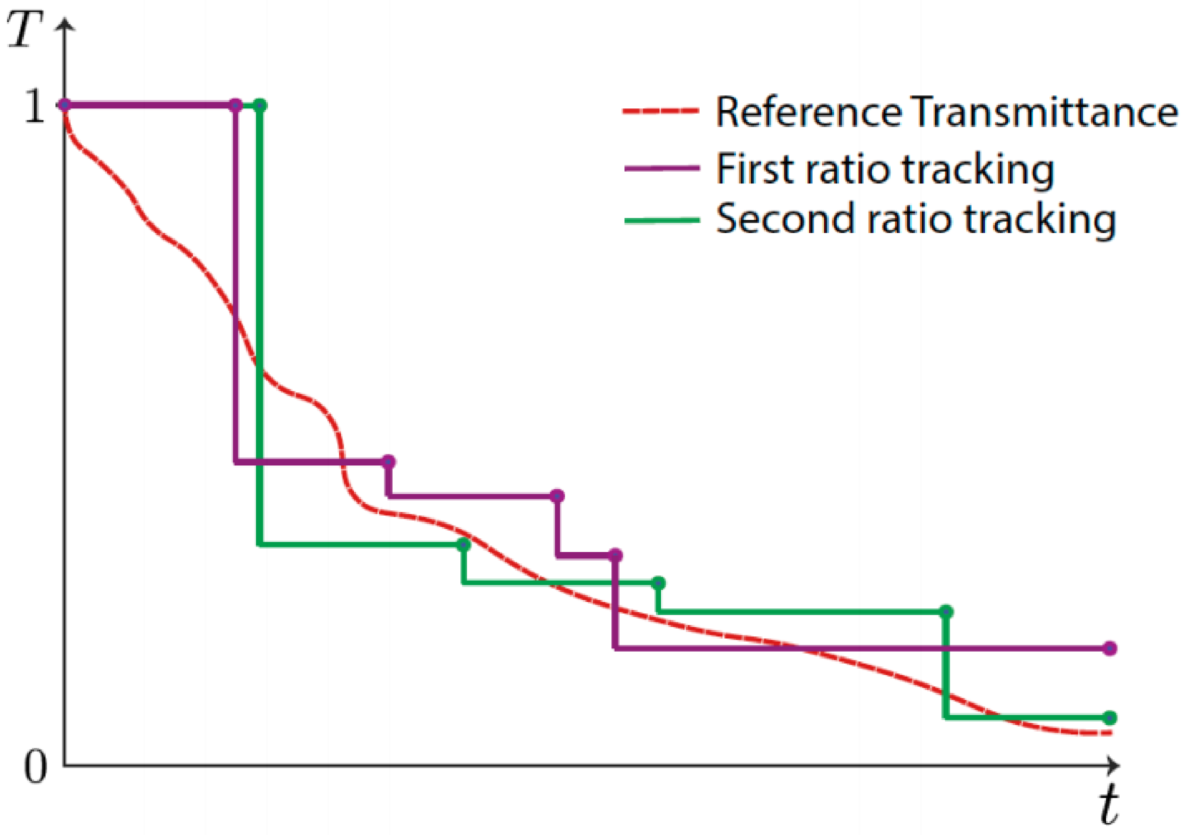

Pseudo-marginal MH: Summary

- MLT with MCMC/MH algorithm from [Veach and Guibas 1997]
  - For each pixel, find several seed paths using traditional path tracing
  - For each seed path, run MH algorithm
- MH algorithm
  - Propose path by $q\left(\bar{x}^{i+1}\middle|\bar{x}^{i}\right)$ from [Pauly et al. 2000], [Jacob 2010]
  - Accept path by $r(\bar{x}^{i}, \bar{x}^{i+1})$
- Need unbiased estimator
  - Build unbiased estimator using unbiased transmittance estimator [this paper] by ratio tracking from [Novak et al. 2014]

Results

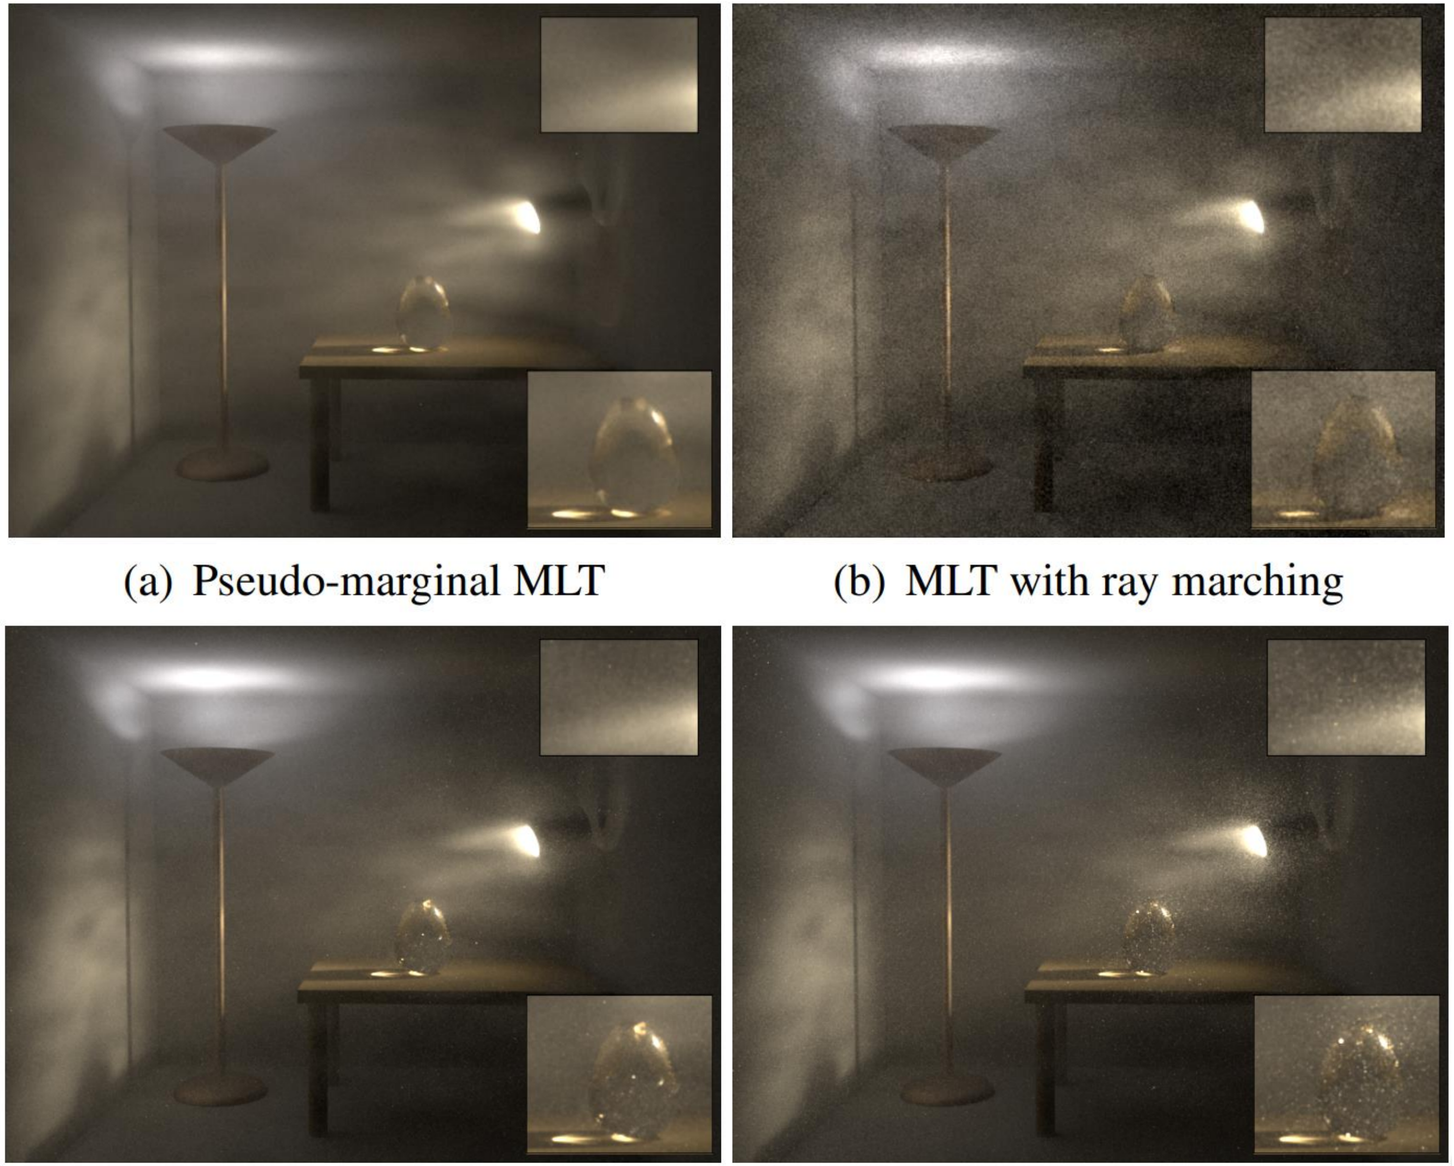

(a) Pseudo-marginal MLT
(b) MLT with ray marching
(c) PSSMLT with ratio tracking
(d) Bidirectional path tracing

Figure 3: *Equal time renderings of the classical scene from Veach filled with anisotropic scattering heterogenous media ($g = 0.85$).*

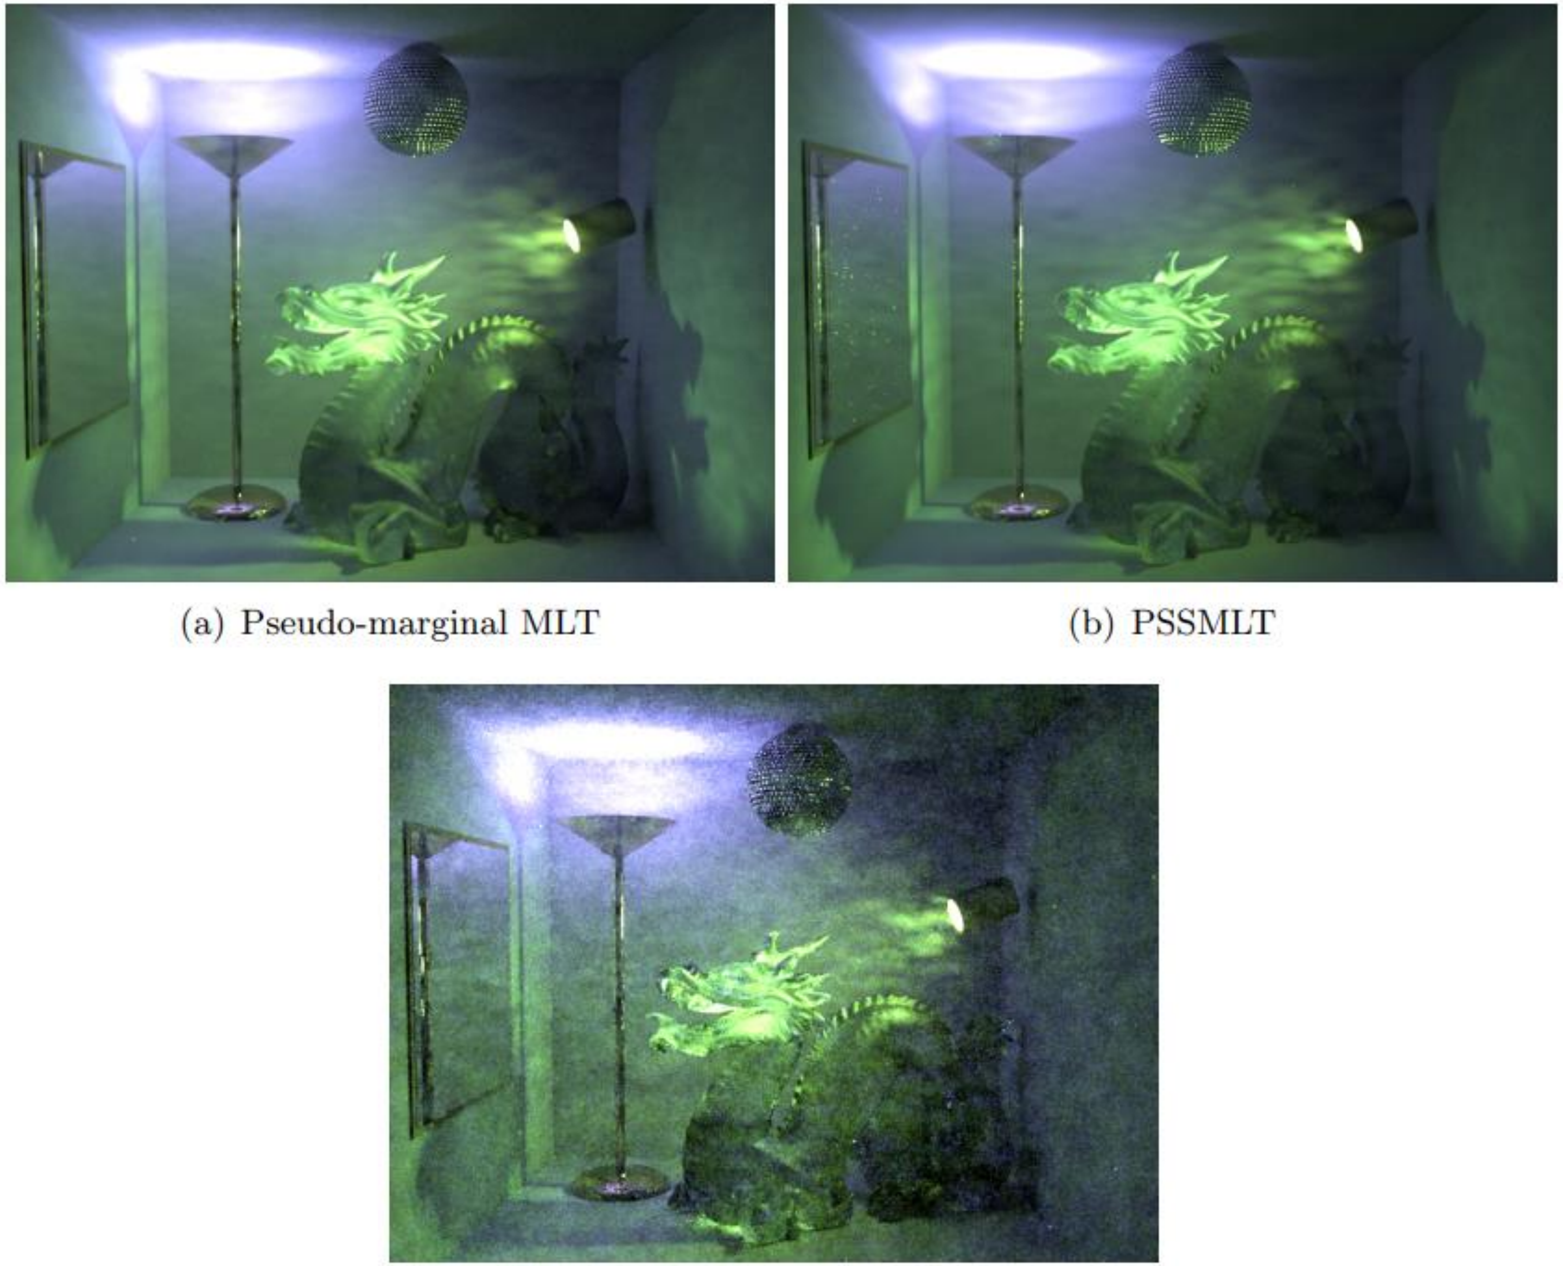

(a) Pseudo-marginal MLT (b) PSSMLT

(c) MLT with ray marching

1.2. **Disco scene.** Figure 2 shows equal time renderings of a scene with heterogenous anisotropic media (resolution= 256^3, $g = 0.85$) and glossy transfer, presenting a very difficult case for previous work.

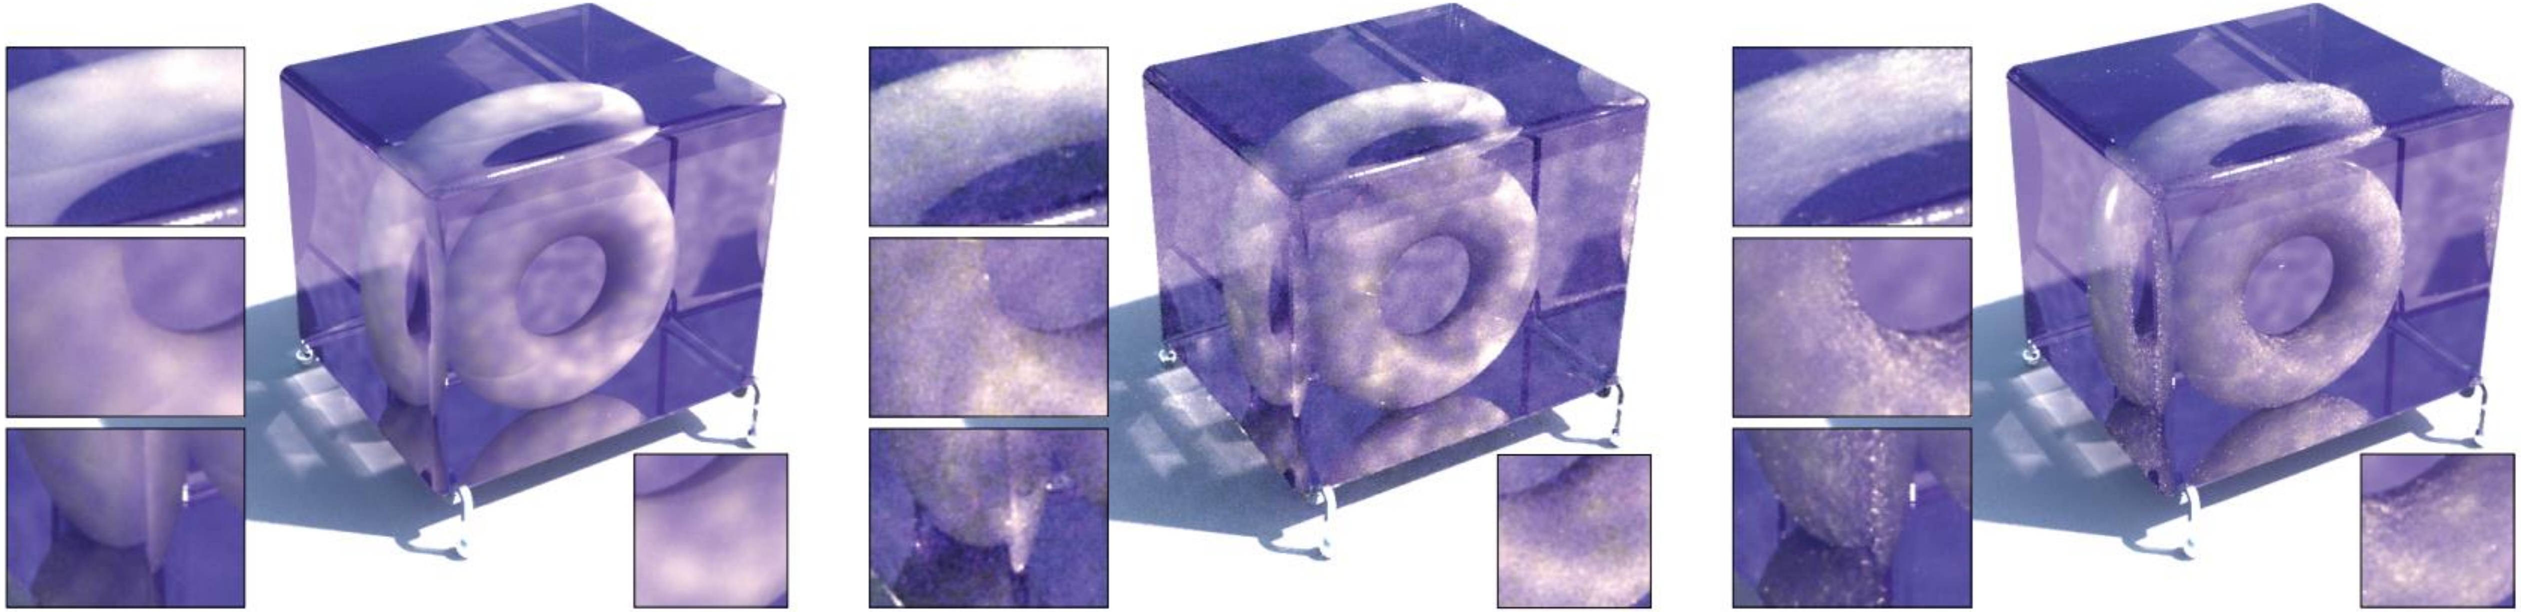

(a) Pseudo-marginal ERPT using ratio tracking (b) ERPT with ray marching (c) PSSMLT with ratio tracking

Figure 1: *Equal time renderings of a refractive glass cube containing isotropic heterogenous participating media and a diffuse torus. a) Our method using ERPT [Cline et al. 2005] with ratio tracking [Novák et al. 2014] to obtain an unbiased estimate of the transmission. b) Same as a) but with deterministic ray-marching used to evaluate the transmission. c) PSSMLT [Kelemen et al. 2002] with ratio tracking.*

Strengths & Weaknesses

Strengths

- Accurate, and also fast in scene with participating media
- Can be extended to other MLT variants (ERPT in example)

Weaknesses

- Only applicable to scenes with participating media
- Rendering parameters should be carefully set