Volumetric Multi-View Rendering – Supplemental document

Basile Fraboni[1,2], Antoine Webanck[2], Nicolas Bonneel[2], and Jean-Claude Iehl[2]

[1]INSA Lyon, CNRS/LIRIS, France [2]Univ Lyon, CNRS/LIRIS, France

Contents

PART ONE

Volumetric shift mappings

Shifting medium interactions consists in transforming the depth t of an interaction sampled using *delta tracking* [NGHJ18, MGJ19] on a base segment that traverses one or several volume sections, to another valid depth location t' on a target segment also traversing volume sections as illustrated in Figure 1. A shift location should ensure that the pdf of the shifted interaction is well defined to be valid – i.e. inside a medium. Note that we discuss shift mappings of null interactions in this document, as the main article application requires null shifts only, but the technique is applicable for scattering and absorption interactions.

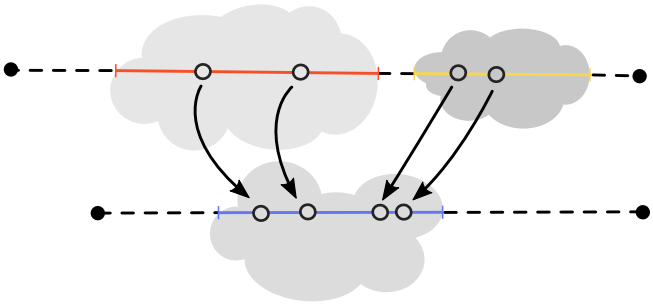

Figure 1: *Volumetric chains shift mapping configuration.*

1. Copy-based shift mappings

We first present copy-based transformations that have been proposed and used multiple times in previous work to shift interactions, such as re-using the same random numbers that generated the base interaction depth. These shift mappings however lead to some failure cases that are depicted in Figure 2.

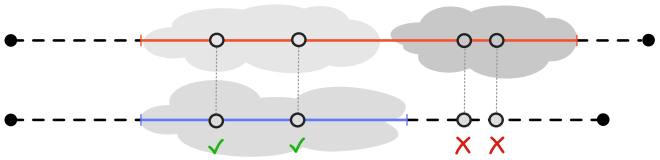

Figure 2: *Volumetric chains shift mapping failures. Given a base path segment (top) composed of several null interactions, the transformed interactions on the target segment (bottom) are constructed by reusing the same free-flight distances – i.e. raw depths. This results in a family of invalid shifts, when shifted interactions fall outside of the medium sections or the segment. Not only depth copy but primary sample copy and optical depth copy shift mapping suffer from the same failure cases.*

1.1. Raw depth copy

First consider the naive approach of replicating the exact free-flight distances t that has been sampled on the base segment for a null interaction x. We hence get the following mapping with its associated Jacobian:

$$t' = S(t) = t \qquad \left|\frac{\mathrm{d}S(t)}{\mathrm{d}t}\right| = |S'(t)| = 1 \tag{1}$$

The pdf of a shifted interaction writes:

$$p(x') = p(t')P_n(x') = p(t)\left|S'(t)\right|^{-1} P_n(x') = p(t)P_n(x') \tag{2}$$

where $P_n(x)$ is the probability of x being a null interaction [NGHJ18]. However, by reusing the same raw sample depth some of the transformed volume interactions may even end up either outside of the medium section or farther than the segment endpoints. This approach has been proposed, and successful, in gradient domain methods [GHV*18], though authors mention that some interactions may be invalid. But we aim at reusing on segments that can be geometrically very different (e.g. two disjoint points of view observing the same pivot), hence this simple shift mapping may lead to a whole family of invalid shifted segment, as illustrated in Figure 2. On the contrary, we want to enforce transformed null scattering interactions to fall inside medium sections, to increase our chances to find a valid shifted segment.

Table 1: *Shift mappings comparison*

Shifts	Copy	Linear Stretch
Primary Sample	$u' = u \Leftrightarrow t' = \frac{\bar{\mu}}{\bar{\mu}'} t$ $\left\lvert \frac{\partial t'}{\partial t} \right\rvert = \frac{\bar{\mu}}{\bar{\mu}'}$	$u' = C \cdot u \Leftrightarrow t' = \frac{-\log(1-C+C\exp(-\bar{\tau}(t)))}{\bar{\mu}'} + R$ $\left\lvert \frac{\partial t'}{\partial t} \right\rvert = \frac{\bar{\mu} \cdot C}{\bar{\mu}'(C-(C-1)e^{\bar{\mu} \cdot t})}$ $C = \frac{1-\bar{T}(t'_{max})}{1-\bar{T}(t_{max})}$
Raw Depth	$t' = t$ $\left\lvert \frac{\partial t'}{\partial t} \right\rvert = 1$	$t' = C \cdot t$ $\left\lvert \frac{\partial t'}{\partial t} \right\rvert = C$ $C = \frac{t'_{\max}}{t_{\max}}$
Majorant Optical Depth	$\bar{\tau}' = \bar{\tau} \Leftrightarrow t' = \frac{\bar{\mu}}{\bar{\mu}'} t + R$ $\left\lvert \frac{\partial t'}{\partial t} \right\rvert = \frac{\bar{\mu}}{\bar{\mu}'}$	$\bar{\tau}' = C \cdot \bar{\tau} \Leftrightarrow t' = C \cdot \frac{\bar{\mu}}{\bar{\mu}'} t + R$ $\left\lvert \frac{\partial t'}{\partial t} \right\rvert = C \cdot \frac{\bar{\mu}}{\bar{\mu}'}$ $C = \frac{\bar{\tau}'_{\max}}{\bar{\tau}_{\max}}$

1.2. Primary sample copy

Another common approach is random sequence replay [KSKAC02, BJNJ17] – i.e. reusing the uniform random numbers that generated the base interaction. Hence, each shifted interaction x' is constructed using the random sample $u' = u$, that has been used to generate the base interaction x. Using the free-flight sampling construction (delta tracking) that samples an homogenized medium of constant majorant (denoted $\bar{\mu}$ on the base segment and $\bar{\mu}'$ on the target segment), the shift mapping is computed as follows:

$$t' = -\frac{\log(1-u')}{\bar{\mu}'} = -\frac{\log(1-u)}{\bar{\mu}'} = -\frac{\bar{\mu}}{\bar{\mu}'}\frac{\log(1-u)}{\bar{\mu}} = \frac{\bar{\mu}}{\bar{\mu}'} t = S(t) \qquad \left\lvert \frac{\mathrm{d}S(t)}{\mathrm{d}t} \right\rvert = |S'(t)| = \frac{\bar{\mu}}{\bar{\mu}'} \tag{3}$$

The shifted depth is thus a stretched version of the base depth, and the interaction PDF reads:

$$p(x') = p(t')P_n(x') = p(t)\left|S'(t)\right|^{-1} P_n(x') = p(t)\frac{\bar{\mu}'}{\bar{\mu}} P_n(x') \tag{4}$$

Again, this way of transforming null interactions may result in invalid shifted interactions, for example when the medium majorant extinction coefficients are equals, $\bar{\mu}' = \bar{\mu} \Leftrightarrow \frac{\bar{\mu}}{\bar{\mu}'} = 1$, we fall back to the invalid case depicted in Figure 2. Another drawback of both depth and PSS copy approaches is there is no obvious way to shift interactions in presence of several disjoint medium sections easily.

1.3. Majorant optical depth copy

A last solution to consider in the family of copy-based shifts, is the majorant optical depth copy $\bar{\tau}'(t') = \bar{\tau}(t)$, where the majorant optical depth of an interaction is copied to find the shifted interaction. We will not detail the complete derivation, but again, this scheme can lead to invalid shifts since shifted interactions are not bounded inside the medium sections on the target segment. For example using the configuration depicted in Figure 2 and given equal majorants in base and target volumes, $\bar{\mu} = \bar{\mu}' = 1$, the majorant optical depth copy is equivalent to raw depth copy and results in the same failure case that precludes the construction of a whole family of shifts.

2. Stretch-based shift mappings

Alternatively, we present three different new shift mappings, summarized in Table 1, that transform base interactions, such that the resulting shifted interactions fall within the medium sections of the target segment. We construct shifted interactions by stretching the cumulated raw depth, the majorant optical depth and the primary samples, by deriving their correct stretching constants and their respective Jacobian determinants.

2.1. Raw depth linear stretch

A simple solution to shift from the base segment to the other by ensuring the shifted interactions to fall inside the target medium sections is to reuse the normalized cumulated distance of the base interactions. This solution corresponds to a linear stretch of the depth traversed in media.

Normalized cumulated distance. We define the *cumulated depth of traversed medium* of an interaction at distance t as follows:

$$l(t) = \int_0^t V(x)\mathrm{d}x \tag{5}$$

where $V(x)$ is an indicator function that equals one if x is within a medium zero otherwise. Note that the cumulated depth can be obtained by summing over each volume interval until reaching the distance t:

$$l(t) = \sum_{j=1}^{n_t} \left(\min\left(t, t_{\max,j}\right) - t_{\min,j}\right) = t - t_{\min,n_t} + \sum_{j=1}^{n_t-1} \left(t_{\max,j} - t_{\min,j}\right) \tag{6}$$

where n_t is the volume section index containing depth t on the base segment, and $t_{\min,j}$ and $t_{\max,j}$ are the entry and the exit of the j-eth volume section of the base segment. Dividing $l(t)$ by the total cumulated depth of traversed medium $l_{\max}$ on a segment gives the *normalized cumulated raw depth*.

Shift formulation. The mapping is then defined by copying the normalized cumulated raw depth of the base interactions:

$$\frac{l'(t')}{l'_{\max}} = \frac{l(t)}{l_{\max}} \Leftrightarrow \quad l'(t') = \frac{l'_{\max}}{l_{\max}} \cdot l(t) \Leftrightarrow \quad l'(t') = C \cdot l(t) \tag{7}$$

and $C = \frac{l'_{\max}}{l_{\max}}$ is a constant defined by the base and the target segment total volume sections cumulated lengths. Plugging the piecewise constant definition of the cumulated depth (Equation 6) into the latter equation leads to the expression of the shifted interaction depth t' as a function of the base interaction depth t:

$$\begin{aligned}
& l'(t') = C \cdot l(t) \\
\Leftrightarrow \quad & \sum_{i=1}^{n_{t'}} \left(\min\left(t', t'_{\max,i}\right) - t'_{\min,i}\right) = C \cdot \sum_{j=1}^{n_t} \left(\min\left(t, t_{\max,j}\right) - t_{\min,j}\right) \\
\Leftrightarrow \quad & t' + a_1 = C \cdot [t + a_2] \\
\Leftrightarrow \quad & t' = C \cdot t + R \\
\Leftrightarrow \quad & t' = S_{\mathrm{raw}}(t)
\end{aligned} \tag{8}$$

where $a_1 = \sum_{i=1}^{n_{t'}-1} \left(t'_{\max,i} - t'_{\min,i}\right) - t'_{\min,n}$, $a_2 = \sum_{j=1}^{n_t-1} \left(t_{\max,j} - t_{\min,j}\right) - t_{\min,m}$ and $R = C \cdot a_2 - a_1$. Applying the transformation for each base null interaction completes the construction of the shifted segment.

Jacobian determinant and PDF. The Jacobian of this mapping is straightforward to compute and is constant – i.e. global – on the whole segment:

$$\left|\frac{\mathrm{d}S_{\mathrm{raw}}(t)}{\mathrm{d}t}\right| = \left|S'_{\mathrm{raw}}(t)\right| = C = \frac{l'_{\max}}{l_{\max}} \tag{9}$$

Finally the PDF of a transformed interaction can be expressed as:

$$p(x') = p(t')P_n(x') = p(t)\left|S'_{\mathrm{raw}}(t)\right|^{-1} P_n(x') = p(t)\frac{l_{\max}}{l'_{\max}} P_n(x') \tag{10}$$

Discussion. In the presence of one medium section on the base and the target segment, this mapping corresponds to a linear stretch of the raw depth. This shift mapping allows a simple construction of the transformed interactions and a simple Jacobian evaluation that is global on the whole segment. However the resulting distribution of null interactions does not account for the properties of the traversed medium sections – e.g. the majorant extinction coefficients of each interval, or the traversed transmittance. We present next two shift mappings that better account for volume properties to better distribute the shifted interactions.

2.2. Majorant optical depth linear stretch

Another solution is to take into account the majorant optical depth $\bar{\tau}$ of each null interaction in the volume sections, to accurately stretch the interactions positions w.r.t the volume density.

Majorant optical depth. The majorant optical depth is obtained by integrating the majorant extinction coefficient from the medium entry to the interaction position:

$$\bar{\tau}(t) = \int_0^t \bar{\mu}(x)\mathrm{d}x \tag{11}$$

As the majorant coefficient is constant over medium sections and zero outside, the majorant optical depth is piecewise linear. It thus can be obtained by summing over the medium sections crossed along the prefix until reaching the interaction at depth t:

$$\begin{aligned}\bar{\tau}(t) &= \sum_j^{n_t} \bar{\mu}_j \left(\min\left(t, t_{\max,j}\right) - t_{\min,j}\right) \\ &= \bar{\mu}_{n_t}\left(t - t_{\min,n_t}\right) + \sum_j^{n_t-1} \bar{\mu}_j \left(t_{\max,j} - t_{\min,j}\right)\end{aligned} \tag{12}$$

where $t_{\min,j}$ and $t_{\max,j}$ are the entry and exit depths of the j-th volume section and n_t is the index of the volume section containing depth t. In case of overlapping media, we replace the unique majorant of a section μ_j by the sum of majorants in overlapping sections $\sum_k \mu_{k,j}$. Dividing $\bar{\tau}(t)$ by the total majorant optical depth $\bar{\tau}_{\max}$ on a segment gives the *normalized majorant optical depth*.

Shift formulation. The mapping is then defined by copying the normalized majorant optical depth of the base interactions:

$$\frac{\bar{\tau}'(t')}{\bar{\tau}'_{\max}} = \frac{\bar{\tau}(t)}{\bar{\tau}_{\max}} \Leftrightarrow \quad \bar{\tau}'(t') = \frac{\bar{\tau}'_{\max}}{\bar{\tau}_{\max}}\bar{\tau}(t) \Leftrightarrow \quad \bar{\tau}'(t') = C \cdot \bar{\tau}(t) \tag{13}$$

where $C = \frac{\bar{\tau}'_{\max}}{\bar{\tau}_{\max}}$ is a constant defined by the base and the target segment volume sections total majorant optical depth. Plugging the piecewise constant definition of the majorant optical depth (Equation 12) into the latter equation leads to the expression of the shifted interaction depth t' as a function of the base interaction depth t:

$$\begin{aligned}
& \bar{\tau}'(t') = C \cdot \bar{\tau}(t) \\
\Leftrightarrow \quad & \sum_{i=1}^{n_{t'}} \bar{\mu}'_i \left(\min\left(t', t'_{\max,i}\right) - t'_{\min,i}\right) = C \cdot \sum_{j=1}^{n_t} \bar{\mu}_j \left(\min\left(t, t_{\max,j}\right) - t_{\min,j}\right) \\
\Leftrightarrow \quad & \bar{\mu}'_{n_{t'}} \cdot t' + a_1 = C \cdot [\bar{\mu}_{n_t} \cdot t + a_2] \\
\Leftrightarrow \quad & t' = C \cdot \frac{\bar{\mu}_{n_t}}{\bar{\mu}'_{n_{t'}}} \cdot t + R \\
\Leftrightarrow \quad & t' = S_{\mathrm{mod}}(t)
\end{aligned} \tag{14}$$

where $a_1 = \sum_{i=1}^{n_{t'}-1} \bar{\mu}'_i \left(t'_{\max,i} - t'_{\min,i}\right) - \bar{\mu}'_{n_{t'}} \cdot t'_{\min,n_{t'}}$, $a_2 = \sum_{j=1}^{n_t-1} \bar{\mu}_j \left(t_{\max,j} - t_{\min,j}\right) - \bar{\mu}_{n_t} \cdot t_{\min,n_t}$ and $R = \frac{C \cdot a_2 - a_1}{\bar{\mu}'_n}$. Applying the transformation for each base null interaction completes the construction of the shifted segment.

Jacobian determinant and PDF. The Jacobian determinant of the complete transformation is:

$$\left|\frac{\mathrm{d}S_{\mathrm{mod}}(t)}{\mathrm{d}t}\right| = \left|S'_{\mathrm{mod}}(t)\right| = C \cdot \frac{\bar{\mu}_{n_t}}{\bar{\mu}'_{n_{t'}}} = \frac{\tau'_{\max}}{\tau_{\max}} \cdot \frac{\bar{\mu}_{n_t}}{\bar{\mu}'_{n_{t'}}} \tag{15}$$

Finally the PDF of a transformed interaction reads:

$$p(x'_i) = p(t_i)\left|S'_{\mathrm{mod}}(t)\right|^{-1} P_n(x'_i) = p(t_i)\frac{\tau_{\max}}{\tau'_{\max}} \cdot \frac{\bar{\mu}'_{n_{t'}}}{\bar{\mu}_{n_t}} P_n(x'_i) \tag{16}$$

Discussion. In the presence of one medium section on the base and the target segment, this mapping corresponds to a linear stretch of the majorant optical depth. Note that in contrary to the raw depth mapping, the Jacobian determinant is local, i.e. it may vary for each interaction due to the presence of the base section majorant $\bar{\mu}_{n_t}$, associated with the base interaction at distance t, and the target section majorant $\bar{\mu}'_{n_{t'}}$ associated with the shifted interaction at distance t'.

2.3. Primary sample linear stretch

The common free-flight sampling approach is unbounded, hence for a given uniform random number u the resulting sampled distance t lie between zero and infinity. However the primary sample can be renormalized over a bounded interval using the integral of the combined transmittance over this interval. Another slightly more complicated working approach is to linearly stretch the primary random sample such that the resulting shifted free-flight distance is bounded in the target volume interval.

Primary random sample. The primary random sample u associated with a free-flight sampled distance t equals the integral of the combined transmittance until t:

$$u = \int_0^t \bar{T}(t) = 1 - \bar{T}(t) = 1 - \exp(-\bar{\tau}(t)) \tag{17}$$

Note that the majorant extinction coefficient is either constant or piecewise constant if several volume sections are crossed. Hence the combined transmittance can be obtained using a product of the transmittance over each interval of constant majorant $\bar{\mu}_i$, or using the inverse exponential of the piecewise optical depth, until reaching the interaction at distance t contained in section n_t:

$$\begin{aligned}\bar{T}(t) &= \exp\left(-\bar{\tau}(t)\right) \\ &= \exp\left(-\sum_j^{n_t} \bar{\mu}_j \left(\min\left(t, t_{\max,j}\right) - t_{\min,j}\right)\right) \\ &= \prod_j^{n_t} \exp\left(-\bar{\mu}_j \left(\min\left(t, t_{\max,j}\right) - t_{\min,j}\right)\right)\end{aligned} \tag{18}$$

where $t_{\min,j}$ and $t_{\max,j}$ are the entry and exit depths of the j-th volume section and n_t is the index of the volume section containing depth t. Dividing the primary sample u by the integral of the combined transmittance $u_{\max}$ along a segment gives the *normalized random sample*.

Shift formulation. The mapping is then defined by copying the normalized random sample of the base interactions:

$$\frac{u'}{u'_{\max}} = \frac{u}{u_{\max}} \Leftrightarrow u' = \frac{u'_{\max}}{u_{\max}} \cdot u \Leftrightarrow u' = C \cdot u \tag{19}$$

where $C = \frac{u'_{\max}}{u_{\max}}$ is a constant defined by the integral of the combined transmittance on the base segment and the target segment. Plugging the piecewise constant definition of the majorant optical depth (Equation 18) into the latter equation leads to the expression the shifted interaction depth t' as a function of the base interaction depth t:

$$\begin{aligned}& u' = C \cdot u \\ \Leftrightarrow & 1 - \bar{T}(t') = C\left(1 - \bar{T}(t)\right) \\ \Leftrightarrow & \bar{T}(t') = 1 - C\left(1 - \bar{T}(t)\right) \\ \Leftrightarrow & \exp\left(-\bar{\tau}(t')\right) = 1 - C\left(1 - \bar{T}(t)\right) \\ \Leftrightarrow & \bar{\tau}(t') = -\log\left(1 - C\left(1 - \bar{T}(t)\right)\right) \\ \Leftrightarrow & \bar{\mu}'_{n_{t'}} \cdot t' = -\log\left(1 - C\left(1 - \bar{T}(t)\right)\right) - a_1 \\ \Leftrightarrow & t' = \frac{-\log\left(1 - C\left(1 - \bar{T}(t)\right)\right) - a_1}{\bar{\mu}'_{n_{t'}}} \\ \Leftrightarrow & t' = \frac{-\log\left(1 - C + C\exp(-\bar{\tau}(t))\right) - a_1}{\bar{\mu}'_{n_{t'}}} \\ \Leftrightarrow & t' = S_{\text{pss}}(t)\end{aligned} \tag{20}$$

where $\bar{\tau}(t')$ is given by Equation 12 and $a_1 = \sum_{i=1}^{n_{t'}-1} \bar{\mu}'_i \left(t'_{\max,i} - t'_{\min,i}\right) - \bar{\mu}'_{n_{t'}} \cdot t'_{\min,n_{t'}}$. Applying the transformation for each base null interaction completes the construction of the shifted segment.

Jacobian determinant and PDF. The Jacobian determinant of the complete transformation is:

$$\left|\frac{\mathrm{d}S_{\text{pss}}(t)}{\mathrm{d}t}\right| = \left|S'_{\text{pss}}(t)\right| = \frac{\bar{\mu}_{n_t}}{\bar{\mu}'_{n_{t'}}} \frac{C}{\left(C - (C-1)\exp\left(\bar{\tau}(t)\right)\right)} \tag{21}$$

Finally the PDF of a transformed interaction is:

$$p(x') = p(t)\left|S'_{\text{pss}}(t)\right|^{-1} P_n(x') \tag{22}$$

Discussion. The scaling factor C of the primary sample is the ratio of transmittance normalization terms over the base and target volume intervals. Although this approach is a valid shift, it requires evaluating several exponential and induces a slightly more complicated Jacobian to compute without significantly improving the repartition of the shifted samples.

3. Results and comparison

This section comments the graphs from Figure 3. A source segment on the domain $[0,9]$ is depicted in the top row (green curves) with its associated majorant $\bar{\mu}$, null μ_n and real μ extinction coefficients (top left), majorant $\bar{\tau}$ and real τ optical depths (top center) and densities of majorant, real and null interactions (top right). Note that we use varying extinctions that simulate heterogeneous media. A second target segment on the domain $[0,3]$ is depicted in the second row (red curves) with its associated extinction coefficient (left), optical depth (center) and distributions of real and null interactions (right). Our shift mappings operators aim at mapping the free-flight distributions of the source segment towards the target segment. The first mapping operates on raw depths (S_{raw} – blue curves), and results in a rescaling of the depth with a constant Jacobian (third row left), and thus in a stretch of the source distributions (fourth to bottom rows left). The second mapping operates on majorant optical depths (orange curves), and results in a rescaling of the majorant optical depth with a piecewise constant Jacobian (S_{mod} – third row center), and thus in a piecewise stretch of the source distributions (fourth to bottom rows center). The last mapping operates on primary samples (brown curves), and results in a rescaling of the primary samples with a more complex Jacobian (S_{pss} – third row right), and thus more complex transformation of the source distributions (fourth to bottom rows right) but still resembling the base shape.

4. Conclusion

We presented three new shift mappings of free-flight samples medium interactions and their Jacobian derivations, that can be used in path reusing application. These mappings allows path reusing within dense heterogeneous participating media.

All shifts allow for remapping distributions from an arbitrary source medium segment, to an arbitrary target medium segment. However, by construction none of them perfectly matches the true distributions of real interactions or null interactions, since we can not construct the ideal mappings from these two distributions to the target ones. Some methods such as optimal transport or dynamic numerical inversion could solve this problem of remapping between complex distributions. However, these methods would be too computationally intensive for our use case where we deal with a large number of medium segments and shift operations.

Instead, our mappings rely on geometric quantities available in close form (e.g. raw distances, combined transmittance and majorant extinction coefficient). Hence, we do not expect them to perfectly match the target distributions (real or null).

Raw depth shift mapping In practice, the raw depth shift mapping S_{raw} results in shifted distributions that closely resemble the source distributions. Although this is not desired for a general purpose mapping as the target distributions may be very different, this mapping works best when the source and target distributions are close to each others.

Primary sample shift mapping The primary sample shift mapping S_{pss} instead better spread the shifted distributions to perfectly match the target majorant density. However, we aim at matching the true distribution of real or null interactions for path reuse purpose (we shift null and real interactions), which is not achieved with this mapping. Furthermore, it sometimes leads to oversampled regions at the end of the segment. Additionally its Jacobian is slightly more complex to evaluate than our other mappings.

Majorant optical depth shift mapping The optical depth shift mapping S_{mod} is balanced between both: it better stretches than the raw depth by taking into account the local majorant coefficient, and most of the time better resembles the true distributions (real and null) than the primary sample shift mapping. Finally, it has a simple evaluation and Jacobian formula. We hence use the latter in our path reuse application as it is a good tradeoff between the quality of shifted real distributions and the complexity of evaluation.

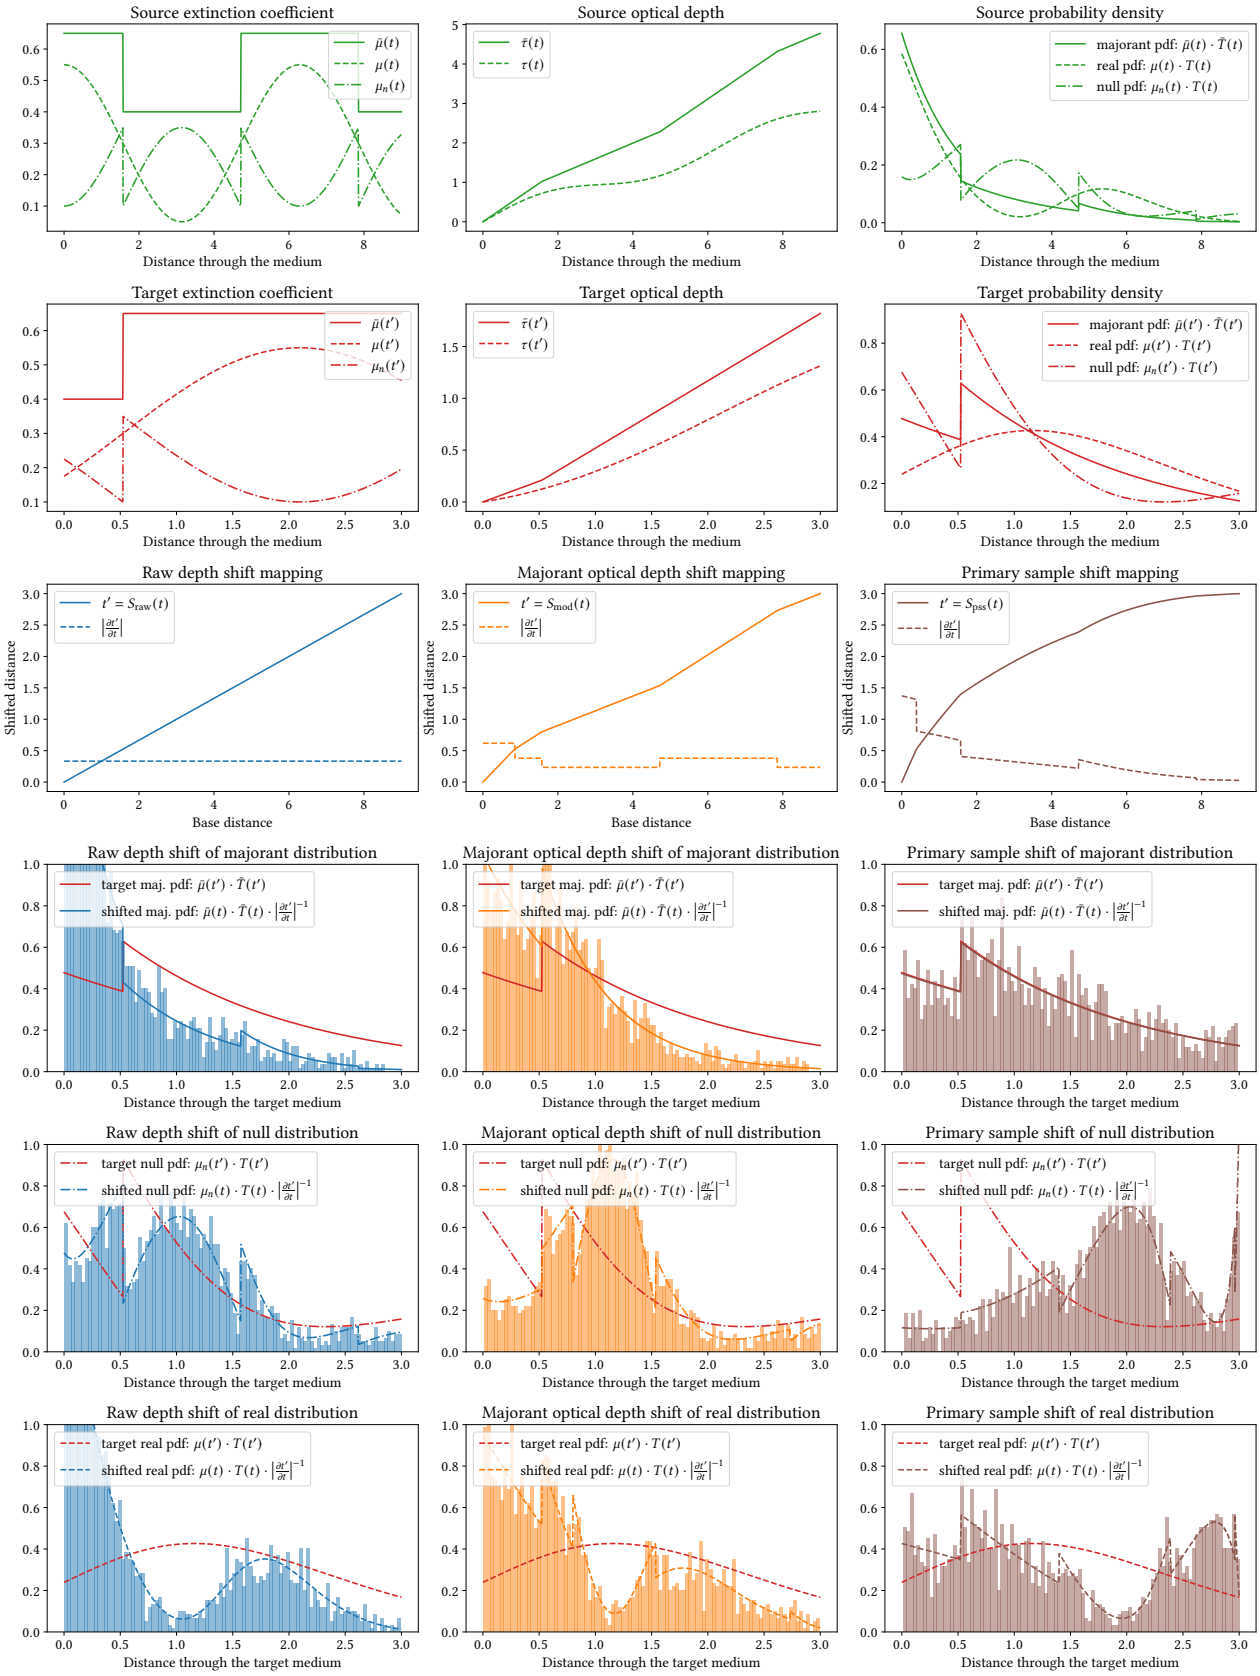

Figure 3: *Comparison of our shift mappings on a test case.*

PART TWO

Computing multi-view weights

In the following section, the keyword *main* is used to reference the equations from the article body. We detail in the following sections the complete derivations of both multiple importance sampling (main, Alg. 1, line 15) and multiple weighted importance sampling (main, Eq. 23 and Alg. 2, line 15) weights.

5. MIS weights

Denoting $w_{i\to j,\text{mis}}(\bar{\mathbf{x}}_j)$ the weight for path $\bar{\mathbf{x}}_j$ reaching pixel j that results from transforming the base path $\bar{\mathbf{x}}_i$ from pixel i, we develop the classic MIS balance heuristic weighting functions [**?**] using the pdfs as defined in the main document (main, Eq. 11 and 12):

$$w_{i\to j,\text{mis}}(\bar{\mathbf{x}}_j) = \frac{p_{i\to j}(\bar{\mathbf{x}}_j)}{\sum_k p_{k\to j}(\bar{\mathbf{x}}_j)} = \frac{\frac{p_{i\to j}(\bar{\mathbf{x}}_j)}{p_{i\to j}(\bar{\mathbf{x}}_j)}}{\sum_k \frac{p_{k\to j}(\bar{\mathbf{x}}_j)}{p_{i\to j}(\bar{\mathbf{x}}_j)}} = \frac{1}{\sum_k \frac{p_{k\to j}(\bar{\mathbf{x}}_j)}{p_{i\to j}(\bar{\mathbf{x}}_j)}} = \frac{1}{\sum_k \frac{c_k q_k(\bar{\mathbf{x}}_k)|S'_{k\to j}(\bar{\mathbf{x}}_k)|^{-1}}{c_i q_i(\bar{\mathbf{x}}_i)|S'_{i\to j}(\bar{\mathbf{x}}_i)|^{-1}}}. \tag{23}$$

Next we show that the above equation reduces to a form that does not depend on j, enforcing the fact the the MIS weight is unique for a given base path $\bar{\mathbf{x}}_i$. We use the chain rule $|S'_{a\to b}(\bar{\mathbf{x}}_a)| \cdot |S'_{b\to c}(\bar{\mathbf{x}}_b)| = |S'_{a\to c}(\bar{\mathbf{x}}_a)|$ to expand the latter equation, which leads to:

$$w_{i\to j,\text{mis}}(\bar{\mathbf{x}}_j) = \frac{1}{\sum_k \frac{c_k q_k(\bar{\mathbf{x}}_k)|S'_{k\to j}(\bar{\mathbf{x}}_k)|^{-1}}{c_i q_i(\bar{\mathbf{x}}_i)|S'_{i\to j}(\bar{\mathbf{x}}_i)|^{-1}} \cdot \frac{|S'_{j\to k}(\bar{\mathbf{x}}_j)|^{-1}}{|S'_{j\to k}(\bar{\mathbf{x}}_j)|^{-1}}} = \frac{1}{\sum_k \frac{c_k q_k(\bar{\mathbf{x}}_k)|S'_{k\to k}(\bar{\mathbf{x}}_k)|^{-1}}{c_i q_i(\bar{\mathbf{x}}_i)|S'_{i\to k}(\bar{\mathbf{x}}_i)|^{-1}}} = \frac{1}{\sum_k \frac{p_{k\to k}(\bar{\mathbf{x}}_k)}{p_{i\to k}(\bar{\mathbf{x}}_k)}} = \frac{1}{S_r} \tag{24}$$

where S_r is the sum of ratios of the base and shifted pdf (main, Eq. 11 and 12) from Alg. 1 in the main document. In the case of the power heuristic, the ratios should be elevated at the desired power α:

$$w_{i\to j,\text{mis},\alpha}(\bar{\mathbf{x}}_j) = \frac{1}{\sum_k \left(\frac{p_{k\to k}(\bar{\mathbf{x}}_k)}{p_{i\to k}(\bar{\mathbf{x}}_k)}\right)^\alpha} = \frac{1}{S_{r,\alpha}} \tag{25}$$

6. MWIS weights

Denoting $w_{i\to j,\text{mwis}}(\bar{\mathbf{x}}_j)$ the weight for path $\bar{\mathbf{x}}_j$ reaching pixel j that results from transforming the base path $\bar{\mathbf{x}}_i$ from pixel i, we develop the balance heuristic MWIS weighting functions (main, Eq. 18) using the pdfs as defined in the main document (main, Eq. 11 and 12) and Eq. 24:

$$w_{i\to j,\text{mwis}}(\bar{\mathbf{x}}_j) = w_{i\to j,\text{mis}}(\bar{\mathbf{x}}_j) \cdot w_{i\to j,\text{wis}}(\bar{\mathbf{x}}_j) = \frac{1}{S_r} \cdot \frac{p_{j\to j}(\bar{\mathbf{x}}_j)}{p_{i\to j}(\bar{\mathbf{x}}_j)} \tag{26}$$

where $p_{i\to j}(\bar{\mathbf{x}}_j)$ is the source pdf and $p_{j\to j}(\bar{\mathbf{x}}_j)$ is the target pdf within the WIS framework.

References

[BJNJ17] BITTERLI B., JAKOB W., NOVÁK J., JAROSZ W.: Reversible jump Metropolis light transport using inverse mappings. *ACM Transactions on Graphics 37*, 1 (2017). 3

[GHV*18] GRUSON A., HUA B.-S., VIBERT N., NOWROUZEZAHRAI D., HACHISUKA T.: Gradient-domain volumetric photon density estimation. *ACM Trans. on Graphics (TOG) 37* (2018). 2

[KSKAC02] KELEMEN C., SZIRMAY-KALOS L., ANTAL G., CSONKA F.: A simple and robust mutation strategy for the Metropolis light transport algorithm. *Computer Graphics Forum 21*, 3 (2002), 531–540. 3

[MGJ19] MILLER B., GEORGIEV I., JAROSZ W.: A null-scattering path integral formulation of light transport. *ACM Trans. on Graphics (SIGGRAPH) 38*, 4 (2019). 2

[NGHJ18] NOVÁK J., GEORGIEV I., HANIKA J., JAROSZ W.: Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum (Eurographics State of the Art Reports) 37*, 2 (2018). 2

[VG95] VEACH E., GUIBAS L. J.: Optimally combining sampling techniques for Monte Carlo rendering. In *SIGGRAPH* (1995), ACM, pp. 419–428. 9