Comparison of null-collision radiance estimators in light transport

Brin Colnar, Matija Marolt, Ciril Bohak, Žiga Lesar

University of Ljubljana, Faculty of Computer and Information Science
E-mail: bc9389@student.uni-lj.si, {matija.marolt, ciril.bohak, ziga.lesar}@fri.uni-lj.si

Abstract

In this paper, we present a comparison of different Monte-Carlo-method-based algorithms for radiance estimation in a heterogeneous medium. The algorithms employ a null-collision approach that acts as fictitious material and homogenizes the medium, supporting the use of closed-form free-path sampling. We give insight into the fundamentals of volumetric light transport simulation, overview the implemented algorithms in a web-based volumetric path tracing framework, and present the evaluation results.

1 Introduction

Monte Carlo methods are widely used to estimate radiative transfer in homogeneous media by approximating the radiative transfer equation by simulating absorption and scattering events. However, in heterogeneous media, the properties like the absorption and scattering coefficients vary spatially, increasing the complexity significantly. To address this, null collisions are used to homogenize the medium and enable closed-form sampling of absorption and scattering events. Galtier et al. [2] have developed an integral formulation that reveals a wide family of null-collision algorithms for estimating the radiance. Kutz et al. [3] propose practical null-collision algorithms for the use in motion picture industry. In this paper, we evaluate the performance of selected algorithms in the context of scientific visualization. Specifically, we evaluate analog delta tracking, weighted delta tracking, and weighted decomposition tracking. Our evaluation will serve as a guide when choosing a suitable algorithm for scientific volume visualization.

2 Radiative transfer equation

To derive the radiative transfer equation (RTE), we must consider how light interacts with matter, which directly influences radiance L. Light can be absorbed or scattered (L_s) when it passes through a participating medium. Additionally, the medium itself can emit light, which we model with a function L_e. We combine these effects in a single differential equation:

$$(\omega \cdot \nabla)L(x,\omega) = -\mu_a(x)L(x,\omega) - \mu_s(x)L(x,\omega) \quad (1)$$
$$+\mu_a(x)L_e(x,\omega) + \mu_s(x)L_s(x,\omega),$$

$$L_s(x,\omega) = \int_{S^2} f_p(x,\omega',\omega)L(x,\omega')\,\mathrm{d}\omega', \quad (2)$$

where μ_a and μ_s are the absorption and scattering coefficients, respectively. Integrating along the direction ω yields the RTE:

$$L(x,\omega) = \quad (3)$$
$$\int_0^\infty T(x,y)\left[\mu_a(y)L_e(y,\omega) + \mu_s(y)L_s(y,\omega)\right]\,\mathrm{d}y,$$

$$T(x,y) = \exp\left(-\int_0^{\|x-y\|}(\mu_a(z)+\mu_s(z))\,\mathrm{d}z\right). \quad (4)$$

3 Algorithms

This section presents delta tracking, weighted delta tracking, and weighted decomposition tracking. The implementation of these algorithms was inspired by Kutz et al. [3], who continue from Galtier et al.'s formulation [2] and produce several sophisticated light estimators.

3.1 Delta tracking

We approximate the integral in Equation (4) by drawing samples of the integrand. By choosing the probability density function of $p(y) \propto T(x,y)$, we can eliminate the evaluation of $T(x,y)$ in the integrand. We achieve this by introducing a fictitious medium that homogenizes the volume without changing the light transport. This medium does not absorb light but only scatters it at a rate $\mu_n(x)$ without changing its direction. We define the majorant as $\mu = \mu_a(x) + \mu_s(x) + \mu_n(x)$, which gives $T(x,y) = \exp(-\int_0^z \mu\,\mathrm{d}z) = \exp(-\mu z)$. We can sample z via the inversion method, $z = -\ln(1-u)/\mu$, where $u \sim U(0,1)$ is uniformly distributed on the unit interval. The one-sample RTE estimator is thus:

$$\frac{\mu_a(y)}{\mu}L_e(y,\omega) + \frac{\mu_s(y)}{\mu}L_s(y,\omega) + \frac{\mu_n(y)}{\mu}L(y,\omega). \quad (5)$$

Delta tracking adds a probabilistic evaluation of the above terms so that only one of them can be evaluated for each sample. This leads to Algorithm 1.

Algorithm 1 Delta tracking

while true **do**
  $t \leftarrow -\frac{\ln(1-u)}{\mu}$
  $x \leftarrow x - t\omega$
  if $u < \frac{\mu_a(x)}{\mu}$ **then**
    return $L_e(x, \omega)$
  else if $u < \frac{\mu_a(x)}{\mu} + \frac{\mu_s(x)}{\mu}$ **then**
    $\omega \leftarrow \text{sample} \propto f_p(x, \omega', \omega)$
  end if
end while

Delta tracking requires a bounding majorant μ for the absorption and scattering probabilities to be physically plausible, which is a significant limitation.

3.2 Weighted delta tracking

Weighted delta tracking mitigates the requirement of finding a bounding majorant μ by redefining the probabilities for evaluating the integrand terms. More precisely, we need to redefine the probabilities in the case when $\mu_n(y) < 0$. Probabilistic evaluation $\langle f(y) \rangle_P$ is a random variable that returns $f(y)/P$ with probability P and 0 otherwise. Note that its expected value is $f(y)$. We can upgrade our one-sample RTE estimator:

$$\left\langle \frac{\mu_a(y)}{\mu} L_e(y, \omega) \right\rangle_{P_a} + \left\langle \frac{\mu_s(y)}{\mu} L_s(y, \omega) \right\rangle_{P_s} + \left\langle \frac{\mu_n(y)}{\mu} L(y, \omega) \right\rangle_{P_n}, \quad (6)$$

where P_a, P_s, and P_n are absorption, scattering, and null-collision probabilities. They can be arbitrary, gracefully handling the case when $\mu_n(y) < 0$. The pseudo-code is shown in Algorithm 2.

Algorithm 2 Weighted delta tracking

$w \leftarrow 1$
while true **do**
  $t \leftarrow -\frac{\ln(1-u)}{\mu}$
  $x \leftarrow x - t\omega$
  if $u < P_a(x)$ **then**
    return $w \times \frac{\mu_a(x)}{\mu} L_e(x, \omega)$
  else if $u < P_a(x) + P_s(x)$ **then**
    $\omega \leftarrow \text{sample} \propto f_p(x, \omega', \omega)$
    $w \leftarrow w \times \frac{\mu_s(x)}{\mu(x)}$
  else
    $w \leftarrow w \times \frac{\mu_n(x)}{\mu(x)}$
  end if
end while

In our implementation in volumetric path tracing framework (VPT) [4], we added a scalar b such that $\mu = b \cdot \max \mu_a + \mu_s$. If $b < 1$, $\mu_n(x)$ can become negative, which can help performance. If $b > 1$, the volume is oversampled, and performance deteriorates. Figure 1 displays our weighted delta tracking implementation results at different values of b. Note the increased variance in Figure 1a. Figure 2 shows another volume in which the true majorant is significantly lower, leading to faster performance when setting $b < 1$.

3.3 Decomposition tracking

Decomposition tracking is an approach that decomposes the original volume into two superimposed sub-volumes: the homogeneous control volume and the heterogeneous residual volume. Each sub-volume has its own absorption and scattering coefficients, listed in Table 1.

Table 1: Coefficients of control and residual volumes.

$\mu_t^c = \mu_a^c + \mu_s^c$	extinction coef. of control volume
$\mu_a^r(x) = \mu_a(x) - \mu_a^c$	absorption coef. of residual volume
$\mu_s^r(x) = \mu_s(x) - \mu_s^c$	scattering coef. of residual volume
$\mu_t^r(x) = \mu_t(x) - \mu_t^c$	extinction coef. of residual volume
$\mu_n(x) = \mu - \mu_t^r(x) - \mu_t^c(x)$	null-collision coefficient

The motivation behind decomposition is that we can reduce the expensive memory look-ups as it allows some events to be evaluated analytically using the control component. Kutz et al. [3] describe two variants of decomposition tracking: analog and weighted decomposition tracking. A downside of analog decomposition is that it requires both the minorant and the majorant to be bounding. Finding these bounds is challenging from a computational standpoint. Weighted decomposition allows non-bounding coefficients to be used. We have implemented the weighted variant. The pseudo-code for weighted decomposition tracking is shown in Algorithm 3.

Algorithm 3 Weighted decomposition tracking

$w \leftarrow 1$
while true **do**
  $t \leftarrow -\frac{\ln(1-\xi)}{\mu(x)}$
  $x \leftarrow x - t\omega$
  $F \leftarrow 0$
  if $u < (F \leftarrow F + P_a^c(x))$ **then**
    return $w \times \frac{\mu_a(x)}{P_a^c(x)\mu} L_e(x, \omega)$
  else if $u < (F \leftarrow F + P_s^c(x))$ **then**
    $\omega \leftarrow \text{sample} \propto f_p(x, \omega', \omega)$
    $w \leftarrow w \times \frac{\mu_s(x)}{P_s^c(x)\mu}$
  else if $u < (F \leftarrow F + P_a^r(x))$ **then**
    return $w \times \frac{\mu_a^r(x)}{P_a^r(x)\mu} L_e(x, \omega)$
  else if $u < (F \leftarrow F + P_s^r(x))$ **then**
    $\omega \leftarrow \text{sample} \propto f_p(x, \omega', \omega)$
    $w \leftarrow w \times \frac{\mu_s^r(x)}{P_s^r(x)\mu}$
  else
    $w \leftarrow w \times \frac{\mu_n(x)}{P_n\mu}$
  end if
end while

A new uniform m, which is also settable in VPT's user interface, allows users to set the fraction of true extinction that will become the minorant. Figure 3 shows 3 decomposition tracking results for different values of m, with $b = 1$.

4 Results

To evaluate the implemented algorithms, we calculated root mean square error (RMSE), which measures the noise

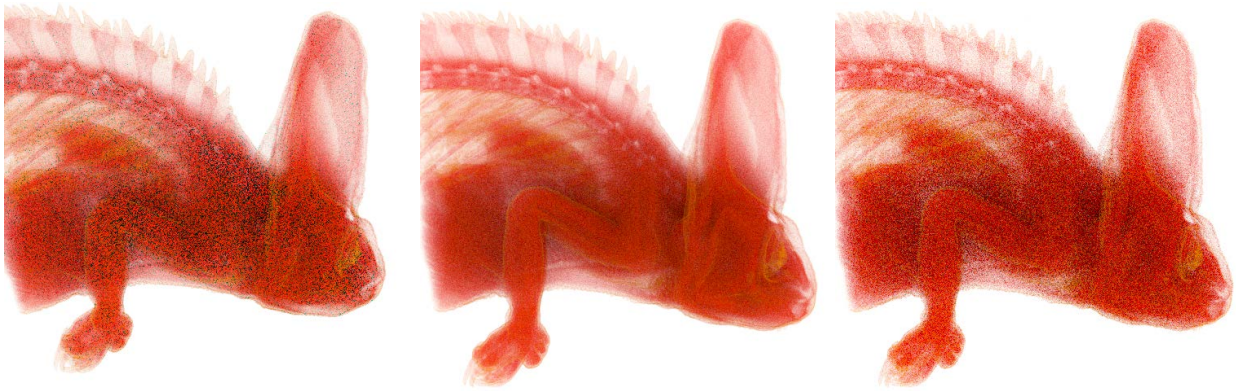

(a) $b = 0.63$. (b) $b = 1$. (c) $b = 10$.

Figure 1: Three different values of parameter b in weighted delta tracking.

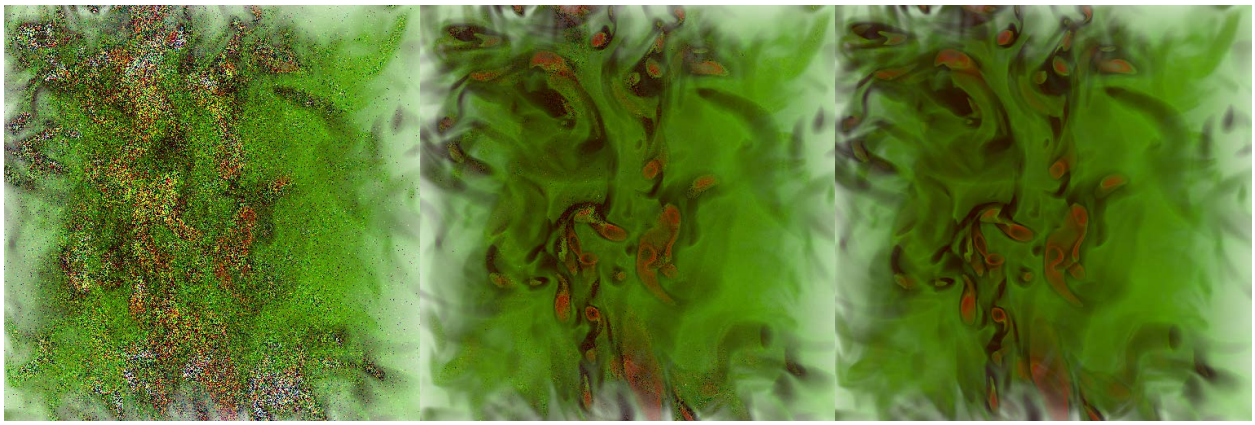

(a) $b = 0.25$. (b) $b = 0.50$. (c) $b = 0.75$.

Figure 2: Three different values of parameter b for weighted delta tracking.

of the given image compared to the converged delta tracking result.

We observe the evolution of RMSE over a 10s time frame. We also compare the time to unit variance (TTUV) values of the converged images. TTUV is defined as the product of render time and RMSE, combining the differences in variance and in computational complexity in a single metric.

4.1 Weighted-delta tracking

First, we compared weighted delta tracking using the values of $b \in \{0.25, 0.5, 0.75, 1.0\}$. Figure 4 shows the evolution of RMSE with respect to the render time. The best RMSE is obtained for $b = 1$ as this is the true majorant. For smaller values of b the implementation achieves larger RMSE, as the variance significantly increases when sampling in the regions where $\mu_n(x) < 0$, which causes the weight w to alternate in sign.

We get a similar insight when analyzing TTUV. Figure 5 shows TTUV values of converged images (taken after 300s) for the given values of b. Similar values are achieved for $b = 1$ and $b = 0.75$, indicating that $\mu_n(x) < 0$ only in small parts of the volume. The increased sampling efficiency counteracts the increase in variance.

4.2 Decomposition tracking

For the decomposition tracking, we observe how different values of m influence RMSE. Figure 5 shows RMSE with respect to $m \in \{0.01, 0.02, 0.03, 0.04, 0.05\}$. In all cases, we have set $b = 1$ (true majorant) so that we could concentrate solely on the influence of m. Figure 5 reveals that TTUV increases with respect to m. This is due to a similar reason as in weighted delta tracking, since increasing m above the true minorant significantly increases the variance due to the alternating sign of w. In this volume, the region where the minorant is non-bounding is large, and the increase in sampling efficiency does not counteract the increase in variance.

4.3 Final comparison

Finally, Table 2 displays TTUV values, comparing the best weighted delta tracking result ($b = 0.75$) with the best weighted decomposition tracking result ($m = 0.01$). The render time was 10s.

5 Conclusion

We evaluated two advanced Monte Carlo methods for volume rendering: weighted delta tracking and weighted

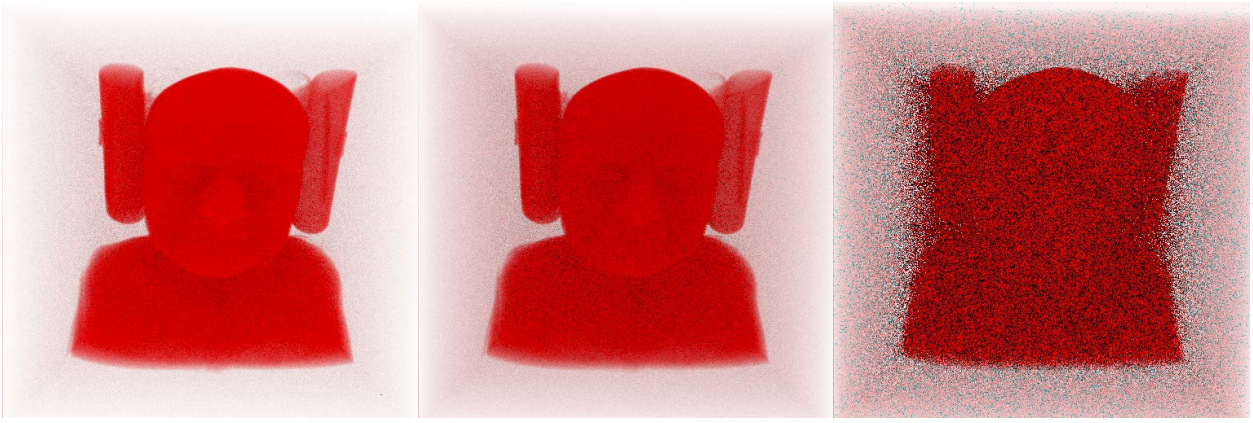

(a) $m = 0.01$. (b) $m = 0.02$. (c) $m = 0.08$.

Figure 3: Three different values of parameter m in weighted decomposition tracking.

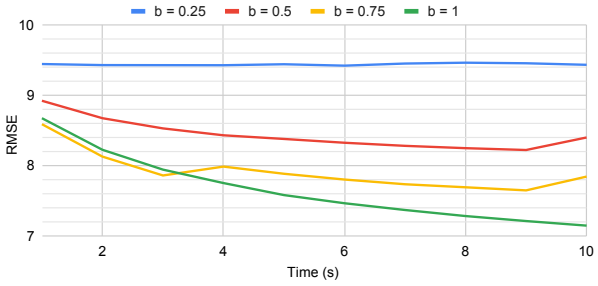

Figure 4: Comparison of RMSE values for different values of parameter b in weighted delta tracking.

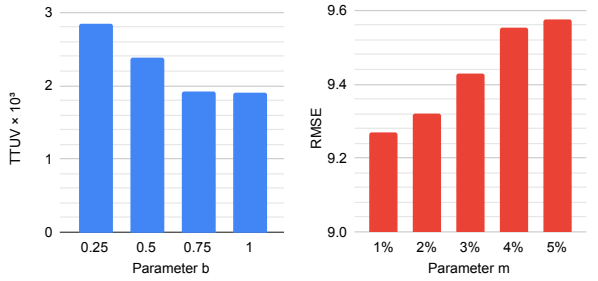

Figure 5: Left: comparison of TTUV values for different values of parameter b in weighted delta tracking. Right: comparison of RMSE values for different values of parameter m in weighted decomposition tracking.

Table 2: TTUV comparison.

Method	Parameter	TTUV
Weighted delta tracking	$b = 0.75$	78.4
Decomposition tracking	$m = 0.01$	95.8

decomposition tracking. Weighted delta tracking exhibited superior performance when the majorant was set close to the true majorant, while weighted decomposition tracking excelled at minorant values close to the true minorant. Although both methods achieved worse quality with non-bounding parameters than with bounding parameters, their advantage lies in the fact that it is not necessary to find the true bounds ahead of time. A possible future extension of the algorithms could include adjusting the parameters while rendering to dynamically optimize convergence.

The source code of this project is available on GitHub [1].

Acknowlegment

This research was conducted as part of the basic research project *Cell visualization of unified microscopic data and procedurally generated sub-cellular structures* [project number J2-50221], funded by the Slovenian Research and Innovation Agency (Javna agencija za znanstvenoraziskovalno in inovacijsko dejavnost RS) from the state budget.

References

[1] Brin Colnar, Matija Marolt, Ciril Bohak, and Žiga Lesar. Comparison of transmittance and distance estimators in light transport. `https://github.com/UL-FRI-LGM/vpt-galtier`. [Online; accessed 8-July-2024].

[2] Mathieu Galtier, Stéphane Blanco, Cyril Caliot, Christophe Coustet, Jérémi Dauchet, Mouna El Hafi, Vincent Eymet, Richard Fournier, Jacques Gautrais, Anaïs Khuong, Benjamin Piaud, and Guillaume Terrée. Integral formulation of null-collision monte carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 125:57–68, 8 2013.

[3] Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Transactions on Graphics*, 36:1–16, 8 2017.

[4] Žiga Lesar, Ciril Bohak, and Matija Marolt. Real-time interactive platform-agnostic volumetric path tracing in webgl 2.0. In *Proceedings of the 23rd International ACM Conference on 3D Web Technology*, New York, NY, USA, 2018. Association for Computing Machinery.