

Handling Fluorescence in a Uni-directional Spectral Path Tracer

M. Mojiżik¹ and A. Fichet^{1,2} and A. Wilkie¹

¹Charles University, Czech Republic

²INRIA, Univ. Grenoble Alpes/LJK, CNRS/LJK, France




Figure 1: A brief summary of the contributions made in this paper. The Cornell box is filled with a fluorescent medium which absorbs at 360nm & re-emits at 620nm, and the two spherical objects have fluorescent surfaces as well. With classical exponential tracking, the fluorescent volume cannot be rendered at all (see figure 10). In image 1a, we use a simple "hotfix" of exponential tracking, which however still performs quite poorly: see section 7.3 for an explanation of the cyan colour cast. Image 1b shows that our proposed fluorescence-aware tracking performs much better for the same number of samples, even in a monochrome tracer. And finally, image 1c shows that we also managed to get Hero wavelength spectral sampling to work with wavelength-shifting media, which again significantly increases convergence speed, and removes any residual colour cast for low sample counts. Image 1d shows a converged version of the scene.

Abstract

We present two separate improvements to the handling of fluorescence effects in modern uni-directional spectral rendering systems. The first is the formulation of a new distance tracking scheme for fluorescent volume materials which exhibit a pronounced wavelength asymmetry. Such volumetric materials are an important and not uncommon corner case of wavelength-shifting media behaviour, and have not been addressed so far in rendering literature. The second one is that we introduce an extension of Hero wavelength sampling which can handle fluorescence events, both on surfaces, and in volumes. Both improvements are useful by themselves, and can be used separately: when used together, they enable the robust inclusion of arbitrary fluorescence effects in modern uni-directional spectral MIS path tracers. Our extension of Hero wavelength sampling is generally useful, while our proposed technique for distance tracking in strongly asymmetric media is admittedly not very efficient. However, it makes the most of a rather difficult situation, and at least allows the inclusion of such media in uni-directional path tracers, albeit at comparatively high cost. Which is still an improvement since up to now, their inclusion was not really possible at all, due to the inability of conventional tracking schemes to generate sampling points in such volume materials.

CCS Concepts

•Computing methodologies → Rendering; Ray tracing;

1. Introduction

At least in principle, it has been known for quite a while how to handle fluorescent substances, i.e. materials which are capable of *shifting* the wavelength of light that interacts with them, in the con-

text of computer graphics computations. But due to technical difficulties with integrating wavelength shifting into modern rendering engines, the feature is still absent from practically all modern production rendering systems. The absence of this feature is an increasingly unsatisfactory situation, as the effect does play quite a

role in the appearance of a number of common materials, such as paper and textiles (both via optical brighteners that are present in a large number of them), and fluorescent warning paint, e.g. on warning signs and emergency vehicles. But as can be seen in figure 2, fluorescence also occurs in volumetric materials, such as some detergents, sodas, minerals, rocks, and in certain transparent plastics.

The omission of the effect from current rendering software is all the more odd, as truly predictive rendering capabilities are now increasingly playing a role in both design workflows, and highly realistic movie productions. Renditions of such materials were, at least up to now, mostly approximative fixes: a fairly common way to represent obviously fluorescent materials such as warning paint in current rendering systems was to simply model them as abnormally reflective, i.e. to set their reflectance to be larger than unity in the part of the spectrum that dominates appearance of the colour. But as the industry is moving towards physically-based workflows which use renderers that require energy-conserving BRDF models, this is increasingly no longer an option [MSHD15].

In this paper, we deal with two of the remaining technical issues that present themselves when wishing to implement fluorescence in a modern spectral renderer that uses Hero wavelength sampling [WND*14], which at the moment seems to be the most efficient way to work with spectra in a Monte Carlo setting.

2. Related work

Here, we briefly review the state of the art with respect to three areas: offline rendering technology that is capable of achieving the highest levels of realism, current practice with regard to handling of fluorescence in rendering, and volume rendering.

2.1. Spectral Path Tracing

If highly accurate results are desired, and if rendering speed is not of prime concern, spectral path tracers are currently the best technology one can use for image synthesis purposes. There are of course also other, increasingly realistic rendering technologies which are capable of interactive frame-rates: these are used for gaming and AR/VR purposes [Tat17]. However, as of today, none of these come even close to spectral Monte Carlo techniques with regard to physical realism. So at least for now, offline techniques are here to stay as a separate technological branch.

In fact, it is only in the space of the last five years that path tracing in its various forms has become the dominant technology for high quality movie VFX [FHF*17]. Interestingly, most actual production work in this area is done with comparatively simple MIS path tracers (simple at least compared to recent advances, such as [KGGH*14]), and by using workflows that are still mainly based on RGB colour [ENSB13]. However, some major movie VFX efforts are already being rendered using spectral path tracers [FHF*17]: one of the reasons given is that it is easier to work with physical measurements and quantities in such a system.

When it comes to representing the full spectrum of light transport, older spectral rendering systems used to work with full representations of spectra [DCWP02], either via a regular sampling


scheme of some sort, or e.g. via hybrid techniques [SFD98]. However, any full representation of spectral data carried along the paths in a modern Monte Carlo renderer has the intrinsic problem that this is only viable if all decisions along the path are made with the same probability for all wavelengths. This issue was long overlooked, as this is not an issue with the simple scenes that many early path tracers were tested on: environments which only featured hard surfaces, and no (or only very simple) scattering effects. But as soon as one works with arbitrary volumetric effects, the only sensible approach is to also randomly sample the wavelength domain, i.e. to use monochrome paths. As this introduces colour noise, which is highly visually objectionable and hard to filter out, Hero wavelength sampling [WND*14], which in turn was based on earlier work that was not yet applicable to production systems [RBA09], was introduced. This technique allows one to run a path tracer with spectral data, while retaining the conceptual simplicity of monochrome path tracing – but without colour noise. The price for this is a comparatively small computational overhead, to keep track of the probabilities that are used. This is the current state of the art with regard to spectral path tracing, so any technique that aims at supporting fluorescence in a modern renderer has to be compatible with it.



(a) Simulated daylight.



(b) Monochrome light, 400nm.



(c) Monochrome light, 550nm.

Figure 2: Photographs of real fluorescent and non-fluorescent objects in a controlled viewing environment. Note the little glass filled with fluorescent dishwashing liquid that becomes totally transparent outside the wavelength range where fluorescence emission occurs in 2c.

2.2. Previous Work on Fluorescence Rendering

Light transport in the presence of fluorescent substances has been investigated to quite some extent in the natural sciences [KM06, KNH05]. The main purpose of these efforts was to build mathematical models that correspond to particular experimental set-ups. The mathematical models which were introduced there of course describe the same phenomenon that we see in renderings: but as the mathematical techniques they were developed for are significantly different from what is being used in graphics, the direct utility of these results for graphics is limited.

Table 1: Nomenclature used in the paper, which is based on the notation used by [FWKH17]. Due to the bi-spectral nature of our work, we explicitly have to label the involved wavelengths λ , and introduce the additional fluorescence-enabled quantities $\sigma_{s\Phi}$, σ_{s_e} , σ_{s_a} , $\widehat{\sigma}_s$ and $\widehat{\sigma}_t$. Figure 3 shows a schematic of the notation elements involved in a fluorescent collision event.

Symbol	Description
\mathbf{x}	Position
t	Ray parameter
ω	Ray direction
\mathbf{x}_t	Parameterised position along ray: $\mathbf{x}_t = \mathbf{x} + t\omega$
λ	Wavelength sampled by a monochrome eye-path
λ_o	Output wavelength for fluorescent collisions
λ_i	Input wavelength for fluorescent collisions
d	Ray length / domain of volume integration: $0 < t < d$
ζ	Random number
$L(\mathbf{x}, \omega, \lambda)$	Radiance at \mathbf{x} in direction ω
Ω	Hemispherical integration domain
Λ	Spectral integration domain
Non Fluorescent parameters	
$\sigma_a(\mathbf{x}, \lambda)$	Absorption coefficient
$\sigma_s(\mathbf{x}, \lambda)$	Non fluorescent scattering coefficient
$\sigma_t(\mathbf{x}, \lambda)$	Non fluorescent extinction coefficient: $= \sigma_a(\mathbf{x}, \lambda) + \sigma_s(\mathbf{x}, \lambda)$
$T(\mathbf{x}, \omega, \lambda)$	Transmission: all radiance that is not absorbed or out-scattered, and which continues unaltered
$f_p(\mathbf{x}, \omega, \omega')$	Phase function
Fluorescent parameters	
$\Phi(\lambda_i, \lambda_o)$	Fluorescence response function, a.k.a. the full re-radiation matrix, or fluorescence efficiency term
$\sigma_{s\Phi}(\mathbf{x}, \lambda_i, \lambda_o)$	Fluorescence cross-talk, i.e. $\Phi(\lambda_i, \lambda_o)$ minus the main diagonal: $\forall \lambda_i = \lambda_o, \sigma_{s\Phi}(\mathbf{x}, \lambda_i, \lambda_o) = 0$
$\sigma_{s_e}(\mathbf{x}, \lambda)$	Fluorescent in-shifting: incoming wavelength-shifted energy transfer (fluorescence re-emission): $= \int_{\Lambda} \sigma_{s\Phi}(\mathbf{x}, \lambda_i, \lambda) d\lambda_i$
$\sigma_{s_a}(\mathbf{x}, \lambda)$	Fluorescent out-shifting: outgoing wavelength-shifted energy transfer (fluorescence absorption): $= \int_{\Lambda} \sigma_{s\Phi}(\mathbf{x}, \lambda, \lambda_o) d\lambda_o$
$\widehat{\sigma}_s(\mathbf{x}, \lambda)$	Fluorescence-aware collision coefficient: $= \sigma_s(\mathbf{x}, \lambda) + \sigma_{s_e}(\mathbf{x}, \lambda)$
$\widehat{\sigma}_t(\mathbf{x}, \lambda)$	Fluorescence-aware extinction coefficient: $= \sigma_t(\mathbf{x}, \lambda) + \sigma_{s_a}(\mathbf{x}, \lambda)$




Figure 3: Schematic of a collision event in the presence of fluorescence. This is a visual representation of the key quantities introduced in table 1.

Within graphics, the first publication to deal with the phenomenon was Glassner [Gla94], who demonstrated not only the inclusion of fluorescence proper into a Whitted ray tracer, but also dealt with time-dependent re-radiation, a.k.a. phosphorescence. A

very interesting extended formulation of the rendering equation capable of describing both effects was introduced: but by modern standards, the practical utility of the rendering technology used is practically non-existent, due to its deterministic nature.

The combination of polarisation and fluorescence within one rendering system was later investigated by [WTP01]: while the renderer used in this work was already a Monte Carlo system, it was still a rather primitive one, which used n -band full spectral representations, and did not run into the wavelength dependency issues mentioned in the previous section; simply because no materials sophisticated enough to expose this flaw were supported. Due to this simplicity, the implementation proposed in that paper got away with using re-radiation matrices that had been re-sampled to the spectral resolution used by the renderer as attenuation elements. While such an approach entirely avoids colour noise, it is slow, suffers from aliasing artefacts due to its limited spectral resolution, and is not extensible to support non-trivial volumes.

Wilkie et al. [WWLP06] investigated a simple layered BRDF model that was capable of handling the conditional dependence of the multi-coloured lobe found in the BRDFs of real world fluorescent surfaces. Hullin et al. [HHA*10] proposed a modern formulation of the modified rendering equation initially introduced by [Gla94], and generally provided a considerably more theoretically sound framework for the handling of the effect. However,

this work did not go as far as to solve all detail problems of supporting the effect in a modern spectral path tracer. In particular, a clear framework for how to efficiently perform volume rendering in fluorescent media was still missing: as figures 2 and 7 show, the pronounced wavelength asymmetry in the light interaction behaviour that some of these media exhibit requires a specialised distance tracking scheme, which has so far been missing.

As part of his Ph.D. thesis which dealt with simulations of neuro-imaging procedures, Abdellah [Abd17] also investigated volumetric fluorescence rendering [ABE*17]. Fluorescent dyes play a crucial role in the marking of neural structures for microscopy and neurosurgery, and they developed a proof of concept rendering system which was capable of handling the strongly fluorescent media typical for these application cases. They tested their system with real fluorophore re-radiation data, but its performance for strongly asymmetric media like that shown in figure 2 was not investigated. Also, the described proof of concept system apparently did not use a modern approach to spectral rendering. The final work about fluorescence in graphics in recent years, by Bendig et al. [BHD*08], investigated real-time simulations of wavelength shifting effects in high performance ray-tracing based rendering systems.

2.3. Volume Rendering

Using volumes in path tracers is now routine in production environments [FWKH17]. Recently, advances were made in efficient handling of inhomogeneous media [KHLN17], and the handling of emissive volumes [SHZD17]. The former work is orthogonal to our proposed scheme, as the results shown in section 7.5 could equally well have been generated with that technique. The latter offers considerable benefits when "genuine" emissive volumes are present: but one of the insidious properties of fluorescent volumes is that they are not actually emissive *per se*, as their emission depends on the intensity and spectral composition of the incoming light. So that technique is unfortunately not applicable to fluo volumes.

3. Fluorescence

Fluorescence is the effect where a molecule re-emits a previously absorbed photon in a lower energy state than the one it was absorbed in: this means the re-emitted light from a collision event with such a molecule has a longer wavelength than the original incoming light. This shift from the absorbed frequency to the re-emitted one is spread around a main frequency as shown in Figure 4. This is referred to as *Stokes shift*.

Such phenomena can be described by a fluorescence response function $\Phi(\lambda_i, \lambda_o)$, which for an excitation wavelength λ_i returns the amount of energy relaxed to another re-emission wavelength λ_o . For $\lambda_i = \lambda_o$, this corresponds to the non-fluorescent reflectance.

3.1. Re-radiation matrix

The fluorescence response $\Phi(\lambda_i, \lambda_o)$ can be represented in a discretised form as a so-called re-radiation matrix. An example dataset is shown in Figure 5. The incoming wavelength λ_i lies on its vertical axis, and the outgoing wavelength λ_o on its horizontal. The diagonal of such a re-radiation matrix describes the amount of light that does not undergo any wavelength shift.




Figure 4: An example of the so-called Stokes shift, here shown for the material found in pink fluorescent 3M pink Post-it notes. Practically all fluorophores exhibit similar characteristics: namely, that there is one main band where energy is absorbed, and one where it is re-emitted. As discussed in the text, this characteristic would make stratified random sampling of wavelength change events problematic. Note that one still needs the full re-radiation matrix of the fluorophore in question to generate such a plot, which only shows the cumulative shifting effect of the material. Figure 5 shows the entire re-radiation matrix of this material.




Figure 5: The full re-radiation matrix $\Phi(\lambda_i, \lambda_o)$ for the material shown in figure 4. Values lower than 0.006 are ignored in this plot, as they are considered measurement noise.

3.2. Shifting probability

Kasha's rule states that the re-emission spectrum distribution should be the same, whatever wavelength within the absorbing area of the spectrum is used for excitation. Only the intensity of this re-emitted spectrum is modified. However, this rule does not always apply – at least not in the strict sense.

While the accumulated energy absorption and re-emission, as shown in Figure 4, represents, respectively, the sum of the columns and the lines of the re-radiation matrix, i.e. $\int_{\Lambda} \Phi(\lambda, \lambda_o) d\lambda_o$ and $\int_{\Lambda} \Phi(\lambda_i, \lambda) d\lambda_i$, the probability for a given wavelength to be re-emitted may differ if the particle does not exactly follow Kasha's rule (and as figure 6 shows, such slight variations are actually the norm). Then, the distribution of possible re-emission wavelengths




Figure 6: Re-emission probability distribution for several excitation wavelengths in the dataset shown in figure 5. This plot amounts to several vertical cuts approximately halfway through the center of the wavelength-shifting activity, for different incident wavelengths. These show that Kasha’s rule is not an exact rule, so that there is some variation: but the overall appearance of the spectra is quite similar overall.

for an excitation wavelength j is given by the corresponding column j in the re-radiation matrix. On the other hand, from a re-emitted wavelength i , the possible spectrum of excitation wavelengths is given by the corresponding line i in the matrix.

4. Bi-spectral BRDF

Even though our goal in this paper is to improve the handling of fluorescent volumes, we start by extending and following up on the work for fluorescent surfaces done by Hullin et al [HHA*10]. Specifically, their description of the *Bi-spectral Bidirectional Reflectance and Re-radiation* (bi-spectral BRDF):

$$f_r((\omega_i, \lambda_i), (\omega_o, \lambda_o)) = \frac{d^2 L(\omega_o, \lambda_o)}{L(\omega_i, \lambda_i) d\omega_i d\lambda_i}$$

which is in turn used within the *bi-spectral rendering equation* (we omit emission for simplicity):

$$L(\omega_o, \lambda_o) = \int_{\Lambda} \int_{\Omega} L(\omega_i, \lambda_i) f_r((\omega_i, \lambda_i), (\omega_o, \lambda_o)) d\omega_i d\lambda_i$$

The wavelength shifting properties of fluorescent surfaces turn out to usually be directionally independent, which would allow:

$$f_r((\omega_i, \lambda_i), (\omega_o, \lambda_o)) = \Phi(\lambda_i, \lambda_o) \cdot f^{\omega}(\omega_i, \omega_o)$$

However, the formalism of bi-spectral BRDF is more general, as it would allow us to also describe more complex surfaces, where e.g. a fluorescent diffuse surface is combined with a non-fluorescent specular surface.

The measurement of bi-spectral BRDFs can yield, thanks to its potential separability, a re-radiation matrix, a discrete version of $\Phi(\lambda_i, \lambda_o)$. Since fluorescent materials one encounters in a computer graphics context only shift energy towards longer wavelengths [DCWP02], the upper right triangle of the matrix is always zero.

5. Wavelength-shifting volumes

5.1. Bi-spectral collision events

Our goal is to introduce support for fluorescent participating media to a uni-directional path tracer. Our starting point is to formulate the integral we will be attempting to evaluate, which we obtain by expanding the one we would be evaluating for the case where no wavelength shifts occur. Specifically, we start with the most common formulation used for typical path tracers, the integral form of the spectral radiative transfer equation (RTE), where we explicitly model a surface event at distance d . For simplicity, we do not include emission and assume a homogeneous volume.

$$A = \int_{\Omega} f_p(\omega_i, \omega, \lambda) L(\mathbf{x}_t, \omega_i, \lambda) d\omega_i$$

$$L(\mathbf{x}, \omega, \lambda) = \int_{t=0}^d T(\mathbf{x}, \mathbf{x}_t, \lambda) \cdot (\sigma_s(\lambda) \cdot A) dt + T(\mathbf{x}, \mathbf{x}_d, \lambda) \cdot L(\mathbf{x}_d, \omega, \lambda) \quad (1)$$

where σ_s represents the scattering coefficient, $\mathbf{x}_t = \mathbf{x} + t \cdot \omega$, as in the point at distance t in direction ω from \mathbf{x} , f_p is the phase function and $T(\mathbf{x}, \mathbf{x}_t, \lambda)$ is the transmittance between \mathbf{x} and \mathbf{x}_t , which is computed as:

$$T(\mathbf{x}, \mathbf{x}_t, \lambda) = e^{-\|\mathbf{x}-\mathbf{x}_t\| \cdot \sigma_t(\lambda)}$$

$$\sigma_t(\lambda) = \sigma_a(\lambda) + \sigma_s(\lambda)$$

where σ_a represents the absorption coefficient, and σ_t the extinction coefficient. Note that the two instances of σ_s in the two formulas represent two different event types: in-scattering (in transmittance computations) and out-scattering (in collision events). In non-fluorescent settings, both of these events end up using the same coefficient, as the amount of light that gets out-scattered from one direction then ends up getting in-scattered into another. In order

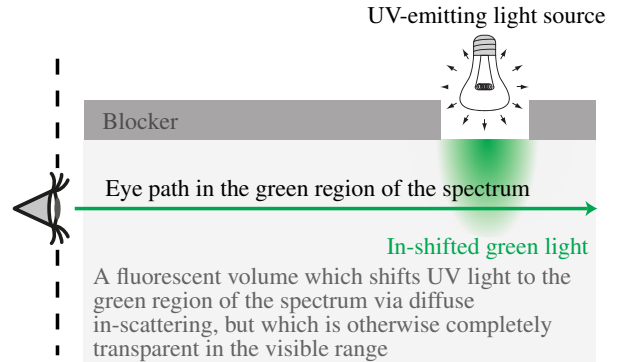


Figure 7: A sketch of what materials with a pronounced wavelength asymmetry in their light interaction behaviour look like from the viewpoint of a monochrome path tracer: the material in question is similar to that seen in the glass of detergent shown in figure 2. If sampling were only done according to the optical properties of the volumetric medium in the green region of the spectrum, no distance tracking events would be generated – the medium is completely transparent there. And randomly sampling fluorescence by potentially jumping up to other wavelengths along the entire path would be extremely inefficient.

to introduce fluorescent volumes, a *fluorescence efficiency term* $\Phi(\lambda', \lambda)$ was introduced by [Gla94]: this describes the efficiency of wavelength shifts when interacting with a fluorophore particle. In our set-up, the so-called *full RTE* would look like this:

$$L(\mathbf{x}, \omega, \lambda) = \int_{t=0}^d T(\mathbf{x}, \mathbf{x}_t, \lambda) \cdot \sigma_s(\lambda) \cdot \int_{\Lambda} \int_{\Omega} f_p(\omega_i, \omega, \lambda_i) \cdot \Phi(\lambda_i, \lambda) L(\mathbf{x}_t, \omega_i, \lambda_i) d\omega_i d\lambda_i dt + T(\mathbf{x}, \mathbf{x}_d, \lambda) \cdot L(\mathbf{x}_d, \omega, \lambda)$$

Although this approach is useful in describing certain fluorescent dyes, it fails to capture the asymmetry in the way light can interact with a volume material as shown in figure 7. For any wavelength λ , into which light gets transported via wavelength shifting collision events, the volume also ends up being translucent (or almost opaque) due to the non-zero $\sigma_s(\lambda)$.

Based on this, we propose certain changes to the integral formulation. We maintain the assumption that wavelength shifts occur only via collision events. To capture the various wavelength shifting behaviours, we parameterise the collision coefficient by both incoming and outgoing wavelengths. Essentially, what in the spectral case was a simple spectrally dependent coefficient, we now assume to be a value taken from a re-radiation matrix. We also parameterise the phase function, turning it into a bi-spectral phase function, for the purpose of symmetry with the bi-spectral BR-RDF [HHA*10], as well as to provide the possibility for expressing more complex phase functions, such as those that have wavelength-shift-dependent properties.

Therefore, to get the incident radiance at a specific wavelength, we also have to integrate over all possible wavelength shifts during collision events, using the *bi-spectral RTE*:

$$\begin{aligned} B &= \int_{\Omega} f_p(\omega_i, \omega, \lambda_i, \lambda) L(\mathbf{x}_t, \omega_i, \lambda_i) d\omega_i \\ C &= \int_{\Lambda} \sigma_{s\Phi}(\lambda_i, \lambda) \cdot B d\lambda_i \\ L(\mathbf{x}, \omega, \lambda) &= \int_{t=0}^d T(\mathbf{x}, \mathbf{x}_t, \lambda) \cdot (\sigma_s(\lambda) \cdot A + C) dt + T(\mathbf{x}, \mathbf{x}_d, \lambda) \cdot L(\mathbf{x}_d, \omega, \lambda) \end{aligned} \quad (2)$$

However, a bi-spectral collision coefficient also affects the computation of transmittance, specifically the extinction coefficient:

$$\hat{\sigma}_t = \sigma_t(\lambda) + \int_{\Lambda} \sigma_{s\Phi}(\lambda, \lambda_o) d\lambda_o \quad (3)$$

Note the different direction used when integrating over the σ_s in extinction coefficient computation, where we integrate over shifts from λ , as opposed to the collision event, where we integrate over shifts into λ . When interpreting σ_s as a re-radiation matrix, the difference is between integrating over a column, or over a row. The difference stems from the aforementioned difference between in-scattering and out-scattering, where in the case of wavelength shifting medium, the light that gets out-scattered does end up getting in-scattered at a different direction, but also at a potentially different wavelength. It is exactly this difference that allows us to properly capture the asymmetry in scattering behaviour shown in figure 7.

It is still possible to express the full RTE (equation 8 in [Gla94])

using this formulation by simply using:

$$\begin{aligned} \sigma_{s\Phi}(\lambda_i, \lambda) &= \begin{cases} \sigma'_s(\lambda) \cdot \Phi(\lambda_i, \lambda) & \lambda_i \neq \lambda \\ 0 & \lambda_i = \lambda \end{cases} \\ \sigma_s(\lambda) &= \sigma'_s(\lambda) \cdot \Phi(\lambda, \lambda) \\ \sigma'_a(\lambda) &= \sigma_a(\lambda) + \int_{\Lambda} \sigma_{s\Phi}(\lambda, \lambda_o) d\lambda_o \end{aligned} \quad (4)$$

where $\sigma'_s(\lambda)$ and $\sigma'_a(\lambda)$ are the coefficients use within the full RTE.

5.2. Importance sampling

Wavelength importance sampling for collision events is equivalent to what is described for bi-spectral BR-RDFs [HHA*10]: the techniques introduced there can be re-used. What requires further investigation are the tracking distances used in volumetric rendering.

The usual approach towards homogeneous media is to sample distances with a probability density proportional to their transmittance, specifically using $t = -\log(1 - \zeta) / \sigma_t(\lambda)$. This approach produces good results for non-fluorescent scattering media, as the same coefficient describes in-scattering and out-scattering, so that $\sigma_s(\lambda) \leq \sigma_t(\lambda)$. However, this no longer has to hold once wavelength shifting gets involved, and distance sampling solely based on the extinction coefficient σ_t can end up skipping all possible collision events, although they would contribute significantly. An example is visible on Figure 10a, where the absorption in the medium is zero, and the collision coefficient is exclusively fluorescent.

A similar problem exists with highly emissive and clear media. In that case, the optimal solution would be to sample the free path distance uniformly between the starting point and the first surface collision. However, unlike fluorescence, any contribution of emission is final, and can therefore be evaluated separately, e.g. via approximative closed form solutions.

5.3. Fluorescence-aware tracking

We propose an improved distance sampling scheme which by default uses uniform sampling for potentially high contributions from non-absorptive media, and exponential sampling for standard scattering media. The core of the technique lies in including the collision coefficient as well as transmittance in the sampling probability derivation. The idea behind it being that instead of only sampling proportionally to transmittance, which by itself may hide how significant fluorescence is, or only sampling uniformly, which is inefficient for higher transmittance, we derive a probability distribution that can naturally choose the most appropriate sampling based on the specific coefficients.

Using the definitions of A and C from equations 1 and 4 respectively, we start with the integral form of bi-spectral RTE, and assume all incoming light to be of the same importance, in order to get what should be proportional to the complete probability.

$$\begin{aligned} 1 &\sim \int_0^d T(\mathbf{x}, \mathbf{x}_t, \lambda) \cdot (\sigma_s(\lambda) \cdot A + C) dt + T(\mathbf{x}, \mathbf{x}_d, \lambda) \\ &= \int_0^d T(\mathbf{x}, \mathbf{x}_t, \lambda) \cdot (\sigma_s(\lambda) + \int_{\Lambda} \sigma_{s\Phi}(\lambda_i, \lambda) d\lambda_i) dt + T(\mathbf{x}, \mathbf{x}_d, \lambda) \\ &= \int_0^d e^{-t \cdot \hat{\sigma}_t(\lambda)} \cdot \hat{\sigma}_s(\lambda) dt + e^{-d \cdot \hat{\sigma}_t(\lambda)} \end{aligned}$$

From this point onward, we will work with two extended volume description coefficients we define, and which are directly used to guide the sampling. These are the *fluorescence-aware extinction coefficient* $\widehat{\sigma}_t$, and the *fluorescence-aware collision coefficient* $\widehat{\sigma}_s$, both of which are formally defined in Table 1. Key to both is that in addition to the "normal" information contained in σ_s and σ_t , respectively, they describe how much light is shed via fluorescence to other wavelengths ($\widehat{\sigma}_t$), or how much is potentially gathered from other wavelengths ($\widehat{\sigma}_s$). This allows the path tracer to make decisions that go beyond what would be possible with the information contained in σ_s and σ_t .




Figure 8: An illustration of the two different types of events that need to be treated separately in equation 5: a collision event, and a surface interaction.

We require the probability density function $p(t)$ of sampling a free-path distance to be proportional to the integrand:

$$p(t) \sim \begin{cases} \widehat{\sigma}_s \cdot e^{-t \cdot \widehat{\sigma}_t} & t < d \\ e^{-d \cdot \widehat{\sigma}_t} \cdot \delta(t-d) & t \geq d \end{cases} \quad (5)$$

We have to differentiate between a collision event occurring, and the case of completely skipping the volume and evaluating the outgoing light from a surface, as illustrated in figure 8.

As the process of deriving them is different for the cases where $\widehat{\sigma}_t > 0$ and $\widehat{\sigma}_t = 0$, we get two different formulas. After computing the normalisation constant, we obtain the CDF as:

$$P(t) = \begin{cases} \frac{1 - e^{-t \cdot \widehat{\sigma}_t}}{1 + \left(\frac{\widehat{\sigma}_t}{\widehat{\sigma}_s} - 1\right) \cdot e^{-d \cdot \widehat{\sigma}_t}} & t < d \\ 1 & t \geq d \end{cases} \quad (6)$$

where $\widehat{\sigma}_t > 0$

$$P(t) = \begin{cases} \frac{t \cdot \widehat{\sigma}_s}{d \cdot \widehat{\sigma}_s + 1} & t < d \\ 1 & t \geq d \end{cases}$$

where $\widehat{\sigma}_t = 0$

By inverting the CDF, we arrive at the formula for computing the

free-path distance sample:

$$t = -\frac{\log\left(1 - \zeta \cdot \left(1 + \left(\frac{\widehat{\sigma}_t}{\widehat{\sigma}_s} - 1\right) \cdot e^{-d \cdot \widehat{\sigma}_t}\right)\right)}{\widehat{\sigma}_t}$$

where $\widehat{\sigma}_t > 0$

$$t = \zeta \cdot \frac{d \cdot \widehat{\sigma}_s + 1}{\widehat{\sigma}_s}$$

where $\widehat{\sigma}_t = 0$

Using this improved distance sampling scheme, it is possible to sample distance in media with $\widehat{\sigma}_s > \widehat{\sigma}_t$, such as fluorescent volumes. Moreover, it is also capable of a more efficient sampling of mostly absorptive surfaces, such as $\widehat{\sigma}_s \ll \widehat{\sigma}_t$, where a purely fluorescence-aware sampling is likely to skip the entire volume.

6. Hero Wavelength Bi-spectral Sampling

The main concept of Hero Wavelength Spectral Sampling (HWSS) is that a single ("Hero") wavelength is used to guide the sampling of the path, and the same path is used to compute the contribution for multiple (usually four) different wavelengths that are stratified over the entire sampled wavelength range. As each wavelength has a chance to be the Hero wavelength, the individual contributions are also MIS weighed. The core advantage of this approach comes from using SIMD operations to compute all four contributions simultaneously, at a usually negligible decrease in computing speed, while effectively removing colour noise from the end result due to sample stratification. As it is the currently best technical option for spectral path tracing, it makes sense to extend it to support for all interesting effects one might encounter in rendering.

6.1. Wavelength shifting

However, the method has not been used in conjunction with wavelength shifting path events so far. The main difference to other path events is that after encountering a wavelength shift, the four wavelengths carried along with the path obviously can no longer remain constant. What happens with the Hero wavelength is sort of obvious: at such an event, it will shift, according to some probability. But what should one do with the other three, if such an event is encountered? Specifically, what should be done about the stratification of the original four wavelengths? Should they all be collapsed onto the Hero wavelength at such an event, effectively turning the path tracer into a monochrome one upon encountering such an event? This would be a valid solution, but inefficient.

Alternatively, the stratification can be retained, by shifting the other three in the same way as the Hero wavelength. However, such a shift probably moves some of them outside of the visible spectral range – and looping back at the other end is not a sensible option in this case. So at least some of the other three wavelengths would be removed from the sample. One last option would be to scale the stratified wavelength set, to always stay within the visible range.

All of these solutions however share the same issue – the majority of energy that is transported via any particular real fluorescence effect is usually localised around two bands: one *absorption*

band, and one emission band, as can be seen in figure 4. This creates problems for the other three wavelengths used by HWSS: if the Hero wavelength is shifted from somewhere in the absorption band to somewhere the emission band (i.e. undergoes a meaningful wavelength shift which will actually end up transporting energy), moving the other three stratified wavelengths by similar distances will, in all probability, not yield meaningful light transport, as no (or very little) shifting happens in those other wavelength brackets.

Which means that although a re-scaled stratified wavelength set that still falls within the boundaries of the visible range (the last option discussed above) would technically be a correct way of doing things, it would still cause HWSS to degenerate to regular monochrome spectral sampling like a straightforward collapsing of wavelengths would: and it would essentially negate any advantages of HWSS after the first fluorescent event. Therefore, it actually makes sense to allow each of the HWSS wavelengths to shift completely independently from each other.

6.2. Bi-spectral HWSS

If the wavelength shift of fluorescent materials is individually importance sampled in HWSS (and as we showed in the previous section, it really ought to be), an issue with sample probabilities arises. A normal HWSS has, in spite of its vector nature, a straightforward sample probability p : namely, $p = p_{d1}$, the probability of the directional decision made for the Hero wavelength. However, if each of the wavelengths λ_i in a HWSS now additionally performs individual importance sampling for wavelength shifting with probability p_{λ_i} , the probability of the entire sample ends up being the product of the involved probabilities: $p = p_{d1} \prod_{i=1}^4 p_{\lambda_i}$.

Which means that if any of the shifting probabilities p_{λ_i} is zero (and this can happen, as there might not be a viable shifting target for some incoming wavelengths), the probability of the entire HWSS vector ends up being zero, which in turn can lead to cases where non-zero radiance samples on other wavelengths end up in a vector sample with overall probability zero. But this sort of thing must never happen in a MC integration: function values of zero can have a non-zero sampling probability, but all non-zero function values must have a non-zero sampling probability. The solution is to ensure that the function value of all HWSS vector components is zero if one of them has a shifting probability of zero. The following derivation presents a way how to do this, while retaining the overall logic of HWSS.

We first introduce a vector notation for equation components. For a function $g(\mathbf{x}, \lambda)$, a vector of n invocations resp. results of invoking $g(\mathbf{x}, \lambda)$ with a different λ for each element would be denoted as:

$$\mathbf{g}^*(\mathbf{x}, \lambda^*) = \begin{bmatrix} g(\mathbf{x}, \lambda_1) \\ \vdots \\ g(\mathbf{x}, \lambda_n) \end{bmatrix}$$

If only normal, non wavelength-shifting materials are present, HWSS at the level of each individual hemispherical integration along a path is straightforward: n -vectorising a specific Monte Carlo sample (which HWSS effectively amounts to, usually with

$n = 4$) has no consequences for the actual integration.

$$L_d(\mathbf{x}, \omega_0, \lambda) = \int_{\Omega} f(\mathbf{x}, \omega_i \rightarrow \omega_0, \lambda) \cdot L(\mathbf{x}, \omega_i, \lambda) d\omega_i$$

The elements of the solution vector are simply a component-wise product of the original vectorised contributions: which is of course the main appeal of the entire idea, as this allows efficient use of SIMD instructions. At a particular recursion level on a path, the local versions of the spectral rendering equation (i.e. the single hemispherical integration at this particular recursion depth) are basically identical in their regular and vectorised forms:

$$\begin{aligned} L_d^*(\mathbf{x}, \omega_0, \lambda^*) &= \int_{\Omega} \mathbf{f}^*(\mathbf{x}, \omega_i \rightarrow \omega_0, \lambda^*) \diamond L^*(\mathbf{x}, \omega_i, \lambda^*) d\omega_i \\ &= \int_{\Omega} \begin{bmatrix} f(\mathbf{x}, \omega_i \rightarrow \omega_0, \lambda_1) \cdot L_d(\mathbf{x}, \omega_0, \lambda_1) \\ \vdots \\ f(\mathbf{x}, \omega_i \rightarrow \omega_0, \lambda_n) \cdot L_d(\mathbf{x}, \omega_0, \lambda_n) \end{bmatrix} d\omega_i \\ L_d^*(\mathbf{x}, \omega_0, \lambda^*) &= \begin{bmatrix} L_d(\mathbf{x}, \omega_0, \lambda_1) \\ \vdots \\ L_d(\mathbf{x}, \omega_0, \lambda_n) \end{bmatrix} \end{aligned} \quad (8)$$

Here, the operator \diamond designates element wise multiplication: such operations can be computed in parallel, e.g. via SIMD instructions. For the following, it is essential to bear in mind that the result obtained via the vectorised form shown in the second equation above is still a *single* Monte Carlo sample: the only difference to the plain form above is that it returns n values for various λ , but all these n values contribute to a single MC sample. The only feature which this vectorisation introduces is the additional sampling of λ , so the spatial integration domain for this vectorised MC sample stays the same as for a non-vectorised sample: Ω .

By contrast, if the n elements of the vectorised form were allowed to also individually pick a sampling direction $\omega_{i,k}$ (which does not happen in HWSS, as the key idea there is to use a single path for all four wavelengths), one would again get n values which contribute to a single MC sample. However, equation 8 would not hold, and the integration domain would actually expand to Ω^n :

$$\begin{aligned} L_d^*(\mathbf{x}, \omega_0, \lambda^*) &= \int_{\Omega^n} \mathbf{f}^*(\mathbf{x}, \omega_i^* \rightarrow \omega_0, \lambda^*) \diamond L^*(\mathbf{x}, \omega_i^*, \lambda^*) \left[d\omega_{i,k} \right]_{1..n} \\ \text{Where : } \left[d\omega_{i,k} \right]_{1..n} &= d\omega_{i,1} \dots d\omega_{i,n} \end{aligned}$$

This, in turn, precludes directly using such a vector sample in the parent MC integration which cast the sample in the first place: simply adding the vector components to the overall estimate is only possible if the integration domains match. But as we will see in the fluorescent case, it is actually possible to derive compensation factors for the vector elements which still allow one to use the elements of such a dimensionally mis-matched vector sample.

Note that the above equation actually does not describe a realistic usage case for MC rendering: in regular path tracers, one specifically does not cast multiple sub-rays for one single Monte Carlo sample, as this was already shown to be inefficient by Kajia in his original paper about the Rendering Equation [Kaj86]. We only use this example to demonstrate that the domain extension effect for

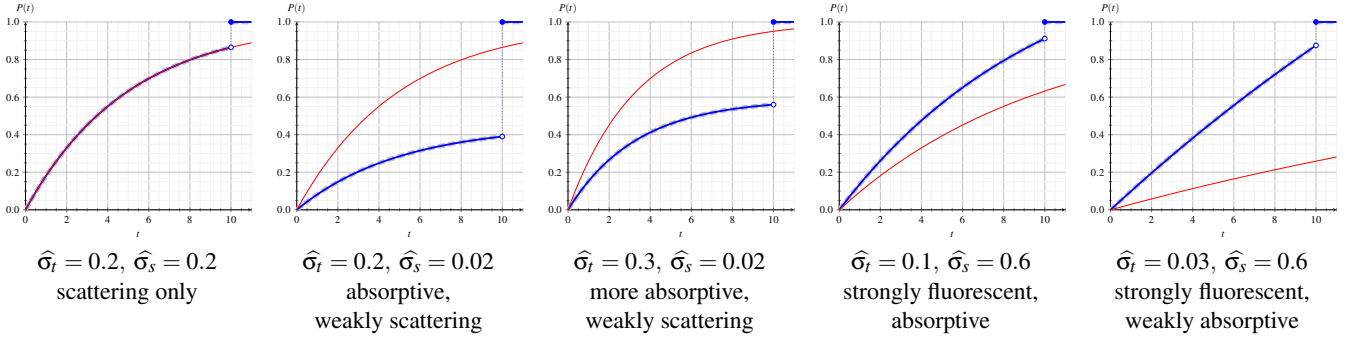


Figure 9: Cumulative distribution functions $P(t)$ of sampling distance t for varying $\hat{\sigma}_t$ and $\hat{\sigma}_s$, and a surface at distance $d = 10$, using standard exponential tracking (red) and the proposed fluorescence-aware tracking (blue). The fluorescence-aware tracking takes the extended collision coefficient into consideration to better match the distribution of the integral that is being sampled, as opposed to standard exponential tracking, which only takes the extinction coefficient into account. The desirable probability distribution to sample the RTE integral with unknown illumination is shown in dashed gray, which in all cases matches the fluorescence-aware tracking plot perfectly. As such, the proposed scheme is able to prefer skipping over the medium for $\hat{\sigma}_s < \hat{\sigma}_t$, where the in-scattered light is more attenuated, which makes the incident light from the surface more valuable. It behaves the same way for $\hat{\sigma}_t = \hat{\sigma}_s$, but prefers sampling collision events for $\hat{\sigma}_t < \hat{\sigma}_s$ (fluorescent medium), where there is a possibility of in-shifting whose out-shifting is not represented in $\hat{\sigma}_t$, which in turn makes exponential tracking incapable of properly sampling it. The last case is especially apparent for $\hat{\sigma}_t \rightarrow 0$, where exponential sampling hardly catches any collision events, while the actual contribution of the integral is linear with respect to distance d . This case is also properly handled by fluorescence-aware tracking, which in this case approaches an almost uniform sampling scheme.

vectorised equations we now also observe for the bi-spectral rendering equation [HHA*10] is actually not unique to it (first the normal form, then vectorised):

$$L_d(\mathbf{x}, \omega_0, \lambda_o) = \int_{\Lambda} \int_{\Omega} f(\mathbf{x}, (\omega_i, \lambda_i) \rightarrow (\omega_0, \lambda_o)) \cdot L(\mathbf{x}, \omega_i, \lambda_i) d\omega_i d\lambda_i$$

$$L_d^*(\mathbf{x}, \omega_0, \lambda_o^*) = \int_{\Lambda^n} \int_{\Omega} f^*(\mathbf{x}, (\omega_i, \lambda_i^*) \rightarrow (\omega_0, \lambda_o^*)) \diamond L^*(\mathbf{x}, \omega_i, \lambda_i^*) d\omega_i [d\lambda_{i,k}]_{1..n}$$

The crucial thing to observe here is that the integration dimension of the outer integral changes to Λ^n for the vectorised form. Intuitively, the reason for this is that each vector element has to integrate over Λ separately – but all of them together still constitute a single MC sample.

The effects of this domain enlargement show up when we expand a single vector element h of the vectorised form, where $1 \leq h \leq n$.

$$L_d^*(\mathbf{x}, \omega_0, \lambda_o^*) = \int_{\Lambda^n} \int_{\Omega} \begin{bmatrix} f(\mathbf{x}, (\omega_i, \lambda_{i,1}) \rightarrow (\omega_0, \lambda_{o,1})) \\ \vdots \\ f(\mathbf{x}, (\omega_i, \lambda_{i,n}) \rightarrow (\omega_0, \lambda_{o,n})) \end{bmatrix} \diamond \begin{bmatrix} L(\mathbf{x}, \omega_i, \lambda_{i,1}) \\ \vdots \\ L(\mathbf{x}, \omega_i, \lambda_{i,n}) \end{bmatrix} d\omega_i [d\lambda_{i,k}]_{1..n}$$

$$= \int_{\Lambda^n} \int_{\Omega} \begin{bmatrix} f(\mathbf{x}, (\omega_i, \lambda_{i,1}) \rightarrow (\omega_0, \lambda_{o,1})) \cdot L(\mathbf{x}, \omega_i, \lambda_{i,1}) \\ \vdots \\ f(\mathbf{x}, (\omega_i, \lambda_{i,n}) \rightarrow (\omega_0, \lambda_{o,n})) \cdot L(\mathbf{x}, \omega_i, \lambda_{i,n}) \end{bmatrix} d\omega_i [d\lambda_{i,k}]_{1..n}$$

$$= \int_{\Lambda} \begin{bmatrix} \int_{\Omega} f(\mathbf{x}, (\omega_i, \lambda_{i,1}) \rightarrow (\omega_0, \lambda_{o,1})) \cdot L(\mathbf{x}, \omega_i, \lambda_{i,1}) d\omega_i \\ \vdots \\ \int_{\Omega} f(\mathbf{x}, (\omega_i, \lambda_{i,n}) \rightarrow (\omega_0, \lambda_{o,n})) \cdot L(\mathbf{x}, \omega_i, \lambda_{i,n}) d\omega_i \end{bmatrix} \diamond \int_{\Lambda^{n-1}} \begin{bmatrix} [d\lambda_{i,k}]_{2..n} \\ \vdots \\ [d\lambda_{i,k}]_{1..n-1} \end{bmatrix} \diamond \begin{bmatrix} [d\lambda_{i,1}] \\ \vdots \\ [d\lambda_{i,n}] \end{bmatrix}$$

$$L_d^*(\mathbf{x}, \omega_0, \lambda_o^*) = \|\Lambda\|^{n-1} \cdot \begin{bmatrix} L_d(\mathbf{x}, \omega_0, \lambda_{o,1}) \\ \vdots \\ L_d(\mathbf{x}, \omega_0, \lambda_{o,n}) \end{bmatrix}$$

Where $\|\Lambda\|$ is the size of the spectral domain.

The additional term $\|\Lambda\|^{n-1}$ is caused by the enlargement of the integration domain, and can be thought of as a purely mathematical artefact that arises from linking the different vector components into a single MC sample. Note that it is a term that does not depend on k because the vector elements in the HWSS do not depend on each other. To compute it, we have to introduce a multiplicative vector term $\mathbf{t}^*((\omega_i, \lambda_i^*) \rightarrow (\omega_0, \lambda_o^*))$ for which holds:

$$\int_{\Lambda^{n-1}} \mathbf{t}^*((\omega_i, \lambda_{i,h}) \rightarrow (\omega_0, \lambda_{o,h})) [d\lambda_{i,k}]_{1..n|k \neq h} = 1$$

for any $\omega_0, \omega_i \in \Omega, \lambda_o^* \in \Lambda^n, \lambda_{i,k} \in \Lambda, h \in \llbracket 1, n \rrbracket$

which allows us considerable freedom to choose this term. One possible solution is a constant expression:

$$\mathbf{t}^*((\omega_i, \lambda_i^*) \rightarrow (\omega_0, \lambda_o^*)) = \left(\frac{1}{\|\Lambda\|^{n-1}} \right)^*$$

This would be mathematically valid, but would just trivially remove

the extra term $\|\Lambda\|^{n-1}$ from each vector element: and the issue with zero probability elements in the HWSS vector which we identified at the beginning of the sector would remain. But the correction term can actually be defined so that it solves just this problem.

In order to preserve our ability to use importance sampling on each wavelength shift separately, we need to introduce a term that specifically eliminates (i.e. zeroes out) cases that cannot be sampled. Which, in turn, means a term that evaluates to non-zero value only when all wavelength shifts are valid ($p_{\lambda_i} > 0, \forall i$). In the setting of Monte Carlo sampling, we can utilise the probabilities of sampling the other wavelength shifts:

$$\mathbf{t}^* ((\omega_i, \lambda_i^*) \rightarrow (\omega_0, \lambda_o^*)) = \begin{bmatrix} \prod_{k=2}^n p(\lambda_{i,k} \rightarrow \lambda_{o,k}) \\ \vdots \\ \prod_{k=1}^{n-1} p(\lambda_{i,k} \rightarrow \lambda_{o,k}) \end{bmatrix}$$

$$\int_{\Lambda^{n-1}} t((\omega_i, \lambda_{i,h}) \rightarrow (\omega_0, \lambda_{o,h})) \left[d\lambda_{i,k} \right]_{1..n|k \neq h} = \prod_{k \neq h=1}^n \int_{\Lambda} p(\lambda_{i,k} \rightarrow \lambda_{o,k}) d\lambda_{i,k} = \prod_{k \neq h=1}^n 1 = 1$$

This approach actually leads to a fairly natural result, where the Monte Carlo sample we would get by dividing by the probability $p(\lambda_i, \omega_i) = \left[p(\lambda_{o,t} \rightarrow \lambda_{i,t}) \right]_{t=1}^n \cdot p(\omega_i)$ ends up canceling out the additional term and leaves only $p(\lambda_{i,k}) \cdot p(\omega_i)$ for each wavelength sample, which is what we would evaluate without using the Hero wavelength approach. The probabilities we use can also be conditional ones, especially in cases where the bi-spectral BRDF [HHA*10] is not separable by wavelength shifts and directions. In such cases, it may be needed to sample wavelengths conditionally: all that matters is that they integrate to one. It is worth noting that in most cases, the two domains should actually be separable. But as long as it cannot be analytically ruled out that no inseparable cases exist, it is a good idea to work on techniques that do not require this property. Especially as this leaves the freedom to derive analytical models which are not separable, which might come in handy for combined models of fluorescent media and surfaces.

The introduced formalism describes the sampling of wavelengths on all samples to be part of a single sampling process, which provides a clear approach as to how to compute MIS weights by computing a single n -sized vector of sampling probabilities. If each of the paths would instead be considered as a separate Monte Carlo sample, the computation of MIS weights starts to get more complicated, especially when considering non-separable events, which could involve the necessity to propagate n^2 probabilities.

7. Implementation & Results

7.1. Distance tracking comparisons

Figure 9 shows characteristic differences in computed distance sampling probabilities between standard exponential tracking and our proposed fluorescence-aware tracking for different combinations of material parameters.

For $\hat{\sigma}_s = \hat{\sigma}_t$, both approaches are identical, apart from the artificial discontinuity at $d = 10$, which technically also occurs with



(a) Standard exponential tracking (b) Exponential tracking using $\max(\hat{\sigma}_t, \hat{\sigma}_s)$ (c) The proposed fluorescence-aware tracking.

Figure 10: A comparison of distance tracking approaches used on a scene that is filled with a fluorescent medium which exhibits the wavelength asymmetry in its interaction with light which we show in figures 2 and 7. Images rendered with 512 samples per pixel. Notice the complete absence of fluorescence events in the case of standard exponential tracking. There are actually two light sources in this box, one in front of the grey dividing polygon, and one behind it. The light nearer to the observer emits no UV light, while the one at the back of the box does. The fluorescent medium in this scene only transfers energy from the UV band to the yellow-green region of the spectrum, and is not activated by light in the visible range. Apart from the fluorescence effect, it is completely clear. The type of medium is similar to that seen in figure 1, although absorption and re-emission wavelengths are different. Figures 18 and 19 show the effect of varying volume absorption and fluorescence, respectively.

exponential sampling in the presence of an object that is present at a finite distance. When $\hat{\sigma}_s < \hat{\sigma}_t$, the fluorescence-aware sampling is more prone to skipping the medium (as in sampling the distance d), as it attempts to prioritise the incoming light from the surface over the weak in-scattered light. For $\hat{\sigma}_s = 0$, it would always skip the medium. However, the higher the $\hat{\sigma}_t$, the higher the influence of transmittance on fluorescence-aware tracking is, and a higher $\frac{\hat{\sigma}_t}{\hat{\sigma}_s}$ ratio is required to induce similar skipping behaviour. The fluorescence aware tracking can therefore be used as a direct replacement for exponential tracking in homogeneous media, where it either functions identically, or takes the wavelength shift/absorption ratio into account, and skips more absorptive media.

Only in fluorescent media can the case of $\hat{\sigma}_s > \hat{\sigma}_t$ occur (and specifically, even then only on receiving wavelengths): there, the fluorescence-aware tracking has an increased probability of sampling a collision event as opposed to exponential tracking. For sufficiently low σ_s and high $\hat{\sigma}_s$, the distribution function approaches that of uniform sampling over the interval $[0, d]$. It even completely turns into uniform sampling for $\hat{\sigma}_t = 0$, as can be seen from the split in section 5.3. This tendency to approach uniform sampling is again influenced by how low $\hat{\sigma}_t$ is, where again for high $\hat{\sigma}_t$, the transmittance takes over for low $\frac{\hat{\sigma}_t}{\hat{\sigma}_s}$ ratios.

7.2. Renderings

Figure 10 showcases the effects of the volume material wavelength asymmetry that is introduced in figures 2 and 7. The set-up shown in figure 10 is a worst-case scenario insofar as the scene is entirely filled with a fluorescent medium that is totally clear in the




Figure 11: The Cornell box contains a floating cube filled with a constant density fluorescent volume that absorbs at 450nm, and re-emits at 550nm. In 11a, note the yellow shadow on the floor caused by absorption at 450nm (blue). This lack of energy is visible on the difference image 11d while 11f shows how the volume acts as a light source, on the re-emission band, and emits light all around due to fluorescence. When illuminating the scene with monochromatic light in the absorption band 11c, the volume is almost opaque, and re-emits strongly (note the yellow ceiling, which would normally be almost black under 440nm mono illumination). When a monochromatic lightsource outside the absorption band 11e is used, the volume is completely clear: this behaviour is the same as the one shown in a real scene in figure 2. Note that the red and green colour of 11d and 11f is falsecolour for signed difference images: red means a negative difference, and green a positive one. The fact that the two images are each a single colour shows that energy is indeed shifted one way between these wavelengths.

visible range: and some parts of the scene are being illuminated by light sources that excite the fluorophores, while others are not. This scene is rendered using both standard exponential tracking, and our proposed new distance tracking approach. Specifically, exponential tracking was tested separately in its naive form, where only σ_{s_e} is used as the coefficient for the inverse mean free path, as well as an altered version that uses $\max(\hat{\sigma}_s, \hat{\sigma}_t)$ as the coefficient (this could be seen as an obvious fix for exponential tracking). As can be seen, naive exponential tracking does not yield any energy transport via fluorescence effects: as expected, its sampling scheme completely misses any down-shifted radiance. While the "quick fix" of modified $\max(\hat{\sigma}_s, \hat{\sigma}_t)$ exponential tracking produces an image that at least shows fluorescence, it is still extremely noisy, even when one uses a very high number of samples per pixel. By comparison, fluorescence-aware tracking produces an image with comparatively low noise levels: it is the only technique that generates an image which looks acceptable.




Figure 12: A comparison of the two distance tracking approaches used on a scene that is filled with a non-fluorescent medium which exhibits wavelength-dependent scattering and absorbance to varying degrees: 12a shows standard exponential tracking, 12b the proposed fluorescence-aware tracking. The run-times were comparable, so this comparison shows that the proposed tracking does not break anything for scenes which do not contain fluorescent media, and can be universally used in a fluorescence-aware renderer. Rendered with approx. 3000 samples per pixel. SNR reference image rendered with approx. 10 thousand samples per pixel.

Figure 11 shows a limited-size homogeneous volume in a Cornell box, and also contains some difference images which demonstrate the effect that the re-emissive volume has on the surrounding environment. Figure 12 shows that the proposed tracking scheme can be left switched on in scenes that contain only normal volumes, while figure 14 demonstrates that fluorescence-aware tracking produces normal and expected results for scenes without any fluorescent volumes, as utilizing the Hero Wavelength Bi-spectral Sampling, since the scene contains a fluorescent sphere. Figure 17 showcases the convergence behaviour of monochrome fluorescence tracing compared to our proposed HWSS solution.

7.3. Monochrome low sample count colour cast

A noticeable feature of some initial low sample count renderings of fluorescent media is a colour cast: examples are figure 1, where

the cast is cyan, and figure 17, where it is purple. In both cases, the effect is only strong in monochrome renderings of media with a comparatively narrow emission band (!), and largely disappears when HWSS is being used. Also, the effect is more pronounced for $\max(\hat{\sigma}_s, \hat{\sigma}_t)$ tracking. In all cases, as convergence goes up with increasing sample count, the colour cast completely disappears.

This peculiar behaviour is a direct consequence of using a uni-directional path tracer to render an asymmetric fluorescent medium. Note that both our proposed "hotfix" and our actual fluorescence tracking have to generate fluo collision events in the emission band of the medium in question: which means that eye-paths which are on these wavelengths stand a non-zero chance of being diverted via a wavelength-changing event, while others do not. As the variance of sampling fluo events is higher than the variance of paths which are not diverted, we initially see a more converged image on the surviving wavelengths: in the case of figure 17, green paths are diverted, while red and blue paths are left alone (and yield an initial purple colour cast). The stronger colour cast of the $\max(\hat{\sigma}_s, \hat{\sigma}_t)$ tracking is due to its less efficient generation of fluo collision events: as a simpler heuristic, it generates a sub-optimal set of tracking events, which increases variance: and this is most noticeable at low sample counts.

7.4. Performance considerations



(a) SNR: 23.063955 dB
Exponential tracking


(b) SNR: 32.277494 dB
Fluo-aware tracking

Figure 13: A comparison of the two distance tracking approaches used on a scene that is filled with a fluorescent medium that exhibits regular scattering alongside fluorescence with an only weak wavelength asymmetry: 13a shows standard exponential tracking, 13b the proposed fluorescence-aware tracking. Rendered with approx. 3000 samples per pixel. SNR reference image rendered with approx. 10 thousand samples per pixel.

All three images in figure 10 were rendered using 512 samples per pixel, with HWSS enabled. A fully converged version of this image still takes around 20k samples, which is significantly more than one would normally need for a converged image of a Cornell box scene. So we freely admit that our proposed sampling scheme is still not as efficient as one might wish. Bear in mind, though, that prior to our work, the best option one would have had would have been case (b) in that figure: and even the $\max(\hat{\sigma}_s, \hat{\sigma}_t)$ "hotfix" used there was, at least to our knowledge, not published before. So essentially, prior to our work, asymmetric fluorescent media in

a uni-directional spectral path tracers simply did not work: all you had to go by was case (a), which failed to capture the effect entirely.

It should also be noted that the scene shown in figure 10 is a replication of a "worst case material" with practically no normal scattering: this does occur in reality, as figure 2 shows, so developing an ability to render such materials is a practical concern. However, figure 13 demonstrates that materials which also exhibit normal scattering alongside volumetric fluorescence are actually considerably less prone to pathological performance when rendered with normal exponential tracking. Our proposed technique still yields a measurable improvement in such cases, though.



(a) Illuminated by an equal energy (b) Illuminated by visible light only. spectrum that covers both the visible and the UV portion of the spectrum.

Figure 14: A Cornell box with a large light and a fluorescent sphere, with fluorescence excitation in the UV band, and emission in the yellow-green band. The image was rendered by using direction sampling only. The fluorescent component of the surface is only visible when UV light is present, and disappears when illuminated with visible light only. Image rendered with a Hero WL tracer that uses our proposed wavelength selection scheme. See figures 16, 17, 18 and 19 in the appendix for image series of other fluorescence scenes at increasing convergence levels: these demonstrate the convergence characteristics of our proposed Hero WL tracing.

7.5. Rendering inhomogeneous fluorescent volumes

As discussed so far, the new tracking scheme we introduced is only applicable to homogeneous volumes which exhibit fluorescence. For media which exhibit extreme properties like the liquid shown in figure 2, the new tracking scheme allows uni-directional rendering techniques to generate images – something which they were not able to do before. However, as figures 13 and 18 demonstrate, the new scheme is only essential for media that are entirely transparent apart from their fluorescent component: homogeneous fluorescent media which exhibit at least moderate absorption and/or scattering can actually be rendered with standard exponential tracking (not efficiently, but they can be rendered). The proposed scheme is still usable in those cases, though, and usually still a bit more efficient.

Extending the proposed tracking scheme in its full form to inhomogeneous media would not be that easy: but fortunately, we found that the simple fix of using $\max(\hat{\sigma}_s, \hat{\sigma}_t)$ during standard Woodcock tracking seems to be already sufficient in a large number of




Figure 15: A heterogeneous medium dataset from Mitsuba in a Cornell box, rendered using a fluorescent medium with an $\widehat{\sigma}_s$ that absorbs at 360nm, and emits at 560nm with no regular absorption coefficient σ_a . Notice the absence of a shadow in the visible range 15a. The shadow appears in the visualisation of the absorption band in 15b (this waveband is already outside the range of human vision, but as our renderer computes spectral images that go into the UV range, we can show this). Outside the absorption or re-emission range, the medium is perfectly clear as shown in 15d.

cases. This actually stands to reason insofar as the main problem with homogeneous media with extreme fluorescence properties is the absence of sampling points at which the tracking can stop: once sampling points are generated, a path tracer will also sample wavelength-changing events. And pretty much by definition, an inhomogeneous medium will have variations in its characteristics which trigger the generation of sample points: which in turn makes them less susceptible to the problem of "rays going all the way through without stopping" – if the $\max(\widehat{\sigma}_s, \widehat{\sigma}_t)$ fix is used, that is. Figure 15 shows that the fix is good enough to yield usable results for volumetric datasets.

7.6. Limitations

The most obvious complication is as to how to deal with infinite volumes, since d is used within the calculations. For most cases, when we examine the formula in section 5.3 and take a limit for $d \rightarrow \infty$, we can see that the case for $\widehat{\sigma}_t > 0$ degenerates to standard exponential sampling, since the contribution of the "surface" ends up attenuated over an infinitely long ray. However, an infinite medium with purely fluorescent properties (as in, for a certain wavelength it holds that $\widehat{\sigma}_s > 0$ and $\widehat{\sigma}_t = 0$) is almost impossible to sample without actively considering the illumination situation, as the RTE integral could evaluate to infinity.

8. Conclusions and Future Work

We made two contributions to enable fluorescence rendering in modern spectral path tracers: we showed how Hero wavelength sampling can be used in conjunction with such materials, and we propose a distance tracking scheme that is capable of handling homogeneous fluorescent volumes. This allows one, for the first time ever, to use asymmetric fluorescent volumes in an uni-directional path tracer. However, the efficiency of our technique still leaves obvious room for improvement: but at least the effect is now basically tractable. In the future, we will mainly attempt to increase the efficiency of the sampling scheme.

Acknowledgements

This work was financed by the Czech Science Foundation under grant number 16-08111S, and also via the European Union Horizon 2020 research and innovation programme, under the Marie Skłodowska-Curie grant agreement No 642841 (DISTRO).

References

- [Abd17] ABDELLAH M.: *In Silico Brain Imaging: Physically-plausible Methods for Visualizing Neocortical Microcircuitry*. PhD thesis, Neuroscience Doctoral School (EDNE), Blue Brain Project (BBP), Ecole Polytechnique Federale de Lausanne, 2017. 4
- [ABE*17] ABDELLAH M., BILGILI A., EILEMANN S., MARKRAM H., SCHÜRMAN F.: A physically plausible model for rendering highly scattering fluorescent participating media. *CoRR abs/1706.03024* (2017). URL: <http://arxiv.org/abs/1706.03024>. 4
- [BHD*08] BENDIG M., HANIKA J., DAMMERTZ H., GOLDSCHMIDT J. C., PETERS M., WEBER M.: Simulation of fluorescent concentrators. In *2008 IEEE Symposium on Interactive Ray Tracing* (Aug 2008), pp. 93–98. doi:10.1109/RT.2008.4634628. 4
- [DCWP02] DEVLIN K., CHALMERS A., WILKIE A., PURGATHOFER W.: Tone reproduction and physically based spectral rendering. In *STAR Proceedings of Eurographics 2002* (Saarbrücken, Germany, Sept. 2002), Fellner D., Scopigno R., (Eds.), Eurographics Association. 2, 5
- [ENSB13] EISENACHER C., NICHOLS G., SELLE A., BURLEY B.: Sorted deferred shading for production path tracing. *Computer Graphics Forum* 32, 4 (2013), 125–132. 2
- [FHF*17] FASCIONE L., HANIKA J., FAJARDO M., CHRISTENSEN P., BURLEY B., GREEN B.: Path tracing in production - part 1: Production renderers. In *ACM SIGGRAPH 2017 Courses* (New York, NY, USA, 2017), SIGGRAPH '17, ACM, pp. 13:1–13:39. URL: <http://doi.acm.org/10.1145/3084873.3084904>, doi:10.1145/3084873.3084904. 2
- [FWKH17] FONG J., WRENNINGE M., KULLA C. D., HABEL R.: Production volume rendering: Siggraph 2017 course. In *SIGGRAPH Courses* (2017), ACM, pp. 2:1–2:79. URL: <http://doi.acm.org/10.1145/3084873.3>, 4
- [Gla94] GLASSNER A.: A model for fluorescence and phosphorescence. In *Fifth Eurographics Workshop on Rendering* (Darmstadt, Germany, June 1994), Eurographics, pp. 57–68. 3, 6

- [HHA*10] HULLIN M. B., HANIKA J., AJDIN B., SEIDEL H.-P., KAUTZ J., LENSCH H. P. A.: Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph.* 29, 4 (July 2010), 97:1–97:7. 3, 5, 6, 9, 10
- [Kaj86] KAJIYA J. T.: The rendering equation. In *Computer Graphics (ACM SIGGRAPH '86 Proceedings)* (Aug. 1986), vol. 20, pp. 143–150. 8
- [KGH*14] KRIVANEK J., GEORGIEV I., HACHISUKA T., VEVOVA P., SIK M., NOWROUZEZHRAI D., JAROZ W.: Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics* 33, 4 (July 2014). 2
- [KHLN17] KUTZ P., HABEL R., LI Y. K., NOVÁK J.: Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Trans. Graph.* 36, 4 (2017), 111:1–111:16. URL: <http://doi.acm.org/10.1145/3072959.3073665>. 4
- [KM06] KIM A. D., MOSCOSO M.: Radiative transport theory for optical molecular imaging. *Inverse Problems* 22, 1 (2006), 23. URL: <http://stacks.iop.org/0266-5611/22/i=1/a=002>. 2
- [KNH05] KLOSE A. D., NTZIACHRISTOS V., HIELSCHER A. H.: The inverse source problem based on the radiative transfer equation in optical molecular imaging. *Journal of Computational Physics* 202, 1 (2005), 323 – 345. URL: <http://www.sciencedirect.com/science/article/pii/S0021999104002888>, doi:<https://doi.org/10.1016/j.jcp.2004.07.008>. 2
- [MSHD15] MENG J., SIMON F., HANIKA J., DACHSBACHER C.: Physically meaningful rendering using tristimulus colours. *Computer Graphics Forum* 34, 4 (2015). 2
- [RBA09] RADZISZEWSKI M., BORYCZKO K., ALDA W.: An improved technique for full spectral rendering. *Journal of WSCG* 17, 1-3 (2009), 9–16. 2
- [SFD98] SUN Y., FRACCHIA F. D., DREW M. S.: *A Composite Model for Representing Spectral Functions*. Tech. Rep. TR 1998-18, School of Computing Science, Simon Fraser University, Burnaby, BC, Canada, Nov. 1998. URL: <ftp://fas.sfu.ca/pub/cs/TR/1998/CMPT1998-18.pdf>. 2
- [SHZD17] SIMON F., HANIKA J., ZIRR T., DACHSBACHER C.: Line integration for rendering heterogeneous emissive volumes. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 36, 4 (2017). 4
- [Tat17] TATARCHUK N.: Advances in real-time rendering, part i. In *ACM SIGGRAPH 2017 Courses* (New York, NY, USA, 2017), SIGGRAPH '17, ACM. URL: <http://doi.acm.org/10.1145/3084873.3096476>, doi:10.1145/3084873.3096476. 2
- [WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero wavelength spectral sampling. *Comput. Graph. Forum* 33, 4 (2014), 123–131. 2
- [WTP01] WILKIE A., TOBLER R. F., PURGATHOFER W.: Combined rendering of polarization and fluorescence effects. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques, London, UK, June 25-27, 2001* (2001), Gortler S. J., Myszkowski K., (Eds.), Springer, pp. 197–204. 3
- [WWLP06] WILKIE A., WEIDLICH A., LARBOULETTE C., PURGATHOFER W.: A reflectance model for diffuse fluorescent surfaces. In *Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia 2006, Kuala Lumpur, Malaysia, November 29 - December 2, 2006* (2006), Lee Y. T., Shamsuddin S. M. H., Gutierrez D., Suaib N. M., (Eds.), ACM, pp. 321–331. URL: <http://doi.acm.org/10.1145/1174429.1174484>. 3

Appendix A: Deriving fluorescence-aware distance sampling.

We have set-up the preliminary sampling probability density function as follows:

$$p'(t) = \begin{cases} \widehat{\sigma}_s \cdot e^{-t \cdot \widehat{\sigma}_t} & t < d \\ e^{-d \cdot \widehat{\sigma}_t} \cdot \delta(t-d) & t \geq d \end{cases}$$

To compute the exact value, we have to compute the normalization function n , so that integrating over $p(t) = p'(t) \cdot n$ evaluates to one:

$$\begin{aligned} n &= \int_0^\infty p(t) dt = \int_0^d \widehat{\sigma}_s \cdot e^{-t \cdot \widehat{\sigma}_t} dt + e^{-d \cdot \widehat{\sigma}_t} \\ &= \left[-\frac{\widehat{\sigma}_s}{\widehat{\sigma}_t} e^{-t \cdot \widehat{\sigma}_t} \right] + e^{-d \cdot \widehat{\sigma}_t} = -\frac{\widehat{\sigma}_s}{\widehat{\sigma}_t} \cdot e^{-d \cdot \widehat{\sigma}_t} + \frac{\widehat{\sigma}_s}{\widehat{\sigma}_t} + e^{-d \cdot \widehat{\sigma}_t} \\ &= \frac{\widehat{\sigma}_s + (\widehat{\sigma}_t - \widehat{\sigma}_s) \cdot e^{-d \cdot \widehat{\sigma}_t}}{\widehat{\sigma}_t} \\ &\text{where } \widehat{\sigma}_t > 0 \end{aligned}$$

$$\begin{aligned} n &= \int_0^\infty p(t) dt = \int_0^d \widehat{\sigma}_s \cdot e^{-t \cdot 0} dt + e^{-d \cdot 0} \\ &= [t \cdot \widehat{\sigma}_s] + 1 = d \cdot \widehat{\sigma}_s + 1 \\ &\text{where } \widehat{\sigma}_t = 0 \end{aligned}$$

The cumulative distribution function contains a discontinuous jump $P(d) = 1$ due to the Dirac function in the second case of $p(t)$. Up until that point, that is for $t < d$, we get:

$$\begin{aligned} P(t) &= \int_0^t p(l) dl = \int_0^t \frac{\widehat{\sigma}_s}{c} \cdot e^{-l \cdot \widehat{\sigma}_t} dl = \frac{\widehat{\sigma}_s \cdot (1 - e^{-t \cdot \widehat{\sigma}_t})}{\widehat{\sigma}_s + (\widehat{\sigma}_t - \widehat{\sigma}_s) \cdot e^{-d \cdot \widehat{\sigma}_t}} \\ &= \frac{1 - e^{-t \cdot \widehat{\sigma}_t}}{1 + \left(\frac{\widehat{\sigma}_t}{\widehat{\sigma}_s} - 1\right) \cdot e^{-d \cdot \widehat{\sigma}_t}} \\ &\text{where } \widehat{\sigma}_t > 0 \end{aligned}$$

$$\begin{aligned} P(t) &= \int_0^t p(l) dl = \int_0^t \frac{\widehat{\sigma}_s}{c} \cdot e^{-l \cdot 0} dl = \frac{t \cdot \widehat{\sigma}_s}{d \cdot \widehat{\sigma}_s + 1} \\ &\text{where } \widehat{\sigma}_t = 0 \end{aligned}$$

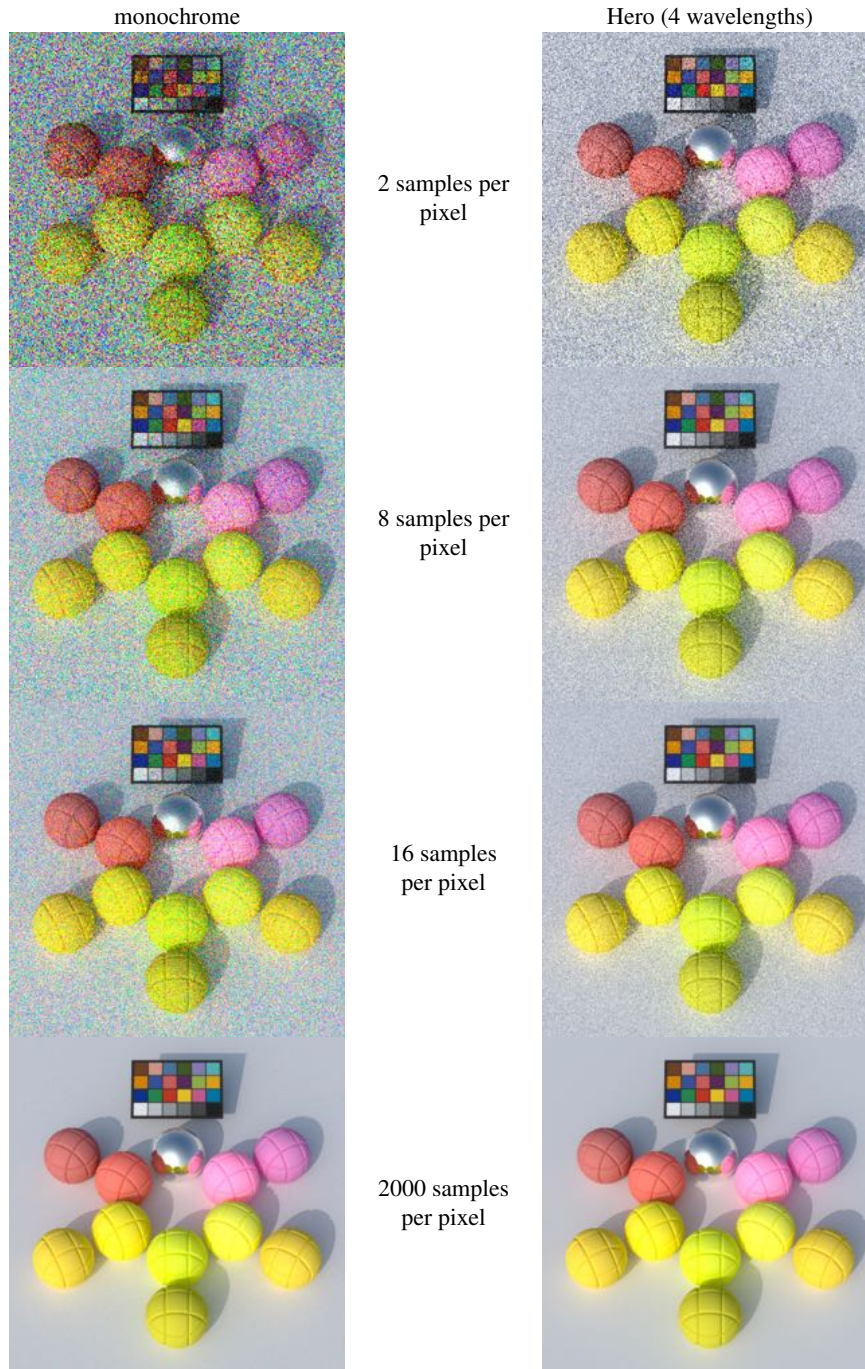


Figure 16: A scene featuring a colour checker and several spheres, illuminated by a physically based sky-dome model which includes a UV component. The spheres in the inner ring all have various fluorescent surfaces, while the corresponding spheres in the outer ring have exactly the same BRDF, just with the re-radiation component set to zero. Rendered with both a monochrome (**left column**) and a Hero (**right column**) path tracer that uses 4 wavelengths. The latter uses our proposed wavelength selection scheme for fluorescent events. The purpose of this figure is to demonstrate the viability of the proposed combination of Hero sampling and fluorescence. As expected, the colour noise is practically eliminated in the right column, even on the fluorescent surfaces. Note, however, that some small amounts of residual colour noise can still be observed on the inner spheres for low sample counts - see also figure 19. The monochrome path tracer eventually catches up and eliminates colour noise as well, but only at a very high sample count. Note that the spatial noise of Hero, 2 samples per pixel (first row on the right) is equivalent to the spatial noise of monochrome, 8 samples per pixel (second row on the left), since the each Hero sample has the possibility of contributing as much as 4 samples.

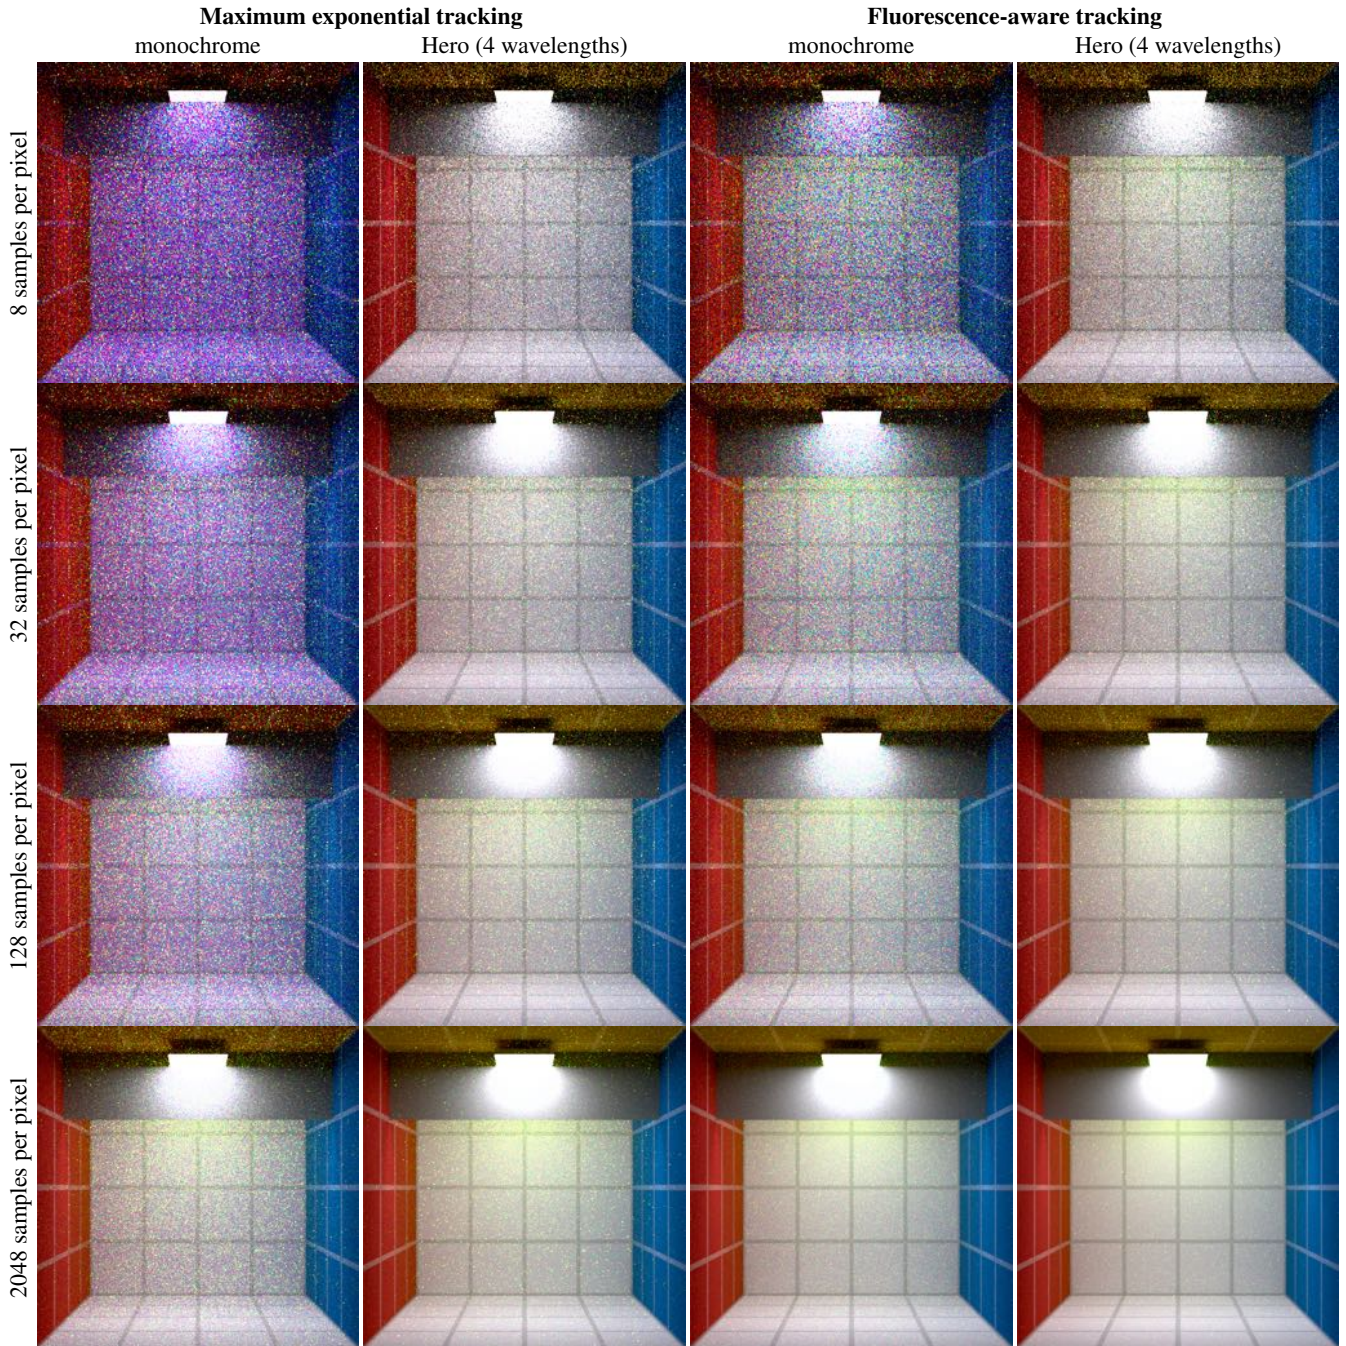


Figure 17: Comparison of monochrome and HWSS for maximum exponential tracking and fluorescence-aware tracking. At 8 samples per pixel, with maximum exponential tracking without HWSS, the colour noise distorts the fluorescent volume colour very strongly, while a good preview of the scene is already possible when using our fluorescent aware tracking combined with HWSS. The purple colour cast of the monochrome initial low sample count images is discussed in section 7.3.

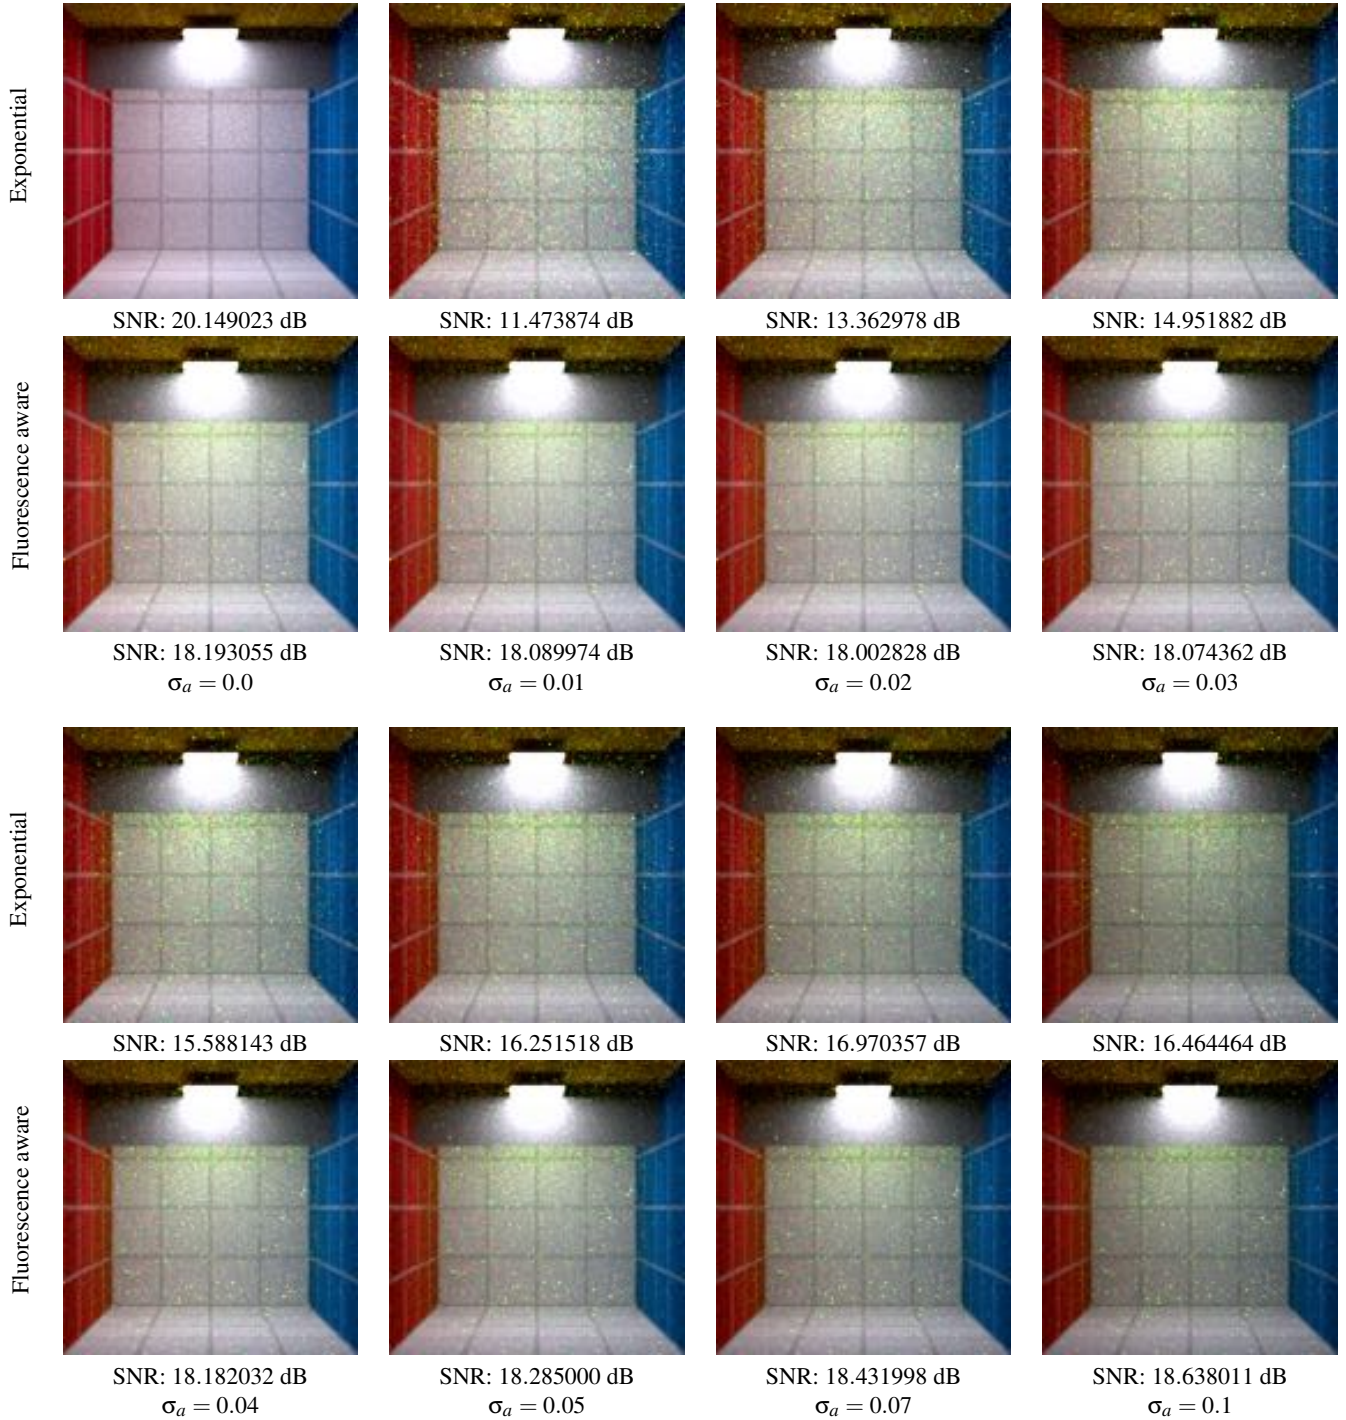


Figure 18: A modified form of the scene from figure 10 where the absorption coefficient $\sigma_a(\lambda)$ in the visible range is continually increased, to provide a comparison of how the two distance tracking schemes perform for media that range from totally clear, to somewhat absorptive. **Top rows:** normal exponential tracking **Bottom rows:** fluorescence-aware tracking. As one can see, normal exponential tracking fails as expected for media which only exhibit fluorescence (top left), but catches up as soon as there is some absorption that can trigger tracking events. Rendered with 128 samples per pixel to give a snapshot of early convergence behaviour for similar sample counts. SNR reference images rendered with 16 thousand samples per pixel.

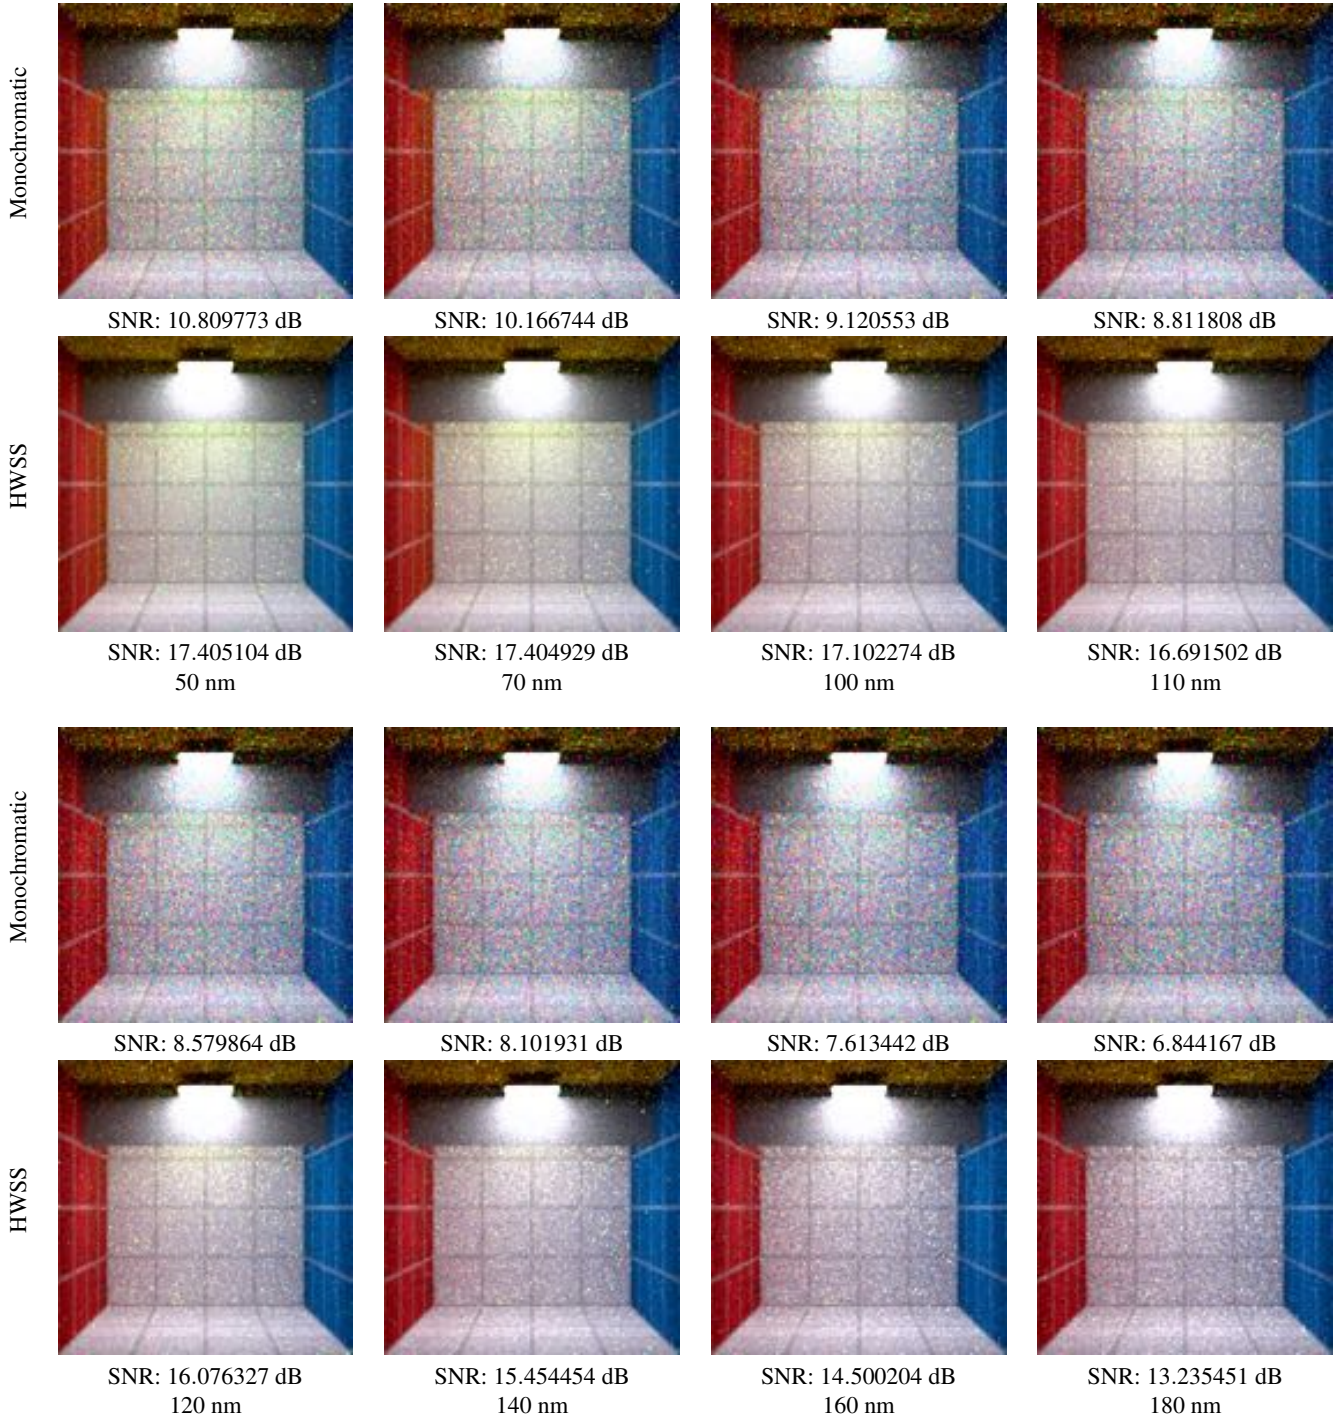


Figure 19: A modified version of the scene from figure 10. Here, the fluorescent medium has a continually widening emission band of the indicated width, centred on 560 nm. This test is aimed at the colour noise performance of the path tracer. **Top rows:** rendered with a monochrome path tracer **Bottom rows:** output of a Hero path tracer (4 wavelengths per path) that uses our proposed scheme for wavelength selection. Rendered with 128 samples per pixel to give a snapshot of early convergence behaviour for similar sample counts. SNR reference images rendered with 16 thousand samples. The noteworthy thing to observe here is how for very narrow emission bands, even the Hero sampler has difficulties finding other wavelengths that still fall within the Stokes shift, which in turn causes colour noise. The wider the emission band becomes, the easier this becomes, and the colour performance improves as a consequence.