

Airborne Three-Dimensional Cloud Tomography

Aviad Levis, Yoav Y. Schechner, Amit Aides
 Dept. Electrical Eng.,

Technion - Israel Inst. Technology, Haifa, Israel

{levisav@tx,yoav@ee,amitibo@campus}.technion.ac.il

Anthony B. Davis

Jet Propulsion Laboratory
 California Inst. of Technology
 Pasadena, California

Anthony.B.Davis@jpl.nasa.gov

Abstract

We seek to sense the three dimensional (3D) volumetric distribution of scatterers in a heterogenous medium. An important case study for such a medium is the atmosphere. Atmospheric contents and their role in Earth's radiation balance have significant uncertainties with regards to scattering components: aerosols and clouds. Clouds, made of water droplets, also lead to local effects as precipitation and shadows. Our sensing approach is computational tomography using passive multi-angular imagery. For light-matter interaction that accounts for multiple-scattering, we use the 3D radiative transfer equation as a forward model. Volumetric recovery by inverting this model suffers from a computational bottleneck on large scales, which include many unknowns. Steps taken make this tomography tractable, without approximating the scattering order or angle range.

1. Introduction

Scattering and refractive media are increasingly considered in computer vision [19, 20, 38, 39, 46, 48], typically for observing background objects [47, 35, 56]. However, in important cases, the medium itself is of interest. For example, remote sensing of the atmosphere seeks to assess the distribution of various airborne scatterers. Image data is used to fit a physical model of light propagation through the medium and recover scatterer properties. Similar efforts use refractive propagation models to recover properties of refractive media [6, 24, 25, 49, 55].

This paper seeks volumetric recovery of a three dimensional (3D) heterogeneous highly scattering medium. Furthermore, this work performs recovery in a very large scale: the atmosphere embedded with clouds. The data comprises images acquired from multiples directions [50], as illustrated in Fig. 1. Such a setup samples the scene's light-field [1, 3, 23, 26, 28]. 3D volumetric recovery is achieved in various domains using tomography, finding wide use in




Figure 1. A multi-angular imager passing over an atmospheric domain enables multi-pixel multi-view acquisition [16, 46]. Incoming solar radiation is a boundary condition. A bidirectional reflectance distribution function characterizes the bottom surface. Light scatters in the medium, generally multiple times, creating a scatter field J . Integrating J and the boundary radiation using corresponding attenuation, yields the radiance $I(x, \omega)$.

biomedical imaging [15, 52] and computational photography [6, 29, 49, 53]. However, this work faces several important challenges. First, due to the large volumes involved, our setup is passive, using the steady, uniform and collimated Sun as the radiation source. This is contrary to most tomography setups, in which the source is controllable. Second, in most tomographic models, as in X-ray, direct-transmission [34] forms the signal, while small-angle scattering has been treated as a perturbation. In contrast, in a medium as a cloud, the source and detector are generally not aligned: *scattering including high orders is the dominant signal component.*

Previous works on scattering tomography focus on limits of the scattering order: either diffusion [9, 21] or single-scattering [2]. However, scattering objects, as clouds, are complex: they can exhibit diffusion in their core, low-order scattering in their boundaries and interaction with their surrounding. To avoid reliance on such scattering order approximations, we use the 3D radiative transfer equation (RTE) as the image formation model. However, avoiding such scattering approximations significantly complicates the model, jeopardizing the prospects of its inversion. To make any type of tomography a practical tool, the inverse problem must be tractable. It is thus necessary to find means to efficiently invert the model, as this paper derives.

The derived mathematical multi-scatter heterogeneous model and algorithm principles can be applied to various media. This paper makes further focus on clouds, which are scattering media made of suspended water droplets. Clouds reflect much of the Sun's radiation out to space, while trapping emitted terrestrial radiation. Clouds thus play a major role in the Earth's radiation budget and the understanding of climate evolution [44], yet their influence over climate change still has a large uncertainty [11]. There is thus motivation to sense clouds in detail.

Lidars and in-situ sampling provide local small-scale data. Current optical remote sensing of Earth's atmosphere generally does not yield 3D volumetric mapping. Rather, remote sensing methodology has largely assumed a plane-parallel atmospheric geometry, with horizontally uniform properties. This crude geometry limits recovery to simple parameters [36, 37, 42]. Moreover, this approximation of atmospheric structure causes biases of the retrieved parameters in most cases [14]. Our 3D volumetric recovery avoids the plane-parallel assumption. The method is demonstrated on images captured from a high altitude aircraft. It is also validated using established atmospheric models coupling fluid dynamics and cloud micro-physics.

2. Theoretical background

2.1. Image Formation Model: Radiative Transfer

Our image formation (forward) model is steady-state 3D radiative transfer. A domain $\Omega \subset \mathbb{R}^3$ has boundary $\partial\Omega$, whose outward facing normal is ϑ (Fig. 1). At position $\mathbf{x} \in \Omega$ and direction of propagation $\boldsymbol{\omega} \in \mathbb{S}^2$ (unit sphere), the radiation field is $I(\mathbf{x}, \boldsymbol{\omega})$. Dependency on wavelength λ is omitted, to simplify the explanation. Let $\boldsymbol{\omega} \cdot \vartheta < 0$ define incoming radiation. The boundary condition is

$$I(\mathbf{x}, \boldsymbol{\omega}) \triangleq I_{BC}(\mathbf{x}, \boldsymbol{\omega}) \quad \text{when } \boldsymbol{\omega} \cdot \vartheta < 0, \quad \mathbf{x} \in \partial\Omega. \quad (1)$$

Radiative transfer satisfies [10]:

$$\boldsymbol{\omega} \cdot \nabla I(\mathbf{x}, \boldsymbol{\omega}) = \beta(\mathbf{x}) [J(\mathbf{x}, \boldsymbol{\omega}) - I(\mathbf{x}, \boldsymbol{\omega})] \quad \mathbf{x} \in \Omega, \quad (2)$$

$$I(\mathbf{x}, \boldsymbol{\omega}) = \int_{\mathbb{S}^2} J(\mathbf{x}', \boldsymbol{\omega}) \beta(\mathbf{x}') e^{-\int \beta(\mathbf{r}) d\mathbf{r}} d\mathbf{x}' + I_{BC} e^{-\int \beta(\mathbf{x}) d\mathbf{x}}$$

$$J(\mathbf{x}, \boldsymbol{\omega}) = \frac{\varpi}{4\pi} \int_{\mathbb{S}^2} p(\mathbf{x}, \boldsymbol{\omega} \cdot \boldsymbol{\omega}') I(\mathbf{x}, \boldsymbol{\omega}') d\boldsymbol{\omega}'$$

Figure 2. The RTE is a recursive interplay between $J(\mathbf{x}, \boldsymbol{\omega})$ and $I(\mathbf{x}, \boldsymbol{\omega})$, given by Eqs. (3,4), thus spanning multiple scattering. Numerically, this forward-model is iterated to convergence.

where β is the \mathbf{x} -dependent extinction coefficient, while

$$J(\mathbf{x}, \boldsymbol{\omega}) = \frac{\varpi}{4\pi} \int_{\mathbb{S}^2} p(\mathbf{x}, \boldsymbol{\omega} \cdot \boldsymbol{\omega}') I(\mathbf{x}, \boldsymbol{\omega}') d\boldsymbol{\omega}', \quad (3)$$

is the *source function* (*in-scattering* term) [10], neglecting visible light emission by the medium. Here ϖ is the single scattering albedo and $p(\mathbf{x}, \boldsymbol{\omega} \cdot \boldsymbol{\omega}')$ is the phase function at \mathbf{x} . The phase function describes the fraction of energy scattered from $\boldsymbol{\omega}'$ to $\boldsymbol{\omega}$ by an infinitesimal volume. Equations (1–3) define a complete radiative transfer forward model for an externally illuminated, non-emitting medium.

Integrating Eq. (2) along a specific direction $\boldsymbol{\omega}$ results in an integral form of the 3D RTE [10, 31]

$$I(\mathbf{x}, \boldsymbol{\omega}) = I_{BC}(\mathbf{x}_0, \boldsymbol{\omega}) \exp \left[-\int_{\mathbf{x}}^{\mathbf{x}_0} \beta(\mathbf{r}) d\mathbf{r} \right] + \int_{\mathbf{x}}^{\mathbf{x}_0} J(\mathbf{x}', \boldsymbol{\omega}) \beta(\mathbf{x}') \exp \left[-\int_{\mathbf{x}}^{\mathbf{x}'} \beta(\mathbf{r}) d\mathbf{r} \right] d\mathbf{x}'. \quad (4)$$

Here \mathbf{x}_0 is a point on the boundary (see Fig. 1). Equation (4) accumulates scattered radiance along a line of sight, weighted by the corresponding extinction. By $\int_{\mathbf{x}}^{\mathbf{x}'} f(\mathbf{r}) d\mathbf{r}$, we mean a line integral over a field $f(\mathbf{x})$ along the segment extending from \mathbf{x} to \mathbf{x}' . Numerically, this is preformed by back-projecting a ray through the medium.

Equations (3,4) express a recursive interplay of the fields J and I , as illustrated in Fig. 2. A recursion effectively amounts to a successive order of scattering, or a Picard iteration [18]. Had $J(\mathbf{x}, \boldsymbol{\omega})$ been computed only through incoming solar radiation, without recursion, the result would have been a single scattering approximation.

2.2. Spherical Harmonics Discrete Ordinates

Numerically solving the radiative transfer is a balance between speed and accuracy. Monte Carlo (MC) methods can handle very complex media, including sharp changes in optical parameters. This makes MC particularly accurate in rendering of surfaces. In MC, radiometric quantities are attained by random sampling the infinite domain of possible light paths. This computational process is very slow, particularly for multi-view images. Faster rendering using less random samples increases noise.

Deterministic RTE solvers discretize the spatial and angular domain. Discretization biases $I(\mathbf{x}, \boldsymbol{\omega})$ towards smooth solutions. Nevertheless, volumetric clouds are smoother than surfaces, thus a computed smooth $I(\mathbf{x}, \boldsymbol{\omega})$ is consistent with the nature of cloud fields. Deterministic solvers are thus prime tools in atmospheric rendering. Moreover, when seeking many radiometric outputs of the same scene, e.g. in multiple viewpoints, a model that solves the RTE directly has favorable computational cost.

A very efficient, deterministic solver is the spherical harmonic discrete ordinates method (SHDOM) [18, 31], which relies on two principles. (i) A spherical harmonics representation for the scatter field allows for efficient computation of angular integrals. (ii) Discrete ordinates models radiation flow along specific directions within the domain. Ref. [41] compares atmospheric MC and SHDOM.

2.3. Air and Cloud Water Droplets

For air molecules and cloud water droplets, $\varpi \simeq 1$ in visible light. Molecules follow the Rayleigh scattering law, which uniquely determines the spatially invariant molecular phase function [2] $p_R(\boldsymbol{\omega} \cdot \boldsymbol{\omega}')$. The Rayleigh scattering coefficient β_R is analytically known per air density. Air density varies mainly with altitude z . With localized tornadoes as exception, air density varies slowly in space and time, and mapped over Earth using long established systems.

The change of refractive index between air and cloud water droplets creates scattering, represented by a cloud scattering coefficient $\beta_c(\mathbf{x})$ and phase function $p_c(\boldsymbol{\omega} \cdot \boldsymbol{\omega}')$. The total extinction and phase function are respectively

$$\beta(\mathbf{x}) = \beta_c(\mathbf{x}) + \beta_R(z) \quad (5)$$

$$p(\mathbf{x}, \boldsymbol{\omega} \cdot \boldsymbol{\omega}') = \frac{p_c(\boldsymbol{\omega} \cdot \boldsymbol{\omega}') \beta_c(\mathbf{x}) + p_R(\boldsymbol{\omega} \cdot \boldsymbol{\omega}') \beta_R(z)}{\beta(\mathbf{x})}. \quad (6)$$

The function p_c is determined through Mie theory by the droplet size distribution. The size varies mainly vertically, within typical air masses. Vertical variations follow curves, whose parameters are measured from satellites [45]. 3D variations are thus mainly attributed here to $\beta_c(\mathbf{x})$.

3. Tomographic Recovery - Inverse Model

We seek to recover $\beta(\mathbf{x})$ within the volume of a cloud. This is equivalent to seeking $\beta_c(\mathbf{x})$, since $\beta_R(z)$ is known. The recovery is based on images, with a complex image-formation model. An efficient approach must be derived for such a complex recovery to be contemplated. Note that Eq. (4) depends on $\beta(\mathbf{x})$ in two ways. One dependency on β is explicit, through line integrals over straight back-projected rays, that are easy to compute. The other dependency is implicit, through J . The implicit dependency is complicated and non-local: a change in $\beta(\mathbf{x}_1)$ can cause a

change in $J(\mathbf{x}_2)$. However, for a *given*, fixed J , it is easy to compute, through (4), the radiance I and $\frac{\partial}{\partial \beta} I$. This observation provides a key for computationally realistic tomographic recovery, which is now explained.

3.1. Operator Notation

Following [5, 7], we use operator notations. For a given boundary condition, the *radiance forward mapping* $I(\mathbf{x}, \boldsymbol{\omega}) = \mathcal{I}(\beta)$ is an operator that transforms an extinction field β into a radiation field I . Let us decompose this mapping into two operators, $\mathcal{I}(\beta) = \mathcal{T}(\beta) \mathcal{J}(\beta)$. Here $J(\mathbf{x}, \boldsymbol{\omega}) = \mathcal{J}(\beta)$ is the *in-scatter forward mapping* from a field β to a field J . The operation $\mathcal{T}(\beta)$ transforms an in-scatter field J to a measurable radiance field by the simple line integrals of Eq. (4).

An aperture function $w \in \Omega \times \mathbb{S}^2$, defines collection of radiance by a detector, over a spatial and angular support. Measurements are thus an operator

$$\mathcal{M}_w \mathcal{I}(\beta) = \langle w, I \rangle_\Omega, \quad \text{where } \langle \cdot, \cdot \rangle_\Omega \equiv \int_\Omega \int_{\mathbb{S}^2} \cdot \cdot \, d\boldsymbol{\omega} d\mathbf{x}. \quad (7)$$

For an idealized single-pixel detector positioned at \mathbf{x}^* , collecting radiation flowing in direction $\boldsymbol{\omega}^*$,

$$\mathcal{M}_w \mathcal{I} = \langle \delta(\mathbf{x} - \mathbf{x}^*) \delta(\boldsymbol{\omega} - \boldsymbol{\omega}^*), I \rangle_\Omega = I(\mathbf{x}^*, \boldsymbol{\omega}^*). \quad (8)$$

Consequently, the forward model is

$$\mathcal{F}_w(\beta) = \mathcal{M}_w \mathcal{I}(\beta) = \mathcal{M}_w \mathcal{T}(\beta) \mathcal{J}(\beta). \quad (9)$$

For numerical recovery, the sought field is discretized

$$\beta(\mathbf{x}) = \sum_{k=1}^{N_{\text{grid}}} \beta_k b_k(\mathbf{x}), \quad (10)$$

where $\{\beta_k\}_{k=1}^{N_{\text{grid}}}$ are discrete parameters, $b_k(\mathbf{x})$ is a unitless interpolation kernel and N_{grid} is the number of grid points. Let \mathbf{y} be a measurement vector and $(\cdot)^T$ denote transposition. Tomographic reconstruction is an estimation of $\beta = (\beta_1, \dots, \beta_{N_{\text{grid}}})^T$, that minimizes a *data fit* cost

$$\hat{\beta} = \arg \min_{\beta} \mathcal{E}[\mathbf{y}, \mathcal{F}(\beta)]. \quad (11)$$

Solving the minimization problem utilizes the gradient of \mathcal{E} with respect to β . In Eq. (5), $\beta_R(z)$ is known, thus $[d\mathcal{E}/d\beta] = [d\mathcal{E}/d\beta_c]$. The gradient is traditionally estimated iteratively, which would require $\mathcal{O}(N_{\text{grid}})$ simulations of the forward model, per iteration. The complexity is exacerbated by the complicated and non-linear form of the forward operator \mathcal{F} . Differentiation using an adjoint RTE was theorized in [32]. A theory and initial results of inverse rendering using P_n approximation are shown in [54]. In




Figure 3. Illustrating a surrogate function. (a) Computing the cost function gradient has very high computational complexity, making gradient based methods expensive. (b) The gradient of a surrogate function $\mathcal{E}[\mathbf{y}, \mathcal{F}(\beta|J_n)]$ is easily computed, driving the minimization forward by iterations. In each iteration a minimum of the surrogate function is attained (Eq. 13), followed by its evolution (Eq. 12), using SHDOM and the newly found β_n .

works that do not attempt 3D volumetric recovery, stochastic gradients [22] are used in a homogeneous medium, while numerical differentiation estimates unidirectional path integrals [13, 30]. Despite seeking non-volumetric recoveries, these works report prohibitive computational complexity.

3.2. Scalable Approach

To turn the hypothetical (11) to a feasible, tractable method, we simplify [27] the computations. Instead of optimizing a function whose gradient is complex to compute, we efficiently optimize β using *surrogate functions* that evolve through iterations (Fig. 3). Let β_n be an estimate of β in iteration n . The first step computes the in-scatter field, which corresponds to the current estimate β_n

$$J_n = \mathcal{J}(\beta_n). \quad (12)$$

This is the computationally complex part, hence we do not estimate its gradient. In fact, we hold J_n constant for a while, despite evolution of β . Let $\mathcal{F}(\beta|J_n) = \mathcal{MT}(\beta)J_n$ serve as a surrogate function in which J_n is fixed. In the second step, keeping J_n fixed, the following optimization finds the next estimate of β

$$\beta_{n+1} = \arg \min_{\beta} \mathcal{E}[\mathbf{y}, \mathcal{F}(\beta|J_n)]. \quad (13)$$

Data-fit using weighted least squares has the form

$$\begin{aligned} \mathcal{E}[\mathbf{y}, \mathcal{F}(\beta|J_n)] \\ = \frac{1}{2} [\mathbf{y} - \mathcal{MT}(\beta)J_n]^T \Sigma_{\text{meas}}^{-1} [\mathbf{y} - \mathcal{MT}(\beta)J_n]. \end{aligned} \quad (14)$$

Here Σ_{meas} is the covariance matrix of the (uncorrelated) measurements. The variance of measurement w is σ_w^2 .

Given measurement vector \mathbf{y} of length N_{meas} and J_n , solving (13,14) is simple using gradient-based methods.

$$\begin{aligned} \frac{\partial}{\partial \beta_k} \mathcal{E}[\mathbf{y}, \mathcal{F}(\beta|J_n)] \\ = \sum_{w=1}^{N_{\text{meas}}} \frac{1}{\sigma_w^2} [\mathcal{F}_w(\beta|J_n) - y_w] \mathcal{M}_w \left[\frac{\partial}{\partial \beta_k} \mathcal{T}(\beta) \right] J_n \end{aligned} \quad (15)$$

For a detector defined in (8),

$$\mathcal{M}_w \left[\frac{\partial}{\partial \beta_k} \mathcal{T}(\beta) \right] J_n = A_{w,k} + B_{w,k}, \quad (16)$$

where

$$\begin{aligned} A_{w,k} &= c_k(\mathbf{x}_0) I_{\text{BC}}(\mathbf{x}_0, \boldsymbol{\omega}^*) \exp \left[- \int_{\mathbf{x}^*}^{\mathbf{x}_0} \beta(\mathbf{r}) d\mathbf{r} \right], \\ B_{w,k} &= \int_{\mathbf{x}^*}^{\mathbf{x}_0} J_n(\mathbf{x}, \boldsymbol{\omega}^*) [b_k(\mathbf{x}) + c_k(\mathbf{x})\beta(\mathbf{x})] \\ &\quad \times \exp \left[- \int_{\mathbf{x}^*}^{\mathbf{x}} \beta(\mathbf{r}) d\mathbf{r} \right] d\mathbf{x}. \end{aligned} \quad (17)$$

The factor $c_k(\mathbf{x}) = - \int_{\mathbf{x}^*}^{\mathbf{x}} b_k(\mathbf{r}) d\mathbf{r}$ can be pre-computed, as it does not depend on the sought field. Given the gradient (15), Eq. (13) can be solved using gradient descent

$$\beta_n^{t+1} = \beta_n^t - \Delta \cdot \frac{\partial}{\partial \beta} \mathcal{E}[\mathbf{y}, \mathcal{F}(\beta|J_n)], \quad (18)$$

where t indexes a gradient-descent step, and Δ is the step size. After every N_{β} gradient descent steps, the field J is updated using SHDOM. Then, gradient-descent is resumed, using the updated J , for another set of N_{β} gradient descents, and so on. Starting with an initial guess β_0 , Eqs. (12,13,18) define an iterative optimization process. Eqs. (12,13) are alternated repeatedly until convergence.

In optimization problems, particularly nonlinear ones, the step size Δ needs to be well set. As a preliminary rule of thumb, we found that for stability, Δ needs to be lower, when the effective optical depth (not a 3D function) of a cloud is lower. Optical depth is currently retrievable using 1D radiative transfer methods [36].

3.3. Computational Efficiency

The surrogate function enables a major reduction in computational complexity, enabling for the first time, feasible 3D recovery. The most expensive part of the process is 3D rendering including multiple scattering, expressed by \mathcal{J} . Hence, \mathcal{J} must be performed sparingly. Let there be N_{DOF} degrees of freedom (unknowns) to recover. In the pioneering work of Ref. [22], focusing on retrieval of a homogeneous anisotropic medium, each gradient descent operation requires $\mathcal{O}(N_{\text{DOF}})$ operations of \mathcal{J} (*rendering operator*).

In our optimization of β , a gradient descent of the surrogate function uses a fixed J_n , hence no operation of \mathcal{J} is made, irrespective of $N_{\text{DOF}} = N_{\text{grid}}$. The operator \mathcal{J} is applied sporadically over time, with frequency that is unrelated to N_{grid} . Hence, our approach accelerates the most expensive part of gradient descent by $\mathcal{O}(N_{\text{grid}})$, facilitating optimization of $\mathcal{O}(10^5)$ unknowns in a short time. Our method is thus scalable.

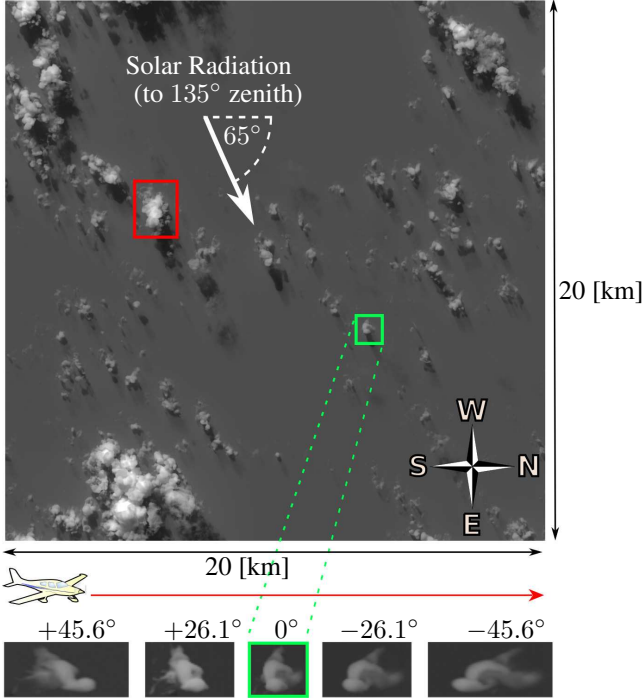


Figure 4. [Top] A cumulus cloud field created by LES [33] is rendered using SHDOM [18]. The clouds in the marked boxes are analyzed by tomography; [Bottom] SHDOM-generated radiance measurements, at five out of the nine viewing angles.

4. Limitations

We have so far not encountered basic limitations that are attributed to the surrogate-function method. Limitations we found in simulations at various scales appear to be due to the basic definition of the problem of passive 3D recovery:

- It is a non-convex problem, having local minima to which gradient descent converges. This can be mitigated by an informed initial guess e.g., using a layered model retrieval or an estimate of where the cloud’s outer boundary lies [8].
- At very high optical depth, the signal coming from deep within the cloud is outweighed by sensor noise. This desensitizes airborne/space-borne signals to cloud densities there. We hypothesize that this problem may be countered using good priors or regularization on cloud structure via fluid dynamics, and ground or in-situ measurements.

5. Simulations

5.1. Scene and Image Rendering

To test the approach in a realistically complex yet controlled scene, we use a *large eddy simulation* (LES) [12, 33] to generate a cloud field (Fig. 4). The LES is a comprehensive tool used by atmospheric scientists to computationally-create physically correct clouds [50]. The key output of the LES is liquid water content over a 3D grid. The clouds here hover above open ocean, whose reflectance [51] is affected by a 10 m/s wind and a typical

chlorophyll concentration of 0.01 mg/m^3 [40]. Here are additional scene parameters.

Atmospheric Constituents: The droplet size is Gamma-distributed, with effective droplet radius $r_{\text{eff}} = 10 \mu\text{m}$ effective variance $v_{\text{eff}} = 0.1$, which are typical values [43]. Mie scattering theory converts these quantities into $\beta_c(\mathbf{x})$ and $p_c(\boldsymbol{\omega} \cdot \boldsymbol{\omega}')$, the cloud phase function. We model molecular scattering using a summer mid-latitude vertical distribution [4], at altitudes ranging within $z \in [0, 20] \text{ km}$. We use $\lambda = 672 \text{ nm}$, where Rayleigh total optical thickness [16, 17] of non-cloudy air is only ~ 0.05 .

Image Rendering: The top of the atmosphere is irradiated by collimated sunlight, directed as described in Fig. 4. An SHDOM code [18] which is popular in atmospheric 3D radiative transfer, emulates¹ measurements similar to those taken by the Multi-angle Spectro-Polarimeter Imager (AirMSPI) [17] at 20 m resolution. The 9 viewing zenith angles are $\pm 70.5^\circ, \pm 60^\circ, \pm 45.6^\circ, \pm 26.1^\circ$, and 0° , where \pm indicates forward/backward along the flight path. Images as viewed from the instrument are rendered in Fig. 4. Poisson and quantization noise are included, according to the specifications of AirMSPI [17].

5.2. Recovery Results

We analyzed two atmospheric volumes, marked by green and red boxes in Fig. 4. Their respective dimensions are $0.72 \times 0.72 \times 1.44 \text{ km}^3$, $1.32 \times 2.22 \times 2.2 \text{ km}^3$. The analysis used $\Delta = 200$ and $N_\beta = 7$ and open horizontal boundaries, expressing observation of an isolated cloud. Updates of the surrogate function stopped when the cost function declines to 1% of its initial value. MATLAB was used on a 2.50 GHz Intel Xeon CPU. The rendering step, implemented in FORTRAN, was parallelized on 8 cores.

The converged reconstructions are displayed in Figs. 5,6 along with the ground-truth and a 3D relative error map. We quantify the recovery error using two measures defined in [2]. For the cloud marked in green (Fig. 4), the relative error in overall recovered mass is $\delta = (5 \pm 0.1)\%$. The relative local error is $\epsilon = (33 \pm 2)\%$. Fig. 7 displays ϵ in three slices. The error is larger at more opaque regions within the cloud. A scatter plot of true vs. estimated values is shown in Fig. 7c: its correlation is $\rho = 0.94$. For the cloud marked in red (Fig. 4), $\delta = 30\%$, $\epsilon = 70\%$, due to the loss of signal. Here $\rho = 0.76$. These results evolve from an initialization of no-cloud, with neither priors nor regularization on cloud structure. Runtime analysis is displayed in Fig. 8.

¹In [33], Fig. 4 is originally rendered using MC.

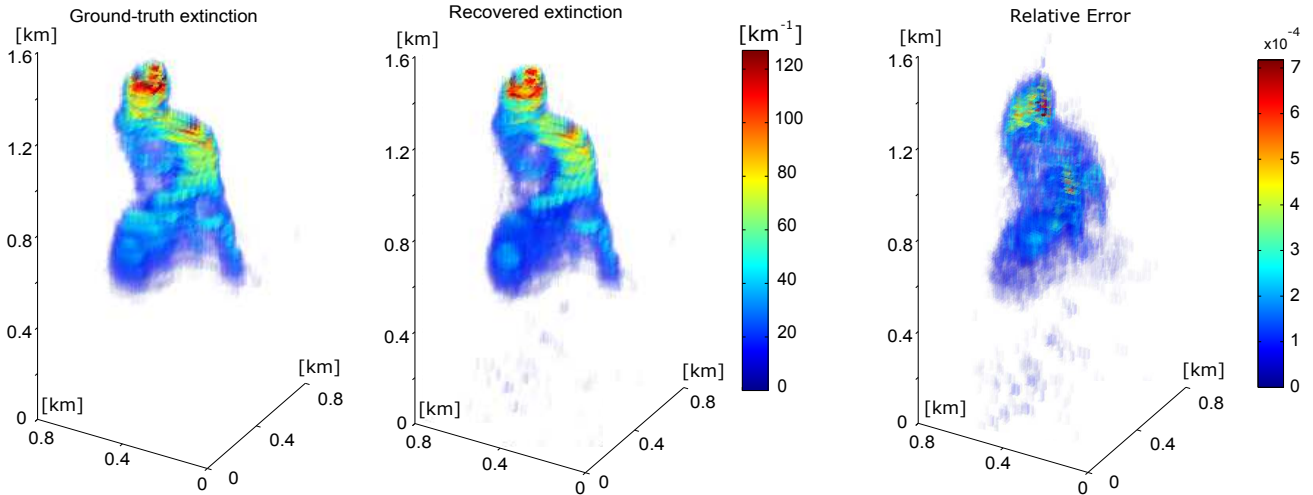


Figure 5. The unknown extinction field is discretized to a $36 \times 36 \times 36$ grid (46,656 unknowns). A volumetric comparison between the true LES-generated cloud and the recovered cloud, based on initialization that assumed no cloud at all. It is evident from the relative error map that the error is larger in the more opaque regions of the cloud.

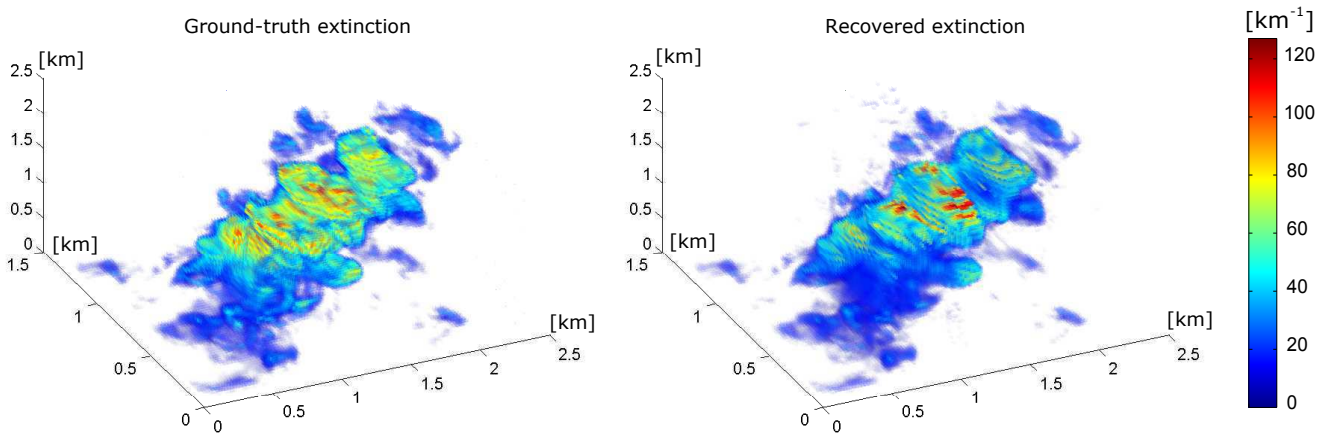


Figure 6. The unknown extinction field is discretized to a $66 \times 111 \times 43$ grid (315,018 unknowns). A volumetric comparison between the true LES-generated cloud and the recovered cloud, based on initialization that assumed no cloud at all. The cloud is extremely optically thick here, completely dissipating the signal (down to sensor noise level) in some areas.

6. Large-Scale Field Experiment

6.1. Real Data and Its Pre-Processing

It is desirable to apply this approach to real data, captured in the huge outdoor field, from multiple remote views [50]. In 2010 NASA initiated field campaigns with AirMSPI [17] at 20 km altitude, aboard NASA's ER-2 aircraft. AirMSPI has an eight-band push-broom camera, mounted on a gimbal for multi-angular observations over a $\pm 67^\circ$ along-track range. AirMSPI had undergone extensive geometric and radiometric calibration, to enable highly accurate quantitative measurements and subsequent products. In a *step-and-stare* mode, the spatial resolution is 10m.

We use the 660nm channel of data from a Pacific flight² done Feb/6/2013 at 20:27 GMT, around global coordinates 32N 123W. The flight path and three out of the nine view angles are displayed in Fig. 9a. We examine an atmospheric portion of $2.6 \text{ km} \times 3.4 \text{ km} \times 2.4 \text{ km}$ in East-North-Up coordinates.

Clouds move due to wind at their altitude, while AirSMPI flies. Motion along-track is difficult to resolve by images, since it aliases as parallax, globally affecting altitude estimation. Motion *across track* was estimated by aligning consecutive frames. This method yielded an assessed cross-track motion of $\approx 37 \text{ km/h}$.

²<https://eosweb.larc.nasa.gov/project/airmspi>




Figure 7. Analysis of the results corresponding to Fig. 5. (a) Slices of the LES generated cloud. (b) The relative error ϵ in the same slices. The error is generally larger for voxels with greater extinction. (c) Scatter plot of the estimated $\hat{\beta}$ vs the ground truth β_{true} (40% of the points were diluted randomly, for display clarity). A red line indicates perfect retrieval.

6.2. Results

The assumed open ocean parameters we used are chlorophyll concentration of 0.01 mg/m³ and surface wind of 15 km/h. The boundary conditions, droplet size distribution, update scheme and convergence criteria are as described in Sec. 5, while here $\Delta = 10$. The extinction field is discretized to $43 \times 56 \times 35$ grid points (86,688 unknowns). Based on the nine views, the converged volumetric reconstruction is displayed in Fig. 9d. The extinction values recovered by our reconstruction indicate an optical mean-free-path of 100-300 meters inside this cloud.

For cross-validation of the method, we excluded the nadir image data from the recovery process, thus using only eight out of the nine raw views in the 3D recovery. Afterwards, we used the recovered cloud to render the missing nadir view. Figure 9b,c compares the raw AirMSPI nadir image to the rendered corresponding view. The same cross-validation process was repeated for the +47° view angle (Fig. 9e,f). The coarseness of the estimated images is due to the cloud-voxel resolution (60m × 60m × 70m), which is lower than the AirMSPI sampling resolution (10m × 10m). Some of the artifacts evident in the volumetric reconstruction are due to the ocean, whose true parameters were not calibrated by us.

7. Summary

This work derives a framework for 3D tomography in scattering media, and in a large uncontrolled environment, where clouds reside. In this type of tomography, the signal is dominated by multiply scattered radiance and the source may be uncontrolled. Neither small angle nor an approximate scattering order limit are assumed by our work. To solve the problem in a tractable way, we develop a surrogate function suitable for the image formation model. This approach aims to enable large-scale tomography of the par-




Figure 8. Convergence rate of the simulation corresponding to Fig. 5, including current runtime in our implementation.

ticulate atmosphere, and thus advance our understanding of the physical processes involved. It can also find use in biomedical tomography, away from the small-angle, single-scattering or diffusion limits.

In computer vision, many recent studies attempt dehazing, i.e., background scene recovery through a scattering medium. The dominant model used so far in dehazing under natural lighting has been based on a single-scattering approximation, ignoring higher-order scattering. Our work herein suggests that an inverse problem including an arbitrary order of atmospheric daylight scattering is solvable. Hence, dehazing may benefit from methods presented here.

Acknowledgements: We are grateful to Dave Diner, Anat Levin, Mike Garay and Daniel Rosenfeld for fruitful discussions. We thank Georgios Matheou, Zheng Qu, Johanan Erez, Ina Talmon, Dani Yagodin for support and Frank Evans for his online code. The work is supported by the US-Israel Binational Science Foundation (BSF grant 2012202), E. and J. Bishop Research Fund and Amazon WSR Grants. AD's work was carried out at the Jet Propulsion Laboratory, California Inst. of Technology, supported by NASA's

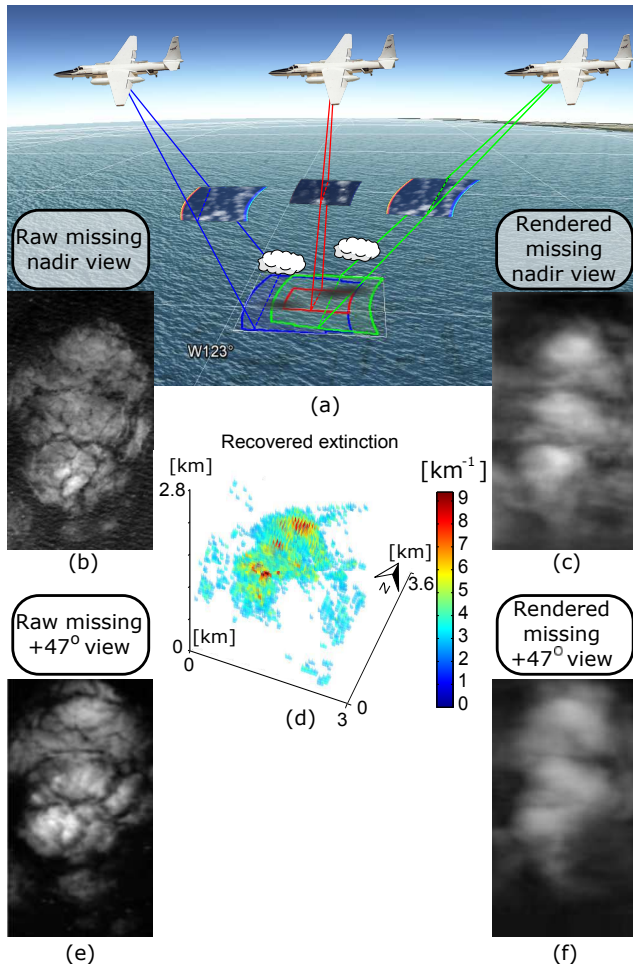


Figure 9. (a) Google Maps illustration of AirMSPI push-broom scan over the Pacific. A domain is viewed at nine angles (three displayed above). (b) Raw AirMSPI measurement at nadir view. (c) Nadir rendered view of a 3D cloud estimated using data that had excluded the nadir. (d) Recovered volumetric extinction field of the sought cloud using all 9 views. (e) Raw AirMSPI measurement at $+47^\circ$ view. (f) The $+47^\circ$ rendered view of a 3D cloud estimated using data that had excluded the $+47^\circ$ raw view.

SMD/ESD/RST and ESTO/AIST programs. YYS is a Landau Fellow - supported by the Taub Foundation. His work is conducted in the Ollendorff Minerva Center. Minerva is funded through the BMBF.

References

[1] E. H. Adelson and J. Y. A. Wang. Single lens stereo with a plenoptic camera. *IEEE TPAMI*, 14:99–106, 1992.

[2] 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.

[3] M. Alterman, Y. Y. Schechner, M. Vo, and S. G. Narasimhan. Passive tomography of turbulence strength. In *Proc. ECCV*, 47–60, 2014.

[4] G. P. Anderson, S. Clough, F. Kneizys, J. Chetwynd, and E. P. Shettle. AFGL atmospheric constituent profiles (0.120 km). Technical report, DTIC Document, 1986.

[5] S. Arridge and J. Schotland. Optical tomography: Forward and inverse problems. *arXiv:0907.2586*, 2009.

[6] 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.

[7] G. Bal. Inverse transport theory and applications. *Inverse Probl.*, 25(5):053001, 2009.

[8] G. Bal, J. Chen, and A. B. Davis. Reconstruction of Cloud Geometry from Multi-Angle Images, Submitted to *Inverse Probl. and Imag.*, 2015.

[9] 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.

[10] S. Chandrasekhar. *Radiative Transfer*. Dover Pub., 1960.

[11] Y.-C. Chen, M. W. Christensen, G. L. Stephens, and J. H. Seinfeld. Satellite-based estimate of global aerosol-cloud radiative forcing by marine warm clouds. *Nature Geo.*, 2014.

[12] D. Chung and G. Matheou. Large-eddy simulation of stratified turbulence. Part I: A vortex-based subgrid-scale model. *J. Atmos. Sci.*, 71:1863–1879, 2014.

[13] C. Cornet and R. Davies. Use of MISR measurements to study the radiative transfer of an isolated convective cloud: Implications for cloud optical thickness retrieval. *J. Geophys. Res. : Atmos.*, 113(D4), 2008.

[14] A. B. Davis and A. Marshak. Solar radiation transport in the cloudy atmosphere: A 3D perspective on observations and climate impacts. *Rep. Prog. Phys.*, 73:026801, 2010.

[15] M. Dierolf, A. Menzel, P. Thibault, P. Schneider, C. M. Kewish, R. Wepf, O. Bunk, and F. Pfeiffer. Ptychographic x-ray computed tomography at the nanoscale. *Nature*, 467(7314):436–439, 2010.

[16] D. J. Diner, J. C. Beckert, T. H. Reilly, C. J. Bruegge, J. E. Conel, R. A. Kahn, J. V. Martonchik, T. P. Ackerman, R. Davies, S. A. W. Gerstl, H. R. Gordon, J. P. Muller, R. B. Myneni, P. J. Stellars, B. Pinty, and M. M. Verstraete. Multi-angle Imaging SpectroRadiometer (MISR) instrument description and experiment overview. *IEEE T. Geosci. Remote. Sens.*, 36:1072–1087, 1998.

[17] D. J. Diner, F. Xu, M. J. Garay, J. V. Martonchik, B. Rheingans, S. Geier, Ab. Davis, B. Hancock, V. Jovanovic, and M. Bull. The Airborne Multiangle SpectroPolarimetric Imager (AirMSPI): A new tool for aerosol and cloud remote sensing. *Atmos. Meas. Tech.*, 6:2007–2025, 2013.

[18] K. F. Evans. The spherical harmonics discrete ordinate method for three-dimensional atmospheric radiative transfer. *J. Atmos. Sci.*, 55:429–446, 1998.

[19] R. Fattal. Single image dehazing. In *Proc. ACM TOG*, 27:72, 2008.

[20] C. Fuchs, M. Heinz, M. Levoy, H.-P. Seidel, and H. Lensch. Combining confocal imaging and descattering. In *Proc. EGSR*, 27:1245–1253, 2008.

[21] A. Gibson, J. Hebden, and S. R. Arridge. Recent advances in diffuse optical imaging. *Phys. Med. Biol.*, 50(4):R1, 2005.

- [22] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin. Inverse volume rendering with material dictionaries. *Proc. ACM TOG*, 32:162, 2013.
- [23] R. Horstmeyer, G. Euliss, R. Athale, and M. Levoy. Flexible multimodal camera using a light field architecture. In *Proc. IEEE ICCP*, 2009.
- [24] I. Ihrke, B. Goidluecke, and M. Magnor. Reconstructing the geometry of flowing water. In *Proc. IEEE ICCV*, 1055–1060, 2005.
- [25] Y. Ji, J. Ye, and J. Yu. Reconstructing gas flows using light-path approximation. In *Proc. IEEE CVPR*, 2507–2514, 2013.
- [26] J. Kim, D. Lanman, Y. Mukaigawa, and R. Raskar. Descattering transmission via angular filtering. In *Proc. ECCV*, 86–99. Springer, 2010.
- [27] A. Levis, Y. Y. Schechner, A. Aides and A. B. Davis. An Efficient Approach for Optical Radiative Transfer Tomography using the Spherical Harmonics Discrete Ordinates Method. *arXiv:1501.06093*, 2015.
- [28] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz. Light field microscopy. *Proc. ACM TOG*, 25:924–934, 2006.
- [29] 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.
- [30] R. Marchand and T. Ackerman. Evaluation of radiometric measurements from the NASA Multiangle Imaging Spectro-Radiometer (MISR): Two-and three-dimensional radiative transfer modeling of an inhomogeneous stratocumulus cloud deck. *J. Geophys. Res. : Atmos.*, 109(D18), 2004.
- [31] A. Marshak and A. B. Davis. *3D Radiative Transfer in Cloudy Atmospheres*. Springer, 2005.
- [32] W. Martin, B. Cairns, and G. Bal. Adjoint methods for adjusting three-dimensional atmosphere and surface properties to fit multi-angle/multi-pixel polarimetric measurements. *J. Quant. Spectrosc. Radiat. Transfer*, 144:68–85, 2014.
- [33] G. Matheou and D. Chung. Large-eddy simulation of stratified turbulence. Part II: Application of the stretched-vortex model to the atmospheric boundary layer. *J. Atmos. Sci.*, 71:4439–4460, 2014.
- [34] 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.
- [35] N. J. Morris and K. N. Kutulakos. Dynamic refraction stereo. In *Proc. IEEE ICCV*, 1573–1580, 2005.
- [36] T. Nakajima and M. D. King. Determination of the optical thickness and effective particle radius of clouds from reflected solar radiation measurements. Part 1: Theory. *J. Atmos. Sci.*, 47:1878–1893, 1990.
- [37] T. Nakajima, M. D. King, J. D. Spinhirne, and L. F. Radke. Determination of the optical thickness and effective particle radius of clouds from reflected solar radiation measurements. Part 2: Marine stratocumulus observations. *J. Atmos. Sci.*, 48(5):728–751, 1991.
- [38] E. Namer, S. Shwartz and Y. Y. Schechner Skyless polarimetric calibration and visibility enhancement, *Optics express* 17:472-493, 2009.
- [39] S. G. Narasimhan, S. K. Nayar, B. Sun, and S. J. Koppal. Structured light in scattering media. In *Proc. IEEE ICCV*, 420–427, 2005.
- [40] NEO - NASA Earth Observations. AQUA/MODIS. <http://neo.sci.gsfc.nasa.gov/>.
- [41] R. Pincus and K. F. Evans. Computational cost and accuracy in calculating three-dimensional radiative transfer: Results for new implementations of Monte Carlo and SHDOM. *J. Atmos. Sci.*, 66:3131–3146, 2009.
- [42] S. Platnick, M. D. King, S. A. Ackerman, W. P. Menzel, B. A. Baum, J. C. Riédi, and R. A. Frey. The MODIS cloud products: Algorithms and examples from Terra. *IEEE T. Geosci. Remote. Sens.*, 41(2):459–473, 2003.
- [43] R. R. Rogers and M. Yau. *A Short Course in Cloud Physics*. International Series in Natural Philosophy, 1989.
- [44] D. Rosenfeld. Aerosol-cloud interactions control of earth radiation and latent heat release budgets. In *Solar Variability and Planetary Climates*, pages 149–157. Springer, 2007.
- [45] D. Rosenfeld, M. O. Andreae, A. Asmi, M. Chin, G. Leeuw, D. P. Donovan, R. Kahn, S. Kinne, N. Kivekäs, M. Kulmala, W. Lau, K. S. Schmidt, T. Suni, T. Wagner, M. Wild and J. Quaas. Global observations of aerosol-cloud-precipitation-climate interactions. *Rev. Geophys.*, 52:750–808, 2014.
- [46] Y. Y. Schechner, D. J. Diner and J. V. Martonchik, Spaceborne underwater imaging, In *Proc. Proc. IEEE ICCP*, 2011.
- [47] Y. Tian and S. G. Narasimhan. Seeing through water: Image restoration using model-based tracking. In *Proc. IEEE ICCV*, 2303–2310, 2009.
- [48] T. Treibitz and Y. Y. Schechner, Polarization: Beneficial for visibility enhancement? In *Proc. IEEE CVPR*, 2009.
- [49] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic reconstruction of transparent objects. *Proc. EGSR*, 2006.
- [50] D. Veikherman, A. Aides, Y. Y. Schechner, and A. Levis. Clouds in the cloud. *Proc. ACCV*, 2014.
- [51] E. F. Vermote, D. Tanré, J.-L. Deuze, M. Herman, and J.-J. Morcette. Second simulation of the satellite signal in the solar spectrum, 6S: An overview. *IEEE T. Geosci. Remote. Sens.*, 35:675–686, 1997.
- [52] L. V. Wang and S. Hu. Photoacoustic tomography: in vivo imaging from organelles to organs. *Science*, 335(6075):1458–1462, 2012.
- [53] G. Wetzstein, D. Roodnick, W. Heidrich, and R. Raskar. Refractive shape from light field distortion. In *Proc. IEEE ICCV*, 1180–1186, 2011.
- [54] S. Wright, M. Schweiger, and S. Arridge. Reconstruction in optical tomography using the P_n approximations. *Meas. Sci. Tech.*, 18:79, 2007.
- [55] 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.
- [56] X. Zhu and P. Milanfar. Stabilizing and deblurring atmospheric turbulence. In *Proc. IEEE ICCP*, 2011.