





Real-Time Underwater Spectral Rendering

Nestor Monzon ¹, Diego Gutierrez ¹, Derya Akkaynak ^{2,3} and Adolfo Muñoz ¹

¹Universidad de Zaragoza - I3A ²University of Haifa ³IUI Eilat

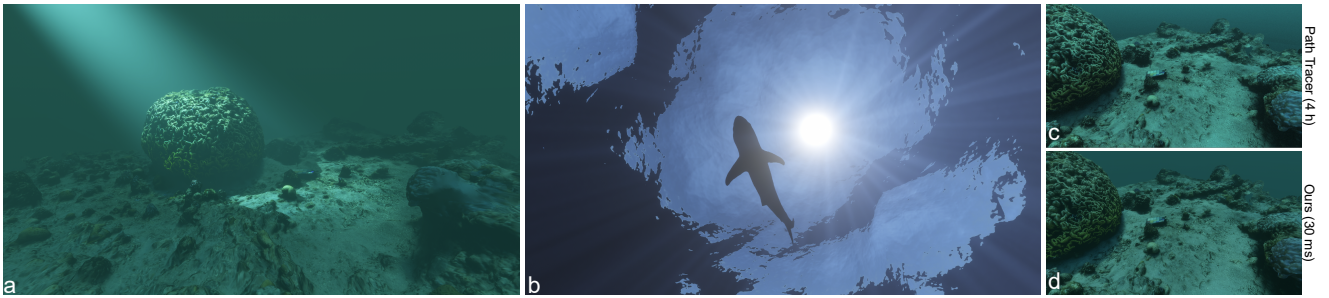


Figure 1: Natural water bodies can show large appearance variations, due to their different optical constituents. For example, coastal, more turbid waters (a) are often greener than open oceanic waters (b). Furthermore, the complex interactions between light and water occur at a great range of scales, from the microscopic scattering events to reflections and refractions at the water surface (see for instance the darker areas in the second image, due to total internal reflection outside of Snell’s window). Accurate simulation of these effects with a spectral path tracer is prohibitively expensive for any interactive application: this reference image (c) was rendered in approximately four hours. In this paper, we propose a spectral method based on oceanographic measurements to achieve real-time performance (around 60 fps) while rendering results (d) close to ground-truth images.

Abstract

The light field in an underwater environment is characterized by complex multiple scattering interactions and wavelength-dependent attenuation, requiring significant computational resources for the simulation of underwater scenes. We present a novel approach that makes it possible to simulate multi-spectral underwater scenes, in a physically-based manner, in real time. Our key observation is the following: In the vertical direction, the steady decay in irradiance as a function of depth is characterized by the diffuse downwelling attenuation coefficient, which oceanographers routinely measure for different types of waters. We rely on a database of such real-world measurements to obtain an analytical approximation to the Radiative Transfer Equation, allowing for real-time spectral rendering with results comparable to Monte Carlo ground-truth references, in a fraction of the time. We show results simulating underwater appearance for the different optical water types, including volumetric shadows and dynamic, spatially varying lighting near the water surface.

1. Introduction

The ability to simulate the appearance of underwater scenes is crucial for a wide number of applications, in particular as an analysis-by-synthesis approach in domains such as underwater imaging and remote sensing. For such applications, beyond a plausible and appealing appearance, the simulation needs to provide physical correctness, considering both the optical properties of water as well as the spectral response curve of the sensors. Furthermore, water absorption and scattering have very characteristic spectral profiles with strong wavelength dependencies, so spectral rendering is a must for accurate color reproduction of underwater scenes.

Spectral light transport simulation of participating media is, on its own, a daunting and time-consuming task, particularly because interactions between light and matter can happen at every differential point of its trajectory, and because of its recursive nature, where every differential scattering interaction generates new light trajectories. The rendering time of offline approaches, on the order of hours, is not practical for many applications for which interactivity is a requirement.

In this paper, we propose a spectral, physically-based, real-time rendering algorithm to simulate underwater light transport at practical frame rates (more than 60 fps in an NVIDIA GeForce GTX

1660 Super). Our method aims to bridge the gap between common real-time solutions in game engines, that are neither spectral nor physically-based enough, and offline simulation techniques, that are prohibitively expensive for interactive applications. Such a speedup in simulations has benefits for a number of applications, including underwater computer vision, visual ecology applications, submarine or remotely operated vehicle training tasks using AR/VR (Augmented/Virtual Reality) setups, or the entertainment industry.

To achieve real-time frame rates, we separate single scattering from multiple scattering, and provide physically-based approximations for each while preserving physical accuracy as much as possible. For **single scattering**, we rely on froxel lighting [Vos14, Hil15, Kov20] for spatially-varying illumination while accounting for the physically-based behavior of light transport (Section 4.1). For **multiple scattering** we rely on the apparent optical property of *diffuse downwelling attenuation* to characterize underwater downwelling irradiance, as usual in ocean optics (Section 3.2). Through some derivations on the Radiative Transfer Equation [Cha50], this enables us to devise an expression that can be evaluated in constant time per wavelength (Section 4.2).

Our results work at interactive frame rates, for a wide number of real-world water types. This simulation of different water types is made possible thanks to their routine characterization in oceanography, and the availability of the real world measurements of the coefficients [SM15, WBF*03]. We analyze the importance of spectral rendering on underwater color reproduction. Our results also provide a rather accurate match compared to a full light transport simulation using path tracing in a minimal fraction of the time.

We also show how our method can be easily integrated with a number of real-time techniques in physically based pipelines, which allows to simulate a variety of underwater phenomena such as moving caustics, the dielectric water surface or the Snell window, as we can see in Figure 1.

2. Related work

Commonly, full light transport simulations of underwater appearance numerically approximate the physics of the real world through Monte Carlo methods, at the cost of extremely long rendering times. These Monte Carlo methods are often improved by bidirectional techniques such as photon mapping [Jen96], which accelerate the convergence of the projected caustic patterns and their volumetric beams and shadows. Frisvad and colleagues [FCJ07] generalized the Lorenz-Mie theory to rendering natural waters, while Gutierrez et al. [GSM08] explored the effects of inelastic scattering (common in oceans due to phytoplankton) in the appearance of water. We refer the reader to a complete review of the most advanced offline techniques recently presented by Droske et al. [DHV*23]. Cerezo et al. [CPP*05] provided a general review on volume rendering techniques.

Real-time techniques. Few works tackle the particularities of physically based rendering of underwater scenes in real time. In game engines, artistic control is often more important than physical accuracy. For example, in *Subnautica* (2018) the rules of physics

are broken to allow for more compelling colors and improved visibility [Dev16]. The most common approach in real-time engines is just a tinted fog approach, where radiance is attenuated exponentially. Ambient light reaching both the surface and the medium is usually made depth-dependant, but without following physical interactions [LKA*17]. These approaches are controllable and computationally inexpensive, but do not allow to render different water types directly based on their coefficients. Moreover, real-time approaches usually adopt fog rendering techniques, which are only suitable for the atmosphere since they assume monochromatic extinction. In water bodies, however, extinction has a strong wavelength dependency [ATS*17, AT18].

Chen et al. [CZSZ22] propose a controllable model for real-time underwater lighting, by making use of a depth-dependant irradiance approximation used in oceanography for ambient illumination. However, they approximate lighting using wideband attenuation coefficients in XYZ space for efficiency, which can lead to large errors in the final images. Furthermore, their approach models ambient illumination as a constant for the whole scene (not accounting for differences in vertical depth), and does not take into account the angular dependence of the *veiling light* (the underwater saturated color, visible when the viewing path is long enough).

Our proposed solution addresses these limitations, increasing accuracy by computing radiance for more wavelength bands, including both single scattering and caustics, and also allowing to simulate the different water types and depths (see Section 3). Furthermore, our method is flexible, capable of reproducing the captured color by different sensors, making it useful for oceanographic studies that employ consumer cameras.

Underwater light diffusion. Diffusion has been a popular approximation for complex multiple scattering problems in fields such as neutron transport [Cas60] and subsurface scattering rendering [JMLH01]. These methods usually rely on reducing these interactions to a one-dimensional problem. In this context, Premoze and Ashikhmin [PA01] developed a model of the diffusion of light in the water volume based on the ocean optics literature and measurements. Although their method was offline and focused on an outside viewer, we propose a similar approach for multiple scattering that allows to simulate the underwater perspective, accounting for the seafloor reflectance, and model different sensors, in real time. We validate our model with respect to full light transport simulations.

3. Theoretical background

3.1. Radiative transfer

The Radiative Transfer Equation (RTE) [Lom87, Cha50] models the scalar radiance transfer (i.e. without polarization or inelastic scattering) in a differential manner. In computer graphics, we usually formulate the RTE in its integral form, also known as the volume rendering equation (VRE) [FWKH17]. This expression models the final radiance $L(\mathbf{o}, \omega)$ reaching the camera (at point \mathbf{o} looking towards ω , see Table 1 for an overview of the symbols), traveling a distance P :

Symbol	Description
λ	Wavelength
σ_s	Scattering coefficient (m^{-1})
σ_a	Absorption coefficient (m^{-1})
σ_t	Extinction coefficient. $\sigma_t = \sigma_s + \sigma_a$ (m^{-1})
f_s	Phase function
K_d	Diffuse downwelling attenuation coefficient
\mathbf{o}	3D position of the sensor
$T(v)$	Transmittance in the medium for a beam travelling v units.
L	Radiance
E	Irradiance (E_D downwelling, E_U upwelling)

Table 1: Main symbols used in the paper.

$$\begin{aligned}
L(\mathbf{o}, \omega) &= T(P)L_s(\mathbf{o} + P\omega, \omega) + L_m(\mathbf{o}, \omega) \\
&= T(P)L_s(\mathbf{o} + P\omega, \omega) + \int_0^P T(v)L_i(\mathbf{o} + v\omega, \omega) dv, \quad (1)
\end{aligned}$$

where $T(P) = e^{-\sigma_t P}$ is the transmittance of the water for distance P . The extinction coefficient, $\sigma_t = \sigma_s + \sigma_a$, accounts for out-scattering and absorption events. L_s is the radiance at $\mathbf{o} + P\omega$, and L_i represents the *in-scattering*: all incident irradiance at any point in the medium scattered towards the sensor. Finally, the integral sums the contribution of every point $\mathbf{o} + v\omega$ along the path. L_m corresponds to the integral below, which accounts for all the medium in-scattering along each differential point in the full path.

The fundamental difficulty of this problem is the computational cost of approximating this integral, as the in-scattering itself at a given point \mathbf{p} is yet another nested integral over all directions Ω :

$$L_i(\mathbf{p}, \omega) = \int_{\Omega} \sigma_s f_s(\mathbf{p}, \omega_i, \omega) L(\mathbf{p}, \omega_i) d\omega_i, \quad (2)$$

where f_s is the phase function, which represents the ratio of light being scattered from ω_i towards ω , and $L(\mathbf{p}, \omega_i)$ the incoming radiance from ω_i . We aim to find an approximation to Equation (1) that does not require the numerical evaluation of it and Equation (2) while respecting the physical behaviour of radiance fields entering the water bodies.

3.2. Ocean optics background

Inherent and apparent optical properties. In the ocean optics literature, the optical properties of water, which determine how it interacts with light, are often classified into Inherent (IOPs) and Apparent (AOPs) Optical Properties. IOPs, such as the absorption and scattering coefficients and the phase function, only depend on the water constituents, and fully describe its light transport properties through Equation (1). AOPs, on the other hand, also depend on the radiance field. For example, the irradiance profile with respect to depth is an AOP, as it depends on the incident irradiance at the surface. However, AOPs are studied because they are easier to measure than IOPs, and many works have been devoted to find relationships between the two.

The so called K functions [Mob94] are the key AOPs in ocean optics, modeling the rate of change of irradiance with respect to

depth. Although they do depend on sun angle and other external factors, they are considered quasi-inherent optical properties due to their robustness to external factors [Mob94, LDC*05, Mob01, Gor89].

Moreover, K functions are widely used in oceanography due to practical aspects, since they are easy to measure in different waters by measuring irradiance at different depths. In many cases, they also show a relationship with biological constituents, and are useful for remote sensing [SB78b, SB78a].

In this work we rely on the *diffuse downwelling attenuation* K-function K_d [Mob22], which is defined for the downward irradiance differential as

$$K_d(y) = -\frac{1}{E(y)} \frac{dE(y)}{dy}, \quad (3)$$

where y is the vertical distance from the surface and $E(y)$ the irradiance at that point.

Asymptotic radiance distribution. All homogeneous, optically deep waters show a characteristic depth-dependent asymptotic radiance distribution. This was first conjectured by Whitney [Whi41b, Whi41a], then given empirical support from actual measurements [Tyl58, Tyl60, JF60, SWOO58], and finally given mathematical proof through Preisendorfer's radiative transfer model [Pre59]. This asymptotic radiance distribution implies that, in homogeneous waters, AOPs such as the K functions become IOPs given enough depth, independent from the boundary conditions at the surface and bottom [Mob94].

Real-world K_d datasets. The earliest datasets of K_d originate from the Swedish Deep Sea Expedition during which oceanographer Nils Gunnar Jerlov made measurements of the downwelling irradiance spectrum for different water bodies as a function of depth, which he later used to calculate K_d values for various locations around the globe. Typically referred to as *Jerlov Water Types* [JER51, Jer76, Jer77], these data categorize the optical conditions of the world's oceans into 10 classes. The waters are divided in five open waters (I-III) (Figure 2, top), and five coastal (1-9C) (Figure 2, bottom). The figure shows their respective wavelength dependent coefficients of absorption σ_a , scattering σ_s , extinction σ_t , and diffuse downwelling attenuation K_d . Note that the absorption and scattering coefficients corresponding to the K_d values of Jerlov water types were obtained by reverse engineering using the Hydrolight software by Solonenko & Mobley [SM15].

While the Jerlov water types offer a compact generalization of most commonly encountered optical conditions, they do not capture the large variety seen across the world's natural bodies of water. Thousands of modern-day measurements of K_d are available in the SeaWiFS Bio-optical Archive and Storage System (SeaBASS) [WBF*03]. Here, we demonstrate our results using the Jerlov water types for simplicity and also because the corresponding IOPs for them are known, but our method can be used with any real-world K_d dataset.

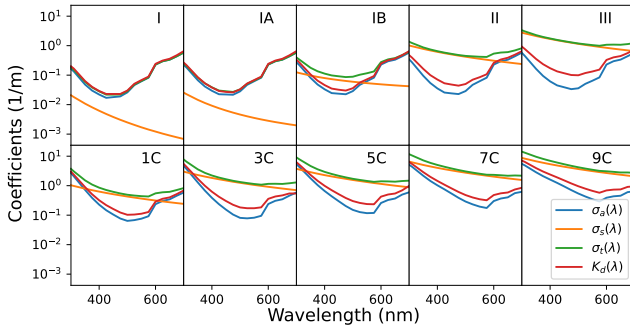


Figure 2: Jerlov water types, five open (top) and five coastal (bottom). They are ordered from lowest (left) to highest (right) scattering to absorption ratio (single-scattering albedo).

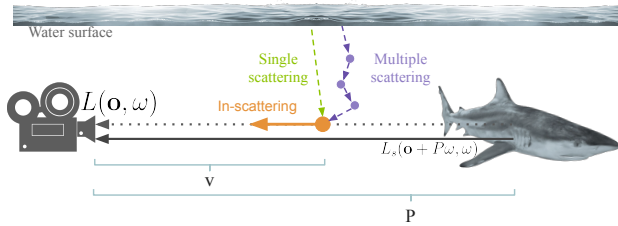


Figure 3: The final radiance L results from the contributions of the surface of the object L_S (black), and path radiance (dotted line). The path radiance is the integral of the medium in-scattering (orange) at each point. Each of the particles in the medium can contribute in-scattering, either directly (single scattering, green), or through several interactions (multiple scattering, purple).

4. Our method

Near the surface of the ocean, single scattering (Figure 3, green) produces sharp, spatially varying caustic patterns, which are important cues for underwater appearance [DHV*23]. On the other hand, lighting due to multiple scattering in the ocean (see Figure 3, purple) is mainly homogeneous horizontally, and slowly varying vertically.

This is related to the asymptotic behaviour of the K functions we mentioned in the previous section [Mob94, Mob22].

From the path radiance $L_m(\mathbf{o}, \omega)$ in Equation (1) we can obtain the following expression, separating the full contributions of single ($L_{SS}(\mathbf{o}, \omega)$) and multiple ($L_{MS}(\mathbf{o}, \omega)$) scattering:

$$\begin{aligned} L_m(\mathbf{o}, \omega) &= \int_0^P T(v) (L_{iSS}(\mathbf{o} + v\omega, \omega) + L_{iMS}(\mathbf{o} + v\omega, \omega)) dv \\ &= \int_0^P T(v) L_{iSS}(\mathbf{o} + v\omega, \omega) dv + \int_0^P T(v) L_{iMS}(\mathbf{o} + v\omega, \omega) dv \\ &= L_{SS}(\mathbf{o}, \omega) + L_{MS}(\mathbf{o}, \omega), \end{aligned} \quad (4)$$

where $L_{iSS}(\mathbf{o} + v\omega, \omega)$ and $L_{iMS}(\mathbf{o} + v\omega, \omega)$ account for the single and multiple scattering at each point, respectively. In the follow-

ing, we detail how we approximate each of these terms. We omit the spectral dependencies of some terms for simplicity, and explain how they are taken into account at the end of the section.

4.1. Single scattering

Single-scattering in-scattering $L_{iSS}(y, \omega)$ results from direct sunlight, refracted by the water surface, and attenuated exponentially by the transmission of the medium $e^{-y \cos(\theta_{sun}) \sigma_t}$, where θ_{sun} is the vertical angle from the refracted sun direction. We emulate this moving light field by the usual approach [DHV*23] of procedurally shadowing a directional light source, due to the great efficiency of this approach.

Then, to approximate the resulting spatially varying illumination, we rely on the *froxel lighting* technique, widely used in game engines [Vos14, Hil15, Kov20], and implemented in Unity's High Definition Rendering Pipeline (HDRP) [Lag18]. Froxel lighting is a two-pass algorithm which first computes a discrete representation of the in-scattering L_{iSS} in the scene, and then traverses this representation to approximate the final single-scattering integral L_{SS} . A summary of this discretization is illustrated in Figure 4, where the view frustum is subdivided in view-space voxels (i.e., froxels).

The second pass approximates the single-scattering integral of each ray by adding the contributions of the corresponding froxels:

$$\begin{aligned} L_{SS}(\mathbf{o}, \omega) &= \int_0^P T(v) L_{iSS}(\mathbf{o} + v\omega, \omega) dv \\ &= d \sum_{\mathbf{o} + v\omega} T(v) L_{iSS}(\mathbf{o} + v\omega, \omega), \end{aligned} \quad (5)$$

where $\mathbf{o} + v\omega \in XV(\omega)$, $XV(\omega)$ is the set of froxels for direction ω , and d the length of each froxel along the camera z axis. The L_{iSS} in-scattering term is able to consider any phase function, as both the viewing and the light direction direction are known when evaluating the expression. We use a forward-scattering Henyey-Greenstein phase function (with anisotropy $g = 0.8$).

To handle the air-water refraction of the sunlight, we use two different lights to represent the sun. The above-surface one has the actual direction and affects the sky rendering, while the second directional light, which affects Equation (5) and the underwater objects, is oriented according to Snell's law.

4.2. Multiple scattering

In general, multiple scattering is prohibitively expensive to simulate. We leverage ocean optics knowledge (and measurements), introduced in Section 3.2, to devise a faster-to-evaluate analytical approximation. We follow a similar approach to Preisendorfer [Pre76, Mob94] and make the approximation that multiple-scattered radiance at any underwater position depends only on vertical depth according to the diffuse downwelling attenuation coefficient K_d as shown in Equation (8). This approximation neglects shadows, that are less prominent in multiple scattering than in single scattering, where we actually account for them (Section 4.1).

We assume that the medium is spatially homogeneous. Therefore, the full downwelling irradiance at any depth y results only

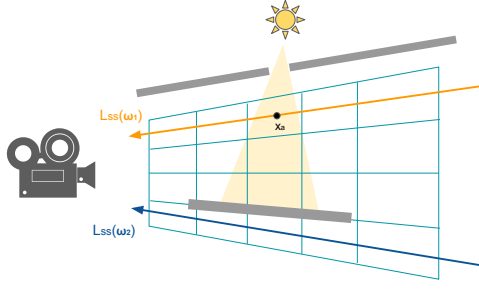


Figure 4: The froxel lighting technique approximates volumetric single scattering in two passes. First: the 3D view frustum is discretized into voxels (blue grid), which store the incident illumination at each point, accounting for occlusions (grey objects in the image). Then the RTE integral of each ray is approximated by adding the in-scattering of each of its voxels. For $L_{ss}(\omega_1)$, in orange, it will thus account for the in-scattering at \mathbf{x}_a , while $L_{ss}(\omega_2)$, in blue, will be zero due to occlusions.

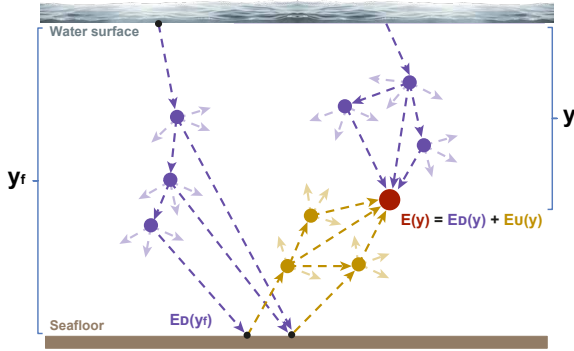


Figure 5: At depth y , the total irradiance ($E(y)$, red) results from adding the descending irradiance ($E_D(y)$, purple) and the upwelling irradiance, previously reflected on the ocean floor ($E_U(y)$, gold).

from the irradiance at the surface $E_D(0)$. We can compute it by integrating the diffuse downwelling attenuation coefficient:

$$E_D(y) = E_D(0)e^{\int_0^y -K_d dv} = E_D(0)e^{-K_d y}. \quad (6)$$

We additionally propose to also account for the radiance being reflected from the seafloor, as illustrated in Figure 5, using a similar approach. We assume a Lambertian horizontal floor at depth y_f , with reflectance coefficient r . Then, the upwelling irradiance $E_U(y)$ being reflected from it can be obtained through Equation (6) as

$$E_U(y_f) = \frac{r}{\pi} E_D(y_f) = \frac{r}{\pi} E_D(0)e^{-K_d y_f}. \quad (7)$$

By adding both, upwelling and downwelling, we get an irradiance profile solely based on the vertical axis:



Figure 6: The water surface radiance L_{WS} accounts for the dielectric nature of the water surface. Snell's window is the circle in the water surface, above the viewer, outside of which there is total internal reflection. Inside it, the whole sky is refracted. Our method naturally simulates this phenomenon, as we can see above the shark.

$$\begin{aligned} E(y) &= E_D(y) + E_U(y_f)e^{-|y-y_f|K_d} \\ &= E_D(0)e^{-K_d y} + \frac{r}{\pi} E_D(0)e^{-K_d y_f} e^{-|y-y_f|K_d} \\ &= E_D(0) \left(e^{-K_d y} + \frac{r}{\pi} e^{-(2y_f-y)K_d} \right). \end{aligned} \quad (8)$$

By assuming an isotropic phase function and substituting the diffuse in-scattering from Equation (4) for this expression we analytically solve the integral over the whole path L_{MS} , as the expression becomes a first-order ODE:

$$\begin{aligned} L_{MS}(\mathbf{o}, \omega) &= \int_0^P T(v) \sigma_s f_s(\mathbf{o} + \omega v, \omega) E(y_v) dv \\ &= \int_0^P e^{-\sigma_t v} \frac{\sigma_s}{4\pi} E_D(0) \left(e^{-K_d y_v} + \frac{r}{\pi} e^{-(2y_f-y_v)K_d} \right) dv \\ &= \frac{\sigma_s E_D(0)}{4\pi(K_d y_{\omega} - \sigma_t)} (e^{(K_d y_{\omega} - \sigma_t)P} - 1) \\ &\quad \cdot \left(e^{-K_d y_o} + \frac{r}{\pi} e^{-K_d(2y_f-y_o)} \right), \end{aligned} \quad (9)$$

where $y_v = y_o + y_{\omega}v$ is the vertical depth v units away of the origin in the direction of the ray. After substituting and integrating, this gets simplified into the final expression.

Note that this approximation has turned the complex computation of multiple scattering into a single $O(1)$ expression (per wavelength, per pixel), accounting for both downwelling irradiance from the sun, as well as the upwelling reflectance from the seafloor.

On top of that expression, we also consider the multiple-scattering radiance coming from the surface $L_s(\mathbf{o}, \omega)$ as:

$$L_s(\mathbf{o}, \omega) = e^{-P\sigma_t} e^{-K_d y_s} \frac{r}{\pi}, \quad (10)$$

which accounts for the diffuse downwelling attenuation until the surface and transmittance attenuation from the surface to the sensor.

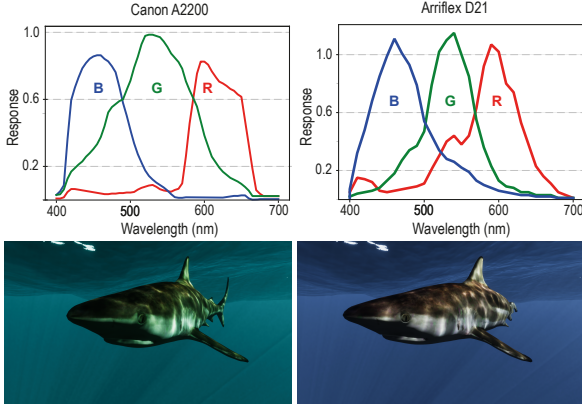


Figure 7: The response functions of two sensors (top), Canon A2200 and Arriflex D21, are used to render the same underwater scene (bottom) of Jerlov 1A, a very clear ocean water, typically of blue appearance to the human eye. The first sensor shows a stronger green channel response than the second one, which results in a clear difference in the RGB values they capture.

Water surface. We also consider the special case of the radiance from the water surface $L_{WS}(\mathbf{p}_{WS}, \omega)$, at a given point p_{WS} , often visible when looking up. We can compute the dielectric color by combining the refracted radiance from the sky L_{sky} and the reflected one (again using L_{MS} Equation (9)), according to the normal \mathbf{n} and the well-known Fresnel equations:

$$L_{WS}(\mathbf{p}_{WS}, \omega) = F(\omega, \mathbf{n}, \eta) L_{sky}(\omega_{fr}) + (1 - F(\omega, \mathbf{n}, \eta)) L_{MS}(\mathbf{p}_{WS}, \omega_{fl}), \quad (11)$$

where $F(\omega, \mathbf{n}, \eta)$ is the ratio of refraction with respect to reflection for the given vectors and the index of refraction of water ($\eta = 1.33$). $L_{sky}(\omega_{fr})$, the radiance from the sky, can be a texture or a physically based approximation. ω_{fr} is the refracted vector, computed according to Snell's law, and ω_{fl} the specular reflection.

This term allows to visualize Snell's window, the circle in the water surface directly above the viewer, approximately 96° in diameter, in which the whole sky is refracted underwater [Mob94], as we can see in Figure 6. Outside this window, delimited by the critical angle, there is total internal reflection.

4.3. Spectral rendering

Until now, we have omitted the wavelength dependency of all coefficients and radiance terms for simplicity. In practice, Equation (9) and Equation (10) are computed per wavelength band, using the corresponding spectral values of the coefficients from the given water type, leading to wavelength-dependent radiance $L(\lambda)$.

In addition, different sensors have different spectral sensitivities, depending on their spectral response function $f(\lambda)$. This function returns the RGB triplet that the sensor outputs when excited by wavelength λ . At the top of Figure 7, we can see two different response functions, for the Canon A2200 and the Arriflex D21, re-

spectively. As we can see in the bottom row, as the green channel of the first camera shows a stronger response, the same scene looks greener. In our work, we simulate a selection of eighty nine response functions from the most widely used sensors, collected from the literature [GN03, JLGS13, SA23].

The final RGB triplet L_{RGB} corresponding to our $L(\lambda)$ captured by response $f(\lambda)$ is obtained as the inner product of both functions:

$$L_{RGB} = \int_{\lambda} L(\lambda) f(\lambda) d\lambda \approx w \sum_{\lambda} L(\lambda) f(\lambda), \quad (12)$$

where w is the width of the bands: $w = (\max(\lambda) - \min(\lambda))/n$, with n the number of wavelength bands evaluated. In practice, we have found $n = 8$ bands to be a good trade-off between accuracy and computational cost, with higher values bringing progressively diminishing returns in color accuracy. In our implementation, Equation (12) and the multiple scattering equations are evaluated per pixel in a single deferred shader.

5. Results

Our method can be integrated into different rendering pipelines. For this work, we have implemented it as a deferred shader in Unity's High Definition Render Pipeline (HDRP) [Lag18], adding support for spectral data (coefficients and sensor response curves), and combining it with the volumetric buffer computed by HDRP's implementation of froxel lighting. In the following results, we used 256-deep slices for the froxels and set the sun spectrum $E_D(0)$ to D65.

Our results show high quality renderings at interactive frame rates even with relatively low-end GPUs: we achieve over 30 frames per second on an NVIDIA GeForce GTX 1050 GPU, and more than 60 frames per second on a more modern NVIDIA GTX 1660 Super. Our physically-based pipeline is able to simulate the appearance of varied Jerlov water types at any depth, captured by a wide variety of sensors. Beyond these results, our method would work with any other water dataset, and any other sensor.

Our video in the supplemental material shows a real-time capture of our interactive application to render different types of ocean waters with different characteristics (see also Figure 9 for a screenshot). First, we start ten meters deep in a Jerlov 1C water (coastal, moderately turbid), and simulate different sensors such as the Nikon D1x or the D220. We then lower depth to 5.6 meters, and change to a Jerlov 3C medium (still coastal, moderately turbid), which presents stronger scattering. We then switch to a Jerlov 1B water, and show how in this case light reaches depths of up to 50 meters. Looking up, Snell's window is visible (as in Figure 6). Last, we modify the water's IOPs, showing how our method can be used to render waters beyond Jerlov types: we first double the scattering coefficient, which increases turbidity, then we increase the absorption coefficient by a factor of three.

Figure 8 shows a reef scene for different open water types, where the 3D model was created from the images in *Sea-thru* dataset D3 [AT19]. Each row shows increasing depth. In the top row ($y_f = 4$ m), we can see how the increasing single-scattering albedo (σ_s/σ_t) from left to right correlates with stronger single-scattering

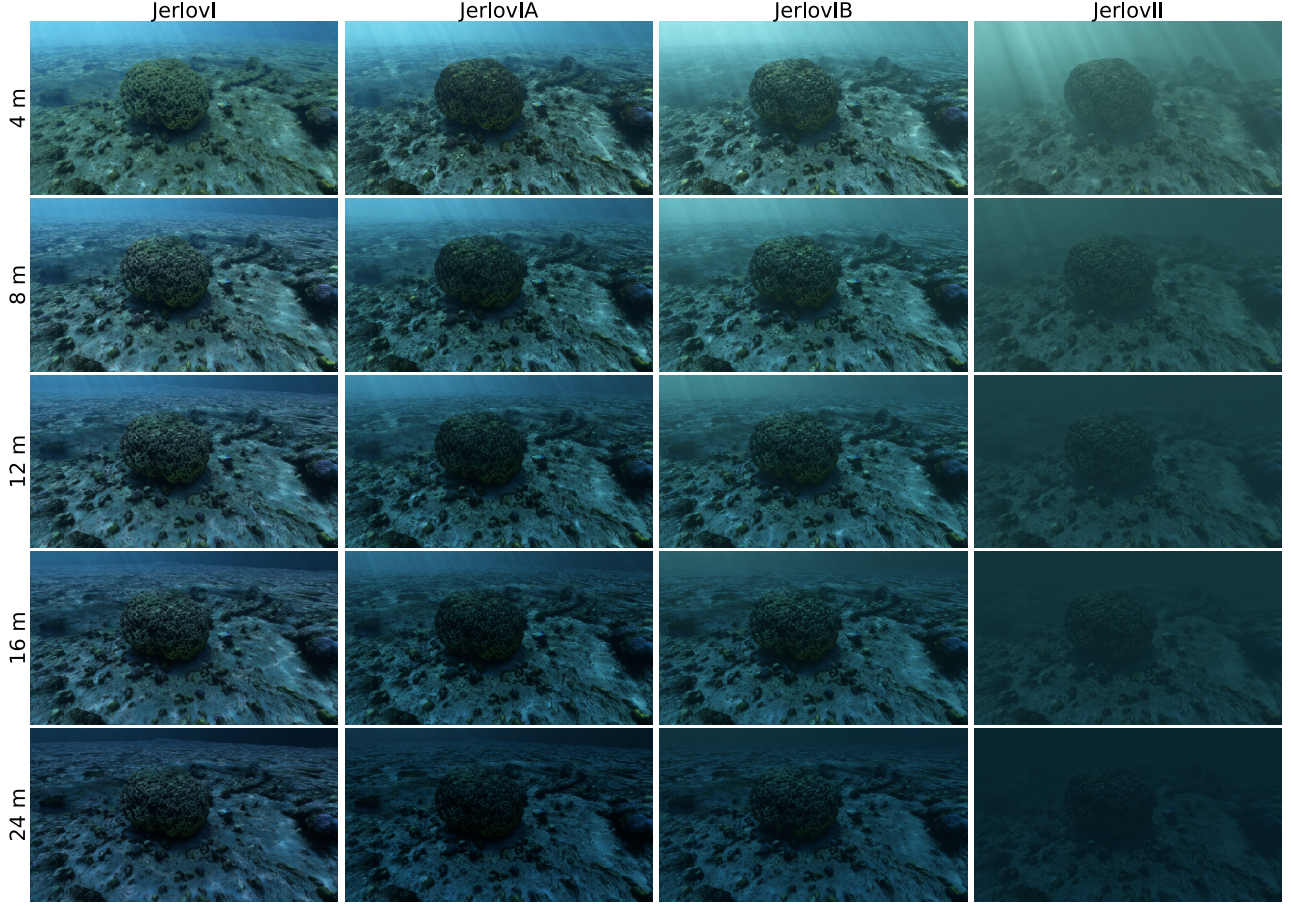


Figure 8: Real-time results of the different Jerlov open waters (I to II) in each column (increasing scattering to the right). Jerlov III is omitted due to low visibility. Each row shows increasing depth. In shallower depths (top), single scattering is much more prevalent.

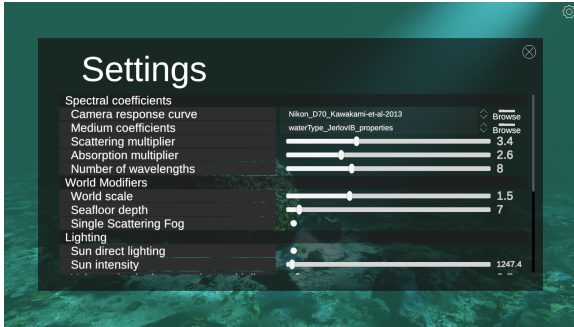


Figure 9: Screenshot of our interactive application, where the user can change multiple parameters defining the scene and its depth, the type of water and its IOPs, the sensor characteristics, and the incoming light. Please refer to the supplemental video.

beams from the water surface. Below, at greater depths, the direct sunlight beams get rapidly attenuated, and the diffuse multiple scattering term becomes dominant.

In Figure 10, we show the equivalent table for different coastal waters, where we can see a similar effect. These coastal waters, with higher scattering due to greater concentrations of suspended particles, rapidly extinguish all light at around ten meters below the surface.

5.1. Comparison with state of the art

The recent USSim method [CZSZ22] is the closest in the literature both in terms of goal and approach: efficient rendering of controllable, physically accurate underwater imagery. The authors also focus on rendering different Jerlov water types based on their IOPs. However, USSim performs its lighting computations in the XYZ space to reduce computation time, and then converts the resulting colors to sRGB. This introduces large errors in the appearance of the final scenes, as shown in Figure 11 for six different Jerlov waters and depths. Moreover, their depth value is constant for each scene. In contrast, our solution is fully spectral and accounts for varying depths within the scene, leading to much more accurate results with respect to the ground-truth, path-traced images in all six cases. For this comparison, in the path tracer and our spectral method we have

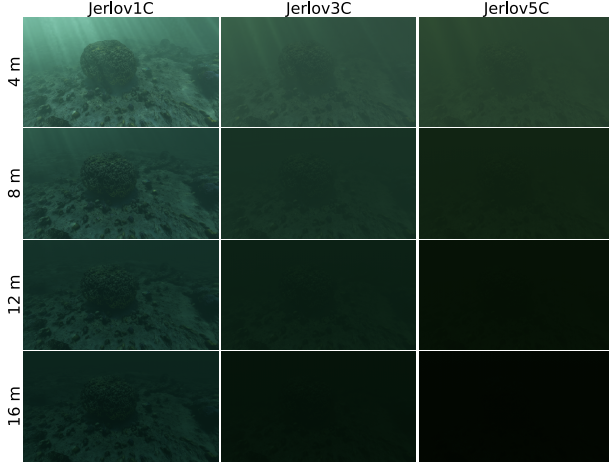


Figure 10: Real-time results of different Jerlov coastal waters (7C and 9C omitted due to low visibility); the rows show increasing depth. In these turbid coastal waters, radiance gets extinguished at around ten meters deep. On the other hand, in shallow depths, single scattering volumetric effects are more clearly visible.

Scene	Chen et al.	Ours
Jerlov 3C d3	0.138	0.059
Jerlov 3C d9	0.057	0.041
Jerlov I d5	0.333	0.117
Jerlov IA d9	0.249	0.107
Jerlov IB d5	0.239	0.107
Jerlov II d3	0.212	0.107
Average	0.205	0.090

Table 2: Average per-pixel RMSE of each method (left: Chen et al. [CZSZ22], right: ours) with respect to the references. These values correspond to the renders shown in Figure 11. Our method consistently produces closer matches to the ground truth, showing on average half the RMSE.

applied the CIE XYZ spectral sensitivity curves, then converted to sRGB to bring all methods to the same color space [RKAJ08].

In Table 2, the average per-pixel Root Mean Square Error (RMSE) of each of the scenes is listed for both methods, compared to the references. On average, we achieve half the RMSE.

5.2. Computational cost

The theoretical per-frame rendering time $t_f(n)$ of our algorithm for n wavelengths is $t_f(n) = t_o + t_{SS} + nt_{MS}(1)$, where t_o is the overhead time mainly due to other pipeline passes (GBuffer, post-processing, antialiasing, etc.), t_{SS} is the cost of single scattering, and t_{MS} the cost to evaluate a single wavelength for multiple scattering. Since only the multiple scattering term depends on the number of wavelengths, the order of the algorithm is $O(n)$ with respect to this parameter.

Figure 12 shows the full frame rendering time (in milliseconds) for the reef scene, varying n between 3 and 250, and rendered

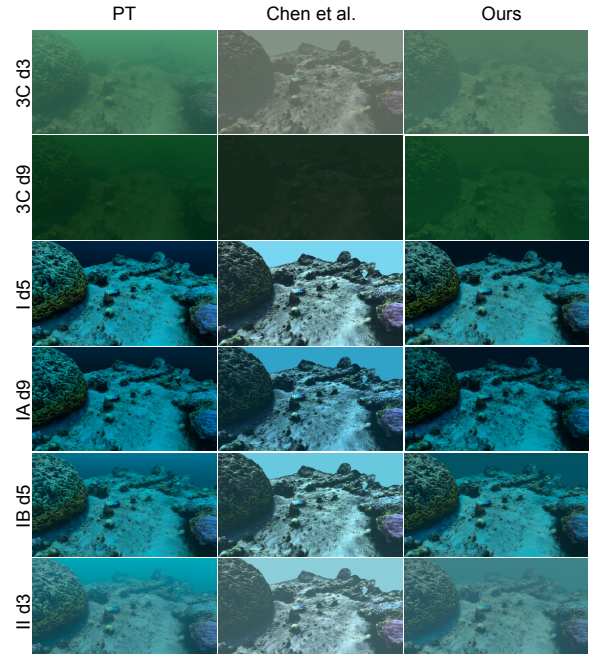


Figure 11: Comparison with the state-of-the-art method of Chen and colleagues (USSim) [CZSZ22], for six different scenes varying water type and depth. PT refers to the path-traced, ground-truth solution. Our method matches the ground-truth references much more closely.

in Full HD with our full single and multiple scattering pipeline, measured in an NVIDIA GeForce GTX 1660 Super GPU. In blue, we show the aggregated cost of the standard rendering passes (t_o), which do not depend on n and add up to 11.06 ms on average. This time is mainly due to the GBuffer pass (50 %), the cascaded shadow maps (10 %), temporal antialiasing and other post-processing effects (such as bloom). In orange, t_{SS} is also independent of n and averages 3.57 ms. In green, we show the cost of the deferred shader evaluating multiple scattering for different wavelength bands. For low values of n , its cost is still lower than single scattering (e.g., at $t_{MS}(32) = 2.5 < 3.57 = t_{SS}$ ms). These low values are the interesting ones in practice, as the spectral resolution of our IOP data is limited to 17 bands. Asymptotically, $t_{MS}(n) \approx 0.056n + 1.17$ ms.

6. Validation

6.1. Qualitative validation

We have compared our multiple scattering approximation using $n = 8$ wavelength bands with a path-traced Monte Carlo reference, implemented in Mitsuba2 [NDVZJ19], which not only does not rely on the downwelling approximation of irradiance, but also considers a forward-scattering Henyey-Greenstein phase function (with mean cosine $g = 0.8$), which is more accurate than our isotropic phase function. We also compare our approximation with an RGB version of our method, evaluating Equations (9) and (10) for the three most representative wavelength bands given a specific sensor response curve.

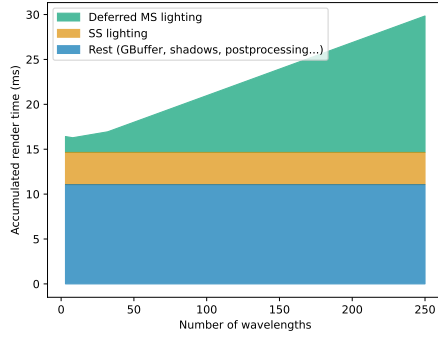


Figure 12: Full rendering time (ms) with respect to the number of wavelengths evaluated, using our single (orange, $t_{SS} \approx 3.57$ ms) and multiple scattering (green, $t_{MS} \approx 0.056n + 1.17$ ms) pipeline. The blue time ($t_o \approx 11.07$ ms) corresponds to standard passes, mainly GBuffer, shadow mapping, antialiasing and postprocessing. Using $n = 8$ as in the rest of the paper, we measured the total time as $t_f(8) = t_{MS}(8) + t_{SS} + t_o = 1.8 + 3.57 + 11.07 = 16.44$ ms.

Figure 13 shows two examples of this experiment for two scenarios: Jerlov IB water captured by the Canon 600D sensor (left), and Jerlov II water captured by the Canon 400D. We show the corresponding renders of these scenes with three methods: a ground-truth volumetric spectral path tracer (top), our real-time spectral method (second row), and a real-time RGB version (bottom). As we can see, the spectral version of our method achieves a very similar appearance to the path-traced references in both cases. Please refer to the supplemental material for more comparisons.

Note that these Monte Carlo spectral renders used as reference take approximately four hours to converge, while our spectral method renders each frame in around 16 ms. Even if the hardware running these methods is not the same (Mitsuba is a multithreaded CPU program, while the real-time solutions run in the NVIDIA GeForce 1660 Super GPU), our method is around 6 orders of magnitude faster. For further analysis of the computational cost, see Section 5.2.

6.2. Quantitative evaluation

We further compare our multiple scattering spectral solution with the RGB implementation using the path tracer simulation as ground truth. We study six combinations of medium coefficients and underwater depths (3, 5 and 9 meters). We include two scattering-dominated waters (Jerlov 3C and Jerlov II), and three absorption-dominated waters (I, IA, IB). Given a scene defined by its medium and depth, we render it using the spectral response curves of 28 different cameras.

We take the pixel-wise difference with respect to the ground truth, using two different metrics: RMSE and the more perceptual Delta E 2000 over CIELAB [Bra03]. The results are summarized in Figure 14 where, for instance, *Jerlov3C_d9* means a Jerlov 3C medium at 9 meters depth. For both metrics, our method achieves a significantly higher accuracy.

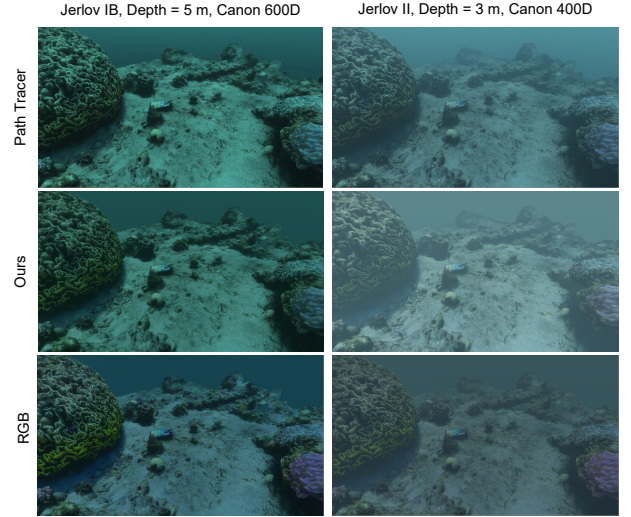


Figure 13: Two underwater scenes (left: Jerlov IB, Canon 600D; right: Jerlov II, Canon 400D) rendered with three methods: a path tracer (4 h), our spectral implementation with 8 wavelengths (16 ms), and an RGB approximation (3 wavelengths) of our method (16 ms).

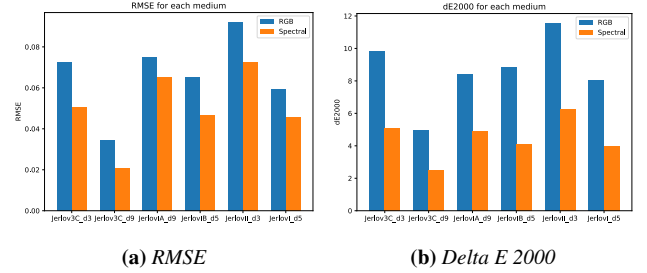


Figure 14: Summary of the average RMSE and Delta E 2000 [Bra03] for RGB (blue) and our spectral rendering method (orange), compared to ground-truth path tracing. For both metrics the error is significantly lower with our method.

6.3. Diffuse downwelling irradiance

The previous results suggest that our choice of an isotropic phase function for the multiple scattering contribution does not introduce as much error as it would in a Monte Carlo simulation, as our references use a more accurate forward-scattering phase function. This can be explained by the fact that the irradiance profile obtained by the downwelling analytical approximation ($e^{-K_d y}$) closely matches the actual irradiance.

In Figure 15, we show a simplified scenario simulating the (monochromatic) irradiance profiles with depth in an infinite, homogeneous water volume. The lines are the Monte Carlo results for different Jerlov water types (each in a color), with their corresponding IOPs. Crosses are the analytical approximation, $E_D(y) = e^{-K_d y}$. Similar experiments can be found in the literature [SM15]. This match also means that K_d can be recovered by a least-squares op-

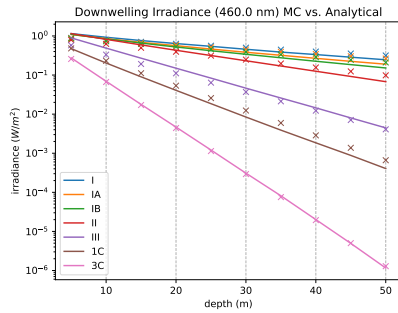


Figure 15: Irradiance profiles with depth of an infinite homogeneous body of various Jerlov water types. Lines: Monte Carlo simulation using the spectral IOPs from the literature [SM15]. Crosses: irradiance $E(y)$ approximation using only K_d , $E(y) = e^{-K_d y}$ (Equation (6)).

timization for other media where it is not routinely measured, as long as their absorption and scattering is known.

7. Conclusions

We have presented a method that achieves real-time performance for rendering a wide variety of spectral underwater scenes, where it naturally reacts to the varying physical coefficients, including the IOPs, the underwater depth, the sun intensity and the response curve of the sensor. We have integrated our method into the Unity engine HDRP to combine it with some of its physically-based rendering features.

Our method accounts for both the downwelling irradiance from the water surface, as well as the diffuse reflectance of the seafloor, improving its accuracy in shallow scenes. Inspired by underwater optics literature, we have proposed an expression for multiple scattering that models downwelling irradiance as a function of the distance to the surface. Our method also models the dielectric nature of the water surface, and reproduces both single scattering volumetric shadows and caustics. By working in the spectral domain, our results consistently achieve a closer match to ground-truth, path-traced references.

Our method assumes homogeneous coefficients and an isotropic phase function. The actual phase function of oceanic water is a combination of its water molecules (isotropic), and suspended particles (strongly forward scattering, with anisotropy $g = 0.924$). In oceanography, it is common to model this with the Fournier-Forand phase function, which for efficiency in Monte Carlo simulations is often approximated with a weighted sum of Henyey-Greenstein phase functions [Mob94, DHV*23]. An interesting avenue for future exploration would be to devise constant-time expressions for both single and multiple scattering considering other analytical phase functions (such as Henyey-Greenstein).

Our method could be improved by tackling its other main limitations: multiple scattering does not consider occlusions or local light sources. Although this limits its applicability, this works well for the usually open geometry of underwater scenes.

Our model has been formally derived from radiative transfer theory and oceanography. Our methodology could be inspiring for other complex media, in which radiative transfer can be combined with heuristic knowledge from a particular field of knowledge (oceanography in our case) to achieve real-time performance. Moreover, another promising research direction would be to apply our model for inverse rendering problems, in order to estimate the physical parameters of water bodies from underwater photographs.

Acknowledgements

Special thanks to the members of the Graphics and Imaging Lab for their help with the manuscript, and to the reviewers for their constructive insight. This work was funded by the Government of Aragon's Departamento de Ciencia, Universidad y Sociedad del Conocimiento through the Reference Research Group "Graphics and Imaging Lab" (ref T34_23R) and by MCIN/AEI/10.13039/501100011033 / FEDER, UE through Project PID2022-141539NB-I00 (IMMERSENSE). DA was funded by the Israel Science Foundation grants #2788/22, #1055/22, Israel Ministry of Science and Technology grant #5052, and the Schmidt Marine Technology Partners. DG was funded by the European Union's European Defense Fund Program through the ENLIGHTEN project under grant agreement No. 101103242. Additionally, NM was supported by a Gobierno de Aragon predoctoral grant. The shark 3D model was created by DigitalLife, licensed under CC BY-NC.

References

- [AT18] AKKAYNAK D., TREIBITZ T.: A revised underwater image formation model. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), pp. 6723–6732. 2
- [AT19] AKKAYNAK D., TREIBITZ T.: Sea-thru: A method for removing water from underwater images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2019), pp. 1682–1691. 6
- [ATS*17] AKKAYNAK D., TREIBITZ T., SHLESINGER T., TAMIR R., LOYA Y., ILUZ D.: What is the space of attenuation coefficients in underwater computer vision? In *CVPR* (2017). 2
- [Bra03] BRAINARD D. H.: 5 - color appearance and color difference specification. In *The Science of Color (Second Edition)*, Shevell S. K., (Ed.), second edition ed. Elsevier Science Ltd, Amsterdam, 2003, pp. 191–216. 9
- [Cas60] CASE K. M.: Elementary solutions of the transport equation and their applications. *Annals of Physics* 9, 1 (1960), 1–23. 2
- [Cha50] CHANDRASEKHAR S.: *Radiative transfer*. 1950. 2
- [CPP*05] CEREZO E., PEREZ F., PUEYO X., SERÓN F., SILLION F.: A survey on participating media rendering techniques. *The Visual Computer* 21 (06 2005). 2
- [CZSZ22] CHEN K., ZHANG L., SHEN Y., ZHOU Y.: Towards controllable and physical interpretable underwater scene simulation. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2022), pp. 2789–2793. 2, 7, 8
- [Dev16] DEVELOPER G.: How subnautica plunges deeper into rendering realistic water, 2 2016. 2
- [DHV*23] DROSKE M., HANIKA J., VORBA J., WEIDLICH A., SABBADIN M.: Path tracing in production: The path of water. In *ACM SIGGRAPH 2023 Courses* (New York, NY, USA, 2023), SIGGRAPH '23, Association for Computing Machinery. 2, 4, 10

- [FCJ07] FRISVAD J. R., CHRISTENSEN N. J., JENSEN H. W.: Computing the scattering properties of participating media using lorenz-mie theory. In *ACM SIGGRAPH 2007 papers*. 2007, pp. 60–es. 2
- [FWKH17] FONG J., WRENNINGE M., KULLA C., HABEL R.: Production volume rendering: Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses* (New York, NY, USA, 2017), SIGGRAPH '17, Association for Computing Machinery. 2
- [GN03] GROSSBERG M., NAYAR S.: What is the space of camera response functions? vol. 2, pp. II – 602. 6
- [Gor89] GORDON H. R.: Can the lambert-beer law be applied to the diffuse attenuation coefficient of ocean water? *Limnology and Oceanography* 34, 8 (1989), 1389–1409. 3
- [GSMA08] GUTIERREZ D., SERON F., MUÑOZ A., ANSON O.: Visualizing underwater ocean optics. *Computer Graphics Forum (Proc. of EUROGRAPHICS)* 27, 2 (2008), 547–556. 2
- [Hil15] HILLAIRE S.: Physically-based & unified volumetric rendering in frostbite, 07 2015. 2, 4
- [Jen96] JENSEN H. W.: Global illumination using photon maps. In *Rendering Techniques '96: Proceedings of the Eurographics Workshop in Porto, Portugal, June 17–19, 1996* 7 (1996), Springer, pp. 21–30. 2
- [JER51] JERLOV N.: Optical studies of ocean waters. reports of the swedish deep-sea expedition, 3. *NO. I* (1951). 3
- [Jer76] JERLOV N. G.: *Marine optics*. Elsevier, 1976. 3
- [Jer77] JERLOV N. G.: Classification of sea water in terms of quanta irradiance. *ICES Journal of Marine Science* 37, 3 (09 1977), 281–287. 3
- [JF60] JERLOV N., FUKUDA M.: Radiance distribution in the upper layers of the sea. *Tellus* 12, 3 (1960), 348–355. 3
- [JLGS13] JIANG J., LIU D., GU J., SÜSTRUNK S.: What is the space of spectral sensitivity functions for digital color cameras? In *2013 IEEE Workshop on Applications of Computer Vision (WACV)* (2013), pp. 168–179. 6
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, Association for Computing Machinery, p. 511–518. 2
- [Kov20] KOVALOV A.: Volumetric effects of the last of us: Part two. In *ACM SIGGRAPH 2020 Talks* (New York, NY, USA, 2020), SIGGRAPH '20, Association for Computing Machinery. 2, 4
- [Lag18] LAGARDE S.: The road toward unified rendering with unity's high definition render pipeline, 8 2018. 4, 6
- [LDC*05] LEE Z.-P., DARECKI M., CARDER K. L., DAVIS C. O., STRAMSKI D., RHEA W. J.: Diffuse attenuation coefficient of downwelling irradiance: An evaluation of remote sensing methods. *Journal of Geophysical Research: Oceans* 110, C2 (2005). 3
- [LKA*17] LIAROKAPIS F., KOUŘIL P., AGRAFIOTIS P., DEMESTICHA S., CHMELÍK J., SKARLATOS D.: 3d modelling and mapping for virtual exploration of underwater archaeology assets. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-2/W3* (02 2017), 425–431. 2
- [Lom87] LOMMEL E.: Die photometrie der diffusen zurückwerfung. *Sitzber. Acad. Wissensch. München* 17 (1887), 95–124. 2
- [Mob94] MOBLEY C.: *Light and Water: Radiative Transfer in Natural Waters*. 01 1994. 3, 4, 6, 10
- [Mob01] MOBLEY C.: Radiative transfer in the ocean*. In *Encyclopedia of Ocean Sciences (Second Edition)*, Steele J. H., (Ed.), second edition ed. Academic Press, Oxford, 2001, pp. 619–628. 3
- [Mob22] MOBLEY C.: The oceanic optics book. 3, 4
- [NDVZJ19] NIMIER-DAVID M., VICINI D., ZELTNER T., JAKOB W.: Mitsuba 2. *ACM Trans. Graph.* 38, 6 (12 2019), 1–17. 8
- [PA01] PREMOŽE S., ASHIKHMIN M.: Rendering natural waters. *Computer Graphics Forum* 20, 4 (2001), 189–200. 2
- [Pre59] PREISENDORFER R. W.: Theoretical proof of the existence of characteristic diffuse light in natural waters. 3
- [Pre76] PREISENDORFER R. W.: *Hydrologic optics*. US Department of Commerce, National Oceanic and Atmospheric Administration ..., 1976. 4
- [RKAJ08] REINHARD E., KHAN E. A., AKYUZ A. O., JOHNSON G.: *Color imaging: fundamentals and applications*. CRC Press, 2008. 8
- [SA23] SOLOMATOV G., AKKAYNAK D.: Spectral sensitivity estimation without a camera. In *IEEE International Conference on Computational Photography (ICCP)* (July 2023). 6
- [SB78a] SMITH R. C., BAKER K. S.: The bio-optical state of ocean waters and remote sensing 1. *Limnology and Oceanography* 23, 2 (1978), 247–259. 3
- [SB78b] SMITH R. C., BAKER K. S.: Optical classification of natural waters 1. *Limnology and Oceanography* 23, 2 (1978), 260–267. 3
- [SM15] SOLOMONENKO M. G., MOBLEY C. D.: Inherent optical properties of jerlov water types. *Applied optics* 54, 17 (2015), 5392–5401. 2, 3, 9, 10
- [SWOO58] SASAKI T., WATANABE S., OSHIBA G., OKAMI N.: Measurements of angular distribution of submarine daylight by means of a new instrument. *Journal of the Oceanographical Society of Japan* 14, 2 (1958), 47–52. 3
- [Tyl58] TYLER J. E.: Radiance distribution as a function of depth in the submarine environment. 3
- [Tyl60] TYLER J. E.: Radiance distribution as a function of depth in the submarine environment. part ii. overcast lighting with tables of radiance-k values for both overcast and sunny lighting conditions. 3
- [Vos14] VOS N.: Volumetric light effects in killzone: Shadow fall. In *GPU Pro 5*, Engel W., (Ed.). CRC Press, 2014, pp. 127–147. 2, 4
- [WBF*03] WERDELL P. J., BAILEY S., FARGION G., PIETRAS C., KNOBELSPIESSE K., FELDMAN G., MCCLAIN C.: Unique data repository facilitates ocean color satellite validation. *Eos, Transactions American Geophysical Union* 84, 38 (2003), 377–387. 2, 3
- [Whi41a] WHITNEY L. V.: The angular distribution of characteristic diffuse light in natural waters. 3
- [Whi41b] WHITNEY L. V.: A general law of diminution of light intensity in natural waters and the percent of diffuse light at different depths*. *J. Opt. Soc. Am.* 31, 12 (Dec 1941), 714–722. 3