

Distributed Sky Imaging Radiometry and Tomography

Amit Aides, Aviad Levis, Vadim Holodovsky, Yoav Y. Schechner, *Member, IEEE*,
Dietrich Althausen, and Adi Vainiger

Abstract—The composition of the atmosphere is significant to our ecosystem. Accordingly, there is a need to sense distributions of atmospheric scatterers such as aerosols and cloud droplets. There is growing interest in recovering these scattering fields in three-dimensions (3D). Even so, current atmospheric observations usually use expensive and unscalable equipment. Moreover, current analysis retrieves partial information (e.g., cloud-base altitudes, water droplet size at cloud tops) based on simplified 1D models. To advance observations and retrievals, we develop a new computational imaging approach for sensing and analyzing the atmosphere, volumetrically. Our approach comprises a ground-based network of cameras. We deployed it in conjunction with additional remote sensing equipment, including a Raman lidar and a sunphotometer, which provide initialization for algorithms and ground truth. The camera network is scalable, low cost, and enables 3D observations in high spatial and temporal resolution. We describe how the system is calibrated to provide absolute radiometric readouts of the light field. Consequently, we describe how to recover the volumetric field of scatterers, using tomography. The tomography process is adapted relative to prior art, to run on large-scale domains and being in-situ within scatterer fields. We empirically demonstrate the feasibility of tomography of clouds, using ground-based data.

Index Terms—Computational Photography, Imaging Radiometry, Computed Tomography, Internet of Things, Sensor Networks

1 INTRODUCTION

SUSPENDED scattering particles are abundant in the atmosphere. These particles, especially aerosols, cloud droplets and ice crystals significantly impact our life, society and economy in different time scales. On the immediate scale, they reduce visibility and solar power, may pollute the air and threaten aviation. The latter impact was exemplified by the 2010 eruption of Iceland’s Eyjafjallajökull volcano, whose ash (aerosols) covered much of Europe’s airspace, grounding all flights. On a short time scale, aerosols affect cloud formation and precipitation. In the longer term, aerosol and cloud processes have a significant impact on the climate system [1].

Be it for conceiving three dimensional (3D) flight paths that minimize hazards, managing solar power [2], [3], understanding cloud physics or improving climate prediction [4], there is a need to sense spatiotemporal distributions of these particles. The distributions express density and micro-physical characteristics [5] in broad 3D fields (Fig. 1). However, most current systems do not support such sensing. Remote sensing methods [6] largely assume a plane-parallel atmospheric model, in which particle characteristics are horizontally uniform over wide spans. In tune with the plane-parallel model [7], [8], atmospheric observational hardware has been designed with coarse spatial resolution, though with sophisticated multi-channel calibrated optical sensing. Moreover, atmospheric-observing satellites have so far been costly and difficult to repair or scale.




Figure 1: We seek scalable multi-view radiometric imaging and tomographic reconstruction of large scale three dimensional atmospheric volumes.

Ground observation systems are easier to place in the field and repair. Nevertheless, they have thus far been challenging to scale. Aerosol lidars [9], [10], specifically, probe vertical profiles of the atmosphere in high vertical resolution, using excellent optics and calibration processes. However, due to their cost, they are only sparsely distributed over the world. Hence they cannot practically provide three-dimensional sensing.

Sensing atmospheric scatterers in 3D over wide fields requires two new technologies. First is a scattering-based computed tomography (CT) approach. CT [11], [12], [13], [14], [15]

- A. Aides, A. Levis, V. Holodovsky, Y. Schechner, and A. Vainiger are with the Viterbi Faculty of Electrical Engineering, Technion-Israel Institute of Technology, Haifa, Israel.
E-mail: amitibo@campus.technion.ac.il
- D. Althausen is with the Leibniz Institute for Tropospheric Research, Leipzig, Germany.

is highly suitable for sensing transparent media [16], [17], [18], [19]. It can benefit from complementary computational advances seen in other problems, such as lighting analysis [20], [21], and even neural networks [22]. CT has thus recently been generalized to recover scattering media using passive and other computational imaging methods [3], [23], [24], [25], [26], [27], [28], [29], [30]. However, hardware to *sense* nature accordingly and thus provide the required data has been missing. As all CT approaches, also passive scattering-tomography requires multi-directional imaging of the atmospheric domain, using a well-calibrated system having high resolution. This raises the second requirement for sensing: a computational imaging system that can work in the field and provide the necessary data. This paper describes a system to meet these needs and complement existing technologies. The system comprises a scalable and reproducible calibrated camera network around a lidar and a sunphotometer. Furthermore, we demonstrate ground-based *tomography* of clouds using this system.

Multi-view imaging [31], [32], [33], [34] is not new to remote sensing. Nevertheless, even high resolution multi-angular orbiting and high altitude platforms such as MISR [35], AirMSPI [36], and MAIA [37] require several minutes in order to capture a domain from multiple views, during which the atmospheric state changes. In contrast, a camera network captures the domain simultaneously, at any time of interest. Furthermore, a camera network on the ground is easy to deploy, upgrade, and scale - provided that it leverages on the economy of scale and spatial breadth associated with cellular communication. Thus, the system is comprised of low-cost untethered units which are part of the internet of things (IoT). This is contrary to prior generations of sky-observing cameras [38], which are tethered and expensive hence operated in solitude or in small unscalable numbers per viewed field.

The scale of the network requires low-priced units. Therefore, special care must be taken when calibrating the network units. This is especially important because the data is intended for quantitative tomographic recovery. Hence the readout of the cameras must be trustworthy as scientific radiometric instruments, contrary to systems that focus on geometry [39], [40], [41], [42], [43], [44], [45], [46]. Hence, in this paper, there is special attention to radiometric aspects, in addition to the scalability of the network system.

In addition to a wide network of calibrated cameras (radiometric imagers), the system uses an aerosol lidar and a sunphotometer. These provide an important calibration anchor and constraints for tomographic recovery.

The rest of this paper is organized as follows. The basics of the image formation model are given in Sec. 2. The interaction of sunlight with the atmosphere is described in Sec. 3. In Sec. 4, we describe the tomographic algorithm we use for recovering the 3D atmospheric medium. Details on our distributed camera network are provided in Sec. 5. Sec. 6 explains processes for calibrating the camera network, including conversion of machine vision cameras into absolute radiometric sensors. In Sec. 7, we present simulations of the camera network. Initial real world experiments show in Sec. 8 the feasibility of using distributed imaging for atmospheric scattering tomography. Finally, Sec. 9 concludes the paper with a brief summary and discussion.

2 IMAGE FORMATION

A ground-based camera samples the light field at a specific location and in a range of directions. We now express the signal at a pixel, for a camera which is calibrated. The cameras and their calibration are described in Secs. 5,6. A pixel in the camera samples a particular line of sight (LOS). The camera location and LOS direction of a specific pixel are jointly referred to as w . Solar radiation propagates through the atmosphere, undergoes scattering and absorption, and eventually reaches a camera on the ground. There, at location and direction w , the corresponding spectral radiance per wavelength λ is denoted $L_{\lambda,w}^{\text{ground}}$, in units of $[\frac{W}{cm^2 \cdot Sr \cdot nm}]$.

Let h be Planck's constant and c the speed of light. The expected number of photons at wavelength λ is the product of light power in Watts and the factor $\frac{\lambda}{hc}$. A sensor on the ground responds to photons in a range of wavelengths. The response of the sensor is given by the sensor's quantum efficiency, QE_{λ} , which measures the probability for an electron to be produced per incident photon. The expected number density of electrons generated by photons at wavelength λ during exposure time Δt is

$$\begin{aligned} N_{\lambda,w}^{\text{sensor}} &= QE_{\lambda} \frac{\lambda}{hc} L_{\lambda,w}^{\text{ground}} \Delta t \\ &= \Gamma_{\lambda} L_{\lambda,w}^{\text{ground}} \Delta t \quad \left[\frac{\text{electrons}}{cm^2 \cdot Sr \cdot nm} \right], \end{aligned} \quad (1)$$

where we define

$$\Gamma_{\lambda} = QE_{\lambda} \frac{\lambda}{hc}. \quad (2)$$

For a sensor having a linear radiometric response, the ratio between $N_{\lambda}^{\text{sensor}}$ and the sensor readout is a fixed-value, denoted γ^{cam} . It depends on the physical and electronic properties of the sensor. In the case of a camera, the readout is a gray level value of a pixel. The ratio γ^{cam} depends on the optics, size of the pixel, and the pixel's receptive solid angle. Let $\gamma_{\Delta t}^{\text{cam}} = \gamma^{\text{cam}} \Delta t$. The expected graylevel is then

$$\hat{I}_w = \gamma_{\Delta t}^{\text{cam}} \int_{\lambda} N_{\lambda,w}^{\text{sensor}} d\lambda = \gamma_{\Delta t}^{\text{cam}} \int_{\lambda} \Gamma_{\lambda} L_{\lambda,w}^{\text{ground}} d\lambda. \quad (3)$$

Eq. (3) describes the image formation process for ground based cameras. It relates radiance to graylevel values, assuming the optical transfer of the camera is uniform.

In reality, a camera has nonuniform optical transfer (vignetting). Section 6 describes how we calibrate and compensate for this nonuniformity. Empirical measurements include random noise, which mainly originates from the discrete nature of photons and electric charges. Typically, this noise is modeled by a Poisson process. Denote incorporation of noise into the expected signal by the operator \mathcal{N} . The noisy measurement in LOS w is then

$$y_w = \mathcal{N}\{\hat{I}_w\}. \quad (4)$$

Direct sunlight should generally be blocked from entering the optical system (including the housing's glass dome). The reason is that direct sunlight creates a strong lens-flare [42], [47] through internal reflection between optical elements. Flare biases image readouts in a spatially varying

manner [48], thus affecting subsequent tomographic recovery. Hence, the camera is equipped with a sunshader tuned to the width of the optics and the solar direction.¹

3 RADIATIVE TRANSFER FORWARD MODEL

One of the possible uses for a wide, dense radiometrically calibrated camera network is scattering tomography, e.g., to estimate the 3D volumetric distribution of the scatterers' density n . The denser a voxel, the stronger the extinction of radiation by this voxel, meaning a larger extinction coefficient, β . Estimation requires a *forward model* for propagation of solar radiation through the atmosphere to the camera network. The solar irradiance at the top of the atmosphere (TOA), L^{TOA} , is flux passing through a square-meter, given in units of $[\frac{W}{m^2 \cdot nm}]$. It is modelled using the ASTM E-490 solar spectral irradiance [49], multiplied by the cosine of the Sun zenith angle, Φ^{SR} :

$$L_{\lambda}^{\text{TOA}} = \cos \Phi^{\text{SR}} L_{\lambda}^{\text{ASTM}}. \quad (5)$$

Solar radiation interacts with the atmosphere by multiple scattering [23] and possibly absorption. The interaction is modeled by radiative transfer (RT). We define light propagation and interactions using an operator RT . This operator depends on the extinction field β_{λ} of the atmospheric medium at wavelength λ . Thus we write

$$L_{\lambda,w}^{\text{ground}} = L_{\lambda}^{\text{TOA}} \text{RT}\{\beta_{\lambda}\}. \quad (6)$$

Let us substituting Eq. (6) into Eq. (3). This yields the forward model for the readout at a pixel that views LOS w

$$\hat{I}_w = \gamma_{\Delta t}^{\text{cam}} \int_{\lambda} \Gamma_{\lambda} L_{\lambda}^{\text{TOA}} \text{RT}\{\beta_{\lambda}\} d\lambda. \quad (7)$$

We use² the spherical harmonic discrete ordinates method (SHDOM) [50] to model RT. This can be stated as

$$\text{RT}\{\beta_{\lambda}\} \approx \text{SHDOM}\{\beta_{\lambda}\}. \quad (8)$$

Substituting Eq. (8) into Eq. (7) yields our numerical forward model,

$$\mathcal{F}_w(\beta_{\lambda}) = \gamma_{\Delta t}^{\text{cam}} \int_{\lambda} \Gamma_{\lambda} L_{\lambda}^{\text{TOA}} \text{SHDOM}\{\beta_{\lambda}\} d\lambda. \quad (9)$$

Ideally, to calculate the integral in Eq. (9), RT would be run at multiple infinitesimally narrow wavelength bands. That would incur high computational resources. To simplify calculations, assume that the optical parameters of the atmospheric medium do not change considerably across the spectral band in which Γ_{λ} is significant. We use the color bands Λ of the camera (Red, Green, Blue, in our case), i.e., $\Lambda \in [\Lambda_R, \Lambda_G, \Lambda_B]$. In each band, define

$$\beta_{\Lambda} = \frac{\int_{\lambda} \beta_{\lambda} L_{\lambda}^{\text{TOA}} d\lambda}{\int_{\lambda} L_{\lambda}^{\text{TOA}} d\lambda}, \quad (10)$$

where each integral is across a spectral band Λ . Then, Eq. (9) is approximated by

$$\mathcal{F}_w(\beta) \approx \gamma_{\Delta t}^{\text{cam}} \text{SHDOM}\{\beta_{\Lambda}\} \int_{\lambda} \Gamma_{\lambda} L_{\lambda}^{\text{TOA}} d\lambda. \quad (11)$$

1. The exception to this mode of operation is extrinsic geometric calibration, described in Sec. 6.1. This process relies explicitly on direct imaging of the sun.

2. The SHDOM code which we use can be downloaded from <https://github.com/aviadleviss/pyshdom>




Figure 2: Simulated haze blobs: (a) Iso-surface visualization of a synthetic 3D aerosol distribution. (b) The simulated haze blobs as synthesized from the viewpoint of one of the real cameras. The surrounding buildings and the sunshader are masked according to the field of view of the real-world camera.

An atmospheric voxel can be comprised of a mixture of several types of scatterers (cloud droplets, aerosols and air molecules). Each type has different optical properties. The overall optical extinction coefficient is the sum of the respective extinction coefficients

$$\beta_{\Lambda} = \beta_{\Lambda}^{\text{cloud}} + \beta_{\Lambda}^{\text{aerosol}} + \beta_{\Lambda}^{\text{air}}. \quad (12)$$

Figure 2 illustrates a simulated 3D aerosol distribution in the form of two dense blobs (the setup is described in Sec. 7) and a corresponding synthetic camera view.

A cloud droplet has optical cross-section $\sigma_{\Lambda}^{\text{cloud}}$, which is computed through *Mie* theory. Using typical parameters taken from the literature [51], the cloud droplet size in our tests is assumed to be Gamma-distributed, with effective droplet radius $r_{\text{eff}} = 10 \mu\text{m}$ and a dimensionless effective variance $v_{\text{eff}} = 0.1$. We rewrite $\beta_{\Lambda}^{\text{cloud}}$ as the product of the (wavelength independent) droplet density n and the pre-calculated optical cross-section

$$\beta_{\Lambda}^{\text{cloud}} = \sigma_{\Lambda}^{\text{cloud}} n. \quad (13)$$

4 TOMOGRAPHIC RECOVERY - INVERSE MODEL

Tomographic reconstruction seeks estimation of the optical densities β_{Λ} . Similarly to [23], this task is formulated as a minimization of a *data fitting* cost \mathcal{E} ,

$$\hat{\beta}_{\Lambda} = \arg \min_{\beta_{\Lambda}} \mathcal{E}[\mathbf{y}, \mathcal{F}(\beta_{\Lambda})]. \quad (14)$$

Here \mathbf{y} is a vector which concatenates all measurements y_w , for all pixels (directions), spectral bands and cameras. Generally, Eq. (14) may seek estimation of the densities of all scatterers mentioned in Eq. (12). However, air density varies mainly with altitude and its extinction profile $\beta_{\Lambda}^{\text{air}}$ can be derived from data transmitted daily by radiosondes launched by nearby meteorological stations. Consequently, $\beta_{\Lambda}^{\text{air}}$ is not considered to be unknown.

When clouds are absent, $\beta_{\Lambda}^{\text{cloud}} = 0$. Then the 3D distribution left to estimate by tomography is that of $\beta_{\Lambda}^{\text{aerosol}}$. However, 3D variations of $\beta_{\Lambda}^{\text{aerosol}}$ are usually by far weaker than that of $\beta_{\Lambda}^{\text{cloud}}$, when clouds exist in the field of view. In the latter case, we assume that aerosol concentrations (far from industrial sources or volcanic plumes) vary mainly



Figure 3: The Raman aerosol lidar used during our experiments. It is the PollyXT system by TROPOS, installed on the roof the Meyer (EE) Building, Technion, Haifa, Israel. The AERONET sunphotometer can be seen on top of the installation cage.

with altitude. Then, it suffices to measure $\beta_{\Lambda}^{\text{aerosol}}$ on a vertical line above the camera network domain. This is achieved using an aerosol lidar (Fig. 3), which is embedded in the sensor network.

For the purposes of demonstration and without loss of generality, the expressions here focus on a single dominant scatterer type, i.e either cloud droplets or aerosols. For simplicity, we refer to estimation of cloud-droplet density, in 3D. Denote by vector \mathbf{y}_c the intensity measurements in all pixels (lines of sight) w of camera c . Correspondingly, $\mathcal{F}_c(\beta_{\Lambda})$ is the vector of modelled intensities in all the pixels and color channels of camera c , according to the forward model (Sec. 3). Then, the cost function in Eq. (14) is set to be

$$\mathcal{E} = \sum_{\Lambda} \sigma_{\Lambda}^{\text{cloud}} \sum_c |\mathcal{M}_c[\mathcal{F}_c(\beta_{\Lambda}) - \mathbf{y}_c]|^2. \quad (15)$$

Here \mathcal{M}_c is a digital mask applied to camera c . This mask was introduced in [52], [53], to attenuate the contribution of pixels around the Sun. Pixels near the sun tend to be noisy and with strong deviation from the model. Moreover, the mask expresses the mechanical sun-shader, which is used to block lens flare and saturation. The mask also marks the surrounding buildings and trees which occlude the sky.

We optimise \mathbf{n} by minimizing Eq. (15). To achieve this, we use a gradient-based method according to the formulation of [53]. The gradient of the fitting cost is

$$\frac{\partial}{\partial \mathbf{n}} \mathcal{E} = 2 \sum_{\Lambda} \sigma_{\Lambda}^{\text{cloud}} \sum_c \mathbf{Q}_c \{ \mathcal{M}_c \mathcal{J}_c(\mathbf{n}) \}^T \mathcal{M}_c [\mathcal{F}_c(\beta_{\Lambda}) - \mathbf{y}_c], \quad (16)$$

where T denotes transposition. Here $\mathcal{J}_c(\mathbf{n}) = \frac{\partial}{\partial \mathbf{n}} \mathcal{F}_c(\beta_{\Lambda})$ is the *Jacobian*. It expresses how the modelled (forward) image in camera c varies in response to variations in \mathbf{n} . The matrix \mathbf{Q}_c is used for conditioning the gradient, as

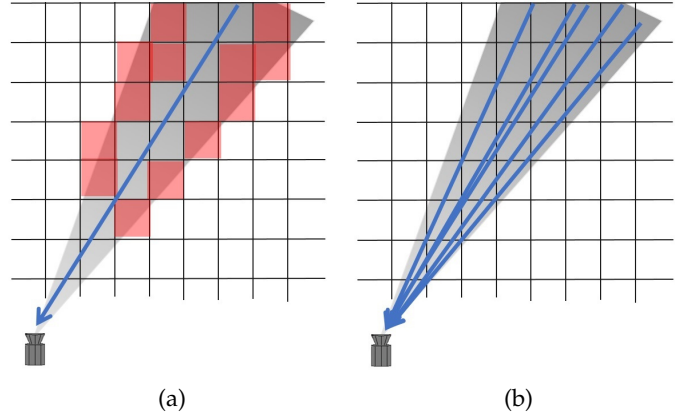


Figure 4: Perturbation of a back-projected ray. (a) When back-projecting a ray through the center of a pixel's instantaneous field of view w , some voxels are overlooked (marked by red rectangles). (b) Perturbation of the back-projected ray through the support of the pixel updates all its voxels.

described in Sec. 4.3. A Jacobian of the RT can be estimated numerically [54], [55], [56], [57], [58], [59]. We opt to use the approximate gradient derived in [23], [53], due to its speed, as reported in the literature.

4.1 Stochastic Gradient Descent

All LOS's w of all pixels in all cameras in the network constitute a set S_{net} . The number of measurements, accounting also for three color channels is $N_{\text{meas}} = 3|S_{\text{net}}|$. Tomography estimates scattering characteristics in N_{grid} voxels. When N_{grid} and N_{meas} are large, the gradient calculation step (Eq. 16) can become computationally expensive.

To speed up the optimization, it is possible to select at each iteration a different subset $S_{\text{sub}} \subset S_{\text{net}}$ of all measurements, and use only this subset when calculating Eq. (16). The subset can be selected independently each time or by weighted sampling from the fit cost (15). In both cases, this method is equivalent to stochastic gradient descent (SGD) [60]. We found that this form of SGD brings significant speedup to the convergence of the cost function, though some experimentation is needed for finding the right size of the subset. If the subset is too large, each iteration takes longer to calculate. If the subset is too small, noise introduced by SGD inhibits practical convergence.

4.2 Back-Projection Perturbation

In perspective cameras, each pixel has a receptive field (angular support) in the form of a cone around w (Fig. 4a), denoted \mathcal{C}_w . A single LOS back-projected from the center of the pixel overlooks some voxels which have (at least partial) overlap with \mathcal{C}_w . Densely sampling w increases N_{meas} and makes calculation of Eq. (16) prohibitively expensive. Instead, in each iteration, per pixel, a random ray $\omega^* \in \mathcal{C}_w$ is selected. In a sense, this method complements the SGD described in section 4.1.

4.3 Conditioning

The algorithm in [23] was developed and tested on images acquired far from the scattering domain. In this




Figure 5: The difference between spaceborne and in-situ setups: (a) The camera is far from the scattering domain. The LOSs of different pixels are approximately parallel. Voxels k and m have similar contribution to the fitted signals. (b) The cameras are near the scattering domain. Voxels close to the camera (e.g. voxel k) affect more LOSs than distant voxels (e.g. voxel m).

case (Fig. 5a), all grid voxels are similarly observed. However when the cameras are near or even inside the scattering medium, as is the case for ground-based cameras (Fig. 5b), some voxels affect the optimization more than others. As a consequence, the voxels near the cameras get updated more frequently and this leads to slow convergence [53] in analogy to ill-conditioned problems. To assist convergence, there is a need to condition the gradient in Eq. (16). This is achieved by multiplying the gradient by a diagonal matrix Q_c . Element (k, k) of Q_c is the reciprocal of the number of random LOS rays [53] which originate from the lens of camera c and pass through voxel k .

5 DISTRIBUTED RADIOMETRIC IMAGERS

Our computational imaging system is an autonomous radiometric imaging sensor network. It is designed for large scale deployment. Each camera is untethered and builds on the infrastructure of cellular communication technologies. Each unit has a lot in common with a smartphone: it has a small built-in computer board (we used the *Hardkernel Odroid U3*), which controls a camera and a cellular communication modem. As such, the units capitalize on the economy of scale of these off-the-shelf devices. The main distinction from commercial cellular phones is a robust structure, sustainable outdoors, with no maintenance for weeks.

The unit is solar powered. Solar energy stored in batteries sustain operation under partially clouded skies and at night. All the electronics reside in a standard housing, which is customized to hold the camera housing, external sunshader motor and an automobile windscreen washer. The washer automatically sprays distilled water each morning directly on the camera glass, to clean it. The washer pump is commanded by the on-board computer, and powered by the unit's batteries. Distilled water leave no residue upon drying. Water is stored in an external tank. Figure 6a shows

a single camera deployed in the field. Figure 6b shows the network infrastructure. Figure 6c shows the mechanical details of a unit. The camera sensor and fisheye lens are encased in a detachable rigid sub-housing. This sub-housing is used during camera calibration in the lab. Intrinsic and vignetting calibrations are then unaffected when the sub-housing is installed inside the overall system housing.

Via cellular communication, users can do two things from anywhere in the world where cellular communication reaches: access the acquired sky images and remotely control the camera settings. These operations are done via a custom software having a graphical user interface.³ This software also provides interactive visualization of the geometric output. These include LOSs of cameras, back-projected in 3D or re-projected to other camera views and 3D visualization of the network overlaid with reconstruction results. Additional information about the systems is given in the *Supplemental Material*.

6 CALIBRATION OF THE CAMERA UNITS

The cameras need to address radiometric requirements. Here we describe the optical aspects of a camera unit and its calibration. The camera sensor converts photons to red-green-blue (*RGB*) values. To use the images as radiometric samples, the *RGB* values are converted to physical quantities. First, the cameras are corrected for geometric and optical nonuniformities. This is described in Sec. 6.1 and 6.3, respectively. Then, the digital *RGB* values translate to scientifically meaningful radiance values, as explained in 6.4.

6.1 Geometric Calibration

Geometric calibration maps each pixel in the camera sensor to its LOS w . This is a two-step process. First, each pixel is mapped to a direction in the camera's coordinate system. This is referred to as *intrinsic calibration*. Each camera has a fisheye lens hence pixels are mapped to a hemisphere. Intrinsic calibration is done in the lab by photographing a checkerboard pattern multiple times while directing the camera in different directions. These images are processed to produce the intrinsic matrix and distortion coefficients of the camera sensor [61].

The second step (*extrinsic calibration*) relates the camera pose to locations and rotations in the Earth coordinate system. Extrinsic calibration of any camera is done after the unit is positioned in the field. The unit's location is known by a GPS receiver carried while laying the unit at its place. Determination of the orientation (yaw, pitch and roll angle) is done by fitting the known Sun azimuth and elevation to its empirical projection pixel in the camera sensor. Here we follow the method of [42]. Figure 7 shows how solar observations are fitted to the model of solar trajectory, according to the extrinsic calibration.

6.2 Dark Current

Dark current is caused by random generation of electrons in the sensor. Dark current depends on the exposure time,

³ The software of the camera network, both for the on-board computer and the graphical user interface, is available for download from <https://github.com/amitibo/cameranetwork>.




Figure 6: The distributed camera network. (a) A camera unit in the field. (b) The camera network can be arbitrarily scaled up, deployed and controlled from any location in the world where there is cellular connectivity. (c) The mechanical layout of the camera unit.

and is mainly noticeable in long exposures taken at night. This effect is calibrated and then compensated for. Per exposure time setting, the camera acquires multiple images at complete darkness in the lab. These images are averaged and stored, per exposure setting, on the unit's on-board computer. Afterwards, when images of targets and scenes are acquired, the dark current image of the corresponding exposure time is subtracted automatically.

6.3 Vignetting Correction

The optical transfer of fisheye lenses is highly nonuniform. To compensate for this effect, a custom calibration setup was built in a dark room. It comprises a fixed, diffuse light source target, about 5cm wide, whose power is regulated. It is based on the Colibri illumination system by ZEISS Microscopy. Its output passes an optical diffuser. The entire camera sub-housing (including sensor array, lens, cover glass dome and the sub-housing itself) is fixed to a gimbal,

far from the target. The gimbal rotates the camera rig in steps, across the entire range of view angles supported by the lens (Fig. 8a). In each step, the camera acquires an image of the fixed target. The target appears as a spot. Per spot and color channel, the central intensity (graylevel) value I and coordinates x of the spot are read. Afterwards, per color channel Λ , a 2D polynomial function f fits the image readout values. The basis for the fit is optimization,

$$\widehat{\eta}_{\Lambda} = \arg \min_{\eta_{\Lambda}} \sum_x |I_{\Lambda}(x) - f(x; \eta_{\Lambda})|^2, \quad (17)$$

where η_{Λ} are the polynomial coefficients (we got best results when using 4th degree 2D polynomials). Large distortions and measurement errors at the edge of the fisheye field of view create outliers. Hence, the RANSAC algorithm is used for the fitting. The function $f(x; \eta_{\Lambda})$ is then normalized such that its maximum value is 1. Once these polynomial functions are calibrated per camera unit, they are stored on the on-board computer of that unit. Then, these functions

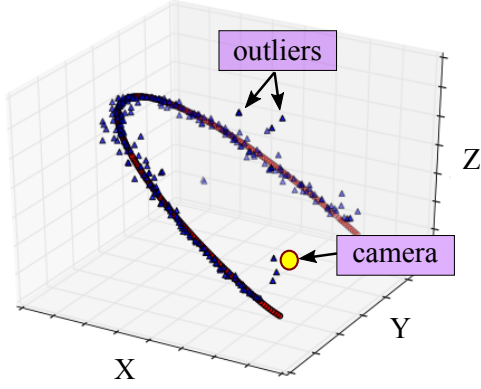


Figure 7: Extrinsic calibration fit. Red circles mark the Sun direction at different hours of the day, projected on the unit sphere. Blue triangles mark the Sun direction measurements at the corresponding times. Outliers are caused by clouds momentarily covering the Sun. The RANSAC algorithm fits the solar trajectory, even in partially clouded days.

are used automatically for vignetting compensation of each acquired image,

$$\hat{I}_\Lambda(x) = f^{-1}(I_\Lambda(x); \eta_\Lambda). \quad (18)$$

Here $\hat{I}_\Lambda(x)$ denotes the graylevel value corrected for nonuniformity. The process is illustrated in Fig. 8.

6.4 Absolute Radiometric Calibration

Let the nonuniformity of the camera be calibrated and compensated, as explained in Sec. 6.3. Then, there is a need to quantitatively relate image graylevels to radiance. This relation requires absolute radiometric calibration. The function Γ_λ (Eq. 2) can be obtained from the specifications of the camera manufacturer. Thus, absolute radiometric calibration amounts to the estimation of γ^{cam} . Here we describe a way to achieve this outdoors using a sunphotometer.

A sunphotometer is a multi-channel, automatic calibrated multispectral sun-and-sky scanning radiometer. The sunphotometer measures the direct solar irradiance and sky radiance at the Earth's surface. It records measurements at pre-determined discrete wavelengths in near-IR and visible spectral bands. Measurements are taken in a couple of dozen directions (Fig. 9). We denote the set of the sunphotometer LOS's (direction and ground location) as S_{S-P} . The sunphotometer is part of an AEROSOL ROBOTIC NETWORK (AERONET) station [62], and we place it among the network cameras (Fig. 3).

During absolute radiometric calibration of a camera, we position the camera in proximity to the sunphotometer. Hence, both instruments are exposed to the same radiance.⁴ Moreover, recall that the camera is geometrically calibrated, as described later in Sec. 6.1. Hence, for each camera pixel, we know its corresponding LOS w . The sunphotometer samples $L_{\lambda,w}^{\text{ground}}$ at $w \in S_{S-P}$. Plugging these ground-truth radiance samples into Eq. (3,4) we get

$$y_w \approx \gamma_{\Delta t}^{\text{cam}} \int_\lambda \Gamma_\lambda L_{\lambda,w}^{\text{ground}} d\lambda. \quad (19)$$

4. On the scale of cloud features (decameters and more), the camera and sunphotometer are essentially co-located.




Figure 8: Vignetting calibration: Each camera sub-housing is set in a custom calibration setup. (a) A distant diffuse light source is acquired while the camera sub-housing pans and tilts in multiple directions. This scan samples and covers the fisheye hemisphere. (b) Per color channel, a 2D polynomial function is fitted to the measured intensity, creating a vignetting correction map. (c) An image before applying vignetting correction. (d) The same image after applying vignetting correction. The effect is most noticeable at the edges of the image (Best viewed in color).

Here y_w is the image grayscale measurement in the pixel corresponding to w . Eq. (19) is a linear constraint on γ^{cam} . Using all scan directions $w \in S_{S-P}$ yields a set of linear constraints. Estimation of γ^{cam} can use least squares,

$$\widehat{\gamma^{\text{cam}}} = \arg \min_{\gamma^{\text{cam}}} \sum_{w \in S_{S-P}} \left| y_w - \gamma^{\text{cam}} \Delta t \int_\lambda \Gamma_\lambda L_{\lambda,w}^{\text{ground}} d\lambda \right|^2. \quad (20)$$

We use RANSAC to handle outliers caused by nearby buildings and the automatic sunshader (Fig. 9b). The value of γ^{cam} is estimated separately for each camera color channel, because each channel has a distinct Γ_λ function.

Beyond camera calibration, we use the sunphotometer for its primary role of assessing aerosol micro-physics, and their optical characteristics: phase function and single scattering albedo. This is done through the AERONET inversion code. The AERONET code does not derive spatially variable aerosol properties in 3D. Instead, it relies on optical readouts integrated over the whole atmospheric column above the station. Hence, the AERONET products are, in a sense, an aggregate of spatially varying aerosol properties. They are used for initialization and a prior for 3D analysis.

An additional way to estimate γ^{cam} does not rely on sunphotometer radiance readouts, but on the AERONET aerosol product and lidar readouts. On a clear day at which the lidar indicates no significant aerosols, the camera takes

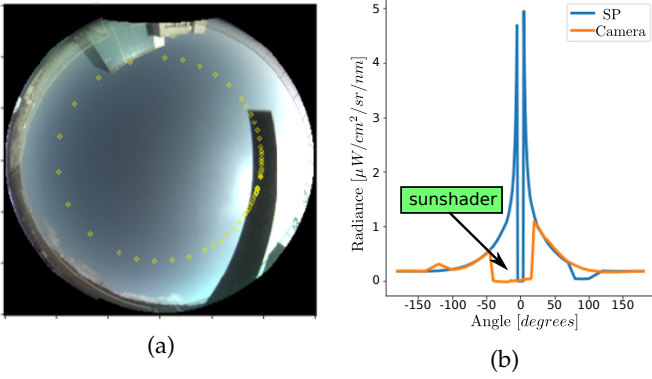


Figure 9: Absolute radiometric calibration of a camera: (a) Almucantar scan directions of an AERONET sunphotometer overlaid on and image taken at the same time, by a camera placed adjacent to the sunphotometer. (b) A fit is found between the sunphotometer Almucantar measurements [blue] and camera gray scale values [orange].

an image of the deep-blue sky. The camera essentially senses sunlight scattered by air and some aerosols. The image is compared to a model image of the clear sky. The latter is obtained by running the image formation forward model (Eqs. 11,12) for $\beta_{\Lambda}^{\text{cloud}} = 0$, while $\beta_{\Lambda}^{\text{aerosol}}$ is set by the lidar and AERONET. We found this approach to be effective.

7 SIMULATIONS

Using simulations, we study some capabilities and limitations of ground-based tomography. The atmospheric domain is defined over a uniform $40 \times 40 \times 60$ grid. We simulate two 3D ellipsoid blobs with uniformly distributed extinction $\beta = 0.1$ [km^{-1}]. They are defined in [52]. One is 6km wide, 2.2km thick and centered at altitude 1.5km; the other is 4km wide, 1.1km thick and centered at altitude 3km (Fig. 2). Care is taken to simulate real network conditions. Hence, the simulated cameras are positioned in 3D following the exact locations and orientations of the true cameras of our network. A mask is applied where real sensor pixels are obscured by the sunshader or surrounding objects. This enables testing algorithms under the constraints of our particular real world setup (Fig. 2). In turn, such simulations can be used to improve the network, so as to optimize the quality of the results.

To allow for realistic simulations, we inject realistic noise to rendered images (Eq. 4). Towards this, we measured shot noise empirically in the lab, by using a camera from our network to capture a set of images of a static scene. Figure 10 is a scatter plot depicting the variance of the grayscale readout in each pixel, $\text{Var}(\hat{I}_x)$, as a function of the mean value per pixel, $E[\hat{I}_x]$. As expected from a Poissonian distribution, the noise variance is linearly dependent on the mean intensity value per pixel. Let a_1, a_2 be the slope and bias, respectively, of a line fitted to the noise variance in Fig. 10. We approximate the noise distribution by a normal distribution around the expectation value. Random noise is added to pixel x of a simulated image \hat{I} using

$$\text{noise}(x) = \mathcal{N}\left(0, \frac{1}{\sqrt{20}}[a_1 \cdot \hat{I}(x) + a_2]\right). \quad (21)$$




Figure 10: Measuring shot noise in our system: Noise variance vs. mean graylevel value per pixel. A linear function (in red) is fitted.

As described in the *Supplemental Material*, each measurement transmitted by any camera in the field relies on 20 raw frames. Therefore we scale down the variance of the normal distribution by $\sqrt{20}$.

Figure 11 shows different visualizations of reconstruction results. The algorithm seems to succeed in reconstructing the distribution both in shapes and in values. However, we also notice that the sunshaders create artifacts. All sunshaders block the view of the camera in the same direction: the direction of the Sun. This implies that cameras in close vicinity share the same directional “blind spot” and do not contribute to the reconstruction of regions occluded by the sunshader. In addition, far away cameras might not be able to compensate the recovery in the blind spots, being either occluded by camera surroundings or viewing at low angles through a thick atmosphere. This can be mitigated by deploying a larger network over a much wider span, enforcing smoothness or another prior on the scene.

8 EXPERIMENTAL DEMONSTRATION

Several experiments were conducted using the camera network. Up to 14 cameras were deployed, around Haifa Bay, Israel, allowing for a maximal baseline of 16 km between the farthest cameras, while keeping a denser distribution around the Technion. Usually, the cameras were set to take measurements every 5 minutes. At specific days, the cameras grabbed measurements every minute.

The domain we are trying to reconstruct is very large and not every part of it is seen from all viewpoints. To counter this, the vertical domain of β^{cloud} is limited to be between 0.5km to 4km. This allows the use of a vertically nonuniform grid. Between 0.5km to 4km, the vertical division of the grid is dense (every 0.1km), while outside this region, the grid is coarser. This speeds up and stabilizes considerably each gradient descent iteration. Second, space carving [42] is used to further constrain the cloud domain that we seek to reconstruct. Smoothness regularization [52] is applied during reconstruction.

Figure 12 shows reconstruction of a cloud field using data taken by our system. As seen from the reconstructed viewpoint (Fig. 12b) most clouds appear to be reconstructed. The algorithm fails to recover the extinction field in shaded

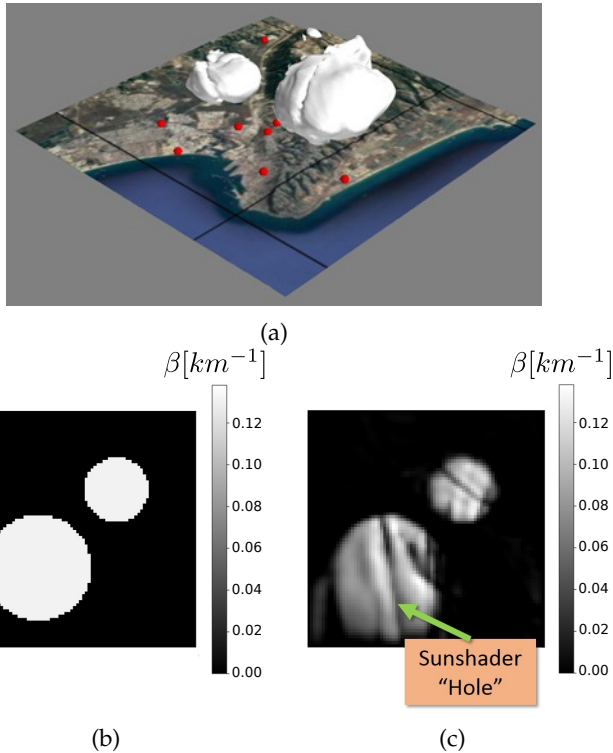


Figure 11: Reconstruction of simulated haze blobs (from Fig. 2): (a) Visualization of the reconstructed aerosol distribution by an iso-surface set to cross $\beta^{\text{aerosol}}(\mathbf{x}) = 0.1[\text{km}^{-1}]$. (b) Cross-section of the original aerosol distribution. (c) Cross section of the reconstructed aerosol distribution. The sunshaders cast “shadows” on the reconstruction distribution, where the algorithm failed to recover the underlying value.

areas (marked by a yellow ellipse). We believe that the reason is that this gray area under a cloud has similar radiance to the sky between clouds. Figure 13 shows iso-surface visualization of the reconstructed cloud field. The iso-surface marks the region where the cloud extinction is higher than 1km^{-1} . In order to reduce clutter, only the inner $6 \times 6 \times 10\text{km}$ domain is visualized.

Figure 14 shows a leave-one-out reconstruction experiment. The viewpoint in Fig. 14a was not used for the tomographic reconstruction (Eq. 14). The synthetic viewpoint, Fig. 14b, and its cross-section, Fig. 14c, demonstrate feasibility to reconstruct a cloud scene. More experimental results and the data for producing these results can be found in <https://webee.technion.ac.il/~yoav/research/aerosol-tomography.html>

9 CONCLUSION

We describe a county-scale (and earth-scalable) camera network. It samples the radiometric lightfield at ground-level observing the sky. It enables exploration of novel remote sensing approaches and ways to sense the atmosphere as it really is: a 3D scattering medium. Through careful calibration and comprehensive analysis, we transform standard RGB cameras into a sophisticated multi-view radiance sampling apparatus on a very large scale. We show how this

network can be used in combination with other instruments. This coordination is useful for calibration and initialization of tomographic reconstruction. Camera networks can supply real time data at high spatial and temporal resolution and complement earth observing systems.

Distributed spaceborne imaging is feasible by formation flying of orbiting nanosatellites. Hence, a mission (CloudCT, funded by ERC) is planned [63] to demonstrate scattering tomography from space, advance spacecraft formation technologies and provide data to reduce climate uncertainties. The relative simplicity of deploying a ground-based network compared to satellites and other remote sensing methods, and its ability to supply unprecedented amount of data on the atmosphere should make networks highly valuable for observing nature quantitatively.

ACKNOWLEDGMENTS

We are grateful to Danny Veikherman, Masada Tzabari, Omer Shubi, Birgit Heese, Julian Hofer, Holger Baars, Karsten Hanbuch, Ronny Engelmann and Ingolf Bernhardt for help in experiments. We thank Orr Cohen, Eyal Zohar, Ron Schneider, Lital Vazana, Kobi Lavi, Ran Zevitch, Johanan Erez, Ina Talmon, Dani Yagodin, Elena Bulkin and Ischack Lozon, for logistic assistance, and Albert Ansmann and Moti Segev for inspiring support. Special thanks to those who hosted our cameras: Israel Shipyards Ltd. (Tamir Neuman), University of Haifa, Israel Ports Company (Shlomo Briman), Gal, Israel and Nelida Aides, Einav Namer, Gideon Bleich, Marine Ltd. (Haim Horesh), Swift garage (Michael), M.M.S. Electric Ltd. (Hanan Brodsky), ‘Igud Arim’ Haifa, Ort Bialik High School (Orit Nitatski), Oranim Academic College, Madatech — Israel National Museum of Science, Technology and Space, and Mei Carmel Ltd. (Nur Eldan).

Yoav Schechner is the Mark and Diane Seiden Chair in Science in the Technion. He is a Landau Fellow - supported by the Taub Foundation. His work is conducted in the Ollendorff Minerva Center, Minvera ifs funded through the BMBF. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 810370: CloudCT), and the German-Israeli Science Foundation (GIF grant I-1262-401.10/2014) with support of the Norman and Helen Asher Fund. The authors acknowledge also funding from the European Union’s Horizon 2020 research and innovation program under grant agreement nos. 654109 (ACTRIS-2).

REFERENCES

- [1] U. Dayan, B. Ziv, T. Shoob, and Y. Enzel, “Suspended dust over southeastern Mediterranean and its relation to atmospheric circulations,” *International Journal of Climatology*, vol. 924, pp. 915–924, Jun. 2008.
- [2] Z. Peng, S. Yoo, D. Yu, D. Huang, P. Kalb, and J. Heiser, “3d cloud detection and tracking for solar forecast using multiple sky imagers,” in *Proceedings of the ACM 29th Annual Symposium on Applied Computing*, 2014, pp. 512–517.
- [3] F. A. Mejia, B. Kurtz, A. Levis, Í. de la Parra, and J. Kleissl, “Cloud tomography applied to sky images: A virtual testbed,” *Solar Energy*, vol. 176, pp. 287–300, 2018.
- [4] K. E. Trenberth, J. T. Fasullo, and J. Kiehl, “Earth’s global energy budget,” *Bulletin of the American Meteorological Society*, vol. 90, no. 3, pp. 311–324, 2009.

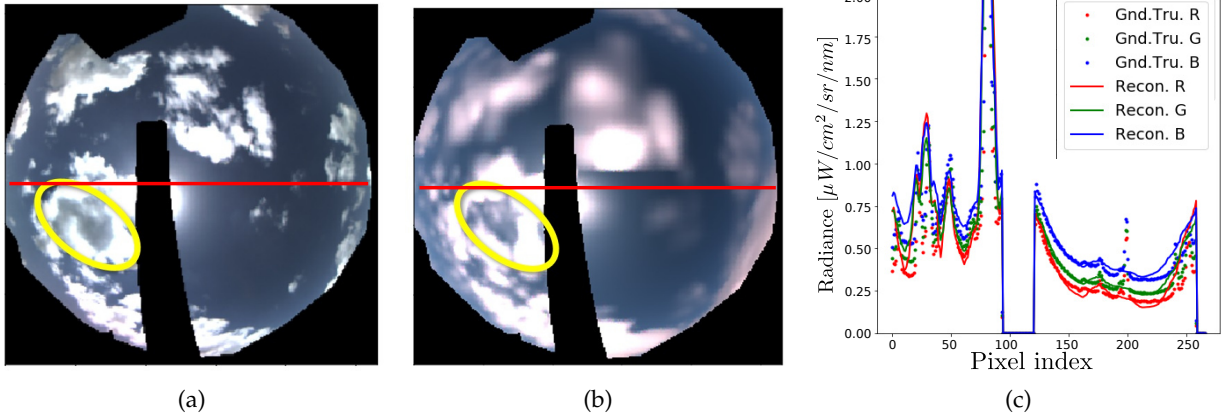


Figure 12: Reconstruction results of a real cloud field: (a) An image taken by the camera network after conversion to radiance measurements and masking of the sunshader, surrounding buildings and vegetation. (b) A reconstructed cloud field image, as synthesized from the same camera viewpoint. Most of the field appears to be reconstructed correctly, in coarse resolution. However, the algorithm fails to recover the extinction field in shaded areas (marked by a yellow ellipse). (c) A horizontal cross-section comparing the true radiance values (denoted “Gnd.Tru”) to the reconstructed values (denoted “Recon.”).




Figure 13: Visualization of cloud field reconstruction, by an iso-surface drawn at $\beta^{\text{cloud}} = 1$. Yellow spheres mark camera locations.

[5] A. Levis, Y. Y. Schechner, and A. B. Davis, “Multiple-scattering microphysics tomography,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[6] T. Nakajima and M. D. King, “Determination of the optical thickness and effective particle radius of clouds from reflected solar radiation measurements. Part I: Theory,” *Journal of the Atmospheric Sciences*, vol. 47, no. 15, pp. 1878–1893, 1990.

[7] Y. Y. Schechner, D. J. Diner, and J. V. Martonchik, “Spaceborne underwater imaging,” in *Proceedings of the IEEE International Conference on Computational Photography (ICCP)*, 2011, pp. 1–8.

[8] 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, 2009, pp. 267–293.

[9] G. P. Gobbi, “Studying Atmospheric Aerosol by Lidar,” in *Exploring the Atmosphere by Remote Sensing Techniques*, pp. 225–241, 2003.

[10] C. Weitkamp, *Lidar: Range-Resolved Optical Remote Sensing of the Atmosphere*. Springer Science & Business, 2006, vol. 102.

[11] D. Gürsoy, Y. P. Hong, K. He, K. Hujsak, S. Yoo, S. Chen, Y. Li, M. Ge, L. M. Miller, Y. S. Chu *et al.*, “Rapid alignment of nanotomography data using joint iterative reconstruction and reprojection,” *Scientific Reports*, vol. 7, no. 1, pp. 1–12, 2017.

[12] I. Ihrke and M. Magnor, “Image-based tomographic reconstruction of flames,” in *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Eurographics Association, 2004, pp. 365–373.

[13] D. Ren, C. Ophus, M. Chen, and L. Waller, “A multiple scattering algorithm for three dimensional phase contrast atomic electron tomography,” *Ultramicroscopy*, vol. 208, p. 112860, 2020.

[14] B. D. Trifonov, “Tomographic reconstruction of transparent objects,” Ph.D. dissertation, University of British Columbia, 2006.

[15] A. Veeraraghavan, A. V. Genkin, S. Vitaladevuni, L. Scheffer, S. Xu, H. Hess, R. Fetter, M. Cantoni, G. Knott, and D. Chklovskii, “Increasing depth resolution of electron microscopy of neural circuits using sparse tomographic reconstruction,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1767–1774.

[16] M. Alterman, Y. Y. Schechner, M. Vo, and S. G. Narasimhan, “Passive tomography of turbulence strength,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. Springer, pp. 47–60.

[17] A. Vainiger, Y. Y. Schechner, T. Treibitz, A. Avni, and D. S. Timor, “Optical wide-field tomography of sediment resuspension,” *Optics Express*, vol. 27, no. 12, pp. A766–A778, 2019.

[18] I. Ihrke, K. N. Kutulakos, H. P. Lensch, M. Magnor, and W. Heidrich, “Transparent and specular object reconstruction,” in *Computer Graphics Forum*, vol. 29, no. 8. Wiley Online Library, 2010, pp. 2400–2426.

[19] C. Ma, X. Lin, J. Suo, Q. Dai, and G. Wetzstein, “Transparent object reconstruction via coded transport of intensity,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 3238–3245.

[20] P. P. Srinivasan, R. Ng, and R. Ramamoorthi, “Light field blind motion deblurring,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 3958–3966.

[21] S. Lombardi and K. Nishino, “Reflectance and natural illumination from a single image,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012. Springer, pp. 582–595.

[22] R. Masuda, H. Iwabuchi, K. S. Schmidt, A. Damiani, and R. Kudo, “Retrieval of cloud optical thickness from sky-view camera images using a deep convolutional neural network based on three-dimensional radiative transfer,” *Remote Sensing*, vol. 11, no. 17, p. 1962, 2019.

[23] A. Levis, Y. Y. Schechner, A. Aides, and A. B. Davis, “Airborne three-dimensional cloud tomography,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 3379–3387.

[24] A. Geva, Y. Y. Schechner, Y. Chernyak, and R. Gupta, “X-ray computed tomography through scatter,” in *Proceedings of The European Conference on Computer Vision (ECCV)*, 2018, pp. 34–50.




Figure 14: Reconstruction results of a real cloud field, in a leave-one-out (cross validation) experiment: (a) An image taken by the camera network after conversion to radiance measurements and masking of the sunshading, surrounding buildings and vegetation. This viewpoint was *not* used in the reconstruction process. (b) Image projection of the reconstructed cloud field, at the left-out viewpoint. (c) Horizontal cross-section (marked by red line in viewpoint images) comparing the true radiance values (denoted “Gnd.Tru.”) to reconstructed values (denoted “Recon.”).

- [25] A. Levis, Y. Y. Schechner, and R. Talmon, “Statistical tomography of microscopic life,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 6411–6420.
- [26] I. Gkioulekas, A. Levin, and T. Zickler, “An evaluation of computational imaging techniques for heterogeneous inverse scattering,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2016, pp. 685–701.
- [27] N. Joshi and M. Cohen, “Seeing Mt. Rainier: Lucky imaging for multi-image denoising, sharpening, and haze removal,” in *Proceedings of the IEEE International Conference on Computational Photography (ICCP)* 2010, pp. 1–8.
- [28] P. Khungurn, D. Schroeder, S. Zhao, K. Bala, and S. Marschner, “Matching real fabrics with micro-appearance models.” *ACM Transactions on Graphics (TOG)*, vol. 35, no. 1, pp. 1–26, 2015.
- [29] H.-Y. Liu, E. Jonas, L. Tian, J. Zhong, B. Recht, and L. Waller, “3D imaging in volumetric scattering media using phase-space measurements,” *Optics Express*, vol. 23, no. 11, pp. 14 461–14 471, 2015.
- [30] G. Satat, B. Heshmat, D. Raviv, and R. Raskar, “All photons imaging through volumetric scattering,” *Scientific Reports*, vol. 6, pp. 1–8, 2016.
- [31] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, B. Curless, S. M. Seitz, and R. Szeliski, “Building rome in a day,” *Communications of the ACM*, vol. 54, no. 10, pp. 105–112, 2011.
- [32] D. Forsyth, J. L. Mundy, A. Zisserman, C. Coelho, A. Heller, and C. Rothwell, “Invariant descriptors for 3d object recognition and pose,” *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 10, pp. 971–991, 1991.
- [33] F. Langguth, K. Sunkavalli, S. Hadap, and M. Goesele, “Shading-aware multi-view stereo,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016, pp. 469–485.
- [34] T. Owens, K. Saenko, A. Chakrabarti, Y. Xiong, T. Zickler, and T. Darrell, “Learning object color models from multi-view constraints,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 169–176.
- [35] J. V. Martonchik, D. J. Diner, R. A. Kahn, T. P. Ackerman, M. M. Verstraete, B. Pinty, and H. R. Gordon, “Techniques for the retrieval of aerosol properties over land and ocean using multi-angle imaging,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, pp. 1212–1227, 1998.
- [36] D. Diner, F. Xu, M. Garay, J. Martonchik, B. Rheingans, S. Geier, A. Davis, B. Hancock, V. Jovanovic, M. Bull *et al.*, “The Airborne Multiangle SpectroPolarimetric Imager (AirMSPI): a new tool for aerosol and cloud remote sensing,” *Atmospheric Measurement Techniques*, vol. 6, no. 8, pp. 2007–2025, 2013.
- [37] D. J. Diner, S. W. Boland, M. Brauer, C. Bruegge, K. A. Burke, R. Chipman, L. Di Girolamo, M. J. Garay, S. Hasheminassab, E. Hyer *et al.*, “Advances in multiangle satellite remote sensing of speciated airborne particulate matter and association with adverse health effects: from MISR to MAIA,” *Journal of Applied Remote Sensing*, vol. 12, no. 4, p. 042603, 2018.
- [38] F. J. Olmo, A. Cazorla, L. Alados-Arboledas, M. A. López-Alvarez, J. Hernández-Andrés, and J. Romero, “Retrieval of the optical depth using an all-sky CCD camera,” *Applied Optics*, vol. 47, no. 34, pp. H182–H189, 2008.
- [39] Y. Y. Schechner and S. K. Nayar, “Generalized mosaicing: High dynamic range in a wide field of view,” *International Journal of Computer Vision*, vol. 53, no. 3, pp. 245–267, 2003.
- [40] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan, “Estimating natural illumination from a single outdoor image,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2009, pp. 183–190.
- [41] —, “Webcam clip art: Appearance and illuminant transfer from time-lapse sequences,” *ACM Transactions on Graphics (TOG)*, vol. 28, no. 5, pp. 1–10, 2009.
- [42] D. Veikherman, A. Aides, Y. Y. Schechner, and A. Levis, “Clouds in The Cloud,” in *Proceedings of the Asian Conference on Computer Vision (ACCV)*. Springer, 2014, pp. 659–674.
- [43] J. T. Kider Jr, D. Knowlton, J. Newlin, Y. K. Li, and D. P. Greenberg, “A framework for the experimental comparison of solar and sky-dome illumination,” *ACM Transactions on Graphics (TOG)*, vol. 33, no. 6, pp. 1–12, 2014.
- [44] E. Bruneton, “A qualitative and quantitative evaluation of 8 clear sky models,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 12, pp. 2641–2655, 2016.
- [45] D. Guimera, D. Gutierrez, and A. Jarabo, “A physically-based spatio-temporal sky model,” in *Proceedings of the XXVIII Spanish Computer Graphics Conference*. Eurographics Association, 2018, pp. 29–37.
- [46] D. M. Romps and R. Öktem, “Observing clouds in 4d with multiview stereophotogrammetry,” *Bulletin of the American Meteorological Society*, vol. 99, no. 12, pp. 2575–2586, 2018.
- [47] F. Koreban and Y. Y. Schechner, “Geometry by deflating,” in *Proceedings of the IEEE International Conference on Computational Photography (ICCP)*, 2009, pp. 1–8.
- [48] E. Sassoon, Y. Y. Schechner, and T. Treibitz, “Flare in interference-based hyperspectral cameras,” in *Proceedings of the IEEE International Conference on Computer Vision (CVPR)*, 2019, pp. 10 174–10 182.
- [49] ASTM, “2000 ASTM Standard Extraterrestrial Spectrum Reference E-490-00,” <https://www.nrel.gov/grid/solar-resource/spectra-astm-e490.html>, 2000, accessed: 2019-09-30.
- [50] K. F. Evans and W. J. Wiscombe, “Improvements to the SHDOM radiative transfer modeling package,” *Thirteenth ARM Science Team Meeting Proceedings*, no. 1d1, pp. 1–6, 2003.
- [51] M. K. Yau and R. Rogers, *A Short Course in Cloud Physics*. Elsevier, 1996.
- [52] A. Aides, Y. Y. Schechner, V. Holodovsky, M. J. Garay, and A. B.

Davis, "Multi sky-view 3D aerosol distribution recovery," *Optics Express*, vol. 21, no. 22, pp. 25 820–25 833, 2013.

- [53] V. Holodovsky, Y. Y. Schechner, A. Levin, A. Levis, and A. Aides, "In-situ multi-view multi-scattering stochastic tomography," in *Proceedings of the IEEE International Conference on Computational Photography (ICCP)*, 2016, pp. 1–12.
- [54] 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," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 144, pp. 68–85, 2014.
- [55] S. Wright, M. Schweiger, and S. Arridge, "Reconstruction in optical tomography using the PN approximations," *Measurement Science and Technology*, vol. 18, no. 1, p. 79, 2006.
- [56] J. Gregson, M. Krimerman, M. B. Hullin, and W. Heidrich, "Stochastic tomography and its applications in 3D imaging of mixing fluids," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, pp. 52:1—52:10, Jul. 2012.
- [57] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin, "Inverse volume rendering with material dictionaries," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, p. 162, 2013.
- [58] 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," *Journal of Geophysical Research: Atmospheres*, vol. 113, no. D4, 2008.
- [59] R. Marchand and T. Ackerman, "Evaluation of radiometric measurements from the NASA Multiangle Imaging SpectroRadiometer (MISR): Two-and three-dimensional radiative transfer modeling of an inhomogeneous stratocumulus cloud deck," *Journal of Geophysical Research: Atmospheres*, vol. 109, no. D18, 2004.
- [60] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proceedings of COMPSTAT*. Springer, 2010, pp. 177–186.
- [61] K. Miyamoto, "Fish eye lens," *Journal of the Optical Society of America A*, vol. 54, no. 8, pp. 1060–1061, 1964.
- [62] A. Setzer, E. Vermote, J. Reagan, Y. Kaufman, T. Nakajima, F. Lacenu, I. Jankowiak, and A. Smirnov, "AERONET-A federated instrument network and data archive for aerosol characterization," *Remote Sensing of Environment*, vol. 66, no. 1, pp. 1–16, Oct. 1998.
- [63] K. Schilling, Y. Y. Schechner, and I. Koren, "CloudCT – computed tomography of clouds by a small satellite formation," in *Proceedings of the 12th IAA symposium on Small Satellites for Earth Observation*. IAA, 2019.



Amit Aides received his Ph.D. degree in Electrical Engineering from the Technion, Israel Institute of Technology, in 2018. He is currently working at Google, Haifa. Before joining Google, he was a deep-learning researcher at IBM Research Labs. His research interests include medical imaging, computational photography, and few-shot learning.



Aviad Levis received his BSc (2013) and Ph.D. (2020) degrees in Electrical Engineering from Ben-Gurion University and Technion - Israel Institute of Technology, respectively. He is currently a post-doctoral researcher at California Institute of Technology at the Computing and Mathematics department. His research interests include remote-sensing and physics-based imaging of natural phenomena.



Vadim Holodovsky received his M.Sc. degree in Electrical Engineering, Technion, Israel, in 2016. He is currently a research team member in the computational-photography lab, Technion, Israel. His research interests include computer-vision algorithms for 3D retrievals and remote-sensing. They also include experiment designs including optics, cameras, and electronic hardware.



Yoav Y. Schechner is a graduate of the Technion - Israel Institute of Technology: BA (Physics 1990), MSc (Physics 1994), PhD (EE 2000). Afterwards, he was a research scientist at Columbia University (CS). Since 2002, he is a faculty member in the Technion's Viterbi Faculty of Electrical Engineering. He is the Diane and Mark Seiden Chair in Science. He is a principal investigator and coordinator of the CloudCT project, funded by ERC. In 2010 and 2011 he was a visiting scientist at Caltech and NASA's

Jet Propulsion Laboratory (JPL). He won the Best Student Paper Award at CVPR in 2017, the Best Paper Awards at ICCP in 2013 and 2018. He is the recipient of the Ray and Miriam Klein Research Award, the Henry Taub Prize for Academic Excellence, the Otto Schwarz Foundation Excellence Award and the Landau Fellowship. His research interests involve outdoor phenomena and all aspects of imaging.



Dietrich Althausen studied physics from 1977 until 1982 at the University of Leipzig. 1992 he got his PhD and since then he is employed at the Leibniz Institute for Tropospheric Research as a senior scientist. Within these 28 years, the focus of his work is the development and application of multi-wavelengths polarization Raman lidars for the vertically resolved observations of particles and clouds in the atmosphere. In 2002 he started the development of small and compact aerosol lidar systems, the "Polly" family. Meanwhile these systems are applied worldwide in the frame of PollyNET. An important result of this research is the observation and characterization of aerosol layers that are transported intercontinentally at different heights - in the troposphere and even in the stratosphere.



Adi Vainiger received her BSc (2012) and MSc (2019) degrees in Electrical Engineering from the Technion - Israel Institute of Technology. Her research interests include computer-vision algorithms for 3D retrievals, remote-sensing, and computational photography.