

# In-Situ Multi-View Multi-Scattering Stochastic Tomography

Vadim Holodovsky,\* Yoav Y. Schechner,\* Anat Levin,<sup>†</sup> Aviad Levis,\* Amit Aides.\*

\* Viterbi Faculty of Electrical Engineering, Technion - Israel Inst. Technology, Haifa, Israel

<sup>†</sup> Dept. Mathematics and Computer Science, Weizmann Inst. Science, Rehovot, Israel

vhodod@tx.technion.ac.il, yoav@ee.technion.ac.il, anat.levin@weizmann.ac.il,

levisav@tx.technion.ac.il, amitibo@tx.technion.ac.il

## Abstract

*To recover the three dimensional (3D) volumetric matter distribution in an object, the object is imaged from multiple directions and locations. Using these images, tomographic computations seek the distribution. When scattering is significant and under constrained irradiance, tomography must explicitly account for off-axis scattering. Furthermore, tomographic recovery must function when imaging is done in-situ, as occurs in medical imaging and ground-based atmospheric sensing. We formulate tomography that handles arbitrary orders of scattering, using a Monte-Carlo model. The model is highly parallelizable in our formulation. This can enable large scale rendering and recovery of volumetric scenes having a large number of variables. We solve stability and conditioning problems that stem from radiative transfer modeling in-situ.*

## 1. Introduction

Scene rendering and recovery involving participating media [1–7] is gaining interest [8–10], while an effort emerges to recover [11–17] three dimensional (3D) heterogeneous scattering media. 3D volumetric recovery is sought in many domains [12, 16, 18–20], including biomedical imaging, based on multi-view computational tomography (CT). However, most CT models (e.g., X-ray) use direct-transmission [21] as a signal, while scattering is often considered to be a disturbance.

Scattering, including high orders,<sup>1</sup> can be an important signal source for tomography. This is especially relevant in the atmosphere, where the uncontrolled source (sun) and detector (wide angle camera) are generally not aligned. Moreover, typical scenes may not comply with limit models of single-scattering [22] or diffusion [23, 24]. General scattering orders are suggested for tomography by [25]. How-

<sup>1</sup>In a *single scattering* regime, each light ray changes direction at most once due to scattering. In a multiple scattering regime, light can change direction by scattering in multiple events (orders).

ever, the medium in [25] is captured from afar: there are no scattering events near the camera.

Sensors may be embedded in a medium. This is the case in several in-vivo medical imaging modalities, as well as ground-based atmospheric sensing [26]. Consequently, this paper derives multi-scattering 3D tomography, where cameras can be arbitrarily close to the medium, and in fact can be in-situ [22]. Such a setup imposes instabilities [27–30] on the image formation model, which strongly affects both rendering and recovery. Moreover, a forward model involving multiple scattering is computationally complex. Attempting an inverse-problem generally magnifies computational complexity, which inhibits its utility.

This work addresses these issues, generalizing [22] to multiples scattering. First, we propose a way to stabilize the forward model, including in-situ image rendering, while being computationally efficient. Efficiency is achieved using a forward Monte Carlo (MC) principle, which parallelizes rendering across multi-view cameras. This gain is in addition to the inherent parallelizable nature of MC, suitable for GPU acceleration [31–34], where each photon packet traverses the medium independent of other packets. Second, in the inverse problem, optimization is made efficient using surrogate functions [25], while overcoming ill-conditioning arising from an in-situ setup.

## 2. Theoretical background

This section describes building blocks of non-emissive radiative transfer. These blocks yield two common MC methods, each having a specific and mutually complementing advantage and disadvantage, in the context of multiview in-situ imaging. Consequently, Sec. 3 derives stable and efficient MC rendering which better addresses the setup. As a case study, we refer to the atmosphere using some expressions that appear in [22], though the formulations apply to scattering media in general.

**Extinction:** Photons propagating through a scattering medium interact with particles. The medium particles have

an *extinction cross section* for interaction with each individual photon. In the atmosphere, per unit volume, the *extinction coefficient* due to aerosols is  $\beta^{\text{aerosol}} = \sigma^{\text{aerosol}}n$ . Here  $\sigma^{\text{aerosol}}$  is the aerosol extinction cross section and  $n$  denotes particle density. The total atmospheric extinction sums the aerosol and molecular contributions,  $\beta = \beta^{\text{aerosol}} + \beta^{\text{air}}$ , where  $\beta^{\text{air}}$  is modeled as a function of the altitude [22] and wavelength  $\lambda$ . The *optical distance* along a photon path  $S$  is

$$\tau = \int_S d\tau = \int_S \beta dl = \int_S \tau^{\text{air}} + \int_S \sigma^{\text{aerosol}} n dl, \quad (1)$$

where  $\tau^{\text{air}} = \int \beta^{\text{air}} dl$ . The fraction of radiation power that gets transmitted through the atmosphere is the *transmittance*  $t$ , which exponentially decays with the optical distance (Beer-Lambert law):

$$t = \exp(-\tau). \quad (2)$$

**Scattering:** A photon interacts with a particle. The unitless *single scattering albedo*  $\varpi$  of the particle determines a probability for scattering. The aerosol single scattering albedo is  $\varpi^{\text{aerosol}}$ . In the visible spectrum, air single scattering albedo is  $\varpi^{\text{air}} \simeq 1$  and emission is negligible.

Let  $\omega, \psi \in \mathbb{S}^2$  (unit sphere) represent photon propagation direction vectors. The relative fraction of energy scattered from  $\psi$  towards  $\omega$  is set by a normalized *phase function*  $P(\omega \cdot \psi)$ . Air molecules yield *Rayleigh* scattering

$$P_{\text{Ray}}(\omega \cdot \psi) = \frac{3}{16\pi} [1 + (\omega \cdot \psi)^2]. \quad (3)$$

For aerosols,  $P(\omega \cdot \psi)$  is often parameterized. The Henyey-Greenstein approximation [22] for aerosols

$$P_{\text{HG}}(\omega \cdot \psi) = \frac{1}{4\pi} \frac{1 - g^2}{[1 + g^2 - 2g(\omega \cdot \psi)]^{\frac{3}{2}}} \quad (4)$$

is parameterized by an *anisotropy parameter*  $-1 \geq g \geq 1$ .

**Radiative Transfer:** The volumetric extinction field at position  $\mathbf{X}$  is  $\beta(\mathbf{X})$ . The radiance field at  $\mathbf{X}$  in direction  $\omega$  is  $I(\mathbf{X}, \omega)$ . At the domain boundary  $\partial\Omega$ , the incoming radiance [25] is  $I_{\partial\Omega}$ . A ray from  $\mathbf{X}$  in direction  $-\omega$  intersects  $\partial\Omega$  at point  $\mathbf{X}_{\partial\Omega}$ . The radiative transfer equation (RTE) [25, 35] in integral form is

$$I(\mathbf{X}, \omega) = I_{\partial\Omega} \exp\left[-\int_{\mathbf{X}}^{\mathbf{X}_{\partial\Omega}} \beta(r) dr\right] + \int_{\mathbf{X}}^{\mathbf{X}_{\partial\Omega}} J(\mathbf{X}', \omega) \beta(\mathbf{X}') \exp\left[-\int_{\mathbf{X}}^{\mathbf{X}'} \beta(r) dr\right] d\mathbf{X}'. \quad (5)$$

Here  $J(\mathbf{X}, \omega)$  is the *in-scattering* [25, 35] volumetric field

$$J(\mathbf{X}, \omega) = \varpi \int_{4\pi} P(\omega \cdot \psi) I(\mathbf{X}, \psi) d\psi. \quad (6)$$

MC is a popular numerical approach to solve Eqs. (5,6).

## 2.1. Monte Carlo Photon Tracking

Given the source radiance and  $\varpi, P(\omega, \psi), \beta$ , MC propagates photons thus realizing Eq. (5). The result is an estimate of the radiance at  $(\mathbf{X}, \omega)$ , across the domain. In our case study, the light source (sun) is effectively located at infinity. A modeled camera  $c$  has center of projection at location  $\mathbf{X}_c$ . Each pixel  $p$  collects radiation flowing from a narrow cone around direction  $\omega_p$ , yielding a raw image  $i_c(\omega_p)$ . Two existing MC approaches [27] render the images  $i_c(\omega_p)$ :

1. Forward Monte Carlo (FMC): photons propagate from the source (sun) to the detector.
2. Backward Monte Carlo (BMC): photons propagate from the detector to the source.

### 2.1.1 Sampling by Inverse Transform

MC treats scattering and extinction as random phenomena sampled from probability distributions. Random sampling at the heart of MC is realized by an *inverse transform* [36] of a specified probability density function. We quote [22] to briefly describe this mathematical process. A random number  $u$  is drawn from a uniform distribution in the unit interval:  $u \sim \mathcal{U}[0, 1]$ . The number  $u$  is transformed to a random variable  $\chi$ , whose cumulative distribution function (CDF) is  $F(\chi)$ . The transform is  $\chi = F^{-1}(u)$ , where  $F^{-1}$  is the inverse of  $F$ . For example, Eq. (2) is interpreted as a probability density function, whose CDF is

$$F(\tau) = \int_0^\tau \exp(-\tau') d\tau' = 1 - \exp(-\tau). \quad (7)$$

There is high probability for a photon packet to keep propagating in a medium, while  $t$  is high. The probability decays as  $t \rightarrow 0$ . Thus a photon packet propagates to a random optical distance

$$\tau^{\text{random}} = F^{-1}(u) = -\ln(1 - u). \quad (8)$$

### 2.1.2 Forward Monte Carlo Photon Tracking

In FMC, photons are generated at the source, illuminating the top of the atmosphere (TOA) uniformly in direction  $\omega_{\text{sun}}$ . Each photon is traced through the medium. Photons that happen to reach camera  $c$  about direction  $\omega_p$  are counted as a contribution to  $i_c(\omega_p)$ . A photon's life cycle is then defined by the following steps (Fig. 1[left]):

(i) A photon-packet having initial intensity  $I_0$  is launched from the TOA in direction  $\omega_{\text{sun}}$ . This is an initial *ray*  $\mathcal{R}_0$ . Per iteration  $s$ :

(ii) On ray  $\mathcal{R}_s$ , sample the place to which the photon-packet propagates. Eq. (8) yields  $\tau^{\text{random}}$ . Hence, based on Eq. (1), a value  $l^{\text{random}}$  is numerically [22] sought to

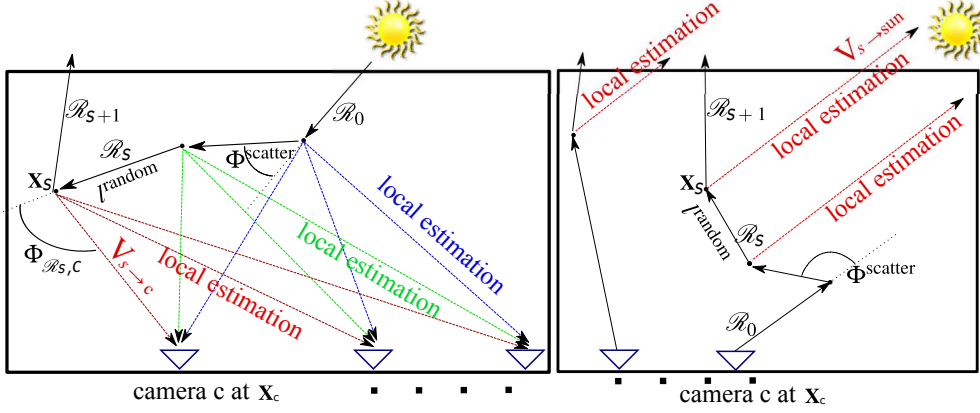


Figure 1. [Left] Multi-view FMC with local estimation. [Right] Multi-view BMC with local estimation.

satisfy

$$\int_0^{l^{\text{random}}} (\sigma^{\text{aerosol}} n + \beta^{\text{air}}) dl = \tau^{\text{random}}. \quad (9)$$

Distance  $l^{\text{random}}$  along  $\mathcal{R}_s$  yields the 3D position  $\mathbf{X}_s$ .

(iii) If  $\mathbf{X}_s$  is outside the domain, the packet is terminated. If  $\mathcal{R}_s$  passes through  $\mathbf{X}_c$ , or a small area around  $\mathbf{X}_c$ , the packet is counted as contributing to the image pixel.

(iv) If  $\mathbf{X}_s$  is inside the medium, then the photon-packet interacts there with a particle. The type of particle (molecule, aerosol) is sampled randomly based on the ratio of extinction coefficients  $\beta^{\text{air}}, \beta^{\text{aerosol}}$  at the voxel containing  $\mathbf{X}_s$ .

(v) If the particle is an aerosol, the photon-packet intensity is attenuated to  $I_{s+1} = \varpi^{\text{aerosol}} I_s$ . For a purely scattering particle, e.g. an air molecule, the photon-packet maintains its intensity. If  $I_{s+1}$  is lower than a threshold, the packet is stochastically terminated, following [28].

(vi) The photon-packet is scattered to a new random direction, determined by inverse transform sampling [37, 38], according to the phase function of the particle (Eqs. 3,4). Let  $\Phi^{\text{scatter}} = \arccos(\boldsymbol{\omega} \cdot \boldsymbol{\psi})$  be the off-axis scattering angle, relative to  $\boldsymbol{\psi}$ . Given a random sample  $u \sim \mathcal{U}[0, 1]$ ,

$$\Phi^{\text{scatter}} = \arccos \left\{ \frac{1}{2g} \left[ 1 + g^2 - \left( \frac{g^2 - 1}{1 + 2gu - g} \right) \right] \right\} \quad (10)$$

for an aerosol. Defining  $\gamma \equiv 4u - 2 + \sqrt{(4u - 2)^2 + 1}$ ,

$$\Phi^{\text{scatter}} = \arccos(\gamma^{\frac{1}{3}} - \gamma^{-\frac{1}{3}}) \quad (11)$$

for a molecule. The scattering azimuth angle around  $\boldsymbol{\psi}$  is sampled from  $\mathcal{U}[0, 2\pi]$ . This scattering event yields a new ray, denoted  $\mathcal{R}_{s+1}$ , thus a new iteration of propagation (ii).

### 2.1.3 Local Estimation In FMC

One way for photons to contribute to a pixel is step (iii) above, which is a rare event. A second way is *local estimation* [27] during step (vi), in every scattering event

(Fig. 1[left]). The local estimation contribution  $W_{\text{le}}$  expresses the probability that a photon scatters towards the camera and reaches the camera without interacting again. Let  $\mathbf{V}_{s \rightarrow c}$  be the vector from the scattering point  $\mathbf{X}_s$  to  $\mathbf{X}_c$ . Let  $t_{s \rightarrow c}$  be the transmittance (2) along  $\mathbf{V}_{s \rightarrow c}$ . Let  $\Phi_{\mathcal{R}_s, c}$  be the angle between  $\mathcal{R}_s$  and  $\mathbf{V}_{s \rightarrow c}$  (Fig. 1[left]). Let  $\varpi$ ,  $P$  be the respective albedo and phase function of the scattering particle: for an aerosol  $\varpi = \varpi^{\text{aerosol}}$ ,  $P = P_{\text{HG}}$ . If the scatterer is a molecule then  $\varpi = \varpi^{\text{air}} \approx 1$ ,  $P = P_{\text{Ray}}$ . Local estimation contributes

$$W_{\text{le}} = \varpi I_s P(\Phi_{\mathcal{R}_s, c}) \frac{t_{s \rightarrow c}}{|\mathbf{V}_{s \rightarrow c}|^2} \quad (12)$$

to pixel  $p$  in camera  $c$ . The factor  $|\mathbf{V}_{s \rightarrow c}|^{-2}$  can be interpreted as if  $\mathbf{X}_s$  is a point radiation source. Due to this factor, FMC is unstable when the camera is *in-situ* i.e. inside the scattering medium [27–30]. Local estimation from scattering points  $\mathbf{X}_s$  close to  $\mathbf{X}_c$  lead to a large increase of image variance. Hence, when  $|\mathbf{V}_{s \rightarrow c}| \rightarrow 0$ ., convergence requires the number of photons to tend to infinity [27].

### 2.1.4 Backward Monte Carlo Photon Tracking

BMC is different from FMC in two major ways. First, from each pixel  $p$  at camera  $c$  a photon packet is separately launched in direction  $-\boldsymbol{\omega}_p$  and then traced back through the medium. If the packet back-traces to the light source (Sun), then the packet counts as contribution to pixel  $p$ . The second difference is how local estimation is done, as detailed below. BMC has the following steps (Fig. 1[right]):

(i) From camera  $c$ , send a photon-packet in direction  $-\boldsymbol{\omega}_p$ . This is the initial ray [22], denoted  $\mathcal{R}_0$ . The packet has an initial intensity  $I_0$ .

Per iteration  $s$ :

(ii) On ray  $\mathcal{R}_s$ , sample the location  $\mathbf{X}_s$  to which the photon-packet propagates, as described in step (ii) of Sec. 2.1.2.




Figure 2. Comparing rendering. In BMC,  $10^7$  initial photons are used, equally divided between all pixels that view the medium. In FMC,  $10^7$  photons uniformly irradiate the domain TOA.

(iii) If  $\mathbf{X}_s$  is outside the domain, the packet is terminated. If  $\mathcal{R}_s \parallel \omega_{\text{sun}}$ , the packet counts as contributing to pixel  $p$ .  
 (iv, v, vi) Sample a photon scattering event: particle type, photon-packet intensity, and scattered direction, as described in steps (iv, v, vi) of Sec. 2.1.2.

Here too, local estimation is performed in conjunction to step (vi), deriving at each scattering event radiance back-traced to the sun. Local estimation expresses the probability that a back propagating photon scatters towards the Sun, then reaches the Sun without interacting again. Let  $\mathbf{V}_{s \rightarrow \text{sun}}$  be the vector from the scattering point  $\mathbf{X}_s$  to the TOA, directed to  $-\omega_{\text{sun}}$  (Fig. 1[right]). Here  $t_{s \rightarrow \text{sun}}$  is the transmittance along  $\mathbf{V}_{s \rightarrow \text{sun}}$ , and  $\Phi_{\mathcal{R}_s, \text{sun}}$  is the angle between  $\mathcal{R}_s$  and  $\mathbf{V}_{s \rightarrow \text{sun}}$ . Local estimation then contributes

$$W_{\text{le}} = \varpi I_s P(\Phi_{\mathcal{R}_s, \text{sun}}) t_{s \rightarrow \text{sun}}. \quad (13)$$

Since the Sun is out of the scattering medium and effectively located at infinity, there is no  $|\mathbf{V}_{s \rightarrow \text{sun}}|^{-2}$  factor at all. Hence sky-images simulated by BMC are stable even in-situ. A comparison is displayed in Fig. 2: FMC rendering is very noisy compared to BMC.

### 3. A Stable and Efficient Forward Model

While having a major stability advantage over FMC, BMC has drawbacks. BMC estimates radiance for one camera and one pixel at a time. In contrast, each single FMC sample trajectory can contribute to multiple viewpoints and pixels simultaneously, using local-estimation (Fig. 1). This is efficient for simulating multiple cameras, which observe the medium from  $N_{\text{views}}$  viewpoints.

We seek to use FMC for several reasons. First, FMC is more efficient for multi-pixel multi-view simulations. Moreover, a gradient-based recovery [25, 39], which we need in Sec. 5 requires the volumetric fields  $I(\mathbf{X}, \omega)$ ,  $J(\mathbf{X}, \omega)$ , not only projected images. Volumetric fields are obtained using FMC without local estimation, hence, are not prone to instabilities. To enable FMC in-situ, however, we need to overcome the  $|\mathbf{V}_{s \rightarrow c}|^{-2}$  instability. In this section, we describe a solution, disposing the problematic  $|\mathbf{V}_{s \rightarrow c}|^{-2}$  factor using voxelization of the field  $J$ . Voxelization of  $J$  biases Eq. (5) to smooth (voxelized) in-scatter

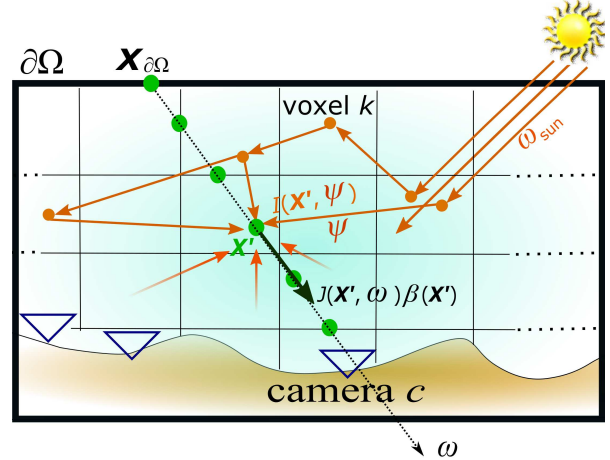


Figure 3. A volumetric scattering medium can be imaged from multiple view, e.g., ground-based cameras observing the atmosphere.

radiance. This bias is generally consistent with smooth 3D scenes, such as the atmosphere. We believe it is a reasonable compromise in such a case, winning both stability and efficiency over simple FMC and BMC, respectively.

As illustrated in Fig. 3, the volumetric domain is discretized into a grid of  $N_{\text{voxels}}$  rectangular cuboid voxels, indexed by  $k$  or  $m$ . As a numerical approximation, assume that within any voxel, the parameters  $\beta(k)$ ,  $\sigma^{\text{aerosol}}$ ,  $\varpi^{\text{aerosol}}$ , and  $g$  are constants, e.g., corresponding to the values at each voxel center. Inspired by volumetric photon mapping [40–42], the solution has three steps: (i) Pre-calculating the geometry of the cameras-grid setup. (ii) FMC calculation of the radiance scattered from voxel  $k$  in the direction of camera  $c$ . (iii) Attenuating radiance along the line of sight (LOS) from voxel  $k$  to camera  $c$ .

#### 3.1. Geometry

At camera  $c$ , the column-stack vector  $i_c$  represents an image having  $N_{\text{pix}}$  pixels. Each pixel collects light from a narrow cone in the domain (Fig. 4). The cone contains or intersects some voxels. The radiant power contributed by voxel  $k$  to camera  $c$  is  $R_c(k)$ , which we define in detail in Section 3.3. Overall, radiance captured at the pixel from all voxels is a weighted sum of  $R_c(k)$  over all voxels  $k$ . This sum is expressed [22] by a sparse  $N_{\text{pix}} \times N_{\text{voxels}}$  matrix operation  $\mathbf{\Pi}_c$ , having reciprocal area units

$$i_c = \mathbf{\Pi}_c \mathbf{R}_c. \quad (14)$$

Here  $\mathbf{R}_c$  is a column-stacked representation of  $R_c(k)$ . The weights of  $\mathbf{\Pi}_c$  represent the relative portion of the radiant power contributing to camera  $c$ , solely due to geometry.

The elements of  $\mathbf{\Pi}_c$  are pre-calculated as follows. Divide pixel  $p$  to  $N_{\text{rays}}$  points, from each of which back-project

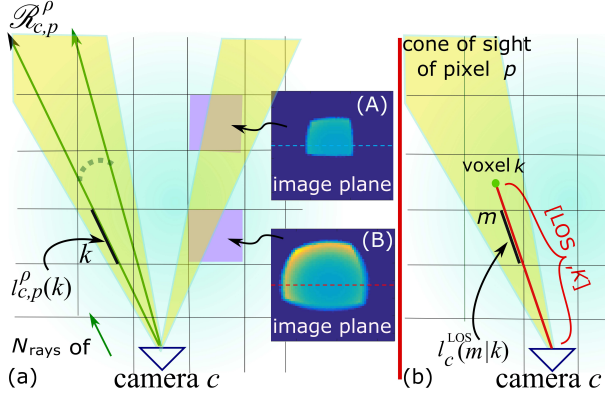


Figure 4. (a) Ray  $\mathcal{R}_{c,p}^\rho$  intersects voxel  $k$ , creating a line-segment  $l_{c,p}^\rho(k)$ . Two isolated voxels (A) and (B) project to the image plane [Insets]. Because of the difference of the voxels' distance from camera  $c$ , their projected support is different. (b) Line-segment  $l_c^{LOS}(m|k)$  is the intersection of voxel  $m$  and line-segment  $[LOS_c, k]$ .

a ray  $\mathcal{R}_{c,p}^\rho$ . The intersection length of  $\mathcal{R}_{c,p}^\rho$  (Fig. 4a) with voxel  $k$  is  $l_{c,p}^\rho(k)$ . Let  $V_{\text{voxel}}$  be a voxel volume. The weight is proportional to a normalized average intersection length,

$$\mathbf{\Pi}_c(p, k) = \bar{l}_{c,p}(k) = \frac{1}{N_{\text{rays}} V_{\text{voxel}}} \sum_{\rho=1}^{N_{\text{rays}}} l_{c,p}^\rho(k) . \quad (15)$$

Eqs. (14,15) express rendering. There is no factor proportional to  $|\mathbf{V}_{s \rightarrow c}|^{-2}$  in Eqs. (14,15). This factor is implicit in the weighted sum matrix  $\mathbf{\Pi}_c$ : each voxel contributes to several pixels, illuminating a spot in the image plane. More rays pass through voxels closer to a camera. Thus, if a scattering event occurs in a voxel for which  $|\mathbf{V}_{s \rightarrow c}|$  is small, the contribution to the image affects more pixels than if the voxel had a large  $|\mathbf{V}_{s \rightarrow c}|$ . This is expressed by a larger spot in the image (Fig. 4a).

### 3.2. Scattered Radiance Calculation with FMC

Define  $L(\mathbf{X}, \boldsymbol{\omega}) = J(\mathbf{X}, \boldsymbol{\omega})\beta(\mathbf{X})$  as *scattered radiance*. Using FMC,  $L(\mathbf{X}, \boldsymbol{\omega})$  can be estimated by caching all the scattering events that occurred at  $\mathbf{X}$  in direction  $\boldsymbol{\omega}$ . Our situation is simpler for two reasons. First, we use a voxelized radiance grid. Hence  $\mathbf{X}$  is discretized to the voxel index  $k$ . Second, as shown below, we only need to store the scattered radiance that contributes to the discrete set of  $N_{\text{views}}$  cameras  $c = 1, \dots, N_{\text{views}}$ . We denote the power scattered from voxel  $k$  in the direction of camera  $c$  by  $L_c(k)$ . For each scattering event [43] in voxel  $k$ , update  $L_c(k)$  by

$$L_c(k) \leftarrow L_c(k) + \varpi I_s P(\Phi_{\mathcal{R}_{s,c}}). \quad (16)$$

Hence  $L_c$  is discretized in space and the relevant directions. Similarly, a discrete version of  $J(\mathbf{X}, \boldsymbol{\omega})$  is

$$j_c(k) \equiv L_c(k)/\beta(k). \quad (17)$$

### 3.3. Optical Transmittance

The transmittance between  $\mathbf{X}$  and  $\mathbf{X}_c$  is

$$t(\mathbf{X}, \mathbf{X}_c) = \exp \left[ - \int_{\mathbf{X}_c}^{\mathbf{X}} \beta(r) dr \right]. \quad (18)$$

Eq. (5) can be re-written as image rendering:

$$I(\mathbf{X}_c, \boldsymbol{\omega}_p) = A + \int_{\mathbf{X}_c}^{\mathbf{X}_{\partial\Omega}} L(\mathbf{X}', \boldsymbol{\omega}_p) t(\mathbf{X}', \mathbf{X}_c) d\mathbf{X}', \quad (19)$$

where  $A$  represents direct solar rays entering the camera. Using notations of [22], denote by  $[LOS_c, k]$  a LOS between camera  $c$  and the center of voxel  $k$  (Fig. 4b). If this LOS intersects voxel  $m$ , then this intersecting line segment has length  $l_c^{LOS}(m|k)$ . From Eq. (1), the optical distance between the center of voxel  $k$  and camera  $c$  is

$$\tau_{LOS_c}(k) = \sum_{m \in [LOS_c, k]} l_c^{LOS}(m|k) \beta(m). \quad (20)$$

Let  $\tau_{LOS_c}$  and  $\beta$  be column-stack vector representations of  $\tau_{LOS_c}(k)$  and  $\beta(k)$ , respectively. Then, Eq. (20) is expressed using matrix notation [22]

$$\tau_{LOS_c} = \mathbf{W}_c \boldsymbol{\beta}, \quad (21)$$

where  $\mathbf{W}_c$  is a  $N_{\text{voxels}} \times N_{\text{voxels}}$  sparse matrix whose element  $(k, m)$  is

$$\mathbf{W}_c(k, m) = \begin{cases} l_c^{LOS}(m|k) & \text{if } m \in [LOS_c, k] \\ 0 & \text{otherwise} \end{cases} . \quad (22)$$

In a voxelized representation, thus, the transmittance between the center of voxel  $k$  to camera  $c$  is

$$T_c(k) = \exp[-\tau_{LOS_c}(k)]. \quad (23)$$

Based on Eqs. (14,16,23),

$$R_c(k) = L_c(k) T_c(k) . \quad (24)$$

Let  $\mathbf{T}_c$ ,  $\mathbf{L}_c$  be the column stack vector representations of  $T_c(k)$  and  $L_c(k)$  respectively. A column-stack vector of all voxel contributions to camera  $c$  is described by

$$\mathbf{R}_c = \mathbf{L}_c \odot \mathbf{T}_c , \quad (25)$$

where  $\odot$  denotes an element-wise (Hadamard) product.

Let  $\mathbf{j}_c$  be a column-stack vector representation of  $j_c(k)$ . From (14,17,25), excluding direct sun light, the image is

$$\mathbf{i}_c(\boldsymbol{\beta}) = \mathbf{\Pi}_c(\mathbf{j}_c \odot \mathbf{Y}_c) = \mathbf{\Pi}_c(\mathbf{j}_c \odot \boldsymbol{\beta} \odot \mathbf{T}_c) , \quad (26)$$

where

$$\mathbf{Y}_c = \boldsymbol{\beta} \odot \mathbf{T}_c. \quad (27)$$

Eqs. (14,26) are a discrete version of Eq. (19), excluding direct solar irradiance of the camera.






Figure 5. The *Haze blobs* aerosol distribution taken from [22]. The density unit is  $10^6$  particles/ $m^3$ .

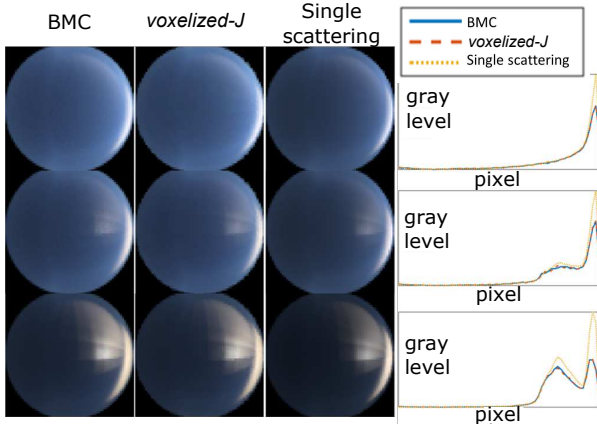


Figure 6. Whole sky (dome) image rendering and a middle horizontal cross-section profile (green channel). [Top] Atm1, [Middle] Atm2, [Bottom] Atm3.

## 4. Rendering Examples

We tested rendering on *Haze blobs* scenes<sup>2</sup> (Fig. 5) and setup parameters described in [22]. The aerosol sea-level density is  $n^{\text{sealevel}}$ . The scenes include an isotropic  $P_{\text{HG}}$  at low density  $n^{\text{sealevel}} \approx 10^6$  (Atm1); anisotropic  $P_{\text{HG}}$  at low density  $n^{\text{sealevel}} \approx 10^6$  (Atm2); and anisotropic  $P_{\text{HG}}$  at high density  $n^{\text{sealevel}} \approx 10^7$  (Atm3).

Fig. 2 demonstrates the improvement achieved by *voxelized-J FMC*, relative to simple FMC (Sec. 2.1.2) on a red channel image of Atm1. Fig. 6 compares whole sky (dome) images rendered using

- BMC, initial  $10^4$  photons per pixel, i.e.,  $\approx 1.5 \cdot 10^9$  photons in total.
- *Voxelized-J FMC* (Sec. 3), using initial  $5 \cdot 10^7$  photons on the TOA and  $N_{\text{rays}} = 10$ . Under these parameters, a pixel readout based on *voxelized-J FMC* has similar variance to rendering based on BMC.

<sup>2</sup>The fields  $\beta^{\text{aerosol}}(\mathbf{X})$  and  $\beta^{\text{air}}(\mathbf{X})$  from [22] are discretized to a  $20 \times 20 \times 40$  voxel grid. For rendering using a *voxelized-J FMC*, the domain is more finely divided into a  $80 \times 80 \times 120$  voxel grid.

- Single scattering approximation [22].

Rendering using *voxelized-J FMC* (Sec. 3) is highly consistent with BMC rendering (Sec. 2.1.4), and similar to single-scattering results [22]. Deviations of MC from single-scattering results are more pronounced where multi-scatter is more significant: near the horizon and generally in Atm3.

Rendering was performed using MATLAB on a 2.50 GHz Intel Xeon CPU, parallelized on 40 cores. Theoretically, for a given pre-calculated geometry (Sec. 3.1), rendering using *voxelized-J FMC* should be  $\approx (1.5 \cdot 10^9)/(5 \cdot 10^7) = 30$  times faster than BMC. However, messaging between cores creates overhead, thus in practice we achieved acceleration of  $\approx 25$ , relative to BMC.

## 5. Inverse Problem

Previous sections rendered images, assuming the field  $\beta(\mathbf{X})$  is known. Now we derive and solve an inverse-problem: given acquired images, what is  $\beta(\mathbf{X})$ ? In addition to air molecules, let there be a single type of aerosol in the domain. Hence, the three-element vector  $[\sigma^{\text{aerosol}}, \varpi^{\text{aerosol}}, g]$  is uniform across the scene and assumed known [44]. The aerosol density is spatially variable and unknown, and so is  $\beta^{\text{aerosol}}(\mathbf{X}) \in \mathcal{C}$ . Here  $\mathcal{C}$  comprises all possible extinction fields that comply with some constraints. Particularly,  $\beta^{\text{aerosol}}$  is non-negative and its spatial support is bounded between the ground and the TOA. A constraint useful for reducing the dimensionality is that  $\beta^{\text{aerosol}}$  is piecewise-constant, following 3D blocks of  $N_x \times N_y \times N_z$  voxels. This constraint is consistent with an assumption that  $\beta(\mathbf{X})$  is generally spatially smooth.

The data are images  $\{i_c^{\text{measured}}\}_{c=1}^{N_{\text{views}}}$ . Recovery is formulated as an optimization of a cost function, to fit the image-formation model to the data [22].

$$\hat{\beta} = \arg \min_{\beta \in \mathcal{C}} \mathcal{E}(\beta) \quad (28)$$

### 5.1. Gradient-based Optimization

Using red-green-blue channels, the RGB extinctions are  $[\beta_{\text{R}}, \beta_{\text{G}}, \beta_{\text{B}}]$ . We solve Eq. (28) using a gradient-based method. Let the variable extinction be  $\beta_{\text{G}}$ . Let  $\mu \in [\text{R}, \text{G}, \text{B}]$  be a channel color. Define  $\tilde{\sigma}_{\mu} = \sigma_{\mu}^{\text{aerosol}}/\sigma_{\text{G}}^{\text{aerosol}}$ . From the known  $\sigma^{\text{aerosol}}$  and  $\beta^{\text{air}}$ , the overall extinction coefficient per channel is determined by  $\beta_{\text{G}}^{\text{aerosol}}$  using

$$\beta_{\mu} = \beta_{\mu}^{\text{air}} + \tilde{\sigma}_{\mu} \beta_{\text{G}}^{\text{aerosol}}. \quad (29)$$

Let  $i_{c,\mu}$  and  $i_{c,\mu}^{\text{measured}}$  denote the modeled and measured image in channel  $\mu$ , respectively. Then, the cost function of Eq. (28) is

$$\mathcal{E}(\beta) = \eta \Psi(\beta_{\text{G}}^{\text{aerosol}}) + \sum_{\mu=\text{R,G,B}} \sum_{c=1}^{N_{\text{views}}} \|\mathcal{M}_c[i_{c,\mu}^{\text{measured}} - i_{c,\mu}(\beta_{\mu})]\|_2^2. \quad (30)$$

Here  $\Psi$  is a regularization term. It expresses the spatial smoothness [22] of  $\beta_G^{\text{aerosol}}$ , while  $\eta$  is a regularization weight. In Eq. (30),  $\mathcal{M}_c$  masks pixels around the Sun. There are two reasons for this masking. Sky-images estimated using MC have high variance in pixels surrounding the Sun due to the stochastic nature of MC. The problem occurs as the phase function is sampled at forward scattering angles.<sup>3</sup> A sun-mask avoids use of noisy modeled pixels. Moreover, a sun occluder is typically applied to real-world cameras, to block lens flare and saturation [26].

Denoting transposition by  $(\cdot)^\top$ , the gradient of (28) is

$$\frac{\partial \mathcal{E}}{\partial \beta} = \eta \frac{\partial}{\partial \beta_G^{\text{aerosol}}} \Psi(\beta_G^{\text{aerosol}}) + 2 \sum_{\mu=R,G,B} \tilde{\sigma}_\mu \sum_{c=1}^{N_{\text{views}}} \mathbf{Q}_c \left\{ \mathcal{M}_c \mathcal{J}_c(\beta_\mu) \right\}^\top \mathcal{M}_c [\mathbf{i}_c(\beta_\mu) - \mathbf{i}_{c,\mu}^{\text{measured}}]. \quad (31)$$

Here the matrix  $\mathcal{J}_c(\beta_\mu)$  is the Jacobian of the vector  $\mathbf{i}_{c,\mu}$  with respect to  $\beta_\mu$ , and  $\mathbf{Q}_c$  is a diagonal weighting matrix which is detailed in Section 6. Element  $(p, k)$  of  $\mathcal{J}_c(\beta_\mu)$  differentiates the intensity of pixel  $p$  in viewpoint  $c$  with respect to the extinction at voxel  $k$ , i.e.,  $\partial i_{c,\mu}(p) / \partial \beta_\mu(k)$ . Estimating  $\mathcal{J}_c(\beta_\mu)$  using Eq. (26) is very complex. The reason is that Eqs. (6,19) express a recursive interconnection of  $I$  and  $J$ , which are functions of  $\beta$  (see 5,19). It is complex to perform recursion per each gradient component.

The Jacobian  $\mathcal{J}_c(\beta_\mu)$  has a closed-form expression under a single scattering approximation [22]. In MC multi-scatter, there is no close-form expression for (31). In [39], FMC fitting of spatially uniform  $\{\beta, \varpi, P(\cdot)\}$  is optimized by a stochastic gradient descent. Gradient computation in [39] is intensive, with a set of three cascade FMC simulations per optimization iteration. Trying a similar formulation in a heterogenous medium would mean  $\mathcal{O}(N_{\text{voxels}})$  FMC renderings per iteration, since there are  $N_{\text{voxels}}$  degrees of freedom.

## 5.2. An Efficient Approach

Instead of a direct estimation of the Jacobian  $\mathcal{J}_c(\beta_\mu)$  and Eq. (31),  $\beta$  can be optimized using a *surrogate function* [25]. Denote by  $\mathbf{j}_{c,\mu}$  a field  $\mathbf{j}_c$  in channel  $\mu$ . Let  $\mathbf{j}_{\text{all}} = \{\{\mathbf{j}_{c,\mu}\}_{c=1}^{N_{\text{views}}}\}_{\mu=R,G,B}$ . For a fixed  $\mathbf{j}_{\text{all}}$ , Eq. (28) is easily minimized as we explain. Gradient-based optimization is iterative. Define  $\beta^{(q)}$  as an estimation of  $\beta$  in the  $q$ 'th iteration. Based on  $\beta^{(q)}$ , the field  $\mathbf{j}_{\text{all}}^{(q)}$  is computed using FMC (Sec. 3.2). This step is *not* an inverse problem but forward-model rendering. Consequently, the computational complexity of this step does not increase with  $N_{\text{voxels}}$ .

<sup>3</sup>To reduce noise in forward scattering, some works [28, 45] approximate the phase function or use *importance sampling*.

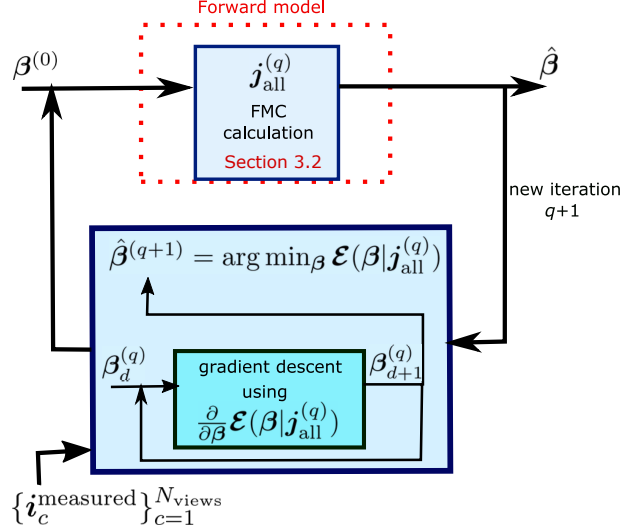


Figure 7. Optimization using a surrogate function  $\mathcal{E}(\beta | \mathbf{j}_{\text{all}}^{(q)})$ .

After  $\mathbf{j}_{\text{all}}^{(q)}$  is derived, it is fixed for a while. Denote  $\mathcal{E}(\beta | \mathbf{j}_{\text{all}}^{(q)})$  as a surrogate function for a fixed  $\mathbf{j}_{\text{all}}^{(q)}$ . Keeping  $\mathbf{j}_{\text{all}}^{(q)}$  fixed,  $\beta$  is evolved. The following iterative optimization process is defined [25]

$$\hat{\beta}^{(q+1)} = \arg \min_{\beta \in \mathcal{E}} \mathcal{E}(\beta | \mathbf{j}_{\text{all}}^{(q)}). \quad (32)$$

Fig. 7 summarizes the iterative optimization process.

When  $\mathbf{j}_{\text{all}}^{(q)}$  is fixed, the Jacobian  $\mathcal{J}_c(\beta_\mu)$  degenerates to a simple calculation. We now derive  $\mathcal{J}_c(\beta | \mathbf{j}_{c,\mu}^{(q)})$ , the Jacobian when  $\mathbf{j}_{\text{all}}^{(q)}$  is a fixed array, i.e.,  $\partial \mathbf{j}_{c,\mu}^{(q)} / \partial \beta_\mu \equiv 0$ . Corresponding to Eqs. (23,27),  $\mathbf{T}_{c,\mu}$  and  $\mathbf{Y}_{c,\mu}$  are the fields  $\mathbf{T}_c$  and  $\mathbf{Y}_c$  in channel  $\mu$ , respectively. Denote by  $\mathbb{D}\{\mathbf{v}\}$  a conversion of vector  $\mathbf{v}$  into a diagonal matrix whose main diagonal elements correspond to the elements of  $\mathbf{v}$ . Using (21,23,26) and expressions from [22] for the gradient of element-wise products, the Jacobian degenerates to

$$\mathcal{J}_c(\beta_\mu | \mathbf{j}_{c,\mu}^{(q)}) = \frac{\partial \Pi_c(\mathbf{j}_{c,\mu}^{(q)} \odot \mathbf{Y}_{c,\mu})}{\partial \beta_\mu} = \Pi_c \mathbb{D}\{\mathbf{j}_{c,\mu}^{(q)}\} \frac{\partial \mathbf{Y}_{c,\mu}}{\partial \beta_\mu}, \quad (33)$$

$$\frac{\partial \mathbf{Y}_{c,\mu}}{\partial \beta_\mu} = \mathbb{D}\{\beta_\mu\} \frac{\partial \mathbf{T}_{c,\mu}}{\partial \beta_\mu} + \mathbb{D}\{\mathbf{T}_{c,\mu}\}, \quad (34)$$

and

$$\frac{\partial \mathbf{T}_{c,\mu}}{\partial \beta_\mu} = -\mathbb{D}\{\exp[-\mathbf{W}_c \beta_\mu]\} \mathbf{W}_c. \quad (35)$$

Given  $\mathcal{J}_c(\beta_\mu | \mathbf{j}_{c,\mu}^{(q)})$ , Eqs. (31,32) can be solved using *gradient descent* as in [25]

$$\beta_{d+1}^{(q)} = \beta_d^{(q)} - \Delta \frac{\partial}{\partial \beta} \mathcal{E}(\beta | \mathbf{j}_{\text{all}}^{(q)}). \quad (36)$$




Figure 8. Example of cost minimization using  $N_{\text{GD}} = 5$ .

Here  $d$  indexes a gradient descent step and  $\Delta$  is the step size. After  $N_{\text{GD}}$  gradient descent steps of  $\beta$ ,  $\mathbf{j}_{\text{all}}$  is updated. Then, gradient-descent of  $\beta$  is resumed, using the updated  $\mathbf{j}_{\text{all}}$  for another set of  $N_{\text{GD}}$  gradient descents [25]. The process reiterates until convergence. Fig. 8 plots an example of cost minimization using  $N_{\text{GD}} = 5$ . A major advantage of the surrogate function is effective gradient calculation, without  $N_{\text{voxels}}$  rendering processes.

## 6. Conditioning the Optimization

Figure 9a illustrates a test atmosphere observed by 25 ground-based cameras. In Fig. 9b, maximum intensity projections (MIP) [46] visualizes  $\beta_1^{(0)}$ . This is the very first step ( $d = 1$ ) in the first iteration ( $q = 0$ ). The field  $\beta_1^{(0)}$  in Fig. 9b stems from  $\mathbf{j}_{\text{all}}^{(0)}$ , which was created by the initialization  $\beta^{\text{aerosol}} = 0$ . The field contains artifacts. The artifacts appear as high values of  $\beta_1^{(0)}$  at voxels near the in-situ cameras. These voxels are unstable. This problem is comparable to ill-conditioned linear optimization.

If all cameras are far from the scattering domain, all voxels are similarly observed [25], and there is no conditioning problem. In-situ, voxels affect unequally the data (Fig. 4a). Due to geometry, if voxel  $k$  projects to more pixels than voxel  $m$  then,  $\partial\mathcal{E}/\partial\beta(k)$  tends to be significantly higher than  $\partial\mathcal{E}/\partial\beta(m)$ . Depending on  $\Delta$ , this imbalance leads to instability in nearby voxels or slow convergence. In the latter, the artifact pattern tends to persist, as seen in Figure 9d.

The remedy is conditioning, achieved using a diagonal weighting matrix  $\mathbf{Q}_c$ . Element  $(k, k)$  of  $\mathbf{Q}_c$  is the number of rays  $\mathcal{R}_{c,p}^p$  that pass through voxel  $k$ . The field  $\beta$  that evolves using this conditioning has weaker artifacts, as visualized in Fig. 9c,e. In addition, Eq. (30) includes a regularization term. Similarly to [22], this paper uses  $\Psi(\beta_{\text{G}}^{\text{aerosol}}) = \|\mathbf{W}\mathcal{L}\beta_{\text{G}}^{\text{aerosol}}\|_2^2$ , where  $\mathcal{L}$  is a matrix representation of the 3D Laplacian operator. The matrix  $\mathbf{W}$  is diagonal: its elements are a function of the altitude of each voxel [22]. The gradient of  $\Psi(\beta_{\text{G}}^{\text{aerosol}})$ , which is used in Eq. (31), is  $2\mathcal{L}^T\mathbf{W}^T\mathbf{W}\mathcal{L}\beta_{\text{G}}^{\text{aerosol}}$ .

## 7. Recovery Simulations

The image formation model has significant nonlinear dependency on the volumetric distribution which we seek to recover. Consequently, analytic derivation of recovery limits and tradeoffs is an open question. We obtain some preliminary insight through numerical simulations.

### 7.1. Images and Noise

We tested recovery on the distributions of Sec. 4. In addition, we ran recovery on the *Haze front* scene of [22], illustrated in Fig. 10. Here  $P_{\text{HG}}$  is anisotropic and  $n^{\text{sealevel}} \approx 10^6$  (Atm4). In all these atmospheres, the aerosol is purely scattering, i.e.,  $\varpi_{\mu}^{\text{aerosol}} \equiv 1 \forall \mu$ . We further tested Atm5, which has the same parameters as Atm2, except that its aerosol is partly absorbing:  $[\varpi_{\text{R}}^{\text{aerosol}}, \varpi_{\text{G}}^{\text{aerosol}}, \varpi_{\text{B}}^{\text{aerosol}}] = [0.69, 0.74, 0.78]$ .

In each scene, a set  $\{i_c^{\text{measured}}\}_{c=1}^{36}$  was rendered using BMC (Sec. 2.1.4) and  $10^4$  photons per pixel. BMC is the most accurate and precise method, despite its slow speed. As in [22], the 36 cameras were placed on a  $6 \times 6$  grid,  $\approx 7\text{km}$  between nearest neighbors. The images are somewhat noisy because MC sampling implicitly induces Poissonian noise. This naturally mimics photon noise in optical imaging. To fully simulate a camera, we incorporate scaling of optical energy to graylevels in a 10-bit camera, white read noise  $o_{\text{noise}}$  and quantization. These operations are expressed by

$$i_c^{\text{measured}} \leftarrow i_c \left( \frac{2^{10}}{\max\{\{\mathcal{M}_c i_c\}_{c=1}^{N_{\text{views}}}\}} \right) + o_{\text{noise}}. \quad (37)$$

The standard deviation of  $o_{\text{noise}}$  is 0.4 graylevels. The values in Eq. (37) are clipped to the range  $[0 \dots 1024]$  and rounded. The resulting images are the input for our reconstruction method and the comparison to [22].

### 7.2. Recovery Results

The analysis used  $\Delta = 10^{-3}$ ,  $N_{\text{GD}} = 5$ ,  $10^7$  photons on the TOA, and  $N_{\text{rays}} = 10$ . Optimization was initialized by  $\beta^{\text{aerosol}} = 0$ , or by a result of a single scattering approximation [22]. The intermediate recovery, which relied on single scattering, was itself initialized by  $\beta^{\text{aerosol}} = 0$ . Convergence was satisfactory after several hundred iterations. The single scattering approximation offers much faster estimation than our stochastic algorithm. The reason is that for single scattering the gradient is expressed in closed form, and the forward model does not require repeated MC sampling. In our implementations, this approximation ran two orders of magnitude faster than MC.

As written in Sec. 2, the aerosol extinction has equivalence to the aerosol density. The density is  $\mathbf{n}^{\text{true}} = \beta_{\text{G,true}}^{\text{aerosol}} / \sigma_{\text{G}}^{\text{aerosol}}$ , where  $\beta_{\text{G,true}}^{\text{aerosol}}$  is the true aerosol extinction in the green channel. Similarly, the estimated corresponding fields satisfy  $\hat{\mathbf{n}} = \hat{\beta}_{\text{G}}^{\text{aerosol}} / \sigma_{\text{G}}^{\text{aerosol}}$ . Recovery



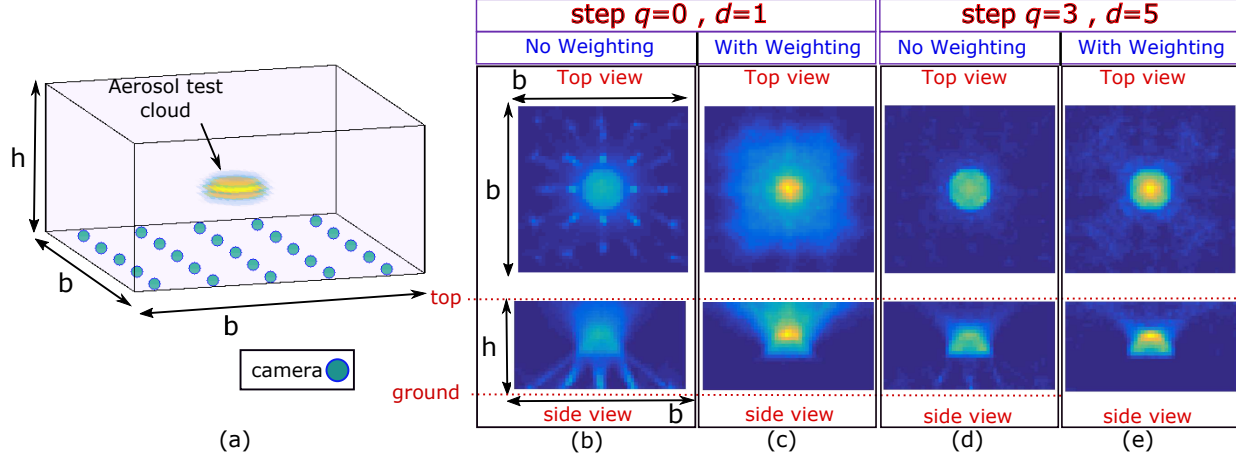


Figure 9. (a) The domain is observed by 25 ground-based cameras. An elliptic aerosol cloud at the center of an air-filled domain. (b,c) Top and side view MIP visualizations of  $\beta_1^{(0)}$ , using a fixed color map scale. (d,e) Top and side view MIP visualizations of  $\beta_5^{(3)}$ ,

Scene	Single scatter; null initialization		Our method; null initialization		Our method; initialized by single scatter	
	$\delta_{\text{mass}}$	$\epsilon$	$\delta_{\text{mass}}$	$\epsilon$	$\delta_{\text{mass}}$	$\epsilon$
Atm1	1.7%	50%	3.4%	26%	3%	23%
Atm2	-6.3%	61%	10%	38%	11%	37%
Atm3	23%	76%	2.4%	71%	5.7%	43%
Atm4	14%	63%	4.1%	27%	7%	28%
Atm5	-5.7%	80%	27%	57%	28%	56%

Table 1. The relative errors resulting from our recovery method, compared to recovery based on a single scattering approximation [22].

quality is quantified as in [22]. Using the  $\ell_1$  norm, the total aerosol relative mass bias is equivalent to

$$\delta_{\text{mass}} = \frac{\|\hat{\mathbf{n}}\|_1 - \|\mathbf{n}^{\text{true}}\|_1}{\|\mathbf{n}^{\text{true}}\|_1} \quad (38)$$

Local errors are quantified by

$$\epsilon = \frac{\|\hat{\mathbf{n}} - \mathbf{n}^{\text{true}}\|_1}{\|\mathbf{n}^{\text{true}}\|_1}. \quad (39)$$

These quantities are summarised in Table 1. Some reconstructions and their errors are illustrated in Figs. 10,11. In our tests, if the standard deviation of  $o_{\text{noise}}$  grows from 0.4 to 5 graylevels, the results are hardly affected. For example, in Atm1,  $\epsilon$  slightly grew from 26% to 28%. We believe the reason is that in daylight, read noise is a minor component relative to fundamental photon noise. Reducing the read noise by using more expensive cameras may thus have only marginal effect on daytime recovery.

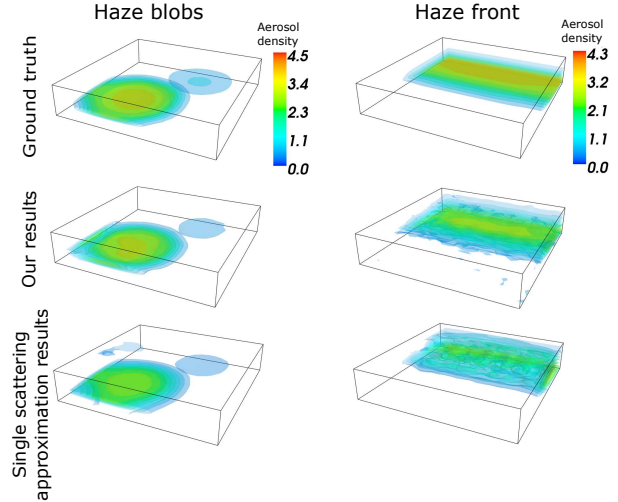


Figure 10. [Top] Ground-truth aerosol distributions: Atm1 [left], and Atm4 [right]. Color encodes aerosol density [ $10^6$  particles/ $\text{m}^3$ ]. [Middle] Our reconstruction. [Bottom] Reconstruction relying on a single scattering approximation [22], ignoring multiple scattering.

Our cameras are spread horizontally, viewing the 3D medium in a range of angles. As the angular range decreases, the recovery suffers from larger vertical smear, particularly away from the cameras. This is analogous to stereo of opaque objects, where a limited baseline increases the axial uncertainty. Fig. 12 shows simulated recoveries of Atm1 when the same 36 cameras are spread over different spans.

Recovery can tolerate use of a smaller  $N_{\text{views}}$ . We recovered Atm1 using different random subsets of the 36 simulated cameras. For each value of  $N_{\text{views}} < 36$ , random photography and recovery was repeated 5 times. Fig. 13




Figure 11. A diagonal slice of the error map,  $|\hat{n} - n^{\text{true}}|$ .

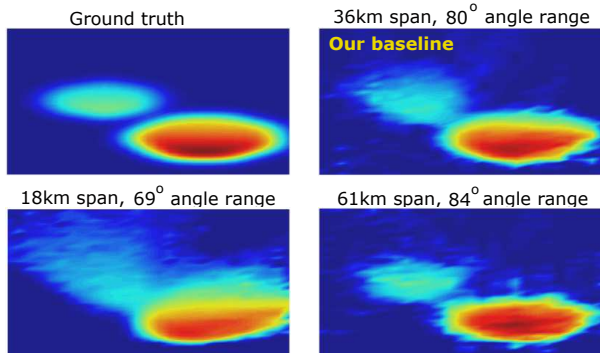


Figure 12. A diagonal slice of the reconstructed domain. In each case, the 36 cameras are spread over a wider or narrower horizontal span, affecting the angular observation range.

plots the errors as a function of  $N_{\text{views}}$ , relying here only the green channel during the optimization.

## 8. Discussion

Our tomographic results have lower error than [22]. This is achieved thanks to explicit use of multiple scattering in the inverse problem. Thus, accounting for multiple scattering is important and significant for recovery, while being computationally feasible. MC can be highly parallelized. Therefore, it will be beneficial to implement the algorithms described here on GPU. MC offers an additional advantage: it can realize scattering events in any location and direction, while discrete ordinate methods [25] are by definition constrained to discretized or band-limited propagation.

What is the recoverable resolution? In estimation problems, avoidance of overfitting of more unknowns requires more data or stronger priors. Our problem is not different. Nevertheless, it may be very difficult to resolve voxels whose length is negligible relative to the mean-free-path, since radiance is hardly affected in such short distances. For example, there may be no benefit using 10cm long voxels in which the mean free path ( $1/\beta$ ) is 10km.

We found that recovery faces difficulties when the opti-




Figure 13. The relative error measures decrease with the number of cameras. Bars represent the standard deviation of errors, based on 5 repeated random tests.

cal distance is very large: then, the signal from deep within the scattering object is low. We also made initial study of other anisotropy values. In addition to the anisotropic aerosol of *Atm2* [22], we tested *particle-type 1* from the aerosol list in [44] (milder anisotropy) and cloud droplets for which  $g \approx 0.85$  (harsher anisotropy). We observed that in these cases, our recovery became worse as anisotropy increases, and this trend appears to exacerbate when  $\beta$  increases. However, the range of parameter variations is too broad to draw definite conclusions from initial tests.

Better priors are a key for advancing multi-scattering tomography. The set of unknowns can be generalized, to include parameters of the phase function and the extinction cross-section. The model itself can be further expanded to include polarized radiative transfer. The model may be adapted to medical imaging. There, the cameras as well as the light sources may be embedded in the scattering domain.

**Acknowledgements:** We are grateful to Anthony Davis, Raanan Fattal and Michael Zibulevsky for fruitful discussions. We thank Mark Sheinin, Johanan Erez, Ina Talmon, Dani Yagodin and Yulia Halupovich for technical support. YYS is a Landau Fellow - supported by the Taub Foundation. His research is supported by the Israeli Ministry of Science, Technology & Space (Grant 51349) and conducted in the Ollendorff Minerva Center. Minerva is funded through the BMBF. AL acknowledges ISF and ERC for financial support.

## References

- [1] R. Fattal. Single image dehazing. In *Proc. ACM TOG*, 27:72, 2008.
- [2] C. Fuchs, M. Heinz, M. Levoy, H.-P. Seidel, and H. Lensch. Combining confocal imaging and descattering. In *Proc. EGSR*, 27:1245–1253, 2008.
- [3] S. G. Narasimhan, S. K. Nayar, B. Sun, and S. J. Koppal. Structured light in scattering media. In *Proc. IEEE ICCV*, 420–427, 2005.
- [4] Y. Y. Schechner, D. J. Diner and J. V. Martonchik, Spaceborne underwater imaging In *Proc. Proc. IEEE ICCP*, 2011.
- [5] R. Kaftory, Y. Y. Schechner, and Y. Y. Zeevi, Variational distance-dependent image restoration In *Proc. IEEE CVPR*, 1–8, 2007.
- [6] T. Treibitz and Y. Y. Schechner, Recovery limits in pointwise degradation In *Proc. IEEE ICCP*, 2009.
- [7] M. Sheinin and Y. Y. Schechner, The next best underwater view In *Proc. IEEE CVPR*, 2016.
- [8] Y. Tian and S. G. Narasimhan. Seeing through water: Image restoration using model-based tracking. In *Proc. IEEE ICCV*, 2303–2310, 2009.
- [9] N. J. Morris and K. N. Kutulakos. Dynamic refraction stereo. In *Proc. IEEE ICCV*, 1573–1580, 2005.
- [10] X. Zhu and P. Milanfar. Stabilizing and deblurring atmospheric turbulence. In *Proc. IEEE ICCP*, 2011.
- [11] J. Kim, D. Lanman, Y. Mukaigawa, and R. Raskar. Descattering transmission via angular filtering. In *Proc. ECCV*, 86–99. Springer, 2010.
- [12] B. Atcheson, I. Ihrke, W. Heidrich, A. Tevs, D. Bradley, M. Magnor, and H.-P. Seidel. Time-resolved 3D capture of non-stationary gas flows. In *Proc. ACM TOG*, 27:132, 2008.
- [13] I. Ihrke, K. N. Kutulakos, H. P. Lensch, M. Magnor, and W. Heidrich. Seidel. State of the art in transparent and specular object reconstruction. In *Eurographics* 2008.
- [14] I. Ihrke, B. Goldluecke, and M. Magnor. Reconstructing the geometry of flowing water. In *Proc. IEEE ICCV*, 1055–1060, 2005.
- [15] Y. Ji, J. Ye, and J. Yu. Reconstructing gas flows using light-path approximation. In *Proc. IEEE CVPR*, 2507–2514, 2013.
- [16] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic reconstruction of transparent objects. *Proc. EGSR*, 2006.
- [17] T. Xue, M. Rubinstein, N. Wadhwa, A. Levin, F. Durand, and W. T. Freeman. Refraction wiggles for measuring fluid depth and velocity from video. In *Proc. ECCV*, 767–782, 2014.
- [18] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz. Light field microscopy. In *Proc. ACM TOG*, 25:924–934, 2006.
- [19] C. Ma, X. Lin, J. Suo, Q. Dai, and G. Wetzstein. Transparent object reconstruction via coded transport of intensity. In *Proc. IEEE CVPR*, 3238–3245, 2014.
- [20] G. Wetzstein, D. Roodnick, W. Heidrich, and R. Raskar. Refractive shape from light field distortion. In *Proc. IEEE ICCV*, 1180–1186, 2011.
- [21] P. Modregger, M. Kagias, S. Peter, M. Abis, V. A. Guzenko, C. David, and M. Stampanoni. Multiple scattering tomography. *Phys. Rev. Lett.*, 113:020801, 2014.
- [22] A. Aides, Y. Y. Schechner, V. Holodovsky, M. J. Garay, and A. B. Davis. Multi sky-view 3D aerosol distribution recovery. *Opt. Express*, 21:25820–25833, 2013.
- [23] D. A. Boas, D. H. Brooks, E. L. Miller, C. A. DiMarzio, M. Kilmer, R. J. Gaudette, and Q. Zg. Imaging the body with diffuse optical tomography. *IEEE Signal Proc. Mag.*, 18(6):57–75, 2001.
- [24] A. Gibson, J. Hebden, and S. R. Arridge. Recent advances in diffuse optical imaging. *Phys. Med. Biol.*, 50(4):R1, 2005.
- [25] A. Levis, Y. Y. Schechner, A. Aides and A. B. Davis. Airborne three-dimensional cloud tomography. In *Proc. IEEE ICCV* 2015.
- [26] D. Veikherman, A. Aides, Y. Y. Schechner, and A. Levis. Clouds in the cloud. In *Proc. ACCV*, 2014.
- [27] A. Marshak and A. B. Davis. *3D Radiative Transfer in Cloudy Atmospheres*. Springer, 2005.
- [28] H. Iwabuchi. Efficient Monte Carlo Methods for Radiative Transfer Modeling *J. Atmos. Sci.*, 63:2324–2339, 2006, and personal communication with H. Iwabuchi.
- [29] T. Kollig, A. Keller. Illumination in the presence of weak singularities *PMonte Carlo and Quasi-Monte Carlo Methods 2004*, Springer Berlin Heidelberg, 2006. 245–257.
- [30] J. Novák, D. Nowrouzezahrai, C. Dachsbacher, W. Jarosz. Virtual Ray Lights for Rendering Scenes with Participating Media In *Proc. ACM TOG*, 31(4), 2012
- [31] J. Křivánek, I. Georgiev, T. Hachisuka, P. Vévoda, M. Šik, D. Nowrouzezahrai, W. Jarosz. Unifying points, beams, and paths in volumetric light transport simulation. In *ACM TOG*, 33(4), 103, 2014.
- [32] Q. Fang, D. A. Boas. Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units. *Opt. Express*, 17(22):20178–20190, 2009.
- [33] M. Ament, C. Bergmann, D. Weiskopf. Ambient Volume Scattering. In *IEEE Trans. Vis. Comput. Graph.*, 19(12):2936–2945, 2013.

- [34] M. Ament, F. Sadlo, D. Weiskopf. Refractive radiative transfer equation. In *ACM Transactions on Graphics*, 33(2):2936–2945, 2014.
- [35] S. Chandrasekhar. *Radiative Transfer*. Dover Pub., 1960.
- [36] L. Devroye. Sample-based non-uniform random variate generation In *Proc. 18th ACM Conf. Winter Simul.*, 260–265, 1986.
- [37] J. R. Frisvad. Importance sampling the Rayleigh phase function *JOSA*, 26(12):2436–2441, 2011.
- [38] T. Binzoni, T. S. Leung, A. H. Gandjbakhche, D. Rufenacht, D. T. Delpy. The use of the Henyey–Greenstein phase function in Monte Carlo simulations in biomedical optics *Phys. Med. Biol.*, 51(17):313–322, 2006.
- [39] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin. Inverse volume rendering with material dictionaries. In *Proc. ACM TOG*, 32:162, 2013.
- [40] H. W. Jensen. Realistic image synthesis using photon mapping. *AK Peters, Ltd*, 2001.
- [41] H. W. Jensen, P. H. Christensen. Efficient simulation of light transport in scenes with participating media using photon maps. In *SIGGRAPH*, 1998.
- [42] W. Jarosz, M. Zwicker, H. W. Jensen. The beam radiance estimate for volumetric photon mapping. In *Computer Graphics Forum*, 27(2), 2008.
- [43] L. S. Kalos, B. Tóth, M. Magdics. Free path sampling in high resolution inhomogeneous participating media. In *Computer Graphics Forum*, 30(1):85–97, 2011.
- [44] J. V. Martonchik, R. A. Kahn, and D. J. Diner, “Retrieval of aerosol properties over land using MISR observations,” in *Satellite Aerosol Remote Sensing over Land*, A. A. Kokhanovsky and G. Leeuw, eds. (Springer Berlin Heidelberg, 2009), pp. 267–293.
- [45] R. Buras, B. Mayer. Efficient unbiased variance reduction techniques for Monte Carlo simulations of radiative transfer in cloudy atmospheres: The solution *J. Quant. Spectrosc. Radiat. Transfer*, 112:434–447, 2011
- [46] J. W. Wallis, T. R. Miller, C. A. Lerner, and E. C. Kleerup. Three-Dimensional Display in Nuclear Medicine In *Proc. IEEE T-MI*, 8(4), 1989

# In-Situ Multi-View Multi-Scattering Stochastic Tomography

## [Supplementary: Notation Table]

Vadim Holodovsky,<sup>\*</sup> Yoav Y. Schechner,<sup>\*</sup> Anat Levin,<sup>†</sup> Aviad Levis,<sup>\*</sup> Amit Aides.<sup>\*</sup>  
<sup>\*</sup> Viterbi Faculty of Electrical Engineering, Technion - Israel Inst. Technology, Haifa, Israel  
<sup>†</sup> Dept. Mathematics and Computer Science, Weizmann Inst. Science, Rehovot, Israel

Acronyms	Description
CDF	Cumulative distribution function
LOS	Line of sight
TOA	Top of the atmosphere

Symbol	Description
$\beta$	Total extinction coefficient
$\beta^{\text{aerosol}}$	Aerosol extinction coefficient
$\beta^{\text{air}}$	Extinction coefficient due to air molecules
$\boldsymbol{\beta}$	Column-stack vector representation of $\beta$
$\beta_\mu$	$\beta$ in color channel $\mu$
$\hat{\beta}$	Estimated $\beta$ (reconstruction)
$\beta_B$	Extinction at the blue channel
$\beta_G^{\text{aerosol}}$	Aerosol extinction at the green channel
$\beta_{G,\text{true}}^{\text{aerosol}}$	Ground truth aerosol extinction at the green channel
$\beta_G$	Extinction at the green channel
$\beta_R$	Extinction at the red channel
$\gamma$	A function of $u$ , used in the Rayleigh inverse transform
$\Delta$	Step size of gradient descent
$\delta_{\text{mass}}$	Aerosol relative mass bias
$\epsilon$	Local aerosol error
$\eta$	Regularization weight
$\lambda$	Wavelength
$\mu$	Color channel index
$\mathbf{\Pi}_c$	Matrix of the portion of the radiant power per voxel, contributing to each pixel in camera $c$
$\rho$	Sub pixel sample
$\sigma^{\text{aerosol}}$	Aerosol extinction cross section
$\sigma_\mu^{\text{aerosol}}$	Aerosol extinction cross section in color channel $\mu$
$\sigma_B^{\text{aerosol}}$	Aerosol extinction cross section in the blue color channel
$\sigma_G^{\text{aerosol}}$	Aerosol extinction cross section in the green color channel
$\sigma_R^{\text{aerosol}}$	Aerosol extinction cross section in the red color channel
$\tilde{\sigma}_\mu$	Aerosol extinction cross section in color channel $\mu$ , normalized by $\sigma_G^{\text{aerosol}}$
$\tau$	Optical distance
$\tau^{\text{air}}$	Optical distance due to air
$\tau^{\text{random}}$	Sampled optical distance
$\tau_{\text{LOS}_c}$	Optical distance between the center of a voxel and camera $c$
$\boldsymbol{\tau}_{\text{LOS}_c}$	Column-stack vector representation of $\tau_{\text{LOS}_c}$
$\Phi_{\mathcal{R}_s,c}$	Angle between $\mathcal{R}_s$ and $\mathbf{V}_{s \rightarrow c}$



$\Phi_{\mathcal{R}_s, \text{sun}}$	Angle between $\mathcal{R}_s$ and $\mathbf{V}_{s \rightarrow \text{sun}}$
$\Phi_{\text{scatter}}$	Scattering angle
$\chi$	Random variable
$\Psi$	Regularization function (i.g., Laplacian)
$\psi$	Direction vector before scattering
$\omega$	Direction vector after scattering
$\omega_p$	Direction vector of pixel $p$
$\omega_{\text{sun}}$	Sun ray direction vector
$\partial\Omega$	Domain boundary
$\varpi$	Single scattering albedo
$\varpi^{\text{aerosol}}$	Aerosol single scattering albedo
$\varpi^{\text{air}}$	Single scattering albedo of air molecules
$A$	Direct solar radiance entering the camera
$B$	Blue
$\mathcal{C}$	Comprises all possible extinction fields that comply with some constraints
$c$	Camera index
$\mathbb{D}\{\mathbf{v}\}$	Conversion of vector $\mathbf{v}$ into a diagonal matrix
$d$	Gradient descent iteration
$\mathcal{E}(\beta)$	Fitting function
$\mathcal{E}(\beta   \mathbf{j}_{\text{all}}^{(q)})$	Surrogate function, for a fixed $\mathbf{j}_{\text{all}}^{(q)}$
$F(\chi)$	Cumulative distribution function (CDF) of random variable $\chi$
$G$	Green
$g$	Anisotropy parameter
$I(\mathbf{X}, \omega)$	Radiance field
$I_0$	Initial photon-packet intensity
$I_s$	Photon-packet intensity at Monte Carlo step $s$
$I_{\partial\Omega}$	Incoming radiance (Sun radiation at the TOA)
$i_c$	Modeled image at camera $c$
$\mathbf{i}_c$	Column-stack vector representation of $i_c$
$\mathbf{i}_c^{\text{measured}}$	Measured image (data) at camera $c$
$\mathbf{i}_{c,\mu}$	$\mathbf{i}_c$ at color channel $\mu$
$\mathbf{i}_{c,\mu}^{\text{measured}}$	$\mathbf{i}_c^{\text{measured}}$ at color channel $\mu$
$J(\mathbf{X}, \omega)$	In-scattering field
$\mathcal{J}_c(\beta_\mu)$	Jacobian of the vector $\mathbf{i}_{c,\mu}$ with respect to $\beta_\mu$
$j_c$	Discretized $J(\mathbf{X}, \omega)$ sampled to the direction of camera $c$
$\mathbf{j}_{\text{all}}$	List of $j_{c,\mu}$ for all $c$ and all $\mu$
$\mathbf{j}_c$	Column-stack representation of $j_c$
$\mathbf{j}_{c,\mu}$	$\mathbf{j}_c$ at color channel $\mu$
$k$	Voxel index
$L(\mathbf{X}, \omega)$	Scattered radiance field
$L_c$	Radiant power scattered by a voxel to camera $c$
$[\text{LOS}_c, k]$	A LOS between camera $c$ and the center of voxel $k$
$\mathbf{L}_c$	Column-stack representation of $L_c$
$\mathcal{L}$	Laplacian operator
$\ell_1$	L1 norm
$l_{c,p}^\rho(k)$	Length of the intersection between Ray $\mathcal{R}_{c,p}^\rho$ and voxel $k$
$\bar{l}_{c,p}(k)$	Weight of voxel $k$ to pixel $p$ in camera $c$
$l_c^{\text{LOS}}(m k)$	Length of the intersection between line $[\text{LOS}_c, k]$ and voxel $m$
$l_{\text{random}}$	Sampled photon distance
$\mathcal{M}_c$	Sun masking operator
$m$	Voxel index
$N_{\text{GD}}$	Number of gradient descent iterations per constant $\mathbf{j}_{\text{all}}^{(q)}$
$N_{\text{pix}}$	Number of pixels in the camera

$N_{\text{rays}}$	Number of ray samples per pixel
$N_{\text{views}}$	Number of viewpoints
$N_{\text{voxels}}$	Number of grid voxels
$N_x, N_y, N_z$	Number of grid voxels per constant $\beta$ block, in the $x$ , $y$ and $z$ axes
$n$	Aerosol particle density
$\mathbf{n}$	Column stack vector representation of $n$
$\mathbf{n}^{\text{true}}$	Column stack vector representation of the ground truth density $n$
$\hat{\mathbf{n}}$	Estimation of $\mathbf{n}$
$n^{\text{sealevel}}$	Aerosol sea-level density
$o_{\text{noise}}$	White read noise
$P$	General phase function
$P_{\text{HG}}$	Henyey-Greenstein phase function
$P_{\text{Ray}}$	Rayleigh phase function
$p$	Pixel index
$\mathbf{Q}_c$	Diagonal weighting matrix
$q$	Optimization iteration index
$R$	Red
$R_c$	Radiant power contributed by a voxel to camera $c$
$\mathbf{R}_c$	Column-stack representation of $R_c$
$\mathcal{R}_0$	Initial ray representation
$\mathcal{R}_s$	Ray representation at Monte Carlo step $s$
$\mathcal{R}_{c,p}^\rho$	Back-projected ray from sub-pixel index $\rho$ in pixel $p$ of camera $c$ ,
$S$	Photon path
$\mathbb{S}^2$	Unit sphere
$s$	Monte Carlo step index
$T_c$	Transmittance between the center of a voxel to camera $c$
$\mathbf{T}_c$	Column-stack representation of $T_c$
$\mathbf{T}_{c,\mu}$	$T_c$ in color channel $\mu$
$t$	Transmittance between two points
$t_{s \rightarrow c}$	Transmittance along $\mathbf{V}_{s \rightarrow c}$
$t_{s \rightarrow \text{sun}}$	Transmittance along $\mathbf{V}_{s \rightarrow \text{sun}}$
$\mathcal{U}$	Unit interval
$u$	Random number drawn from a uniform distribution in the unit interval
$\mathbf{V}_{s \rightarrow c}$	Vector from $\mathbf{X}_s$ to $\mathbf{X}_c$
$\mathbf{V}_{s \rightarrow \text{sun}}$	Vector from $\mathbf{X}_s$ to the TOA, directed to $-\boldsymbol{\omega}_{\text{sun}}$
$V_{\text{voxel}}$	Voxel volume
$\mathbf{W}$	Diagonal weighting matrix of $\mathcal{L}$
$\mathbf{W}_c$	Sparse matrix representing the lengths $l_c^{\text{LOS}}(m k)$
$W_{\text{le}}$	Local estimation contribution
$\mathbf{X}$	Position vector
$\mathbf{X}_{\partial\Omega}$	Ray boundary intersection
$\mathbf{X}_c$	Position vector of camera $c$
$\mathbf{X}_s$	Position vector at Monte Carlo step $s$
$\mathbf{Y}_c$	Column-stack representation of the Hadamard product between $\beta$ and $T_c$
$\mathbf{Y}_{c,\mu}$	$\mathbf{Y}_c$ in color channel $\mu$

---