Atmospheric
Measurement
Techniques

Retrieving 3D distributions of atmospheric particles using Atmospheric Tomography with 3D Radiative Transfer – Part 1: Model description and Jacobian calculation

Jesse Loveridge[1], Aviad Levis[2], Larry Di Girolamo[1], Vadim Holodovsky[3], Linda Forster[4], Anthony B. Davis[4], and Yoav Y. Schechner[3]

[1]Department of Atmospheric Sciences, University of Illinois, Urbana, IL 61801, USA
[2]Computer and Mathematical Sciences Department, California Institute of Technology, Pasadena, CA 91125, USA
[3]Viterbi Faculty of Electrical and Computer Engineering, Technion – Israel Institute of Technology, Haifa 3200003, Israel
[4]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

Correspondence: Jesse Loveridge (jesserl2@illinois.edu)

Abstract. Our global understanding of clouds and aerosols relies on the remote sensing of their optical, microphysical, and macrophysical properties using, in part, scattered solar radiation. These retrievals assume that clouds and aerosols form plane-parallel, homogeneous layers and utilize 1D radiative transfer (RT) models, limiting the detail that can be retrieved about the 3D variability in cloud and aerosol fields and inducing biases in the retrieved properties for highly heterogeneous structures such as cumulus clouds and smoke plumes. To overcome these limitations, we introduce and validate an algorithm for retrieving the 3D optical or microphysical properties of atmospheric particles using multi-angle, multi-pixel radiances and a 3D RT model. The retrieval software, which we have made publicly available, is called Atmospheric Tomography with 3D Radiative Transfer (AT3D). It uses an iterative, local optimization technique to solve a generalized least squares problem and thereby find a best-fitting atmospheric state. The iterative retrieval uses a fast, approximate Jacobian calculation, which we have extended from Levis et al. (2020) to accommodate open and periodic horizontal boundary conditions (BCs) and an improved treatment of non-black surfaces.

We validated the accuracy of the approximate Jacobian calculation for derivatives with respect to both the 3D volume extinction coefficient and the parameters controlling the open horizontal boundary conditions across media with a range of optical depths and single-scattering properties and find

that it is highly accurate for a majority of cloud and aerosol fields over oceanic surfaces. Relative root mean square errors in the approximate Jacobian for a 3D volume extinction coefficient in media with cloud-like single-scattering properties increase from 2 % to 12 % as the maximum optical depths (MODs) of the medium increase from 0.2 to 100.0 over surfaces with Lambertian albedos < 0.2. Over surfaces with albedos of 0.7, these errors increase to 20 %. Errors in the approximate Jacobian for the optimization of open horizontal boundary conditions exceed 50 %, unless the plane-parallel media providing the boundary conditions are optically very thin (~ 0.1).

We use the theory of linear inverse RT to provide insight into the physical processes that control the cloud tomography problem and identify its limitations, supported by numerical experiments. We show that the Jacobian matrix becomes increasing ill-posed as the optical size of the medium increases and the forward-scattering peak of the phase function decreases. This suggests that tomographic retrievals of clouds will become increasingly difficult as clouds become optically thicker. Retrievals of asymptotically thick clouds will likely require other sources of information to be successful.

In Loveridge et al. (2023a; hereafter Part 2), we examine how the accuracy of the retrieved 3D volume extinction coefficient varies as the optical size of the target medium increases using synthetic data. We do this to explore how the increasing error in the approximate Jacobian and the increas-

ingly ill-posed nature of the inversion in the optically thick limit affect the retrieval. We also assess the accuracy of retrieved optical depths and compare them to retrievals using 1D radiative transfer.

1   Introduction

Cloud and aerosol properties retrieved from the inversion of remote sensing measurements (Stephens and Kummerow, 2007; Dubovik et al., 2011) are a critical source of information for understanding and testing the closure of the Earth's radiation budget (Raschke et al., 2005; McFarlane et al., 2016; Zhou et al., 2016), validating dynamical atmospheric models of varying complexity (Hack et al., 2006; Endo et al., 2015; Bodas-Salcedo et al., 2016), and developing parameterizations in large-scale models (Hill et al., 2012; Xie and Zhang, 2015). Cloud radiative feedbacks and aerosol–cloud interactions (ACIs; Bellouin et al., 2020), particularly in cumuliform clouds (Sherwood et al., 2014; Vial et al., 2018), are key sources of uncertainty in projections of future climate (Sherwood et al., 2020) and the forecasting of weather (Van Weverberg et al., 2018) and solar energy (Jimenez et al., 2016).

As dynamical modeling of the atmosphere and climate becomes more and more complex, there is a greater demand for high-quality observations to constrain the uncertain processes within the models and inform model development (Morrison et al., 2020). New observational techniques are required that can provide robust statistics of small-scale, spatially resolved cloud and aerosol microphysical parameters so that their controlling processes can be constrained in both high- and low-resolution modeling. We describe a novel remote sensing retrieval technique with the potential to meet these needs by providing 3D instantaneous snapshots of volumetric properties of the atmosphere, thus making complete use of the resolution of the sensors.

Scattered solar radiation is one of the best candidates for providing high-resolution constraints on aerosol and cloud microphysics, due to its well-documented sensitivity to the properties of particles in those size ranges (Dubovik et al., 2002; King and Vaughan, 2012; Dzambo et al., 2021; Ewald et al., 2021) and our ability to design narrowband sensors that can cost-effectively reach a high spatial resolution (tens of meters) from space. Typical cloud and aerosol remote sensing retrieval algorithms do not utilize realistic 3D radiative transfer (RT) models to interpret measured scattered solar radiation (e.g., Dubovik et al., 2011; Grosvenor et al., 2018). Instead, they make use of two key simplifying assumptions when interpreting radiance measurements. First, they assume that the media (e.g., clouds) form horizontally homogeneous, plane-parallel layers within the field of view of each radiance measurement. Second, they assume that there is no radiative interaction between the regions within the field of view of each radiance measurement in what is known as the independent pixel approximation (IPA).

These assumptions dramatically reduce the computational complexity of the retrieval process, but they also compromise retrieval accuracy (Marshak et al., 2006; Zhang et al., 2012; Kato and Marshak, 2009), leading to errors in, for example, cloud optical depth that can have domain biases of $\sim 35\%$ in cumuliform clouds (Seethala, 2012). These errors contribute to the large inter-instrument inconsistencies in retrievals (Di Girolamo et al., 2010; Lebsock and Su, 2014; Ahn et al., 2018; Fu et al., 2019; Painemal et al., 2021). These errors arise because clouds and aerosols can have strong horizontal gradients in their physical and optical properties (Marshak et al., 1997; Gerber et al., 2001; Zhao and Di Girolamo, 2007; Kahn et al., 2007), which breaks the assumption of the IPA and necessitates the use of 3D RT to accurately model the transport of radiation (Davies, 1978; Cahalan et al., 2005).

These errors distort into our climate records and also affect retrievals made using a combination of passive and active sensors (Saito et al., 2019). The retrieval assumptions also limit the amount of detail about the cloud and aerosol fields that can be retrieved using passive sensing as, once the IPA is imposed, there is no way to identify the vertical geometric variability in cloud microphysics within a layer without active instruments, which in the case of cloud radar requires strong assumptions about the shape of the particle size distribution. An algorithmic advance in operational retrievals is required to better extract the information about 3D variability in cloud and aerosol microphysics contained within high-resolution solar radiance measurements to provide the novel observations required for advancing cloud and aerosol science.

Many algorithms have been demonstrated that utilize 3D RT to improve remote sensing retrievals of cloud properties but have been limited to using radiometric information from a single mono-angle imager at a time (Marshak et al., 1998a; Marchand and Ackerman, 2004; Zinner et al., 2006; Cornet and Davies, 2008) or a single zenith radiance measurement (Fielding et al., 2014). This restriction limits the amount of information obtainable from the radiance field, so that only column integrals or horizontal averages of cloud properties can be inferred, rather than their three-dimensional variability. As a result, several of these methods have relied on external sources of information, such as scanning radar or in situ data (Marchand and Ackerman, 2004; Fielding et al., 2014), or strong assumptions that are not generally applicable (Marshak et al., 1998a; Zinner et al., 2006; Cornet and Davies, 2008). It is clear that a new source of information is required to relax the strong assumptions required by these retrieval algorithms to enable the retrieval of the 3D spatial structure of cloud and aerosol microphysics.

Multi-angle imagery is a promising source of information to constrain the 3D structure of the atmosphere. Inverse problems, where multi-angle boundary measurements are used to

infer internal structure, are commonly known as tomography. In atmospheric science, tomographic methods have been applied to retrieve cloud properties using non-scattering, multi-angle microwave emission (Huang et al., 2008, 2010); water vapor, using microwave attenuation (Jiang et al., 2022); and aerosol, using scattering measurements (Garay et al., 2016; Zawada et al., 2017, 2018). Multi-angle imaging has been utilized to systematically retrieve the geometric properties of clouds (Muller et al., 2002, 2007) and aerosols (Kahn et al., 2007) using stereoscopic methods.

In the field of medical imaging, multiple detectors are routinely utilized to retrieve spatially varying optical properties of the human body from multiply scattered near-infrared radiation in a process known as diffuse optical tomography (Arridge and Schotland, 2009; Bal, 2009). Taking this as inspiration, a similar tomographic approach has been proposed and formalized for the retrieval of spatially varying atmospheric constituents from multi-angle imagery of multiply scattered solar radiation in the atmospheric context (Martin et al., 2014). These methods must make use of 3D RT models, as 1D RT models are unable to reproduce the angular variations in the observed radiation field (Di Girolamo et al., 2010). Tomographic methods have the potential to provide remote sensing retrievals of volumetric cloud and aerosol properties, such as the 3D distribution of volume extinction coefficient, and possibly even microphysical quantities.

Tomography problems are commonly solved using iterative, physics-based optimization procedures similar to state-of-the-art methods in aerosol remote sensing (Xu et al., 2019; Gao et al., 2021). They can also be solved using statistical methods (Zhang and Zhang, 2019; Ronen et al., 2022) or heuristic methods, which have been explored recently in the atmospheric science context (Alexandrov et al., 2021). Tomographic methods in diffuse optical tomography typically make use of computationally efficient forward-adjoint methods to linearize a 3D RT model and calculate the cost function gradients (Arridge and Schotland, 2009). Similar methods have been employed in plane-parallel retrievals of aerosol properties (Hasekamp and Landgraf, 2005). So far, the tomographic retrievals utilizing forward-adjoint methods have only been demonstrated in atmospheric sciences in 2D (Martin and Hasekamp, 2018).

Similar tools have been developed elsewhere. The Monte Carlo 3D RT equation (RTE) solver McArtim (Deutschmann et al., 2011) is specialized for radiance derivative calculations using forward-adjoint methods but uses a backward Monte Carlo technique, similar to other work (Loeub et al., 2020), with importance sampling, which will not scale well to the multi-angle imagery required for tomography. Other Monte Carlo techniques for derivative calculations typically use forward methods with path recycling (Langmore et al., 2013; Yao et al., 2018; Czerninski and Schechner, 2021); however, most implementations of such models lack the variance reduction methods required for the efficient modeling

of radiances with sharply peaked phase functions (Buras and Mayer, 2011; Wang et al., 2017) and have not been benchmarked on atmospheric problems.

Of the available 3D RTE solvers that are benchmarked on atmospheric scattering problems (Cahalan et al., 2005), the deterministic (i.e., explicit) spherical harmonics discrete ordinates method (SHDOM; Evans, 1998) is the most computationally efficient for tomography. This is due to the need to simulate many radiometric quantities. SHDOM is almost 2 orders of magnitude more computationally efficient than Monte Carlo on CPU for multi-angle imagery (Pincus and Evans, 2009). Monte Carlo solvers specialized for 3D atmospheric scattering problems have been slow to adopt a GPU-based computation, which is anticipated to give a reduction in the wall time of between 1 and 2 orders of magnitude (Efremenko et al., 2014; Ramon et al., 2019; Wang et al., 2021; Lee et al., 2022), thereby making Monte Carlo competitive against SHDOM in the future.

At the time of writing, there is no publicly available adjoint to a deterministic (i.e., explicit) 3D RTE solver appropriate to the atmospheric context like SHDOM. A forward-adjoint linearization of the SHDOM method has been developed (Doicu and Efremenko, 2019), and an SHDOM solver has been extended so that general adjoints appropriate for tomography can be computed (Doicu et al., 2022b). This forward-adjoint linearization, following the theory of Martin et al. (2014), is also technically able to compute the Jacobian matrix of partial derivatives of the forward model. However, this is computationally inefficient for tomography problems where the number of measurements is very large (Martin et al., 2014). Unfortunately, the software implementing the forward-adjoint linearization of SHDOM is not publicly available, which means we are unable to build upon these advances. Fortunately, a computationally efficient approximation to the adjoint of SHDOM has been developed and used to demonstrate the success of fully 3D retrievals of the volume extinction coefficient of clouds using multi-angle, mono-spectral imagery in a first for atmospheric remote sensing (Levis et al., 2015). This method of approximate linearization has been extended to utilize multi-spectral (Levis et al., 2017) and polarized (Levis et al., 2020) observations. This approximate linearization opens up computationally efficient access to an approximate Jacobian matrix. However, so far, only gradient-based optimization methods have been used, and it is unclear how robust or efficient the approximate linearization will be when combined with optimization methods which make direct use of the Jacobian matrix. Interestingly, the forward-adjoint method of cloud tomography using SHDOM suffered from slow convergence, and the authors only found success in their synthetic tomographic retrievals when utilizing the approximate linearization of Levis et al. (2020) and Doicu et al. (2022a) in combination with their adjoint method (Doicu et al., 2022b).

The method of Levis et al. (2020) is the most mature and successful remote sensing retrieval using solar radiances and

3D RT available in the atmospheric sciences. The method is still restricted in that its implementation is limited to isolated 3D domains and Lambertian surfaces, and the approximate linearization is a poor approximation for non-black surfaces. Despite these limitations, the method's maturity makes it the ideal starting point for developing retrievals of 3D volumetric microphysical parameters at similar resolutions to those used in large-eddy simulations. The future spaceborne CloudCT mission (Schilling et al., 2019) will provide the required simultaneous multi-angle imagery for tomographic retrievals. Existing airborne instruments such as AirMSPI (Airborne Multi-angle Spectro Polarimetric Imager; Diner et al., 2013) and AirHARP (Airborne Hyper-Angular Rainbow Polarimeter; McBride et al., 2020) and the space-borne MISR (Multi-angle Imaging SpectroRadiometer) and MAIA (Multi-Angle Imager for Aerosols) also have the potential for tomographic retrievals, though they must additionally deal with the effects of cloud evolution (Ronen et al., 2021), as they do not acquire their observations simultaneously. The availability of these measurements makes the continued development of tomographic algorithms especially timely. Retrievals of this sort have the potential to provide the robust statistics of small-scale cloud and aerosol properties required for constraining cloud processes (Morrison et al., 2020), especially when extended to include information from other instruments such as cloud radar.

In this two-part series of papers, we present and validate an extension to the retrieval framework of Levis et al. (2020), which we have implemented and made publicly available in the software package Atmospheric Tomography with 3D Radiative Transfer (AT3D; Loveridge et al., 2022). This paper, which is Part 1, is devoted to the description of the retrieval methodology and the underlying theory of the retrieval, along with supporting numerical evidence. Part 2 of this study is devoted to tomographic retrievals on synthetic data to validate the method.

In Sect. 2, we describe the retrieval software AT3D, which is quite general in that it is designed for retrieving the 3D microphysical properties of external mixtures of atmospheric particles using multi-angle, multi-pixel, and possibly multi-spectral polarized radiances. In Sect. 3, we describe extensions to the method of Levis et al. (2020) to include an improved treatment of non-black surfaces and the retrieval of a plane-parallel medium in which the 3D domain is embedded, thereby improving the realism of the method. The Appendix documents the discrete implementation of the algorithm and model verification.

Despite the successful demonstrations of the tomographic retrieval (Levis et al., 2015, 2017; Martin and Hasekamp, 2018; Levis et al., 2020; Doicu et al., 2022a, b), it is still unclear how the effectiveness of tomographic techniques will vary with scattering regime. Previous studies have shown that success is not uniform, with poorer performance in optically thick clouds (Levis et al., 2015). It is not clear whether this is a result of a limitation in the approximate linariza-

tion method or a physical limitation. In Sect. 4, we present the theory of linear inverse transport problems and use it to explain the limitations of tomography in general, to provide insight into the physical processes that control the cloud tomography problem. This theory is drawn from both the wider literature on inverse problems (Bal and Jollivet, 2008; Bal, 2009; Chen et al., 2018; Zhao and Zhong, 2019) and the relevant literature in the atmospheric sciences that have studied the loss of information about spatial detail of cloud properties in multi-pixel radiances due to multiple scattering (Marshak et al., 1995, 1998b; Davis et al., 1997; Forster et al., 2020).

In Sect. 5, we perform a detailed quantitative validation of the approximate linearization to SHDOM that is utilized in AT3D, following Levis et al. (2020). Despite the success of the method in several test cases in Levis et al. (2015, 2017, 2020), no validation of the approximation itself has yet been performed, which has made it as yet unclear how the method will generalize to the wider variety of scattering regimes present in the cloudy atmosphere. We present an alternative derivation of the approximation that places it in the context of the forward-adjoint formalism developed by Martin et al. (2014). We can then explain the success of the approximate method using the theory presented in Sect. 4. In Sect. 6, we briefly quantitatively contrast stratiform and cumulus cloud geometries, in terms of the well-posedness of the tomography problem from a linear perspective, using the approximate Jacobian. We summarize our results in Sect. 7. In Part 2 of this study, we explore how these issues affect the fully nonlinear retrieval problem and demonstrate the effectiveness of the method described here.

2 Atmospheric Tomography with 3D Radiative Transfer (AT3D)

Atmospheric Tomography with 3D Radiative Transfer (AT3D) is a software package designed to perform tomographic retrievals of atmospheric properties. It poses the inverse problem as a nonlinear, generalized least squares problem that is solved using iterative local optimization techniques. The solution procedure is physics based and uses the 3D RT model SHDOM (Evans, 1998) as its forward model to connect retrieved quantities to measured radiance. SHDOM is an explicit solver of the polarized 3D RTE that is well established in atmospheric science. During the intercomparison of 3D radiative transfer codes (I3RC; Cahalan et al., 2005), SHDOM was well within the consensus results. This is also true of intercomparisons of polarized RT (Emde et al., 2015, 2018).

In brief, SHDOM solves the integral form of the monochromatic vector RTE on a Cartesian grid using a fixed-point iteration scheme for collimated solar or thermal emission sources of radiation. A spherical harmonic expansion representation of the radiation field is used for computing the SOURCE function of the RTE, while a discrete ordinate

representation is used for the streaming of radiation. At each iteration, the SOURCE function is transformed to discrete ordinates, new radiances are computed at each grid point using a short characteristic scheme, and a new spherical harmonic representation of the SOURCE function is computed. An adaptive spatial grid is employed so that grid cells with a variation in the SOURCE function larger than a threshold are split in half, generating new grid points. The number of spherical harmonics kept at each grid point is also adaptively truncated. SHDOM uses delta-M scaling (Wiscombe, 1977) and the truncated multiple-scattering (TMS) approximation (Nakajima and Tanaka, 1988) to treat problems involving highly anisotropic scattering. The truncation fraction used in the delta-M scaling of the optical properties is set by the angular resolution of the SHDOM solver and the Legendre expansion of the phase function.

AT3D builds on the software implemented by Levis et al. (2020), which itself builds upon the work of Evans (1998) in the publicly available Fortran implementation of SHDOM. AT3D is also a Python wrapper for the SHDOM RTE solver developed using the F2PY (Fortran to Python interface generator) tool (Peterson, 2009); this enables easy interfacing with external optimization libraries from SciPy (Virtanen et al., 2020). The use of Python also enables interactivity, even in high-performance-computing (HPC) environments, which accelerates data exploration and code prototyping. The key features of the SHDOM software are preserved in AT3D, with the only notable exception being that AT3D does not yet implement the message-passing interface (MPI)-based parallelization of SHDOM, so it is not yet able to efficiently utilize HPC resources to solve large-scale forward or inverse problems.

AT3D's strength is as a provider of a physics-based, and therefore flexible, method for solving the inverse problem of atmospheric tomography. It is therefore perfectly suited for performing sensitivity tests to changes in the measuring instrument's configuration (e.g., number of view angles and sensor resolution). AT3D supports the retrieval of multiple external mixtures of 3D distributions of scattering particles with solar and thermal sources, using arbitrary combinations of possibly polarized, multi-wavelength monochromatic radiances. Each unknown can be retrieved on a 3D grid or on a user-specified simplified spatial basis (e.g., column averages). AT3D does not yet support non-simultaneous measurements and the corresponding retrieval of a time-varying cloud (Ronen et al., 2021). Flexibility with the configuration of any retrieval problem is supported using object-oriented and functional programming in the Python wrapper. Currently, AT3D includes just the Rayleigh and Mie scattering particle models for homogeneous spheres distributed with SHDOM (https://nit.coloradolinux.com/shdom.html, last access: 20 March 2023; Evans, 1998). The parameterization of the size distribution and accompanying selection of unknowns (e.g., droplet number concentration or liquid water content) for retrievals is flexible.

We note that the software is far from a black-box tool and operates more as a library of high-level objects and functions that can be combined in short Python scripts according to user specifications. This level of flexibility is good for a research tool that is under active development, though it can lead to a steeper learning curve than in a well-defined executable with a fixed set of options common that is in other RT software packages. The code is well documented, and several tutorials are included with the code to mitigate this. Users or potential developers are welcome to contact the corresponding author to discuss their potential use case.

Just like the SHDOM software, AT3D is not a complete RT package and does not include detailed spectroscopic or particle scattering data that can be found elsewhere (Emde et al., 2016; Gordon et al., 2022; Saito et al., 2021). Interfacing with these packages is relatively simple, as the data are represented in AT3D using the xarray package (Hoyer and Hamman, 2017), which supports a variety of file formats such as NetCDF.

AT3D supports all surface bidirectional reflectance distribution functions (BRDFs) available in SHDOM but does not yet include linearization with respect to the parameters describing surface BRDFs. AT3D supports the inversion of optical or microphysical properties with solar, thermal, or combined sources. This is a helpful extension over Levis et al. (2020) for the far shortwave infrared (SWIR) and infrared (IR). However, linearization with respect to atmospheric temperature is not yet supported. These aspects are under development. The radiance calculations in AT3D have been generalized from SHDOM to support more realistic sensor geometries and sensor spatial response functions. However, as noted above, observables are currently monochromatic, and some small extension to the software would be required to accommodate observables requiring multiple monochromatic RTE solutions during inversion. The representation of phase functions within the SHDOM solver has been modified so that the SHDOM model is differentiable (Appendix B). The package also includes several useful tools, including a stochastic generator for making synthetic clouds (described in Part 2), a space-carving algorithm (Lee et al., 2018) for performing volume cloud masking for retrieval initialization (Sect. 5.3), and several basic regularization schemes.

While many practical extensions have been made to the retrieval software since Levis et al. (2020), including a comprehensive verification (Appendix C and E), which is critical to establish the veracity of the scientific results (Kanewala and Bieman, 2014), the novel elements of the retrieval software are the extensions of the linearization to non-homogeneous surfaces and open boundary conditions (BCs), which are documented in the following section.

3 Retrieval methodology

We now present an overview of the iterative retrieval process. There are many mathematical terms introduced in this section. A glossary is provided in Appendix A for reference, and a flowchart of the retrieval process is shown in Fig. 1. The solution of the tomography problem involves the selection of a state vector (a) that parameterizes a discrete representation of the atmospheric optical or physical properties and best fits the available measurements (y) and any prior knowledge of the unknown state. The total size of the state vector (a) depends on the domain size and discretization scheme, but even for the smallest problems with a 3D gridded representation, it will typically range upwards of 10 000. We note that the state vector does not necessarily need to consist of physical variables and may instead consist of, for example, linear combinations of physical variables. This is a generalization that is preferable to ensure the associated optimization problem is well scaled (Nocedal and Wright, 2006).

The measurement vector (y) contains multi-angle, multi-pixel radiances, which may also be multi-spectral and polarized. The number of observations will also typically range above 10 000 and exceed the dimension of the state vector to avoid ill-posedness. The forward model $F(a)$ provides the mapping from the state vector to the measurement space by producing synthetic measurements equivalent to y, based on the unknown state vector (a) and any fixed ancillary data. The forward model consists of multiple components, including the mapping from the state vector (a) to the 3D optical properties at all required wavelengths; the solution of the required 3D RT problems; and, finally, the sampling of the radiance field at the required positions, angles, wavelengths, and polarization states to produce synthetic measurements. We describe each of these components of the forward model in the following subsection.

We select a best-fitting state by minimizing a scalar cost function. The scalar cost function χ^2 is chosen to penalize the misfit from the measurements in a generalized, least squares sense.

$$\chi^2 = (y - F(a))^T S_\epsilon^{-1} (y - F(a)) + R(a), \tag{1}$$

where the error covariance matrix of the residual between the measurements (y) and the forward model $F(a)$ is denoted by S_ϵ and accounts for both measurement uncertainty and forward-model uncertainty. AT3D currently supports a block diagonal S_ϵ, where error correlations are allowed between different Stokes components measured at each pixel, which supports the inclusion of certain types of forward-modeling error and instrumental noise (van Harten et al., 2018). It does not yet support systematic error correlations between pixels due to, for example, uncertainties in the flat-fielding operation, camera-to-camera intercalibration, band-to-band calibration error, or absolute calibration error. $R(a)$ is a differentiable regularization term that reflects prior knowledge about the structure of the unknown state. Note that there is no re-

quirement here that the regularization term takes the form of an a priori distribution. As such, the formulation in AT3D is more general than the optimal estimation frameworks utilized widely in atmospheric remote sensing (Sourdeval et al., 2013; Wang et al., 2016). We note that the cost function in Eq. (1) can take modified forms if transforms (e.g., logarithmic) are used to stabilize (pre-condition) the optimization process (Chance et al., 1997).

The solution to the inverse problem is given by the minimizer of the cost function, subject to the box constraints described by vectors of lower bounds (l) and upper bounds (u) on each element of the state vector. Formally,

$$\widetilde{a} = \operatorname*{argmin}_{a} \chi^2, \text{ subject to } l \le a \le u. \tag{2}$$

For example, these bounds can be used to ensure positivity of the liquid water content or volume extinction coefficient. This optimization problem can be solved efficiently, using a local minimization technique, such as the limited-memory Broyden–Fletcher–Goldfarb–Shanno method for bounded minimization (L-BFGS-B; Byrd et al., 1995), when the cost function is differentiable. The local minimization proceeds through the selection of an initial guess a_0 and iteratively updating the state through repeated evaluation of the cost function and its gradient. At the mth iteration, we then have the following:

$$a_{m+1} = a_m + \Delta a_m. \tag{3}$$

The retrieval method requires the selection of an initial guess, which can have a significant influence on the optimization. In Part 2 of this study, we describe some of the ways that initial guesses can be generated in AT3D.

The L-BFGS-B method is a quasi-Newton method which selects the update to the state vector (Δa_m) by first selecting a search direction through the minimization of an approximate, local, quadratic model of the cost function. The quadratic model uses an approximation to the Hessian of the cost function, which is formed by analyzing how the cost function gradient changes over the most recent M iterations. M is a hyperparameter of the optimization. In essence, the approximation to the Hessian of the cost function is formed from the finite differencing of successive gradient vectors. As such, it only reflects curvature information along the directions that the optimization trajectory is currently exploring. After the selection of a search direction, an inexact line search, obeying the Wolfe–Armijo conditions, which ensure stability and convergence, is then used to select the final update Δa_m along the search direction. The implementation of L-BFGS-B from the SciPy library is used (Virtanen et al., 2020). The L-BFGS-B algorithm has a better convergence rate than simple gradient descent but still only requires the evaluation of the cost function and its gradient. The computational cost of the update is modest once the gradient vector is computed. The storage requirement of L-BFGS-B is
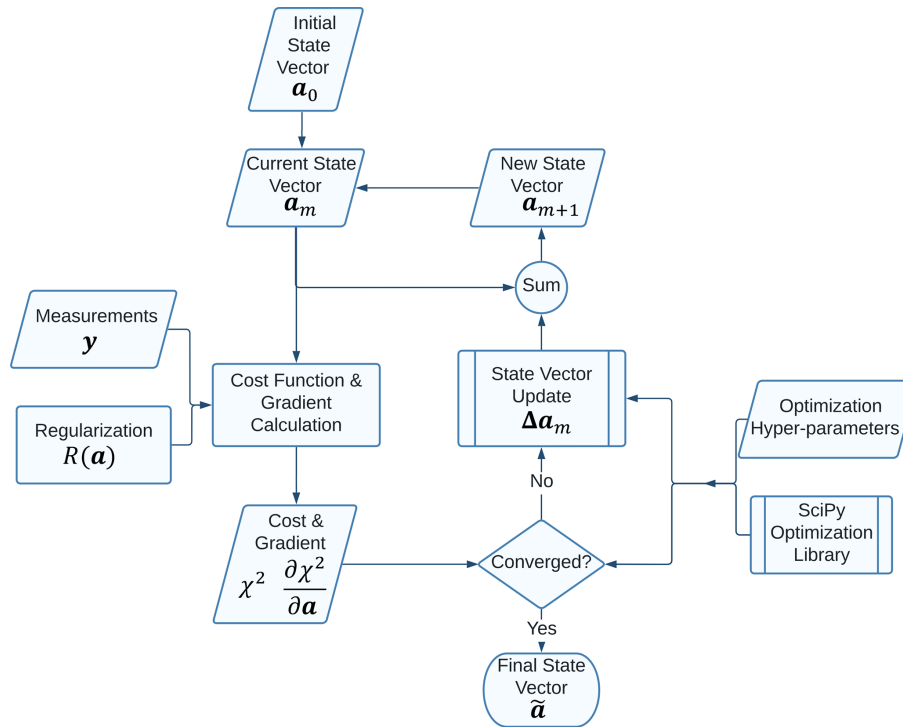
Figure 1. A flowchart depicting the overall iterative retrieval methodology of AT3D.

also modest and is limited to storing the state updates and gradient changes over the past M iterations, so it scales linearly with the size of the state vector. This makes the method appropriate for the large-scale optimization problem of cloud tomography.

The expression for the gradient of the data fit term required by the L-BFGS-B algorithm is as follows:

$$\frac{\partial \chi^2}{\partial \boldsymbol{a}} = 2(\boldsymbol{y} - \boldsymbol{F}(\boldsymbol{a}))^{\mathrm{T}} \mathbf{S}_\epsilon^{-1} \mathbf{K}, \tag{4}$$

where $\mathbf{K}$ is the Jacobian matrix containing the partial derivatives of the ith output of the forward model with respect to the jth component of the state vector. This is determined as follows:

$$\mathrm{K}_{ij} = \frac{\partial F_i(\boldsymbol{a})}{\partial a_j}. \tag{5}$$

It is in our interest to understand the information content of our measurements and the factors that control the convergence rate of the retrieval so that these can be maximized. The linear information content in the measurements in the vicinity of a cost function minimum can be determined by using the Fisher information matrix, which, in the case of Gaussian errors, takes the simple form $\mathbf{K}^{\mathrm{T}} \mathbf{S}_\epsilon^{-1} \mathbf{K}$. The singular-value spectrum of this matrix describes the magnitude of retrieval uncertainties across different directions in state space, and this distribution is largely controlled by the singular-value spectrum of $\mathbf{K}$. Strong off-diagonal error covariances in

the measurements can also affect the spectrum of the Fisher information matrix through $\mathbf{S}_\epsilon^{-1}$, though multi-angle imagers typically have block-diagonal error covariances that limit this effect. The largest physical values of the singular-value spectrum of $\mathbf{K}$ are physically bounded (Chen et al., 2018), and a linearization of RT tends to have rapidly (i.e., exponentially) decaying singular values (Culver et al., 2001). As such, we summarize the singular-value spectrum using just the condition number of the Jacobian matrix $\kappa(\mathbf{K})$, which is the ratio of its largest (s_{largest}) and smallest (s_{smallest}) singular values.

$$\kappa(\mathbf{K}) = \frac{s_{\mathrm{largest}}}{s_{\mathrm{smallest}}} \tag{6}$$

The L-BFGS-B method can suffer in systems where the condition number becomes large, where poor accuracy and slow convergence can occur (Zhu et al., 1997), even for quadratic cost functions. A slow convergence rate is particularly troubling, as it will determine the computational feasibility of performing cloud tomography. In Sect. 4, we relate the structure of the Jacobian matrix, specifically the condition number, to the properties of the medium through the principles of RT theory and support these principles with numerical simulations. We then examine the ability of the approximate Jacobian to accurately represent the true Jacobian matrix and to correctly capture the information content of the measurements about the state vector. Before presenting these results, we must first introduce the formulation of the forward model and the exact calculation of the Jacobian matrix.

3.1 Forward-model description

We now describe the formulation of the forward model and calculation of its Jacobian matrix. The theory for such calculations has already been presented for general 3D problems (Martin et al., 2014) and also in the specific context of the SHDOM model with periodic BCs (Doicu and Efremenko, 2019). In the work of Levis et al. (2015, 2017, 2020), open BCs were used without any incoming radiance at the domain edges (i.e., vacuum BCs). Neither of these configurations fully represents the realistic case of retrieving a heterogeneous 3D domain embedded within a horizontally infinite medium such as might be performed when we retrieve a field of cumulus clouds embedded in a cloud-free atmosphere. In this section, we describe a forward model for this scenario and its linearization, focusing on the specific context of the SHDOM solver. We describe the linearization of the forward model with respect to the parameters that control both the 3D domain and the embedding horizontally infinite medium, as implemented in AT3D.

Note that the principles of the forward model itself remain unchanged from the implementation of open BCs in the original SHDOM code. In SHDOM, when open horizontal BCs are selected, incoming radiances may be prescribed at the horizontal boundaries of the primary domain of 3D RT through the solution of auxiliary RT problems that describe the radiance field in the embedding medium. These auxiliary RT problems are 2D and 1D, describing the plane-parallel embedding medium. The coupling of the BCs between the RT problems, and the fact that a sampled radiance measurement contains contributions from both the primary 3D RTE solution and also the auxiliary RTE problems, introduces some complexity in the mathematical formulation of this model and its linearization. Our description below builds upon the formalism presented by Martin et al. (2014) and applies it to describe existing behavior of the SHDOM model. This description includes features that were not described in Evans (1998), such as the influence of the open BCs on the calculation of the radiance. This detailed description is necessary so that we can differentiate between the exact calculation of the Jacobian matrix and the approximations used in AT3D in Sect. 3.2 and 3.3. For a more pragmatic description of the essence of the approximate Jacobian calculation that does not include the treatment of the boundary conditions presented here, readers may refer to Levis et al. (2020). Our treatment focuses on the continuous problem rather than the details of numerical implementation, such as the delta-M scaling of the optical properties, except where conceptually necessary. Pertinent details on the numerical implementation related to the delta-M scaling and TMS correction in SHDOM can be found elsewhere (Evans, 1998; Doicu and Efremenko, 2019) and in Appendix B and D. Section 3.1.1 presents the definitions and geometry of the model. Section 3.1.2 describes the RT solution procedure. Section 3.1.3 describes the radiance calculation.

3.1.1 Problem setting

We begin by first defining the spatial domain of interest in which the monochromatic vector 3D RTE will be solved. The SHDOM model adopts a Cartesian geometry, and the physical domain $D \subset R^3$ is the horizontally infinite slab of thickness L_z, where $z \in [0, L_z]$. This domain is broken up into nine cuboids, and a different RT problem is solved in each. A top-down view of the arrangement of the domains is shown in Fig. 2. The primary physical domain of 3D RT is the cuboid D_1, described by the position vector $\mathbf{r}$, with smooth boundary ∂D_1.

$$D_1 = \left\{ \mathbf{r} \in R^3 : 0 < x < L_x, 0 < y < L_y, 0 < z < L_z \right\} \qquad (7)$$

The auxiliary RTE domains, D_2 through D_9, are around this primary domain, which we may consider to also be cuboids, but they have horizontal boundaries at some very large distance from the primary domain. We will denote the position of these horizontal boundaries as being at infinity; for example, $x = -\infty$. Practically speaking, the absolute extent of these auxiliary cuboids only needs to be greater than the position of any considered sensor. The absolute size only comes into play during the calculation of radiances in SHDOM and not the solution of the RTE itself, as will soon become apparent. In our description, we only consider a subset of the auxiliary domains and their coupling to the primary domain, as the rest can be treated similarly, following symmetry. D_2 and D_4 are examples of corner auxiliary domains, which share 1D edges with the primary domain (D_1), while D_3 is an example of a side auxiliary domain, which shares a horizontal plane with the primary domain (D_1). Specifically,

$$D_2 = \left\{ \mathbf{r} \in R^3 : -\infty < x < 0, y > L_y, 0 < z < L_z \right\}, \qquad (8)$$

$$D_3 = \left\{ \mathbf{r} \in R^3 : -\infty < x < 0, 0 < y < L_y, 0 < z < L_z \right\}, \qquad (9)$$

$$D_4 = \left\{ \mathbf{r} \in R^3 : -\infty < x < 0, y < 0, 0 < z < L_z \right\}. \qquad (10)$$

Optical properties are defined at all positions in D and are allowed to vary in 3D within D_1. Within the auxiliary domains, the optical properties are simplified, depending on whether the auxiliary domain shares a corner (e.g., D_2 and D_4) or a horizontal side (e.g., D_3) with the primary domain D_1. Optical properties are homogeneous in the x and y directions in the corner domains and are homogeneous in the direction normal to the boundary of the side domain that is shared with the primary domain. In the case of D_3, this is the x direction.

To describe the RT problems and their coupling, we need to distinguish between the domains and their boundaries, for which we must define some relevant sets. These definitions are also illustrated in Fig. 3. The set of directions for the RTE is the unit sphere $\mathbb{S}^2$ scanned by the propagation direction $\mathbf{\Omega}$. For any domain D_n, if we define the unit outward normal to
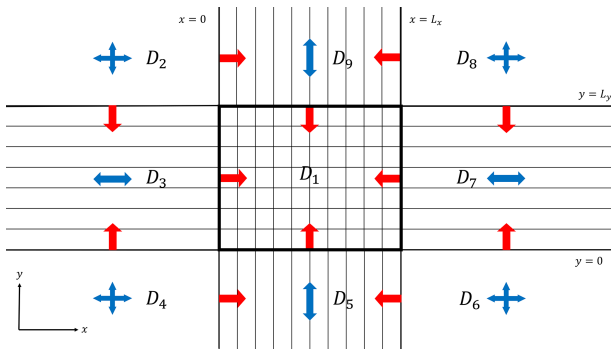
Figure 2. Top-down view of the geometry of the system of RT problems. D_1 denotes the domain of 3D RT. The blue arrows denote the directions in each domain for which there are periodic horizontal BCs. The red arrows denote which domains supply incoming radiances at each boundary. For example, the RTE problem solved on D_3 supplies the horizontal radiance BCs to the 3D RTE problem solved on D_1 at the common plane shared between them.
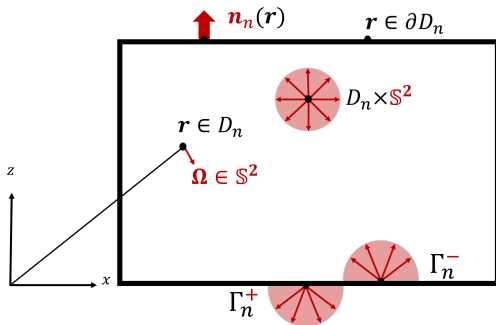

Figure 3. A side view of the nth RT domain illustrating several key definitions, such as the internal, incoming, and outgoing sets and the normal vector. Directional quantities are shown in red, while positional quantities are shown in black. See Eq. (11) and the associated discussion in the main text for more details.

the boundary of domain D_n as $\boldsymbol{n}_n(\boldsymbol{r})$, then we can define the incoming Γ_n^- and outgoing sets Γ_n^+ on the boundary of the domain as follows:

$$\Gamma_n^{\pm} = \left\{ (\boldsymbol{r}, \boldsymbol{\Omega}) \in \partial D_n \times \mathbb{S}^2 : \pm \boldsymbol{n}_n(\boldsymbol{r}) \cdot \boldsymbol{\Omega} > 0 \right\}. \quad (11)$$

These two sets allow the separation of radiance at the boundary of the domain into the sets of directions that are entering the domain ($\boldsymbol{\Omega}^-$) and those directions where radiance is leaving the domain ($\boldsymbol{\Omega}^+$).

For each domain ($n = 1, \ldots, 9$), we solve the monochromatic vector RTE for the polarized radiance field (i.e., Stokes vector) $\boldsymbol{I}_n(\boldsymbol{r}, \boldsymbol{\Omega}) = [I, Q, U, V]_n^{\mathrm{T}}$ on the internal set $D_n \times \mathbb{S}^2$. In the following, we omit the domain subscript unless relevant, i.e., when considering the coupling between domains. The union of the internal set and boundary is written as $D_n \times \mathbb{S}^2 \oplus \Gamma_n^{\pm}$. The RTE can be written, in terms of the trans-

port operator, as follows:

$$\mathcal{L}[\boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega})] = \boldsymbol{\Omega} \cdot \nabla \boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}) + \boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}) \sigma(\boldsymbol{r})$$
$$- \omega(\boldsymbol{r}) \sigma(\boldsymbol{r}) \int_{\mathbb{S}}^2 \boldsymbol{Z}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}') \boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}') \mathrm{d}\boldsymbol{\Omega}', \quad (12)$$

where $\sigma(\boldsymbol{r})$ is the volume extinction coefficient, $\omega(\boldsymbol{r})$ is the single-scattering albedo, and $\boldsymbol{Z}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}')$ is the phase matrix. The RTE is then simply

$$\mathcal{L}[\boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega})] = \boldsymbol{f}(\boldsymbol{r}, \boldsymbol{\Omega}) \text{ on } D_n \times \mathbb{S}^2, \quad (13)$$

where $\boldsymbol{f}(\boldsymbol{r}, \boldsymbol{\Omega})$ is the volume source vector. For example, thermal emission is an isotropic unpolarized volume source that reads as $\boldsymbol{f}(\boldsymbol{r}, \boldsymbol{\Omega}) = (1 - \omega(\boldsymbol{r})) \sigma(\boldsymbol{r}) \boldsymbol{B}(\boldsymbol{r})$, where $\boldsymbol{B}(\boldsymbol{r}) = [B_\lambda(T), 0, 0, 0]^{\mathrm{T}}$ with $B_\lambda(T)$ being the Planck blackbody radiance function. Equation (13) is paired with appropriate BCs, which constrain the solution through enforced continuity of the radiance field at the boundary. The BCs of the primary, side, and corner domains are different, and this introduces significant complexity in the formulation of the forward model. Corner domains (e.g., D_2 and D_4) have periodic BCs in both the x and y direction, while side domains (e.g., D_3) have periodic boundaries only in the direction normal to the boundary of the primary domain. The periodic boundary conditions in the auxiliary domains ensure that the RT solution in each auxiliary domain is independent of the 3D domain so that the system of RTs is solvable. This approximation neglects multiple-scattering interactions between the heterogeneous medium (D_1) and the auxiliary domains. As a result, open horizontal boundary conditions are an approximate treatment of the RT solution for a heterogeneous medium embedded in a horizontally homogeneous medium. Features like cloud and surface adjacency effects cannot be modeled unless the domain of 3D radiative transfer is sufficiently large enough to resolve them. The periodic boundary conditions can be expressed as follows. In D_3, a side domain, we have the following in the x direction:

$$\boldsymbol{I}_3(x = -\infty, y, z, \boldsymbol{\Omega})|_{\Gamma_3^+} = \boldsymbol{I}_3(x = 0, y, z, \boldsymbol{\Omega})|_{\Gamma_3^-}. \quad (14)$$

Directions of periodicity in the RT in each domain are denoted by the blue arrows in Fig. 2. For the horizontal directions that do not have periodicity, we have open BCs, with incoming radiance prescribed by the boundary source vector $\boldsymbol{g}(\boldsymbol{r}, \boldsymbol{\Omega})$ and the reflection of outgoing radiance by the reflection operator $\mathcal{R}$:

$$\boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-} - \mathcal{R}\left[\boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^+}\right] = g(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-} \text{ on } \Gamma^-. \quad (15)$$

The reflection operator $\mathcal{R}$ is defined as the integral over the hemisphere of incoming directions weighted by $\boldsymbol{R}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}')$, which is a polarized bidirectional reflectance distribution function (BRDF), as follows:

$$\mathcal{R}\left[\boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}')|_{\Gamma^+}\right](\boldsymbol{r}, \boldsymbol{\Omega})$$
$$= \int_{\Omega^+} |\boldsymbol{n} \cdot \boldsymbol{\Omega}'| \boldsymbol{R}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}') \boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}') \mathrm{d}\boldsymbol{\Omega}', \text{ on } \Gamma^+. \quad (16)$$

We note that, in the implementation of SHDOM, the BRDF function $\mathbf{R}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}')$ is non-vanishing only on the lower boundary of all the domains ($z = 0$). Within SHDOM, the boundary source vector $\boldsymbol{g}(\boldsymbol{r}, \boldsymbol{\Omega})$ is also restricted to being composed of four components. The first three are thermal emissions from the surface $\boldsymbol{g}^{\text{BOT}}(\boldsymbol{r}, \boldsymbol{\Omega})$, a horizontally homogeneous, isotropic emission from the domain top $\boldsymbol{g}^{\text{TOP}}(\boldsymbol{r}, \boldsymbol{\Omega})$, and a unidirectional collimated source due to solar illumination, $\boldsymbol{g}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})$, with intensity $\boldsymbol{F}_0$ incident on the domain top as follows:

$$\boldsymbol{g}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega}) = \boldsymbol{F}_0 \delta(\boldsymbol{\Omega} - \boldsymbol{\Omega}_{\text{sun}}) \delta(z - L_z). \quad (17)$$

The fourth component is the most complex and is only defined on the horizontal sides, which we will denote by $\boldsymbol{g}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})$. This boundary source vector on the sides is the solution of the neighboring auxiliary RTE problem and thus represents the incoming light into each domain due to the embedding medium. Taking the side domain D_3 as an example, the incoming boundary source at each of the boundaries shared with corner domains D_2 and D_4 is simply the outgoing radiance fields in those corner domains:

$$\boldsymbol{g}_3^{\text{SIDE}}(x, y = L_y, z, \boldsymbol{\Omega})\Big|_{\Gamma_3^-} = \boldsymbol{I}_2(x, y = L_y, z, \boldsymbol{\Omega})\Big|_{\Gamma_2^+}; \quad (18)$$

$$\boldsymbol{g}_3^{\text{SIDE}}(x, y = 0, z, \boldsymbol{\Omega})\Big|_{\Gamma_3^-} = \boldsymbol{I}_4(x, y = 0, z, \boldsymbol{\Omega})|_{\Gamma_4^+}. \quad (19)$$

Similarly, for the primary domain, D_1, the incoming radiance at all horizontal sides will be prescribed by the four auxiliary, side RT solutions. For example, for the side shared between D_1 and D_3, we have the following:

$$\boldsymbol{g}_1^{\text{SIDE}}(x = 0, y, z, \boldsymbol{\Omega})\Big|_{\Gamma_1^-} = \boldsymbol{I}_3(x = 0, y, z, \boldsymbol{\Omega})|_{\Gamma_3^+}. \quad (20)$$

This flow of incoming radiance is denoted by the red arrows in Fig. 2. We can see that there is a one-way propagation of radiance from the corner domains to the side domains and finally to the primary domain. This interaction is one-way due to the periodic BCs used in the solution of the auxiliary problems and ensures that the system of RTEs is solvable.

3.1.2 Radiative transfer solutions

With this basic setup, we can now describe the solution procedure of the system of RTEs. The solution to the RTEs is the first step in modeling specific instrument observables, i.e., the radiance at a particular pixel on a sensor. This forms the essence of the forward model that connects the unknown state to the measurements in AT3D. First, the corner RTEs must be solved to provide BCs to the side problems, and then these side problems must be solved to provide the BCs for the primary 3D problem. At this point, we go into some detail below about the solution procedure, as the concepts introduced here are necessary for describing the approximate Jacobian calculation that is actually used in AT3D.

In SHDOM, the radiance field is decomposed into the direct solar radiance $\boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})$ and the diffuse, scattered radiance field $\boldsymbol{I}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})$:

$$\boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}) = \boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega}) + \boldsymbol{I}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega}). \quad (21)$$

This is done so that the anisotropic scattering of the angular singularity of the solar source can be treated more accurately. We must also separate the incoming boundary source at the horizontal sides into its direct $\boldsymbol{g}_{\odot}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})$ and diffuse components $\boldsymbol{g}_{\text{d}}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})$:

$$\boldsymbol{g}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega}) = \boldsymbol{g}_{\odot}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega}) + \boldsymbol{g}_{\text{d}}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega}). \quad (22)$$

The direct component of the boundary source at the horizontal sides is due to the propagation of the solar beam through an auxiliary RT domain to the incoming boundary of the domain under consideration. The direct radiance field is simply the propagation of the boundary solar and direct sources into the domain, which are attenuated by the transmission along their optical path of propagation. On the other hand, the specification of the diffuse radiation requires the solution of multiply scattering RT problems whose source vectors are no longer angular singularities. To solve for the direct radiance field, we define the streaming operator $\mathcal{T}$ which propagates boundary radiances into the domain with the appropriate attenuation. This streaming operator maps Γ_n^- to $D_n \times \mathbb{S}^2 \oplus \Gamma_n^{\pm}$ and is defined as follows:

$$\mathcal{T}\left[\boldsymbol{I}(\boldsymbol{r}', \boldsymbol{\Omega}')\right](\boldsymbol{r}, \boldsymbol{\Omega})$$
$$= \delta\left(\frac{\boldsymbol{r}' - \boldsymbol{r}}{\|\boldsymbol{r}' - \boldsymbol{r}\|} - \boldsymbol{\Omega}'\right) \delta(\boldsymbol{\Omega} - \boldsymbol{\Omega}') \boldsymbol{I}(\boldsymbol{r}', \boldsymbol{\Omega}') \text{T}(\boldsymbol{r}, \boldsymbol{r}'), \quad (23)$$

where we have made use of the transmission between two points. The transmission is defined as follows:

$$\text{T}(\boldsymbol{r}, \boldsymbol{r}') = \exp\left(-\int_0^{\|\boldsymbol{r}' - \boldsymbol{r}\|} \sigma\left(\boldsymbol{r}' - l\frac{\boldsymbol{r}' - \boldsymbol{r}}{\|\boldsymbol{r}' - \boldsymbol{r}\|}\right) \text{d}l\right). \quad (24)$$

The streaming operator provides the direct solar radiance solution, namely

$$\boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega}) = \mathcal{T}\left[\boldsymbol{g}_{\odot}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega}) + \boldsymbol{g}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})\right]. \quad (25)$$

The RTE for the diffuse radiance then takes the following form:

$$\mathcal{L}[\boldsymbol{I}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})] = \boldsymbol{f}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})$$
$$= \sigma(\boldsymbol{r})\omega(\boldsymbol{r}) \int_{\mathbb{S}^2} \boldsymbol{Z}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}') \boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega}) \text{d}\boldsymbol{\Omega}'$$
$$+ \sigma(\boldsymbol{r})(1 - \omega(\boldsymbol{r})) \mathbf{B}(\boldsymbol{r}), \quad (26)$$

where $\mathbf{B}(\boldsymbol{r})$ is an optional isotropic unpolarized blackbody emission vector. The BCs of this equation are straightforward modifications to Eq. (15):

$$\boldsymbol{I}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-} - \mathcal{R}\left[\boldsymbol{I}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^+}\right] = \boldsymbol{g}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-}$$
$$= \mathcal{R}\left[\boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^+}\right] + \boldsymbol{g}^{\text{TOP}}(\boldsymbol{r}, \boldsymbol{\Omega}) + \boldsymbol{g}^{\text{BOT}}(\boldsymbol{r}, \boldsymbol{\Omega})$$
$$+ \boldsymbol{g}_{\text{d}}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega}), \quad (27)$$

where we have defined the diffuse boundary source vectors $g_d^X(r, \Omega)$ for top (X = TOP), bottom (X = BOT), and side (X = SIDE) boundaries. This system of RTEs is solved using a solution operator $\mathcal{U}$, which is an abstract representation of the SHDOM solver, as described in Martin et al. (2014):

$$I_d(r, \Omega) = \mathcal{U} \begin{bmatrix} f_d \\ g_d \end{bmatrix}. \tag{28}$$

In SHDOM, all of the RTE problems are solved jointly so that the BCs of the primary 3D problem evolve over the iterative solution procedure, along with the solution for the primary diffuse radiance field.

3.1.3 Observable evaluation

Now that we have outlined the solution procedure for the system of RTEs, the final step in the evaluation of the forward model is the sampling of the radiance fields by the sensors. With this final step, we will have described how observables (elements of the forward model) are modeled in AT3D. We will then be able to evaluate the cost function that measures the misfit between our modeled state and the measurements we are using in a given tomography problem. We will then also be able to present, in Sect. 3.2, the derivatives of these observables with respect to elements of the state vector which form the Jacobian matrix. These are used in AT3D to perform the tomographic retrieval.

The sampling operation to calculate observables can be expressed as the inner products between a sensor response function and the radiance fields. Let us define the inner products on each domain in terms of test fields v and w:

$$\langle v, w \rangle_{D_n \times \mathbb{S}^2} = \int_{D_n} \int_{\mathbb{S}^2} v(r, \Omega) \cdot w(r, \Omega) \, dS_\Omega dV_r; \tag{29}$$

$$\langle v, w \rangle_{\Gamma_n^-} = \int_{D_n} \int_{\Omega^-} v(r, \Omega) \cdot w(r, \Omega) \, dS_\Omega dV_r; \tag{30}$$

$$\langle v, w \rangle_{\Gamma_n^+} = \int_{D_n} \int_{\Omega^+} v(r, \Omega) \cdot w(r, \Omega) \, dS_\Omega dV_r. \tag{31}$$

For fields that are defined on both the boundaries and interior of each domain, we also define a composite inner product over the union of the domains $D_n \times \mathbb{S}^2 \oplus \Gamma_n^\pm$, which is simply the sum of the right-hand sides of Eqs. (29), (30), and (31).

As an example, an observable could be radiance exiting a domain, which could be modeled as an inner product between an as-yet-unspecified sensor response function $P_i(r, \Omega)$ and the radiance field, $\langle P_i(r, \Omega), I(r, \Omega) \rangle_{\Gamma_n^+}$. There are several complications with this simplistic picture, due to both the coupling between the different RT domains and the numerical representation of the radiance field in SHDOM. To address these issues completely, we proceed by first defining the particular sensor response functions used in AT3D. Then we describe how to evaluate the radiances at particular positions within each domain with SHDOM. Finally, we describe how these components combine to provide a radiance representative of the coupled system.

The sensor response functions $P_i(r, \Omega)$ in the original SHDOM software (Evans, 1998) take the form of an idealized, singular sampling at position r_i and angle Ω_i, with polarization analyzer O_i, which is a vector that weights the contribution of the different Stokes components to each observable. For example, $O_i = [1, 0, 0, 0]$ for an intensity measurement. In AT3D, we have generalized the sensor response function to a weighted sum over k singular samplings, which we refer to as sub-pixel rays. Each sub-pixel ray has its own position within the domain, $r_{ik} \in D$, and angle, Ω_{ik}, with weights, w_k, that sum to unity. In this way, we can more accurately model the field of view of sensors with a resolution much coarser than the resolution of the RT grid. This addition to AT3D enables the straightforward modeling of more realistic imagers, unlike in the SHDOM software. The sensor response function is then

$$P_i(r, \Omega) = O_i \sum_k w_k \delta(r - r_{ik}) \delta(\Omega - \Omega_{ik}), \tag{32}$$

where the weights are determined by a quadrature scheme over the sensor's detectors, e.g., a 2D Gauss–Legendre quadrature over each pixel. Given the linearity of the inner product, there is an easy generalization from a single quadrature point to a sum over several points, as in Eq. (32), so we simply consider sampling by just one sub-pixel ray in the following descriptions. This removes the need for summation over the k sub-pixel rays, but we keep the k index as a subscript so that sub-pixel quantities are clearly differentiated from pixel quantities in the following discussion.

The form of Eq. (32) indicates that we only need to be able to evaluate radiances at a set of singular positions r_{ik} and angles Ω_{ik} to evaluate the inner product. This is not as simple as it appears as, while we already have the diffuse radiance solution from SHDOM defined at every position and direction in Eq. (28), it is insufficiently accurate for radiances at particular positions and angles. This is because the diffuse radiance is represented in SHDOM on an angularly smooth basis of spherical harmonics that is appropriate for fluxes but not for highly anisotropic radiances.

To sample accurate radiances at position r_{ik} and angle Ω_{ik} from the RT solution, the formal solution of the RTE is used instead. To formalize this procedure, we define some quantities and operators. First, there is the effective volume source and the effective boundary source of the RTE, which are defined, respectively, as follows:

$$\hat{f}_d(r, \Omega) = \sigma(r) \omega(r) \int_{\mathbb{S}^2} Z(r, \Omega, \Omega') I_d(r, \Omega') \, d\Omega' + f_d(r\Omega), \tag{33}$$

and

$$\hat{g}_d(r, \Omega)|_{\Gamma^-} = \mathcal{R}[I_d(r, \Omega)|_{\Gamma^+}] + g_d(r, \Omega)|_{\Gamma^-}. \tag{34}$$

The effective volume source is also commonly known as the SOURCE function of the RTE problem (e.g., Evans,

1998), though we have adopted our more precise nomenclature to avoid ambiguity with other sources such as boundary sources. The two effective sources are the arguments of the formal solution of the RTE, which is stated below in Eq. (36). The effective volume source is represented on a basis of spherical harmonics in SHDOM, which is quite accurate, as the radiance fields have been angularly smoothed by convolution with the phase matrix and are therefore far less affected by truncation error.

The final definition required for the radiance calculation is that of the volume streaming operator, which integrates the effective volume source along a characteristic. It maps from $D_n \times \mathbf{S}^2$ to $D_n \times \mathbf{S}^2 \oplus \mathbf{\Gamma}_n^+$ and is defined in terms of the test field $\boldsymbol{v}$ as follows:

$$\mathcal{S}\left[\boldsymbol{v}\left(\boldsymbol{r}', \boldsymbol{\Omega}\right)|_{D_n \times \mathbf{S}^2}\right]_{D_n \times \mathbf{S}^2} \mathbf{\Gamma}^+\left(\boldsymbol{r}, \boldsymbol{\Omega}\right)$$

$$= \int_0^{\|\boldsymbol{r}'-\boldsymbol{r}\|} \boldsymbol{v}(\boldsymbol{r} - l\boldsymbol{\Omega}, \boldsymbol{\Omega}) T\left(\boldsymbol{r} - l\boldsymbol{\Omega}, \boldsymbol{r}\right) \mathrm{d}l. \tag{35}$$

The radiance at a given position $\boldsymbol{r}_{ik}$ and angle $\boldsymbol{\Omega}_{ik}$ is then given by the following:

$$\langle \delta\left(\boldsymbol{\Omega} - \boldsymbol{\Omega}_{ik}\right) \delta(\boldsymbol{r} - \boldsymbol{r}_{ik}), \boldsymbol{I}\left(\boldsymbol{r}, \boldsymbol{\Omega}\right)\rangle_{D_n \times \mathbf{S}^2 \oplus \mathbf{\Gamma}_n^{\pm}}$$

$$= \boldsymbol{I}_{\odot}\left(\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik}\right) + \mathcal{T}\left[\hat{\boldsymbol{g}}_{\mathrm{d}}\left(\boldsymbol{r}, \boldsymbol{\Omega}\right)|_{\mathbf{\Gamma}^-}\right]\left(\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik}\right)$$

$$+ \mathcal{S}\left[\hat{\boldsymbol{f}}_{\mathrm{d}}\Big|_{D_n \times \mathbf{S}^2}\right]\left(\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik}\right). \tag{36}$$

In our formulation, the direct component of the radiance $\boldsymbol{I}_{\odot}$, which is singular in direction, is assumed to never be observed, and so we actually neglect the first term on the right-hand side of Eq. (36). This is not a strong limitation for the back- or side-scattering observation geometries of Earth-viewing remote sensing instruments deployed on airborne and satellite platforms. The TMS method used for radiance calculation in SHDOM can have significant inaccuracies near the solar direction (Nakajima and Tanaka, 1988). Therefore, further extension of this retrieval method to include measurements of the direct solar radiance should also include improvements to the SHDOM solver itself.

The evaluation of an element of the forward model involves contributions from these streaming operations from each domain along the line of sight of the sensor. Not all domains contribute similarly. This is because we would like to treat the coupled system as one cohesive approximation to the horizontally infinite atmosphere in the evaluation of the forward model by applying the formal solution of the RTE across all unified domains, ignoring all horizontal boundaries. This means that all domains that intersect the line of sight will contribute their effective volume source to the observable. Expressing this precisely takes some complexity, as we wish to express the radiance as a sum over inner products over each domain so that we can easily express differentiation of the observables with respect to the state vector, following Martin et al. (2014).
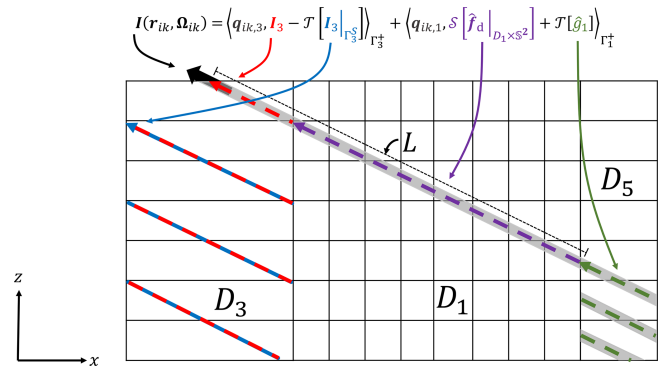

Figure 4. A side view of the system of RTE domains, illustrating the differences in how radiances are calculated during the solution of the RTEs in each domain and during the evaluation of the forward model. Consider the calculation of the radiance at the position $\boldsymbol{r}_{ik}$ and angle $\boldsymbol{\Omega}_{ik}$ at the upper boundary of D_3 denoted by the large black arrow. During the solution of the RTE in D_3, the radiance is calculated by integration along the red characteristic that follows the periodic BCs at the edge of D_3. On the other hand, when the forward model is evaluated, the radiance is calculated through integration along the characteristic denoted by the gray shading which passes through D_3 into D_1 and D_5. The correspondence between the mathematical expressions for the evaluation of the forward model (Eq. 41), in this case, and the graphical illustration are shown. See the main text for additional details.

To begin, let us consider the example shown in Fig. 4 before introducing the general mathematical description, which follows in Eq. (41). The mathematical expression specific to this example is shown in Fig. 4. We want to calculate a radiance at the position $\boldsymbol{r}_{ik}$ and direction $\boldsymbol{\Omega}_{ik}$ of a sub-pixel ray. This sub-pixel ray has a line of sight that points in the opposite direction to the propagation of the radiance that is sampled by the ray, $-\boldsymbol{\Omega}_{ik}$. In Fig. 4, this line of sight is presented as gray shading. We can see that the line of sight of the sub-pixel ray overlaps the three domains, D_3, D_1, and D_5, and so they will all contribute to the radiance calculation.

The line of sight first intersects D_3. All we actually want from D_3 is the contribution to the radiance from the effective volume source in the region of overlap between the dashed red line and the gray shading. Let us call this contribution the "overlap contribution". If we were to just evaluate the radiance at position $\boldsymbol{r}_{ik}$ and direction $\boldsymbol{\Omega}_{ik}$ using only the RT solution of D_3 using Eq. (36), then we would be following the periodic boundary conditions of D_3, with the integration domain of the streaming operator denoted by the red arrow in Fig. 4. Let us call this radiance the "D_3 radiance". To isolate the overlap contribution, we must additionally calculate the periodic boundary radiance which is in blue in Fig. 4. This boundary radiance, attenuated along the line of sight to $\boldsymbol{r}_{ik}$, must then be subtracted from the D_3 radiance to give the overlap contribution, which is a portion of the observable.

We then need to add the contributions to the observable from D_1 by evaluating Eq. (36) in that domain, which is

shown in purple in Fig. 4. The line of sight intersects the horizontal boundary of D_1. The evaluation of the boundary term in Eq. (36), which is the second term on the right-hand side, requires a contribution from D_5. Following the open BCs (e.g., Eq. 18), the incoming radiance at the horizontal boundaries of D_1 is simply the radiance at the boundary of D_5, which is calculated using Eq. (36). This calculation requires integration along the green chord.

To write this mathematically, we first need to isolate the incoming set on the horizontal side, which is a subset of the incoming set defined in Eq. (11):

$$\Gamma_n^S = \left\{ (\boldsymbol{r}, \boldsymbol{\Omega}) \in \Gamma_n^- : 0 < z < L_z \right\}. \tag{37}$$

We then need to define an indicator function for the domains for which we need to calculate direct radiance contributions. For the example above, the domains that directly contribute are D_1 and D_3. D_5 does not directly contribute; it contributes instead through the open BC of D_1. This distinction is important for linearization of the forward model. We use an indicator function to separate between these cases, which takes the form of a Heaviside function along the line of sight.

To define the indicator, we need to define a length denoted by L, which is the distance until the intersection of the line of sight with a boundary of a domain with a lower dimensionality of radiative transport than the current domain. This intersection marks the point at which the domains contribute indirectly through the open BCs. Recall that the corner domains are 1D, the side domains are 2D, and the primary domain is 3D. For the example in Fig. 4, we would find the intersection at the boundary between D_1 and D_5. The corresponding length L is illustrated in Fig. 4. We can then define the indicator function as follows:

$$H_L(x) = \begin{cases} 1, & 0 \le x < L \\ 0, & \text{otherwise} \end{cases}. \tag{38}$$

Note that the indicator function is not inclusive. Now, we can define the volume response functions $\boldsymbol{p}_{ik,n}(\boldsymbol{r}, \boldsymbol{\Omega})$ and the attenuated boundary response function $\boldsymbol{q}_{ik,n}(\boldsymbol{r}, \boldsymbol{\Omega})$ for the subpixel ray for each of the n domains.

$$\boldsymbol{p}_{ik,n}(\boldsymbol{r}, \boldsymbol{\Omega}) = \mathbf{O}_i \delta(\boldsymbol{r}_{ik} - \boldsymbol{r}) \delta(\boldsymbol{\Omega} - \boldsymbol{\Omega}_{ik}) \quad \text{in } D_n \times \mathbb{S}^2, \tag{39}$$

and

$$\boldsymbol{q}_{ik,n}(\boldsymbol{r}, \boldsymbol{\Omega}) = H_L(\|\boldsymbol{r}_{ik} - \boldsymbol{r}\|) \mathbf{O}_i \delta\left(\frac{\boldsymbol{r}_{ik} - \boldsymbol{r}}{\|\boldsymbol{r}_{ik} - \boldsymbol{r}\|} - \boldsymbol{\Omega}_i \right)$$
$$\delta(\boldsymbol{\Omega} - \boldsymbol{\Omega}_{ik}) \, \mathrm{T}(\boldsymbol{r}, \boldsymbol{r}_{ik}) \quad \text{on } \Gamma_n^+. \tag{40}$$

These two response functions will sample the radiance within each domain if the sensor is internal to a domain and also at the boundaries of all of the domains along the line of sight between the sensor and the primary domain. The boundary sample is weighted by the transmission to the sensor. In the example in Fig. 4, $\boldsymbol{q}_{ik,1}$ and $\boldsymbol{q}_{ik,3}$ are non-zero, while $\boldsymbol{q}_{ik,5}$ and all other attenuated boundary response functions

are zero. Additionally, all $\boldsymbol{p}_{ik,n}$ are zero in the example in Fig. 4, as $\boldsymbol{r}_{ik}$ is not interior to any domain.

The forward model is then expressed as follows:

$$F_i(\boldsymbol{x}) = \langle \boldsymbol{p}_{ik,1}, \boldsymbol{I}_{1,\mathrm{d}} \rangle_{D_1 \times \mathbb{S}^2} + \langle \boldsymbol{q}_{ik,1}, \boldsymbol{I}_{1,\mathrm{d}} \rangle_{\Gamma_1^+}$$
$$+ \sum_{n=2}^{9} \left(\left\langle \boldsymbol{p}_{ik,n}, \boldsymbol{I}_{n,\mathrm{d}} - \mathcal{T}\left[\boldsymbol{I}_{n,\mathrm{d}} \big|_{\Gamma_n^S} \right] \right\rangle_{D_n \times \mathbb{S}^2} \right.$$
$$\left. + \left\langle \boldsymbol{q}_{ik,n}, \boldsymbol{I}_{n,\mathrm{d}} - \mathcal{T}\left[\boldsymbol{I}_{n,\mathrm{d}} \big|_{\Gamma_n^S} \right] \right\rangle_{\Gamma_n^+} \right). \tag{41}$$

The inner products over the internal sets will be non-zero only if the sensor is internal to that domain and will therefore be non-zero only for one domain. The inner products over the boundaries will be non-zero only when the line of sight of the ray intersects the outgoing boundary of a domain. The first two terms on the right-hand side of Eq. (41) describe the sampling of the radiance field in the primary domain of 3D RT (D_1), which follows the formulation of Martin et al. (2014).

The second set of inner products shows the contributions of the auxiliary domains. We have removed the component from the radiance field that originates from the periodic horizontal BCs. The resulting expressions are still differentiable, though the existence of an adjoint formulation has not been proven. We do not consider the adjoint formulation of this system here. This forward model is valid for all radiances, except those that are measured exactly in the horizontal plane, where it is undefined just like for periodic horizontal BCs. We now have a full description of the forward model $F(\boldsymbol{a})$ used in AT3D to connect radiometric observables to the state of the atmosphere.

The derivatives of those inner products with respect to the state vector can be expressed using tangent-linear or forward-adjoint principles, as long as care is taken to address the sensitivity of this RT solution to changes in the incoming boundary radiance. This extension to the optimization of the BCs is not discussed in Martin et al. (2014) but is relatively straightforward and will be described below.

3.2 Linearization of the forward model

We can now calculate derivatives of the forward model, with respect to a component of the state vector a_j, which form the essence of the tomographic retrieval employed in AT3D. These derivatives, which form the Jacobian matrix ($\mathbf{K}$, Eq. 5), tell us how to optimally adjust the state vector to better match the measurements. We compute these derivatives for the first set of inner products in Eq. (41) over the primary 3D domain, following Martin et al. (2014). Let $F_{i,n}(\boldsymbol{a})$ refer to the contribution of the nth RTE domain to the forward-model output. For $n = 1$, we have the following:

$$\frac{\partial F_{i,1}(\boldsymbol{a})}{\partial a_j} = \left\langle \boldsymbol{p}_i, \frac{\partial \boldsymbol{I}_{1,\mathrm{d}}}{\partial a_j} \right\rangle_{D_1 \times \mathbb{S}^2} + \left\langle \boldsymbol{q}_i, \frac{\partial \boldsymbol{I}_{1,\mathrm{d}}}{\partial a_j} \right\rangle_{\Gamma_+^1}. \tag{42}$$

These inner products are evaluated by solving a modified RTE problem for the derivatives of the diffuse radiance. The

following expression holds for an arbitrary RTE domain and is formed by differentiating the RTE (Eq. 26) and regrouping terms to find the volume source vector of the modified RTE for the derivatives of the radiance field, $\Delta \boldsymbol{f}_j$:

$$
\begin{aligned}
\mathcal{L}\left[\frac{\partial \boldsymbol{I}_{\mathrm{d}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}\right] &= -I_{\mathrm{d}} \frac{\partial \sigma}{\partial a_j} \\
&+ \left(\frac{\partial \sigma}{\partial a_j} \omega + \frac{\partial \omega}{\partial a_j} \sigma\right) \int_{\mathbb{S}^2} \boldsymbol{Z} \boldsymbol{I}_{\mathrm{d}} \mathrm{d}\boldsymbol{\Omega}' + \sigma \omega \int_{\mathbb{S}^2} \frac{\partial \boldsymbol{Z}}{\partial a_j} \boldsymbol{I}_{\mathrm{d}} \mathrm{d}\boldsymbol{\Omega}' \\
&+ \boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})\left[\frac{\partial \sigma}{\partial a_j} \omega \boldsymbol{Z} + \frac{\partial \omega}{\partial a_j} \sigma \boldsymbol{Z} + \frac{\partial \boldsymbol{Z}}{\partial a_j} \sigma \omega\right] \\
&+ \left(\frac{\partial \sigma}{\partial a_j}(1-\omega) - \frac{\partial \omega}{\partial a_j} \sigma\right) B + \sigma(1-\omega) \frac{\partial B}{\partial a_j} \\
&+ \frac{\partial \boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} \omega \sigma \boldsymbol{Z} = \Delta \boldsymbol{f}_j.
\end{aligned}
\tag{43}
$$

The spatio-angular structure of the volume source of this modified RTE ($\Delta \boldsymbol{f}_j$) is controlled by the radiance field, which varies with the direction of solar illumination, for example.

We also differentiate the BCs of the RTE (Eq. 27) to obtain the associated boundary source vector of the modified RTE, $\Delta \boldsymbol{g}_j$, as follows:

$$
\begin{aligned}
\left.\frac{\partial \boldsymbol{I}_{\mathrm{d}}}{\partial a_j}\right|_{\boldsymbol{\Gamma}_-} &- R\left[\left.\frac{\partial \boldsymbol{I}_{\mathrm{d}}}{\partial a_j}\right|_{\boldsymbol{\Gamma}_+}\right] = \mathcal{R}\left[\left.\frac{\partial \boldsymbol{I}_{\odot}}{\partial a_j}\right|_{\boldsymbol{\Gamma}_+}\right] \\
&+ \frac{\partial \mathcal{R}}{\partial a_j}\left[\boldsymbol{I}_{\odot} + \boldsymbol{I}_{\mathrm{d}}\right] + \frac{\partial \boldsymbol{g}^{\mathrm{TOP}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} + \frac{\partial \boldsymbol{g}^{\mathrm{BOT}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} \\
&+ \frac{\partial \boldsymbol{g}_{\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} = \Delta \boldsymbol{g}_j.
\end{aligned}
\tag{44}
$$

The derivative of the reflection operator (Eq. 16) is

$$
\frac{\partial \mathcal{R}}{\partial a_j}[\boldsymbol{I}] = \int_{\Omega^+}\left|\boldsymbol{n} \cdot \boldsymbol{\Omega}'\right| \frac{\partial \mathbf{R}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}')}{\partial a_j} \boldsymbol{I}(\boldsymbol{r}, \boldsymbol{\Omega}') \mathrm{d}S_{\boldsymbol{\Omega}'}.
\tag{45}
$$

The solution to these modified systems is then

$$
\frac{\partial \boldsymbol{I}_{\mathrm{d}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} = \mathcal{U}\left[\begin{array}{c} \Delta \boldsymbol{f}_j \\ \Delta \boldsymbol{g}_j \end{array}\right].
\tag{46}
$$

Equation (46) constitutes a tangent-linear model to SHDOM (Eq. 28). To evaluate $\Delta \boldsymbol{f}_j$, we just need the solution of the forward RTE problem and the optical property derivatives. We also have to calculate the derivatives of the direct solar beam; that is,

$$
\begin{aligned}
\frac{\partial \boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} &= \frac{\partial T}{\partial a_j}\left[\boldsymbol{g}_{\odot}^{\mathrm{S}}(\boldsymbol{r}, \boldsymbol{\Omega}) + \boldsymbol{g}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})\right] \\
&+ \mathcal{T}\left[\frac{\partial \boldsymbol{g}_{\odot}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}\right].
\end{aligned}
\tag{47}
$$

The derivative of the streaming operator (Eq. 23) simply follows from the derivative of the transmission function

(Eq. 24):

$$
\begin{aligned}
&\frac{\partial \mathcal{T}}{\partial a_j}\left[\boldsymbol{I}(\boldsymbol{r}', \boldsymbol{\Omega}')\right](\boldsymbol{r}, \boldsymbol{\Omega}) \\
&= -\mathcal{T}\left[\boldsymbol{I}(\boldsymbol{r}', \boldsymbol{\Omega}')\right](\boldsymbol{r}, \boldsymbol{\Omega}) \int_0^{\|\boldsymbol{r}'-\boldsymbol{r}\|} \frac{\partial \sigma}{\partial a_j} \\
&\quad \left(\boldsymbol{r} - l\frac{\boldsymbol{r}'-\boldsymbol{r}}{\|\boldsymbol{r}'-\boldsymbol{r}\|}\right) \mathrm{d}l'.
\end{aligned}
\tag{48}
$$

The remaining terms to be specified are the derivatives of the boundary source vectors on the domain top, $\frac{\partial \boldsymbol{g}^{\mathrm{TOP}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$; bottom, $\frac{\partial \boldsymbol{g}^{\mathrm{BOT}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$; and sides, $\frac{\partial \boldsymbol{g}_{\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$ and $\frac{\partial \boldsymbol{g}_{\odot}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$. The first two terms are determined by local analytic relationships with the state, so that their derivatives can be readily calculated. For instance, the surface emission source on the right-hand side of Eq. (15) depends on temperature, and the surface BRDF in Eq. (16) on the left-hand side may also be parameterized by elements of the state vector. The second two terms involve the side boundary sources, which are nonzero only for the open boundaries and involve coupling between the different RTE solutions. The direct component, $\frac{\partial \boldsymbol{g}_{\odot}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$, is just the direct radiance derivative, $\frac{\partial \boldsymbol{I}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$, on the adjoining boundary and can therefore be calculated through recursive application of Eq. (46) across each domain until the direct solar source is only the domain top, where $\frac{\partial \boldsymbol{g}_{\odot}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} = 0$. We do not need to differentiate the domain-top solar source, $\boldsymbol{g}_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})$, which is fixed and known. The derivatives of the diffuse boundary source are more complex to calculate, requiring their own RTE solutions. Specifically, the derivatives of the diffuse boundary source on the 3D primary domain D_1 require RTE solutions on the side domains (e.g., D_3). The derivatives of the diffuse boundary source on the side domains require RTE solutions on the corner domains (e.g., D_2 and D_4). We can then evaluate the following:

$$
\begin{aligned}
\frac{\partial \boldsymbol{g}_{1,\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} &= \left\langle\delta(x), \frac{\partial \boldsymbol{I}_{3,\mathrm{d}}}{\partial a_j}\right\rangle_{\boldsymbol{\Gamma}_+^3} + \left\langle\delta(y), \frac{\partial \boldsymbol{I}_{5,\mathrm{d}}}{\partial a_j}\right\rangle_{\boldsymbol{\Gamma}_+^5} \\
&+ \left\langle\delta(x - L_x), \frac{\partial \boldsymbol{I}_{7,\mathrm{d}}}{\partial a_j}\right\rangle_{\boldsymbol{\Gamma}_+^7} \\
&+ \left\langle\delta(y - L_y), \frac{\partial \boldsymbol{I}_{9,\mathrm{d}}}{\partial a_j}\right\rangle_{\boldsymbol{\Gamma}_+^9}.
\end{aligned}
\tag{49}
$$

For each of the side domains, we then have, for example,

$$
\begin{aligned}
\frac{\partial \boldsymbol{g}_{3,\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} &= \left\langle\delta(y - L_y), \frac{\partial \boldsymbol{I}_{2,\mathrm{d}}}{\partial a_j}\right\rangle_{\boldsymbol{\Gamma}_+^2} \\
&+ \left\langle\delta(y), \frac{\partial \boldsymbol{I}_{4,\mathrm{d}}}{\partial a_j}\right\rangle_{\boldsymbol{\Gamma}_+^4}.
\end{aligned}
\tag{50}
$$

We can now evaluate all of the derivatives of the forward model as follows:

$$
\frac{\partial F_i(\boldsymbol{a})}{\partial a_j} = \left\langle \boldsymbol{p}_{i,1}, \frac{\partial \boldsymbol{I}_{1,\mathrm{d}}}{\partial a_j} \right\rangle_{D_1 \times \mathbb{S}^2} + \left\langle \boldsymbol{q}_{i,1}, \frac{\partial \boldsymbol{I}_{1,\mathrm{d}}}{\partial a_j} \right\rangle_{\Gamma_1^+}
$$

$$
+ \sum_{n=2}^{9} \left(\left\langle < \boldsymbol{p}_{i,n}, \frac{\partial \boldsymbol{I}_{n,\mathrm{d}}}{\partial a_j} - \mathcal{T}\left[\frac{\partial \boldsymbol{I}_{n,\mathrm{d}}}{\partial a_j} \bigg|_{\Gamma_n^\mathrm{S}} \right] \right. \right.
$$

$$
\left. - \frac{\partial \mathcal{T}}{\partial x_j}\left[\boldsymbol{I}_{n,\mathrm{d}} \big|_{\Gamma_n^\mathrm{S}} \right] \right\rangle_{D_n \times \mathbb{S}^2}
$$

$$
\left. + \left\langle \boldsymbol{q}_{i,n}, \frac{\partial \boldsymbol{I}_{n,\mathrm{d}}}{\partial a_j} - \mathcal{T}\left[\frac{\partial \boldsymbol{I}_{n,\mathrm{d}}}{\partial a_j} \bigg|_{\Gamma_n^\mathrm{S}} \right] - \frac{\partial \mathcal{T}}{\partial a_j}\left[\boldsymbol{I}_{n,\mathrm{d}} \big|_{\Gamma_n^\mathrm{S}} \right] \right\rangle_{\Gamma_n^+} \right). \quad (51)
$$

This is an exact treatment of the derivatives of the forward model and is what would be numerically approximated by performing finite differencing of the forward model. This is not the approach directly employed in AT3D.

3.3 Approximate Jacobian calculation

In AT3D, we evaluate approximate derivatives of the forward model instead of exactly evaluating Eq. (51). This is done for reasons of algorithmic simplicity and computational efficiency but naturally can have consequences for the accuracy of the retrieval. This is quite common in optimization problems (Ye et al., 1999; Eppstein et al., 2003; Dwight and Brezillon, 2006), as convergence is the key criterion to measure success of an optimization-based retrieval, rather than a high-accuracy solution to any particular linearized problem. The approximate derivatives use approximate solutions to the RTEs for the radiance derivatives. Specifically, the approximation to the derivatives uses a no-scattering assumption in the solution of the tangent-linear model and resulting evaluation of the radiance fields, following Levis et al. (2020). The no-scattering assumption is equivalent to a zeroth order approximation to a successive order of the scattering solution to Eq. (46). It does not involves setting the single-scattering albedo to zero. The result is that the multiple-scattering tangent-linear model in Eq. (46) does not need to be evaluated, which saves the computational expense of evaluating an additional 3D RT model at each iteration to calculate derivatives. The approximate radiance derivatives require only the evaluation of the formal, integral solutions, which are just line integrations with the same geometry as the forward model (see Fig. 4). This gives us an easy way to adjust the unknown state to better match the measurements and, from there, perform the tomographic retrieval. This key approximation is the essence of AT3D and makes it computationally suitable for solving practical tomography problems.

Let us formalize the approximate derivative calculation. We define the effective volume source and the effective boundary source of the radiance derivative RTE analogously to in the forward solution (Eqs. 33 and 34) as follows:

$$
\Delta \hat{\boldsymbol{f}}_j(\boldsymbol{r}, \boldsymbol{\Omega}) = \sigma(\boldsymbol{r})\omega(\boldsymbol{r}) \int_{\mathbb{S}^2} \boldsymbol{Z}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}') \frac{\partial \boldsymbol{I}_{\mathrm{d}}}{\partial a_j}(\boldsymbol{r}, \boldsymbol{\Omega}') \, \mathrm{d}\boldsymbol{\Omega}'
$$

$$
+ \Delta \boldsymbol{f}_j(\boldsymbol{r}, \boldsymbol{\Omega}), \quad (52)
$$

and

$$
\Delta \hat{\boldsymbol{g}}_j(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-} = \mathcal{R}\left[\frac{\partial \boldsymbol{I}_{\mathrm{d}}}{\partial a_j}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^+} \right] + \Delta \boldsymbol{g}_j(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-}. \quad (53)
$$

The inner products in Eq. (51) may then be evaluated, following the same rule for singular sampling as in the forward model (Eq. 36):

$$
\left\langle \delta(\boldsymbol{\Omega} - \boldsymbol{\Omega}_{ik})\delta(\boldsymbol{r} - \boldsymbol{r}_{ik}), \frac{\partial \boldsymbol{I}_{n,\mathrm{d}}}{\partial a_j} \right\rangle_{D_n \times \mathbb{S}^2 \oplus \Gamma_n^\pm}
$$

$$
= \mathcal{T}\left[\Delta \hat{\boldsymbol{g}}_{\mathrm{d}}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma_n^-} \right](\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik})
$$

$$
+ \mathcal{S}\left[\Delta \hat{\boldsymbol{f}}_j \big|_{D_n \times \mathbb{S}^2} \right](\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik}). \quad (54)
$$

In AT3D, Eq. (54) is approximated by invoking the no-scattering assumption, so the effective volume and boundary sources contain no recursive dependence on $\frac{\partial \boldsymbol{I}_{\mathrm{d}}}{\partial a_j}$. In particular, we approximate the effective volume source for the radiance derivatives as follows:

$$
\Delta \hat{\boldsymbol{f}}_j \big|_{D_n \times \mathbb{S}^2} \approx \Delta \boldsymbol{f}_j(\boldsymbol{r}, \boldsymbol{\Omega}). \quad (55)
$$

We also neglect the first term on the right-hand side of Eq. (53), which is the reflection of diffusely scattered radiance derivatives at the boundary. These two no-scattering approximations are applied for a given domain for the evaluation of integrals of the form of Eq. (54). However, there is also an effect through the BCs that further approximates $\Delta \boldsymbol{g}_j(\boldsymbol{r}, \boldsymbol{\Omega})|_{\Gamma^-}$, due to the application of the no-scattering approximation to all domains. The diffuse horizontal boundary sources for the radiance derivatives in Eq. (44), $\frac{\partial g_{n,\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$, are calculated by applying the same singular sampling of the radiance derivative expressed in Eq. (54) on the other domains. For a side domain (e.g., D_3), the horizontal diffuse boundary radiance comes from only corner domains (e.g., D_2 and D_4), which themselves do not have open horizontal BCs. As such, the diffuse horizontal boundary sources for radiance derivatives in Eq. (50) are approximated as follows:

$$
\frac{\partial \boldsymbol{g}_{3,\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} = \left\langle \delta(y - L_y), \frac{\partial \boldsymbol{I}_{2,\mathrm{d}}}{\partial a_j} \right\rangle_{\Gamma_+^2}
$$

$$
+ \left\langle \delta(y), \frac{\partial \boldsymbol{I}_{4,\mathrm{d}}}{\partial a_j} \right\rangle_{\Gamma_+^4} \approx \frac{\partial \boldsymbol{g}_{3,\mathrm{d}}^{\mathrm{SIDE},0}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}
$$

$$
= \left\langle \delta(y - L_y), \mathcal{T}\left[\Delta \boldsymbol{g}_j |_{\Gamma^-} \right](\boldsymbol{r}, \boldsymbol{\Omega}) + \mathcal{S}\left[\Delta \boldsymbol{f}_j |_{D_2 \times \mathbb{S}^2} \right](\boldsymbol{r}, \boldsymbol{\Omega}) \right\rangle_{\Gamma_+^2}
$$

$$
+ \left\langle \delta(y), \mathcal{T}\left[\Delta \boldsymbol{g}_j |_{\Gamma^-} \right](\boldsymbol{r}, \boldsymbol{\Omega}) + \mathcal{S}\left[\Delta \boldsymbol{f}_j |_{D_4 \times \mathbb{S}^2} \right](\boldsymbol{r}, \boldsymbol{\Omega}) \right\rangle_{\Gamma_+^4}. \quad (56)
$$

Note that, when compared to Eq. (54), we have made the approximation that $\boldsymbol{\Delta}\hat{\boldsymbol{g}}_j(\boldsymbol{r}, \boldsymbol{\Omega})|_{\boldsymbol{\Gamma}^-} \approx \boldsymbol{\Delta}\boldsymbol{g}_j(\boldsymbol{r}, \boldsymbol{\Omega})|_{\boldsymbol{\Gamma}^-}$ in Eq. (56), in addition to Eq. (55). The simpler form of the boundary term is because the corner domains have no horizontal sides, so $\frac{\partial \boldsymbol{g}_{n,\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} = 0$, without requiring further approximation. For the side domains (e.g., D_2 and D_4) and primary domain (D_1), the approximation of $\frac{\partial \boldsymbol{g}_{n,\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$, with $\frac{\partial \boldsymbol{g}_{n,\mathrm{d}}^{\mathrm{SIDE},0}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$, leads to the approximation of the effective boundary source as follows:

$$\boldsymbol{\Delta}\hat{\boldsymbol{g}}_{n,j}(\boldsymbol{r}, \boldsymbol{\Omega})|_{\boldsymbol{\Gamma}_n^-} \approx \boldsymbol{\Delta}\hat{\boldsymbol{g}}_{n,j}^0(\boldsymbol{r}, \boldsymbol{\Omega})|_{\boldsymbol{\Gamma}_n^-} = \mathcal{R}\left[\left.\frac{\partial \boldsymbol{I}_\odot}{\partial a_j}\right|_{\boldsymbol{\Gamma}_+^n}\right]$$
$$+ \frac{\partial \mathcal{R}}{\partial a_j}\left[\boldsymbol{I}_{n,\odot} + \boldsymbol{I}_{n,\mathrm{d}}\right] + \frac{\partial \boldsymbol{g}_n^{\mathrm{TOP}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j} + \frac{\partial \boldsymbol{g}_n^{\mathrm{BOT}}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}$$
$$+ \frac{\partial \boldsymbol{g}_n^{\mathrm{SIDE},0}(\boldsymbol{r}, \boldsymbol{\Omega})}{\partial a_j}. \tag{57}$$

With this, the approximate radiance derivatives in any domain at position $\boldsymbol{r}_{ik}$ and direction $\boldsymbol{\Omega}_{ik}$ are calculated as follows:

$$\left\langle \delta(\boldsymbol{\Omega} - \boldsymbol{\Omega}_{ik})\,\delta(\boldsymbol{r} - \boldsymbol{r}_{ik}), \frac{\partial \boldsymbol{I}_{n,\mathrm{d}}}{\partial a_j}\right\rangle_{D_n \times \mathbb{S}^2 \oplus \boldsymbol{\Gamma}_n^\pm}$$
$$\approx \mathcal{T}\left[\boldsymbol{\Delta}\hat{\boldsymbol{g}}_{n,j}^0(\boldsymbol{r}, \boldsymbol{\Omega})|_{\boldsymbol{\Gamma}_n^-}\right](\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik})$$
$$+ \mathcal{S}\left[\left.\boldsymbol{\Delta}\boldsymbol{f}_j\right|_{D_n \times \mathbb{S}^2}\right](\boldsymbol{r}_{ik}, \boldsymbol{\Omega}_{ik}), \tag{58}$$

with the appropriate recursive calculation of $\boldsymbol{\Delta}\hat{\boldsymbol{g}}_{n,j}^0$. The numerical implementation of these integrals in Eq. (58) is described in Appendix D. The computational cost of this calculation relative to the evaluation of radiances (Eq. 41) is presented in Appendix G.

As this series of approximations follows from a no-scattering assumption in the tangent-linear model, the error in the resulting radiance derivatives, and therefore forward-model derivatives, will depend on the relative contribution of higher orders of scatter to the radiance derivative at the positions and angles sampled by the sensor. When the single-scattering albedo is near unity, these contributions will be most significant. They will also be relatively large when the optical path between the source and sensor is large, as the radiance reaching the sensor will necessarily have undergone many scattering events. The approximation is clearly appropriate for emission problems without scattering or in the single-scattering limit, where it is also exact. For highly scattering solar transport problems, this approximation requires some justification. The approximation of the volume source of the radiance derivative term in Eq. (55) is the same as described in Levis et al. (2020). The treatment of the surface source of the radiance derivative term here is an extension from Levis et al. (2020), as $\boldsymbol{\Delta}\hat{\boldsymbol{g}}_j$ was assumed to vanish in

that formulation. We note that the treatment of the volume source of the radiance derivative term is identical to the formulation of Zawada et al. (2017), where it is appropriate in the quasi-single-scattering context of limb scattering.

In this section, we have described the tomographic retrieval methodology, including a full description of the forward model and its approximate linearization, which is utilized in AT3D for computationally efficient solutions to the retrieval problem. In the following section, we present a linear analysis of the conditioning of the inverse radiative transfer problem and other properties of the exact Jacobian matrix. This provides us with a more detailed framework for understanding the limitations of the cloud tomography method and how it will generalize to the full range of scattering regimes in the atmosphere. We will then apply this framework in Sect. 5 to understand, in more detail, the quantitative behavior of the approximate linearization described here.

4   Linearized inverse radiative transfer

Our analysis in this section is focused on describing the information content about the spatial variability in the optical properties contained within multi-angle measurements and therefore is focused on the inversion of the RT within a linearized context. We begin this analysis by first describing the conditioning of the exact Jacobian matrix, which we measure by its condition number (Eq. 6). Studying the exact Jacobian matrix provides a point of reference for understanding the effects of using approximate Jacobian matrix. Through the comparison, we can identify which limitations of tomography are physical and which are a result of using an approximate Jacobian. This section contains some results that have wider implications for tomographic retrievals and also some which are specific to the SHDOM solver. The implications of these results are discussed in Sect. 4.3.

The structure of the inverse problem and the nature of the Jacobian approximation used in AT3D can most easily be conceptualized if we use the equivalent adjoint formulation of the inner products, for reasons that we make clear below. This formulation will provide the basis for our presentation of the qualitative theory of inverse RT. Here we give only a brief summary of the adjoint formulation of the Jacobian calculations. More details on the formulation of the adjoint problem can be found in Martin et al. (2014) for 3D vector RT or in similar formulations for plane-parallel media (Hasekamp and Landgraf, 2005). In the adjoint formulation, the inner products in Eq. (51) are expressed in terms of an adjoint radiance field, whose sources are now the sensor response functions, $\boldsymbol{p}_i$ and $\boldsymbol{q}_i$ (instead of the Sun). The adjoint RT problem is solved using an adjoint solution operator $\mathcal{U}^*$, which is adjoint to the forward solution operator $\mathcal{U}$. In practice, the adjoint RTE is not directly solved. Instead, a pseudo-forward RTE problem is solved, whose source is the direction-reversed and polarization-flipped sensor response

functions $p_i^{\dagger}$ and $q_i^{\dagger}$:

$$p_i^{\dagger} = Q^{\dagger} p_i(r, -\Omega), \tag{59}$$

and

$$q_i^{\dagger} = \mathbf{Q}^{\dagger} q_i(r, -\Omega), \tag{60}$$

where

$$\mathbf{Q}^{\dagger} = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & -1 \end{pmatrix}. \tag{61}$$

The pseudo-forward problem is solved by the forward-solution operator $\mathcal{U}$ so it can, in principle, be solved by a forward model like SHDOM as long as its implementation supports sufficiently general volume and boundary source vectors. In practice, this pseudo-forward problem is much more numerically challenging due to the spatio-angular singularity of the sources $p_i^{\dagger}$ and $q_i^{\dagger}$ in 3D (Doicu and Efremenko, 2019; Martin and Hasekamp, 2018). The adjoint radiance field can then be found by direction reversal and polarization flipping of the pseudo-forward radiance $I^{\dagger}(r, \Omega)$.

$$I^{\dagger}(r, \Omega) = \mathcal{U} \begin{bmatrix} p_i^{\dagger} \\ q_i^{\dagger} \end{bmatrix} \tag{62}$$

We can then express the inner products for the forward-model derivatives which form the elements of the Jacobian matrix as follows:

$$\begin{aligned} \frac{\partial F_{i,n}(a)}{\partial a_j} &= \left\langle p_i, \frac{\partial I_d}{\partial a_j} \right\rangle_{D_n \times \mathbb{S}^2} + \left\langle q_{i,1}, \frac{\partial I_d}{\partial a_j} \right\rangle_{\Gamma_n^+} \\ &= \left\langle Q^{\dagger} I_i^{\dagger}(r, -\Omega), \Delta f_j \right\rangle_{D_n \times \mathbb{S}^2} \\ &+ \left\langle Q^{\dagger} I_i^{\dagger}(r, -\Omega), \Delta g_j \right\rangle_{\Gamma_n^+}. \end{aligned} \tag{63}$$

The spatial variations in the forward-model derivatives are controlled by the optical distance to the sensor, through $I_i^{\dagger}$, and also by optical distance to the solar source, through Δf_j. Given the singular form of the sensor response functions, the RTE problem for the pseudo-forward radiance is a pencil beam illumination problem. This is a well-studied RT problem (Doicu et al., 2020; Liemert and Kienle, 2013; Martelli et al., 2016), as its solution is Green's function for 3D RT. The approximate Jacobian in the 3D RT domain is equivalent to approximating the pseudo-forward solution by its direct-beam solution.

$$\begin{aligned} & \left\langle p_i, \mathcal{T} \left[\Delta g_j^0(r, \Omega)|_{\Gamma_1^-} \right](r, \Omega) + \mathcal{S} \left[\Delta f_j|_{D_1 \times \mathbb{S}^2} \right] \right\rangle_{D_1 \times \mathbb{S}^2} \\ & + \left\langle q_i, \mathcal{T} \left[\Delta g_j^0(r, \Omega)|_{\Gamma_1^-} \right](r, \Omega) + \mathcal{S} \left[\Delta f_j|_{D_1 \times \mathbb{S}^2} \right] \right\rangle_{\Gamma_1^+} \\ &= \left\langle Q^{\dagger} \mathcal{T}[p_i^{\dagger}](r, -\Omega), \Delta f_j \right\rangle_{D_1 \times \mathbb{S}^2} \\ & + \left\langle Q^{\dagger} \mathcal{T}[q_i^{\dagger}](r, -\Omega), \Delta g_j^0 \right\rangle_{\Gamma_1^+} \end{aligned} \tag{64}$$

With the forward-adjoint formulation described above, each of the columns of the Jacobian matrix corresponds to a different observing geometry and therefore to a different pencil beam pseudo-forward RTE solution. We can use this to express the independence of columns in the Jacobian matrix in terms of the independence of the pseudo-forward-radiance fields, weighted by the radiance derivative source vectors Δf_j and Δg_j. In this way, we link the properties of the Jacobian matrix, such as its condition number, which controls the difficulty of the retrieval problem, to the optical properties of the medium using arguments from RT theory. To be clear, we use the linearized framework to classify the difficulty of the iterative optimization process based on the current state vector. The linearization around the ground truth does provide some information on the difficulty of the inverse problem, but we must bear in mind that, given the non-convex nature of the inversion, the iterative retrieval may encounter a difficult cost function structure far from the ground truth, depending on the optimization trajectory. This linearized framework provides us with some hypotheses about the behavior of AT3D's tomographic retrieval that can be used to explain the behavior of fully nonlinear retrievals in terms of physical principles.

4.1 Numerically evaluating the Jacobian matrix

We support the physical arguments presented in this section with quantitative evidence from numerical experiments. The quantitative evidence is produced through the numerical calculation of reference Jacobian matrices using a two-point central difference around a wide range of media or base states. These reference Jacobian matrices are also used to quantitatively evaluate the approximate Jacobian calculation in Sect. 5. The numerical experiments that we use are as follows.

Each Jacobian matrix is calculated around a reference configuration of the state vector, which is referred to as a base state. Each base state is a 3D Gaussian extinction field embedded in a uniform extinction field evaluated on a grid with 50 m resolution and 21 grid points in each direction, hence primary domain dimensions $L_x = L_y = L_z = 1$ km. Open horizontal BCs are used. This configuration is chosen as a simple analogue of an isolated cloud with a sufficiently high resolution so that we can study how the elements of Jacobian matrix vary with position within the cloud. The extinction field at each grid point is given by the following:

$$\sigma(x, y, z) = \sigma_{\mathrm{bg}} + A \exp\left(\frac{(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2}{r^2} \right), \tag{65}$$

where $\sigma_{\mathrm{bg}} = 0.1$ km^{-1} is a uniform background extinction, $(x_0, y_0, z_0) = (L_x, L_y, L_z)/2$ is the center of the computational domain, and $r = 0.2$ km (i.e., four grid cells $= L_z/5$). The constant A (in km^{-1}) is varied so that the maximum vertical optical path of the Gaussian is 0.1, 5.0, 40.0, and 100.0. When we include the uniform background, which contributes an optical depth of $\sigma_{\mathrm{bg}} L_z$, the maximum vertical op-

tical paths of the medium $\tau_{\max} = \int_0^{L_z} \sigma(x_0, y_0, z) \, dz$ are 0.2, 5.1, 40.1, and 100.1. But, in the following results, we label our base states by the maximum optical path of the Gaussian.

Each base state has a spatially uniform single-scattering albedo and phase function and a uniform Lambertian surface albedo. The single-scattering albedos of the base states are 0.9, 0.99, and 1.0, and the surface albedos are 0.0, 0.2, and 0.7. Three different combinations of phase function and angular resolution are tested. The first uses an isotropic phase function with 16 zenith discrete ordinate bins and 32 azimuthal discrete ordinate bins. The second uses a Mie phase function equivalent to a gamma droplet size distribution with effective radius $r_e = 10 \, \mu m$ and effective variance $v_e = 0.1$ evaluated at a wavelength of $0.86 \, \mu m$ at the same angular resolution. The third configuration also uses the Mie phase function but uses a reduced angular resolution of only two zenith discrete ordinate bins and four azimuthal discrete ordinate bins. The different combinations of these parameters give a total of 108 different base states.

For each base state, we numerically evaluate the Jacobian matrix. We choose the observational sampling to consist of 33 different imaging sensors described by perspective projections that image the domain simultaneously from nadir and also 32 combinations of four zenith angles [75.0°, 60.0°, 45.6°, 26.1°] and eight relative azimuthal angles [0°, ±45°, ±90°, ±135°, 180°]. Each imaging sensor has a field of view of 6° and 26×26 pixels and points at the center of the domain from a distance of 10 km. The cosine of the solar zenith angle is set to 0.3 for all simulations. These observations thus capture both the diffuse radiance escaping the shadowed side of the cloud and the backscattering radiance from its illuminated side. We choose the state vector to be the extinction at every second grid point in each dimension to reduce the computational expense in the finite-differencing calculations. We also exclude the grid points at the domain boundaries that parameterize the open BCs from the state vector. The state vector is therefore of length 1100 ($11 \times 11 \times 10$). It is important to remember that ill-posedness (or instability or ill-conditioning) in an inverse problem is dependent on the choice of representation for the retrieved quantity. Here, we tackle a fully 3D representation with equidistant grid points throughout the cloud. We use this setup to demonstrate the issues that arise from a generalization of remote sensing to 3D that would not arise in, for example, a retrieval of only optical depth using 3D RT (e.g., Marchand and Ackerman, 2004).

Increments to the state vector for numerical evaluation of the Jacobian matrix are set as follows:

$$\Delta a_j = \mathrm{sgn}(a_j) \max(0.01 |a_j|, 0.01). \qquad (66)$$

This decision, and other numerical considerations of the finite differencing, is described and justified in Appendix F. We also describe our procedure for accelerating the finite-differencing calculations, which is based on a method discussed in Evans (1998) for accelerating multi-spectral SHDOM solutions. This method can lead to an acceleration of up to ~ 100 times for finite-differencing calculations with optically thick base states.

4.2 Results

For each of the base states, we evaluate the condition number (Eq. 6) of the reference Jacobian, which is shown in Fig. 5. We see increasing condition numbers with larger optical thickness and with a reduced phase function anisotropy and higher angular resolution. Larger condition numbers indicate the increasing instability of the linearized inverse problem. There is little variation in the condition number with single-scattering albedo over this range and almost no dependence on surface albedo. The condition number ranges from well conditioned $\kappa(\mathbf{K}) \sim 10^1$ to very ill-conditioned $\kappa(\mathbf{K}) \gg 10^5$. The maximum values reached $\kappa(\mathbf{K}) \sim 10^{16}$, which is extremely ill-posed, likely due to noise from the finite differencing. This transition towards ill-posedness is not a deficiency of the observing geometries, which can occur and have been documented elsewhere (Holodovsky et al., 2016), or resolution. The problem setting employed here has hemispherical observations, so all of the cloud is well observed. This is confirmed by the low condition number in the optically thin setting. Instead, the cause of the larger condition numbers is traceable to the nature of the radiation transport itself; i.e., it is largely inherited from the continuous RTE problem (Bal and Jollivet, 2008; Chen et al., 2018; Zhao and Zhong, 2019). As we increase the order of scattering, we increase the spatio-angular smoothness of the radiation field, and it therefore loses information about the spatial detail of the extinction field. This smoothing operation of the scattering leads to increasing ill-posedness under inversion, as quantified by the condition number.

Since the condition number of the Jacobian is expected to affect the accuracy and convergence rate of the tomographic retrievals (Sect. 3), it is important to understand the physical and numerical principles that lead to this behavior. We can explain the results in Fig. 5, using the forward-adjoint framework, by considering both the structure of the reference RTE solutions used to calculate Δf_j and the pseudo-forward RTE solutions, which together form the Jacobian elements. To characterize the reference and pseudo-forward RTE solutions, we will make use of the Knudsen number (Kn), which is the ratio of the mean free path (Davis and Marshak, 2004) to the domain length scale. This quantity can be used to define the ballistic ($Kn \gg 1$) and diffusion ($Kn \ll 1$) regimes of transport. In a homogeneous medium, Kn is simply the reciprocal of the optical depth but tends to be larger in a heterogeneous medium as transmission becomes more efficient (Davis and Marshak, 2004). Our explanations in the following paragraphs are simply descriptions of the qualitative behavior of scattering, supported by rigorous mathematical or numerical results where appropriate.
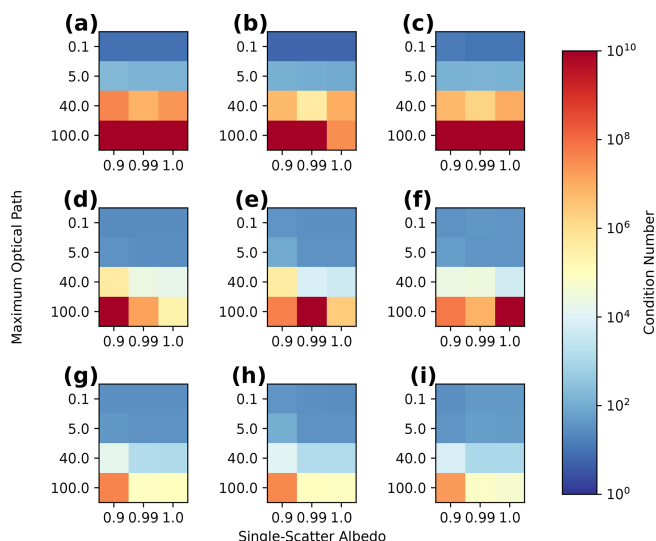
Figure 5. The condition number (see text for details) of the finite-differenced Jacobian for 3D Gaussian clouds with different combinations of single-scattering albedo and maximum optical path. Panels **(a)**, **(b)**, and **(c)** have isotropic phase functions. Panels **(d)**, **(e)**, and **(f)** have Mie phase functions, and panels **(g)**, **(h)**, and **(i)** have Mie phase functions with a reduced angular resolution (see text for details). Panels **(a)**, **(d)**, and **(g)** have a black surface. Panels **(b)**, **(e)**, and **(h)** have Lambertian surface albedos of 0.2, while panels **(c)**, **(f)**, and **(i)** have Lambertian surface albedos of 0.7. The crimson squares exceeding the color scale reach 10^{16}.

Figure 6. A conceptual diagram illustrating the decrease in magnitude of the Jacobian elements in the solar direction through a cross section of a homogeneous spherical cloud. The mismatch in the magnitude of the Jacobian elements between the illuminated and shadowed sides increases as the optical depth increases.

The Knudsen number depends on the definition of the system. In any partly cloudy system, we will have $Kn \sim 1$, due to the dominance of ballistic trajectories from the Sun to the surface. On the other hand, within individual clouds, we may have $Kn \ll 1$. We consider individual clouds as the system under consideration, and the mean free path used in the definition of the Kn is the domain average rather than at any particular position (Davis and Marshak, 2004). Kn is a classification of the optical properties, not of the transport, so it is common to both the reference and the pseudo-forward RTE and can be used to describe both. $Kn \ll 1$ does not guarantee that the diffusion approximation for radiative transport holds everywhere in the system but rather that there is an interior region where it holds to order Kn in the appropriately non-dimensionalized RTE problem (Chen et al., 2018).

In the ballistic limit ($Kn \gg 1$), both the reference RT problem, which controls the radiance derivative source vector Δf_j, and the pseudo-forward problems, which sample Δf_j to calculate the elements of the Jacobian matrix in Eq. (63), are dominated by their direct beams. The scattering is limited to low order. This means that Δf_j does not have a strong spatial dependence, as the optical paths are small enough that the exponential nature of transmission has not yet come into play. The pseudo-forward radiance, which has a pencil beam source, has a highly singular spatial distribution that is largely restricted to the line of sight of each sensor pixel. Be-
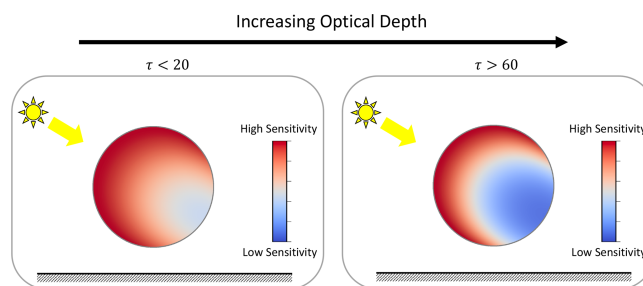
cause of this, the overlapping spatial support between Δf_j and the pseudo-forward radiance is highly localized in space. This means that the pseudo-forward radiances corresponding to different sensor pixels are highly independent and so are the columns of the Jacobian matrix. As a result, the measurements can resolve spatial variability at the pixel scale. This leads to the low condition numbers in the optically thinnest base states in Fig. 5, as the measurement resolution and the grid resolution are similar.

The pseudo-forward radiances and Δf_j are also highly anisotropic. The angular variation in the measurements therefore contains information about both the 3D variation in the phase matrix and also the 3D variation in the scattering coefficient. In the continuous inverse transport problem, this high degree of spatio-angular singularity in the pseudo-forward radiances in the ballistic limit enables the unique and stable inversion of both quantities simultaneously (Bal, 2009; Bal and Jollivet, 2008). The downside of the ballistic regime is that the scattering coefficient and phase matrix are not separable, and even high angular moments of the phase matrix must be retrieved to avoid errors in the retrieval of the scattering coefficient.

On the other hand, in the limit of $Kn \ll 1$, the radiation transport degenerates to a diffusion problem (Chen et al., 2018; Davis and Marshak, 2001). The radiance derivative sources Δf_j now span several orders of magnitude, with large values optically close to the Sun and much smaller values optically far from the Sun. In a homogeneous medium, the behavior in the diffuse limit is an exponential decay of the actinic flux as a function of the diffusion length from the source (Davis and Marshak, 2001). This leads to a scale mismatch in the magnitude of the Jacobian elements between the illuminated and shadowed sides of a cloud, as illustrated in Fig. 6. There is a corresponding trend towards decreasing anisotropy of Δf_j, with increasing optical distance from the Sun.

An example of the pseudo-forward-radiance solution for one of the optically thick base states examined here is shown
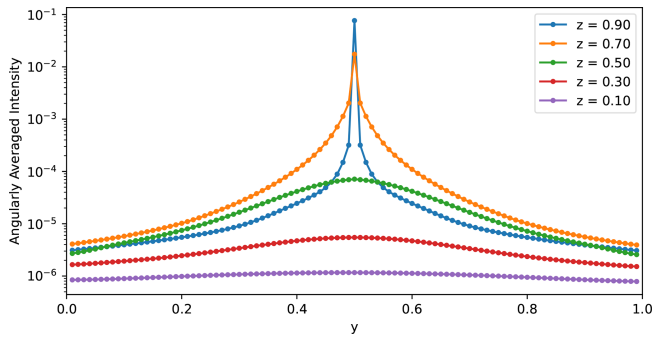
Figure 7. Cross sections of the angularly averaged intensity (actinic flux) for a pencil beam problem. The medium is the Gaussian extinction field with the maximal optical depth of 40.0, single-scattering albedo of 1.0, and the Mie phase function with full angular accuracy (see text for details). This simulation is done with SHDOM using an increased 101 points in each dimension (10 m resolution) to resolve the volumetric radiance field. The pencil beam source is located at the top of the center of the domain ($x = 0.5$, $y = 0.5$, and $z = 1.0$) and pointed at nadir. Each colored line shows a single transect of the angularly averaged intensity at a different altitude.

Figure 8. A conceptual diagram illustrating the decrease in magnitude of the Jacobian elements with distance from the sensors through a transect through a homogeneous spherical cloud. The mismatch in the magnitude of the Jacobian elements between the outer edges and the interior of the cloud increases as the optical dimension of the cloud increases.

in Fig. 7. For the pseudo-forward-radiance fields, near the pencil beam sources, the direct beam still dominates. With increasing optical depth from the source, the diffuse pseudo-forward radiance, which has undergone many more scattering events, begins to dominate. The diffuse radiance field is angularly smoothed, through repeated convolution with the phase matrix, and spatially smoothed through the streaming of the angularly smooth source fields. In a homogeneous medium, this manifests itself as an exponential decay of the angularly averaged intensity with distance from the source, which also causes a scale mismatch in the Jacobian elements, though this time between regions optically close and optically far from the sensors. The increasing spatial smoothness of the pseudo-forward radiance with optical depth indicates that the measurements become sensitive to only increasingly wide averages of the optical properties with increasing optical distance from the sensor. This causes ill-conditioning under inversion (Chen et al., 2018; Zhao and Zhong, 2019). The spatial smoothing effect of the wide pseudo-forward radiance is the three-dimensional or depth-resolved manifestation of the 2D radiative smoothing effect that has been widely studied, based on the consideration of a plane-parallel, homogeneous base state (Marshak et al., 1995; Davis et al., 1997).

The exponential decay of the pseudo-forward radiance is a feature that is unique to the tomographic problem of retrieving three-dimensional or depth-resolved spatial variability. This feature has been recognized in other tomographic applications utilizing diffuse light (Tian et al., 2010; Niu et al., 2010). This feature means that the sensitivity of measurements to changes in optical properties is rapidly lost with increasing optical distance from the sensor. While the pseudo-
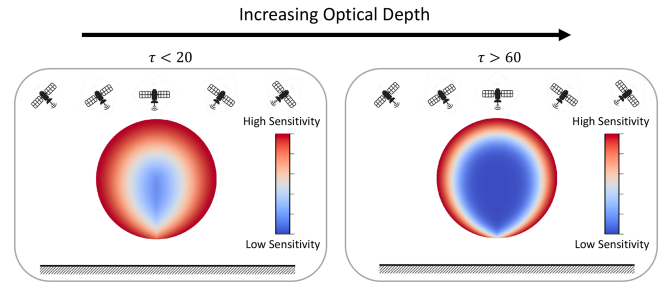
forward radiances corresponding to adjacent sensor pixels remain quite independent in the region optically close to the sensors where the radiance fields are localized, they become indistinguishable in the region optically far from the sensors due to smoothing and their decay to zero. The condition number measures this worst-case loss of independence (Chen et al., 2018), which occurs in the region which is optically far from the sensors and also from the Sun (through Δf_j). As the optical dimension of the medium increases, this worst-case loss of independence is exacerbated, causing the rapid growth of the condition number (Fig. 5). The exponential decay to zero of optical depth sensitivity causes a scale mismatch in the Jacobian elements between regions optically close and optically far from the sensors, similar to that which occurs in the forward problem, as illustrated in Fig. 8.

The two scale mismatches related to the sensors and the Sun that develop in the Jacobian matrix as the medium becomes optically thick are a significant source of instability in the Jacobian, as measurements will be orders of magnitude more sensitive to changes in optical properties in regions close to the Sun and sensors than those further away. This can be seen quantitatively by binning the absolute magnitude of elements of the reference Jacobian matrices by the solar delta-M transmission to each grid point and the delta-M transmission of the minimum optical path from all sensors to each grid point (Fig. 9). This latter quantity is referred to as the minimum sensor transmission. The delta-M transmission is derived from the path integral of the delta-M-scaled extinction, consistent with the calculation of direct transmission used in the SHDOM solver. This latter metric is a measure of overall optical distance from the sensors to each point in the cloud.

The optical path has also been used similarly to define a region of the cloud in which measurements are no longer sensitive to rearrangements of the small-scale features of the extinction field, known as the "veiled core" of the cloud (Forster et al., 2020). Forster et al. (2020) used a threshold of at least $\tau > 5$ along the line of sight of all sensors inter-

secting each volume element to define the volumetric extent of the veiled core of the cloud. This optical depth threshold is roughly equivalent to Delta-M-scaled minimum sensor transmissions of less than 0.1 for the simulations employed here using the Mie phase function and full angular accuracy. Given that the set of sensors covers the upper hemisphere, the minimum sensor transmission is loosely equivalent to minimum transmission to the edge of the domain. Figure 9 shows the increasing scale mismatch in derivative magnitudes between the optical interior and exterior as the total optical depth of the medium increases and the Knudsen number decreases. Figure 9d, h, and l show that derivatives of radiances with respect to changes in extinction within the veiled core can be almost 2 orders of magnitude smaller than those with respect to extinction changes in the exterior.

Note that it is the mismatch of scales between sensitivity to the interior and exterior that is important here, not the absolute smallness of the pseudo-forward radiance in the veiled core of the cloud (up to numerical precision). If the state vector were restricted to describing a region that only included similar transmissions from all sensors, then the problem of ill-conditioning would be much reduced. Such a partitioning is not generally available. The interior regions must be optically thin enough that the nonlinearity of the transmission causing the scale mismatch is small. In an optically thick cloud, this would require extremely detailed knowledge about the extinction field in the edge regions of the cloud, which is not generally available. For moderately opaque clouds, lidar may provide valuable information to constrain the outer portions of the cloud and thereby mitigate this issue.

The arguments presented so far explain the dependence of the difficulty of the inverse problem (measured by condition number of the Jacobian matrix) on the extinction field. They also explain the sensitivity of the condition number to low-order single-scattering properties such as the asymmetry parameter and single-scattering albedo, which modulates the diffusion lengths and transport mean free paths that control the spatial smoothing and exponential decay in the diffuse limit (Davis et al., 2021). They do not, however, explain the strong sensitivity of the condition number or spatial structure of the Jacobian matrix to the angular resolution of the SHDOM solver present in Figs. 5 and 9, respectively. We now consider this behavior.

In the diffusion limit, Davis et al. (2021) investigate the role of the asymmetry parameter in controlling the spatial smoothing and exponential decay of the forward and pseudo-forward problems using the theory of random walks. We can see from our results, which show the sensitivity of the Jacobian to angular resolution in Figs. 5 and 9, that, in addition to what is discussed in Davis et al. (2021), there is also a role for the higher-order moments of the phase function beyond the asymmetry parameter and that these interact with the spatial moments of the pseudo-forward radiance. In particular, we know that, in a medium with a large forward-scattering peak,

more radiance will stay angularly close to the direct beam after several scattering events, even if the average direction of propagation is lost rapidly due to backscattering. We can hypothesize that, even if the asymmetry parameter remains unchanged, a larger forward peak will skew the pseudo-forward radiance near to the direct beam, increasing its localization property and thereby reducing ill-conditioning.

This hypothesis is borne out in the approximate numerical framework of SHDOM, where the forward-scattering peak of a phase function is treated by the delta-M approximation. The larger the forward-scattering peak and the lower the angular resolution of the SHDOM solver, the more of the phase function is angularly unresolved by the model and is lumped together into the direct transmission. As such, a larger forward-scattering peak or lower angular resolution act to transform the medium to one with a higher effective Knudsen number and thereby improve its conditioning. For the Mie phase function considered here, the use of low angular accuracy causes an almost halving of the extinction compared to the high angular accuracy, which is itself also roughly a halving of the true extinction. The change in the condition number with a change in angular resolution of the SHDOM solver is substantial and indicates that the stability of the inverse problem depends on the discretization of the system.

4.3 Discussion

4.3.1 Physical implications

In Sect. 4, we have documented the behavior of a linearized tomography problem. A number of these results have general implications that are not specific to the SHDOM model or the use of the approximate Jacobian described here. We now take the time to consider the wider implications of these results. We have shown that the condition number of the inverse problem largely depends on the Knudsen number or optical size of the medium, as supported by theory. We should therefore expect the convergence rate of an iterative retrieval to decrease in optically thicker media, as discussed in Sect. 3. As such, the tomographic retrieval of optically thicker media is expected to be computationally more expensive due to both RTE solutions becoming more expensive in optically thicker media and the need for more iterations of optimization to achieve a user-specified level of accuracy. High levels of retrieval accuracy may not be obtainable in optically thick media due to the extreme ill-conditioning, possibly causing slower convergence than the stopping condition of an optimization procedure, while still far from the optimal solution.

We have also shown the existence of substantial spatial variability in the linear sensitivity of radiances to changes in the extinction field in optically thick clouds. In particular, linear sensitivity decreases exponentially with optical depth from the Sun and from the sensors, likely causing slow convergence of the extinction field in these regions. This feature was also noted in Levis et al. (2015). We have shown
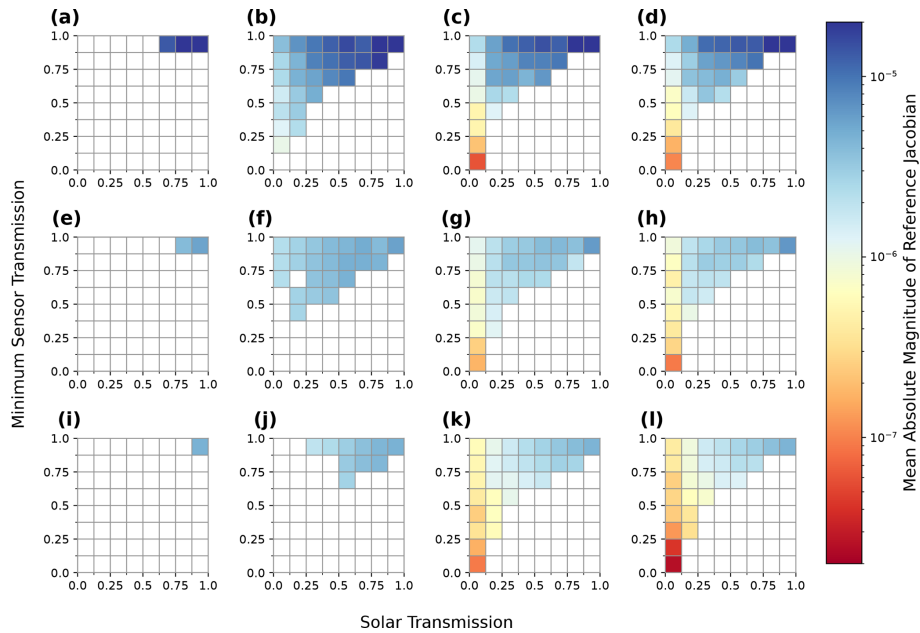
Figure 9. The mean absolute magnitude of finite-difference Jacobian for 3D Gaussian clouds in each bin of solar and minimum sensor delta-M transmission. Each column corresponds to a base state, with a maximum cloud optical depth from 0.1, 5.0, 40.0, and 100.0, increasing from left to right. Each row corresponds to a different combination of phase function and angular accuracy with isotropic phase function, Mie with full angular accuracy, and Mie with a reduced angular accuracy in descending order. All base states have a black surface with conservative scattering.

here that it is, at least in part, a result of a physical limitation and not just the approximations used within their paper. The fastest way to decrease the misfit with the measurements will be to change the extinction field optically close to the sensors and also the Sun. If iterative retrievals of optically thick clouds are unconstrained apart from measurements, then the retrieved extinction field may approach a local minimum with little change from the initialization in the cloud center and on the shadowed side of the cloud. This behavior may introduce a solar-zenith-angle dependence to the retrieval error, despite the use of 3D radiative transfer.

The issues apparent in optically thick clouds appear to substantially limit the applicability of the method, but we must bear in mind that, in terms of both number and area, a large portion of trade cumulus cover comes from small clouds (Zhao and Di Girolamo, 2007). Many trade cumulus tend to be smaller than 800 m in geometric depth (Chazette et al., 2020; Guillaume et al., 2018), and the average adiabatic fractions for these clouds can be significantly less than unity (Eytan et al., 2022). Most of these clouds will have maximum optical depths less than 40, which suggests that prior information or regularization will not be essential for ensuring high-fidelity retrievals of these clouds. We have not so far considered the feasibility of tomography in other, more stratiform, cloud types. We explore the sensitivity to such differences in Sect. 6, using the more computationally efficient approximate Jacobian calculation.

4.3.2　Numerical implications

We have also identified an important sensitivity of the ill-conditioning of the retrieval to the numerical discretization of the method. Of course, ill-conditioning is always sensitivity to discretization choices. For example, if we were to only retrieve a single unknown per column with parameterized vertical variability, then the condition number of the corresponding Jacobian matrices in the optically thickest cases considered here would be of the order of 10, which is the same as the optically thinnest tomography problems. This distinction is not unique to our study.

The important behavior that we have documented here for the first time is the importance of the angular discretization of the forward model for determining the conditioning of the model, rather than standard properties such as the number and spacing of measurements or the spatial resolution. The sensitivity of ill-conditioning to angular discretization arises from the presence of strongly peaked phase functions and the use of the delta-M approximation, which reduces the effective optical depth of the medium to account for strong forward scattering that is unresolved by the angular discretization.

The large decrease in condition number observed when decreasing angular accuracy suggests that using low angular resolution may be beneficial in inversions. The tradeoff for forming a better-conditioned inverse problem in this way is that the forward model is a poorer approximation for re-

ality and may in fact have significant biases (Evans, 1998). Even if a benefit in the convergence rate is not apparent, then these results indicate that retrieval results should not be expected to generalize to other angular resolutions, even when the angular resolution is high enough that the forward model has converged in accuracy. In particular, inversions may actually decrease in fidelity as the angular resolution is increased. Additionally, results should not be expected to generalize between phase functions with substantially different forward-scattering peaks such as Mie vs. Henyey–Greenstein phase functions, even if they have the same asymmetry parameter.

The conditioning results presented here are likely only representative of other explicit RTE solvers like SHDOM utilizing the delta-M approximation. It is unclear how the results will generalize to other widely used methods of solving the RTE, such as Monte Carlo. Those state-of-the-art Monte Carlo solvers in atmospheric radiative transfer that do use the delta-M or phase function truncation approximations may have some similar dependence of their conditioning on the truncation fraction, as documented here for SHDOM. However, modern implementations of phase function truncation tend to be dependent on scattering order (Wang et al., 2017) to avoid strong bias and so may not express the behavior demonstrated by SHDOM.

5 Validating the approximate Jacobian calculation

With the theory introduced in the previous section in place, we can examine the consequences of the approximation to the Jacobian (Eq. 58) used in AT3D and how that approximation exhibits itself as quantitative errors. We examine how well the approximate Jacobian reproduces the behavior described in Sect. 4. As described in Sect. 3.3, we are approximating the pseudo-forward radiance by its direct beam (Eq. 61). In this case, we can see that the entire diffuse pseudo-forward-radiance profile (Fig. 7) will be neglected. This means that the approximate Jacobian does not represent the non-local sensitivity of measurements to changes in optical properties outside of their field of view, other than through changes to the direct solar transmission. This may seem a rather extreme approximation, but we must bear in mind that the most stable information to extract is contained in the highly localized direct-beam component (Bal and Jollivet, 2008). The success of the approximation requires that the pixel-to-pixel smoothing only becomes important and significant when the medium is optically thick enough so that there will be significant decay in the sensitivity with distance from the sensor. In this case, the loss of sensitivity to the cloud interior from all pixels (in the linearized setting) is much more important than neglecting the pixel-to-pixel smoothing. In this section, we test the extent to which this is true quantitatively.

From the theory developed in the previous section, we also expect a sensitivity of the approximate Jacobian to

the dimensionality of the RT problem. In 3D, the pseudo-forward-radiance field is highly singular in space (Fig. 7) and anisotropic. On the other hand, as the dimensionality of the transport decreases, the pseudo-forward solution will become increasingly less singular. For example, in 1D, the pseudo-forward source is a plane illumination, just like the Sun (Hasekamp and Landgraf, 2005). The symmetry of the source means that the pseudo-forward radiance does not disperse perpendicular to the collimated source, but instead, it substantially modifies the depth dependence of the pseudo-forward-radiance field and reduces its anisotropy. This means that the direct-beam approximation will perform worse, and the performance of the approximation to the Jacobian calculation adopted here cannot be expected to generalize from 3D to 2D or, especially, 1D problems. This has an important implication for our extension of the approximate Jacobian to the linearization of the system of RTEs described in Sect. 3 to model a 3D domain embedded in a plane-parallel atmosphere. As such, we separate our error analysis between the elements of the approximate Jacobian corresponding to the primary 3D domain D_1 (Fig. 2) and the other elements controlling the boundary radiances through the auxiliary RTE problems.

Note that, in the following analysis, we quantitatively validate derivatives of radiances only with respect to volume extinction coefficient at different grid points rather than other optical properties such as the single-scattering albedo or components of the phase matrix. The approximation to the Jacobian only approximates the pseudo-forward radiance and therefore is common to derivatives of radiances with respect to all optical properties. However, each optical property interacts differently with the pseudo-forward-radiance field. For example, the single-scattering albedo interacts with all spherical harmonics of the radiance field, while the extinction coefficient interacts with all except the isotropic component. As such, errors in radiance derivatives with respect to single-scattering albedo may be disproportionately affected by error in the pseudo-forward radiance that occur optically far from the sensors compared to derivatives with respect to extinction. In contrast, the higher-order expansion coefficients of the phase matrix are only sensitive to the high-order spherical harmonics of the radiance field and will therefore be disproportionately affected by errors in the pseudo-forward radiance that occur optically close to sources. As a result, Jacobian errors may be distributed slightly differently in space and observation angle as the angular structure of Δf_j changes with the nature of the optical (or microphysical) unknown. We focus solely on extinction to illustrate how errors in the approximate pseudo-forward radiance propagate to the Jacobian as the retrieval of extinction is fundamental for retrievals of spatial structure.

To quantify the agreement between the finite-differenced Jacobians defined in Sect. 4.1 and the equivalent approximate Jacobians ($\widetilde{\mathbf{K}}$), we use the relative Frobenius error, which is

an element-wise relative root mean square error (RMSE):

relative Frobenius error

$$= \frac{\left\|\widetilde{\mathbf{K}} - \mathbf{K}\right\|_F}{\left\|\mathbf{K}\right\|_F} = \frac{\sqrt{\sum_{ij}\left(\widetilde{\mathbf{K}}_{ij} - \mathbf{K}_{ij}\right)^2}}{\sqrt{\sum_{ij}\left(\mathbf{K}_{ij}\right)^2}}. \quad (67)$$

Our validation approach of comparing the approximate Jacobian to the finite-differenced Jacobians measures the consistency of the forward model with the approximate Jacobian. The consistency between these two is what is important for the robustness of the local optimization method used in the tomographic retrieval implemented in AT3D. As such, the relative Frobenius error includes the effects of computational noise in the forward model that will emerge in the finite-differenced derivatives. When evaluating derivatives calculated using a forward-adjoint principle, the relative Frobenius error was around 0.04, reflecting uncertainties in the finite differencing and inconsistencies between the forward and adjoint models (Doicu et al., 2020). We expect a similar degree of accuracy in the optically thin limit for the approximate Jacobian. We performed extensive verification of the approximate Jacobian. We identified a number of instances where the computational noise in the finite-differenced derivatives is actually the dominant source of error, based on comparison against an analytic solution (Appendix E). For that reason, it is important to keep in mind that this error metric is a measure of the consistency of the approximate Jacobian with the forward model and that the stability of the forward model itself is also contributing.

5.1 Primary domain

For the primary domain, we are analyzing the accuracy of the approximate Jacobian across the same base states as in Sect. 4. We also define the state vector in the same way, so we are not analyzing derivatives with respect to the open BCs. We can see in Fig. 10 that the error rapidly grows beyond a benchmark value of 0.02 as the clouds become optically thicker, scattering is closer to conservative, and surface albedos become larger. The small change in the error between a black surface and a Lambertian surface, with an albedo of 0.2, in Fig. 10 shows that there is little sensitivity of the error to the surface albedo when it is not too reflective. This indicates good suitability of the approximate Jacobian for media over oceanic or other dark surfaces. We also see greater agreement for more forward-scattering media, especially at lower angular resolution. This is because the approximate Jacobian treats the direct beam exactly; hence, the more energy within the direct beam due to delta-M scaling, the more accurate the approximation. A lower angular accuracy therefore gives the benefit of more consistent derivatives with the forward model, with the downside being that the forward model will have larger errors against reality. To better understand these systematic differences between the approximate and
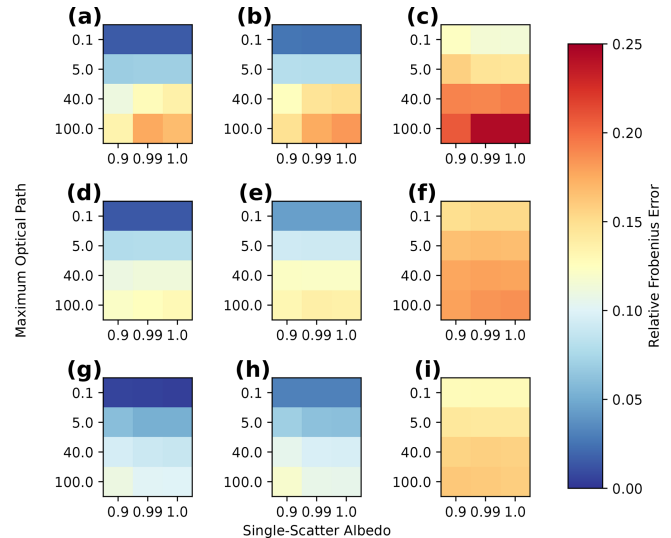

Figure 10. The relative Frobenius error (Eq. 67) of the approximate Jacobian with respect to the finite-differenced Jacobian for the same 3D Gaussian clouds, as used in Fig. 5. Panels **(a)**, **(b)**, and **(c)** have isotropic phase functions. Panels **(d)**, **(e)**, and **(f)** have Mie phase functions, and panels **(g)**, **(h)**, and **(i)** have Mie phase functions with a reduced angular resolution (see text for details). Panels **(a)**, **(d)**, and **(g)** have a black surface. Panels **(b)**, **(e)**, and **(h)** have Lambertian surface albedos of 0.2, while panels **(c)**, **(f)**, and **(i)** have Lambertian surface albedos of 0.7.

reference Jacobian matrices, we also calculate how the errors are distributed in space and angle, i.e., in state and measurement space.

5.1.1 Spatial variation in error

We again categorize the Jacobian elements according to the delta-M transmission from the Sun to each grid point and the minimum delta-M transmission from the sensors to each grid point. For Jacobian elements in each bin, we calculate the RMSE and normalize it by the root mean square magnitude of the entire reference Jacobian matrix calculated by finite differencing. This indicates which grid points produce approximate Jacobian elements with errors large enough to significantly change the overall direction of the gradient. Figure 11 shows that these errors are largest in the regions at the exterior of the cloud close to both the sensor and Sun. These Jacobian entries are also the largest in magnitude and have the largest higher-order derivatives, due to the curvature of the transmission function, and are also expected to have the largest errors in the finite differencing.

We also show the ratio of the mean absolute magnitude of the approximate Jacobian to the mean absolute magnitude of the reference Jacobian in each bin (see Fig. 12). We see that the typical magnitude of the approximate Jacobian decays much quicker with optical depth from the Sun and sensors than the reference. This means that the approximate Ja-

cobian has an enhanced scale mismatch in sensitivity to the properties of the cloud at the exterior and interior, thus exacerbating the ill-conditioning problem outlined in Sect. 4.2. Physically, this is due to the faster exponential decay of direct transmission than diffuse radiance. We could then hypothesize that the condition number of the approximate Jacobian would be higher than that of the reference in the optically thickest cases. This is not borne out in Fig. 13, which shows that condition numbers for the approximate Jacobian are not appreciably larger than for the reference Jacobian (see Fig. 5) and are actually smaller in the optically thick limit. Thus, the condition numbers in Fig. 5 for the reference Jacobian are likely larger in the optically thick limit due to numerical noise in the derivative calculations. The patterns of error illustrated in Figs. 11 and 12 are similar for the non-zero surface albedos and non-conservative scattering except with overall larger errors (not shown).

5.1.2 Angular variation in error

We also examined the sensitivity of the Jacobian errors to the set of scattering angles in the observations. In Fig. 14, we show the relative RMSE in the Jacobian elements corresponding to each viewing angle, grouped by scattering angle. The shape of the error depends on the phase function. In general, the larger errors occur in the backscattering directions for the isotropic phase functions (Fig. 14a, b, c, d). These observation geometries also include large sensitivity to grid points that are optically close to the Sun and therefore have the largest truncation errors, consistent with Fig. 11. These truncation errors may account for a substantial portion of the scattering angle dependence. When a Mie phase function is used (Fig. 14i, j, k, l), the error is also angularly dependent, with a minimum around the rainbow direction and a maximum at scattering angles of around 100°.

We performed a further investigation of these angular patterns in the error using much simpler clouds with hyper angular observations. We defined cloud base states that consisted of just a single, optically thin, cloudy grid point and observations with 1° zenith angle spacing both along and across the solar plane. We confirmed that the dip in error observed around the rainbow scattering angles in Fig. 14 is due to the Mie phase function, as this feature does not appear when a featureless Henyey–Greenstein phase function is used with the same asymmetry parameter. There is further angle dependence in the error that is not correlated with scattering angle, which we can attribute to multiple-scattering effects. We do not expect a consistent angular dependence for these features. For example, a peak in error at scattering angles of around 100° appears for our single grid point clouds, regardless of the choice of a Mie or Henyey–Greenstein phase function. Changing the solar zenith angle for our simplified single grid point clouds revealed that this error is actually dependent on zenith, peaking near nadir. We should expect the angular dependence of the error to vary with the base state

cloud structure and that it cannot easily extrapolate from our simplified cloud to the more complex base states quantified in Fig. 14. However, we have identified that there are angular features in the error that depend only on the phase function due to the increasing relevance of single-scattered radiation in the rainbow scattering angles. Further work with a wider diversity of cloud base states would be needed to reliably isolate multiple-scattering signals in the error that are independent of the phase function details.

5.2 Auxiliary domains

Here, we assess the accuracy of derivatives of radiances with respect to changes in the extinction fields in the auxiliary RTE domains that control the open horizontal BCs. We use the same set of base states and observations as in Sects. 4.1 and 5.1. However, we focus only on the appropriateness of the approximate Jacobian for open horizontal BCs. We compute derivatives with respect to extinction along the horizontal boundaries for every fourth point in the vertical and every fifth point in the horizontal. This set of base states has a low optical depth of 0.1 in the auxiliary domains and is therefore analogous to the case of an optically thin embedding medium, such as the cloud-free atmosphere.

These errors in the approximate Jacobian are displayed in Fig. 15 and show much larger overall errors than for the internal extinction derivatives shown in Fig. 9 (note the difference in the color scale). The larger errors in the optically thin cases, when compared to Fig. 10, are driven by the poorer applicability of the Jacobian approximation to the 2D and 1D RT in the auxiliary domains. The much weaker dependence of the errors in the optical thickness of the primary domain of 3D RT than in Fig. 10 indicates that the neglect of multiple scattering within the primary domain is much less important than the errors due to the application of the Jacobian approximation to the lower-dimensional auxiliary domains. There is a much greater sensitivity to the phase function and angular accuracy, with nearly a halving of the error when moving from the isotropic phase function and Mie phase function to the two-stream Mie phase function.

We also examine another set of base states which are just plane-parallel cloud layers (Fig. 16). The domain and discretization for 3D RT are the same as in Sects. 4.1 and 5.1, and the horizontally infinite portion is modeled using the open horizontal BCs rather than periodic assumptions. The extinction is distributed homogeneously within the $(1\,\text{km})^3$ domain. These base states are analogous to stratiform clouds. We calculate derivatives only with respect to the domain boundary extinction to examine the appropriateness of the approximate Jacobian for these situations. We vary the optical depth of the plane-parallel layers across the same values used previously for the maximum optical depth in the 3D Gaussian extinction fields. The errors in the approximate Jacobian increase much faster with optical depth than for the 3D Gaussian extinction fields (Fig. 10) and are much larger.
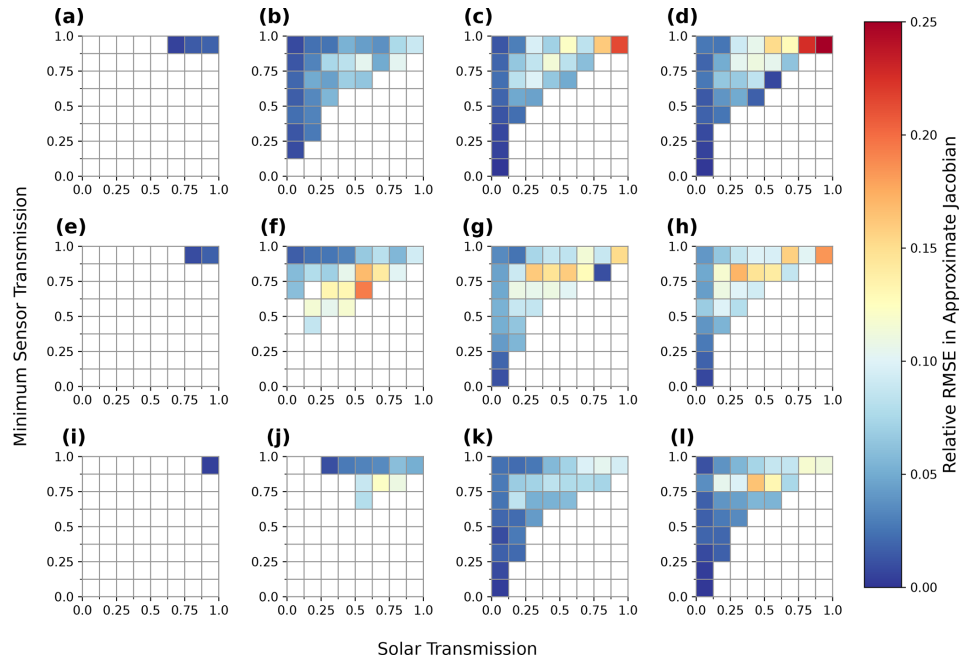
Figure 11. The relative RMSE error in the approximate Jacobian for 3D Gaussian clouds in each bin of solar and minimum sensor delta-M transmission (see Fig. 9 and associated text for definitions). The normalization of the RMSE in each bin is the root mean square magnitude of the entire Jacobian matrix (see Eq. 67). As in Fig. 9, each column corresponds to a base state with maximum cloud optical depth from 0.1, 5.0, 40.0, and 100.0, increasing from left to right. Each row corresponds to a base state with different phase function and angular resolution, with isotropic phase function (**a, b, c, d**), Mie phase function with high angular resolution (**e, f, g, h**), and Mie phase function with low angular resolution (**i, j, k, l**).
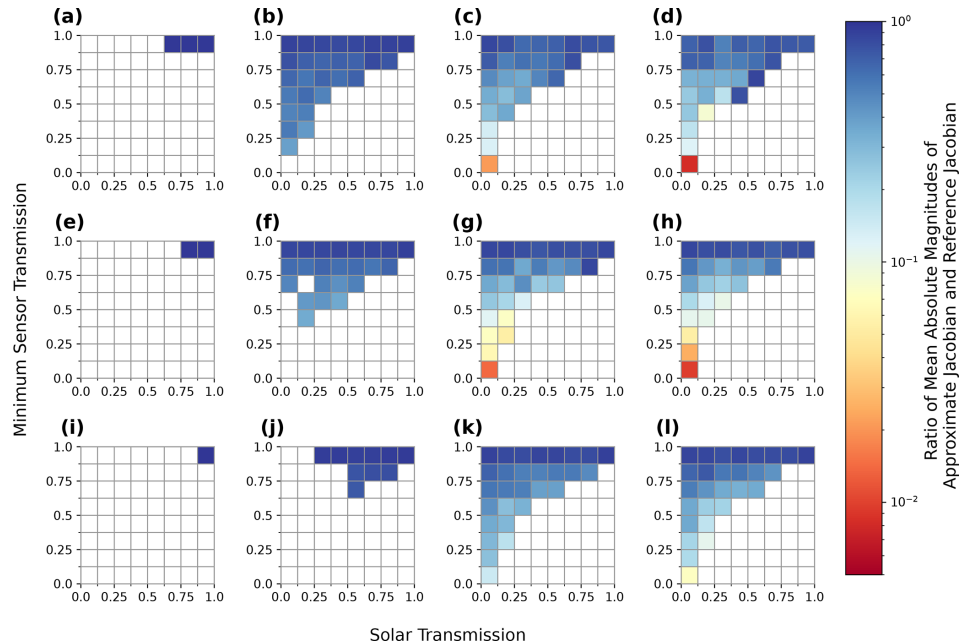
Figure 12. The ratio of the mean absolute magnitude of the approximate Jacobian to the mean absolute magnitude of the finite-difference Jacobian in each bin of solar and minimum sensor delta-M transmission for 3D Gaussian clouds with conservative scattering over a black surface. This figure is the same as Figs. 9 and 11. As in Fig. 9, each column corresponds to a base state, with maximum cloud optical depth from 0.1, 5.0, 40.0, and 100.0, increasing from left to right. Each row corresponds to a base state with different phase function and angular resolution, with isotropic phase function (**a, b, c, d**), Mie phase function with high angular resolution (**e, f, g, h**), and Mie phase function with low angular resolution (**i, j, k, l**).
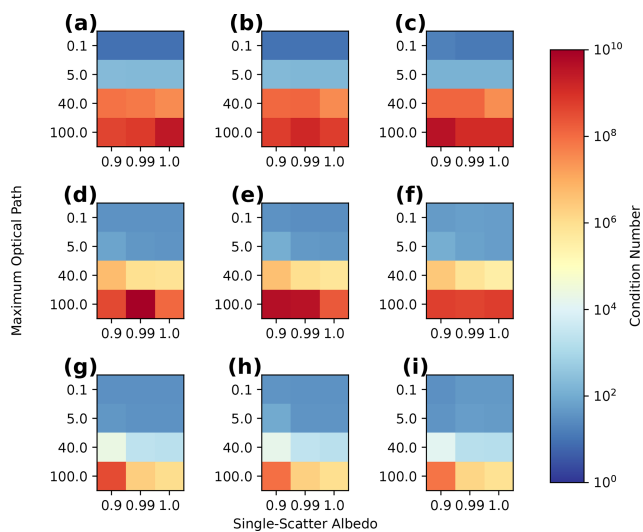
Figure 13. The condition number (see text for details) of the approximate Jacobian for 3D Gaussian clouds with different combinations of single-scattering albedo and maximum optical path. Same as Fig. 5 but for the approximate Jacobian. Panels **(a)**, **(b)**, and **(c)** have isotropic phase functions. Panels **(d)**, **(e)**, and **(f)** have Mie phase functions, and panels **(g)**, **(h)**, and **(i)** have Mie phase functions with a reduced angular resolution (see text for details). Panels **(a)**, **(d)**, and **(g)** have a black surface. Panels **(b)**, **(e)**, and **(h)** have Lambertian surface albedos of 0.2, while panels **(c)**, **(f)**, and **(i)** have Lambertian surface albedos of 0.7.

This indicates that the approximate Jacobian has no skill for BC optimization in conditions of stratiform cloud.

5.3 Discussion

Here, we discuss the implications of the errors in the approximate Jacobian for the iterative retrieval. The relative Jacobian errors documented above bound the relative errors that can occur in the calculation of cost function gradients, which additionally depend on the structure of the measurement residuals. The theoretical examination of the consequences of gradient errors in the L-BFGS-B method have been only recently investigated (Shi et al., 2021), due to the interest in developing stochastic variants for deep learning and other similar applications. The L-BFGS method uses finite differencing to approximate the Hessian. When there is noise in the gradients, the approximate Hessian can become corrupted. With bad curvature information, typically only very small step sizes will be valid, or there may be a complete failure to select a valid search direction. This would result in the early termination of the optimization, possibly far from a local optimum, and will become more significant as the errors in the approximate Jacobian increase, i.e., as clouds become optically thicker. In this sense, both the inherent ill-conditioning of the inverse problem and the approximate Jacobian errors should have a similar deleterious effect on retrieval performance. Moreover, it will not be possible to

disentangle these two effects in nonlinear retrievals without comparison against a reference method that uses an unapproximated method to linearize the forward model, whether it is a forward or forward-adjoint method. As such, we cannot make quantitative statements about the consequences of using the approximated Jacobian without performing the optimization with a known ground truth. The combined effects of the approximate Jacobian and ill-conditioning on retrieval accuracy can be examined in idealized circumstances where other sources of uncertainty in the retrieval are minimized. We perform such simulations in Part 2 of this study.

The typical solution for gradients with errors is to perform a step-lengthening procedure (Shi et al., 2021). However, in systems that are highly ill-conditioned, even without noise, such as the optically thick clouds discussed here, a step-lengthening procedure may cause significant difficulty in the selection of an update vector satisfying the stabilizing Wolfe–Armijo line search conditions. We note that the errors in the gradients induced by the approximate Jacobian are deterministic, and not stochastic noise, so there is an opportunity for them to be highly correlated from one base state to the next. If this occurred, then it would result in a certain amount of cancellation of the errors and better approximation of the Hessian of the cost function through the L-BFGS method. We have not examined this quantitatively here due to the computational expense of numerically calculating second-order derivatives, but this should be kept in mind when considering differences in performance of the approximate Jacobian for tomographic retrievals when used with first-order optimization methods (e.g., gradient descent) vs. quasi-Newton methods such as L-BFGS-B.

For the auxiliary domains, overall errors are much larger, reaching relative Frobenius errors far in excess of 100 %, even for optically thin atmospheres (with isotropic phase functions). Errors are relatively independent of the scattering regime of the internal medium but are very sensitive to the optical depth in the auxiliary domains. These results indicate that, while the approximate Jacobian proposed in Levis et al. (2020) is appropriate for 3D media, it is much less so for lower-dimensional transport. This indicates that a retrieval of stratiform cloud properties using the approximate Jacobian and open horizontal BCs is ill-advised. The ability to optimize BCs using AT3D may still be useful for retrieving a best-fitting cloud-free atmosphere jointly with the retrieval of a 3D cloud field, as the cloud-free atmosphere is optically thin.

6 The importance of heterogeneity for tomography

We can make use of the computational efficiency of the approximate Jacobian to explore the dependence of the condition number of the Jacobian on the spatial structure and optical thickness of the cloud field in more detail than in Figs. 5 and 13. In particular, we illustrate the critical importance of
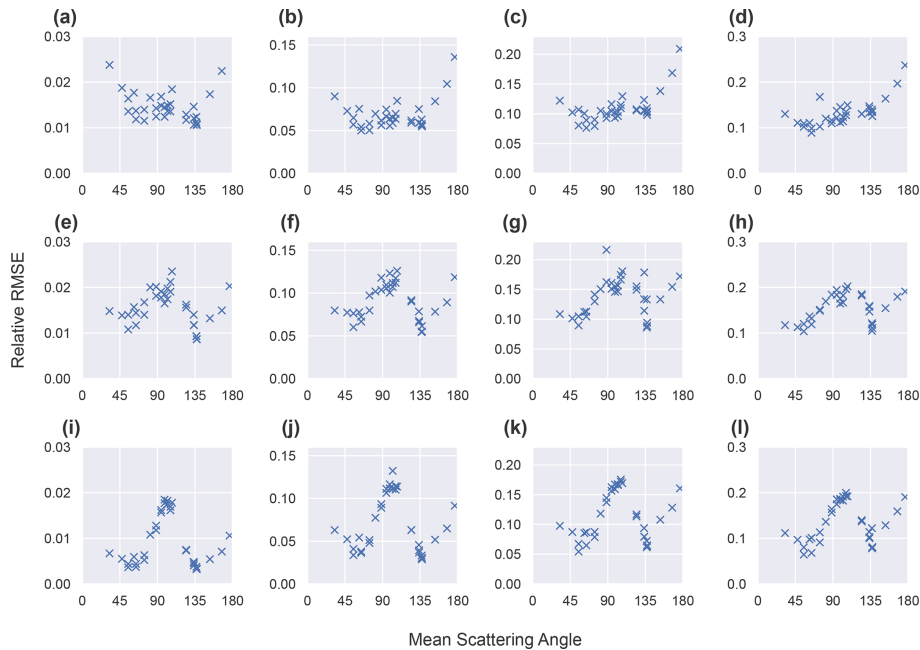
Figure 14. The dependence of the relative RMSE in the approximate Jacobian for each image on the mean scattering angle of each image for cloud base states with conservative scattering and a black surface. As in Figs. 9, 11, and 12, each row corresponds to a base state with different phase function and angular resolution, with isotropic phase function (**a, b, c, d**), Mie phase function with high angular resolution (**e, f, g, h**), and Mie phase function with low angular resolution (**i, j, k, l**). Columns correspond to base states with maximum optical depths of 0.1 (**a, e, i**), 5 (**b, f, j**), 40 (**c, g, k**), and 100 (**d, h, l**).

the spatial structure of the extinction field for determining the feasibility of tomography. We contrast the behavior of plane-parallel homogeneous clouds and 3D Gaussian extinction fields, which are restricted to a $(1 \text{ km})^3$ domain under inversion. Note that the plane-parallel homogeneous cloud is still resolved in 3D, and the radiation transport is in 3D. We are examining our ability to retrieve inhomogeneities embedded within the cloud. As such, the dimensionality effects described in Sect. 5 are not at play here. We use the same hemispheric set of observations and the same grid geometry used in Sects. 4.1 and 5.1. The clouds are conservatively scattering and use the Mie phase function with a black surface. We calculate the derivatives with respect to all internal grid points in each horizontal grid point but only every second grid point in the vertical. This is done to avoid the need to adopt special techniques for the singular-value decomposition (SVD) of large matrices to calculate the condition number.

We classify the plane-parallel clouds by their vertical optical depth and the 3D Gaussian extinction fields by their vertical optical depth to the center of the cloud (half of their diameter). We refer to this as the "optical dimension". This classification puts the two types of extinction field on the same footing in terms of the minimum transmission from any sensor to the grid point that is optically furthest from all sensors. They are therefore equivalent in terms of the presence of a veiled core, as defined by Forster et al. (2020) and investigated in Sect. 4.2.

6.1 Results

In Fig. 17, we can clearly see the exponential growth of the condition number of the approximate Jacobian at larger optical dimensions, consistent with theory, indicating exponential growth with the inverse Knudsen number (Zhao and Zhong, 2019). The plane-parallel clouds are notable in that the condition number increases at a faster rate. This shows how much larger the Knudsen number (and hence mean free path) is in a heterogeneous cloud like the 3D Gaussian extinction field and how much more information about spatial detail is preserved for a similar optical thickness. The importance of finite cloud edges for sensing cloud vertical structure from multi-angle radiances has also been demonstrated in nonlinear retrievals in a 2D setting (Martin and Hasekamp, 2018). The purely geometric part of this effect is partially reflected in the use of the optical radius, rather than diameter of the 3D Gaussian extinction field, when comparing to the plane-parallel clouds. This reflects the ease with which oblique sensors can constrain the cloud base and edges when clouds have aspect ratios around 1, even when they have large vertical optical depths.

In Fig. 17b, we see that the condition number of the approximate Jacobians in the 3D Gaussian clouds remains roughly invariant before the onset of exponential growth. This is likely due to the fact that the Knudsen number remains large enough that a diffusion regime has not developed
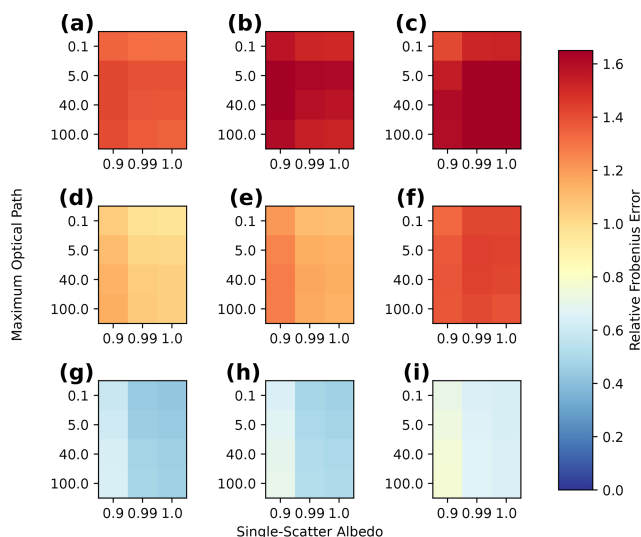
Figure 15. Similar to Fig. 10 but for derivatives of radiances with respect to the extinction within the auxiliary domains that parameterize the open horizontal BCs. Note the difference in color scale from Fig. 10. Panels (a), (b), and (c) have isotropic phase functions. Panels (d), (e), and (f) have Mie phase functions, and panels (g), (h), and (i) have Mie phase functions with a reduced angular resolution (see text for details). Panels (a), (d), and (g) have a black surface. Panels (b), (e), and (h) have Lambertian surface albedos of 0.2, while panels (c), (f), and (i) have Lambertian surface albedos of 0.7.

within the cloud. The reduced angular accuracy results indicate that the clouds are equivalent to those operating with almost halved extinction, as expected. This includes the transition point from slow scaling to exponential scaling for the 3D Gaussian extinction field.

6.2 Discussion

The results in Fig. 17 indicate that plane-parallel clouds are a lot more ill-posed under inversion than finite clouds and will require stronger regularization of prior information. This fact highlights the fact that, from a fundamental perspective, highly heterogeneous cloud fields are actually much simpler targets for remote sensing than homogeneous, stratiform clouds, as the vertical variations within the cloud can be inferred much more easily from passive imagery. Additionally, the results in Fig. 17 also indicate that tomographic retrievals of optically thin ($\tau \leq 3$) stratiform media (e.g., cirrus) should also be effective, as these optically thin cirrus clouds are better conditioned than the Gaussian clouds with optical depths typical of thin cumulus (optical dimension ~ 10). Tomographic retrievals of cumuliform clouds in this optical depth range have already demonstrated some success (Levis et al., 2020). This analysis alone does not demonstrate that tomographic retrievals will be successful but does indicate that the proposed retrieval algorithm will be most effective for broken trade cumulus, thin cirrus, and also pos-
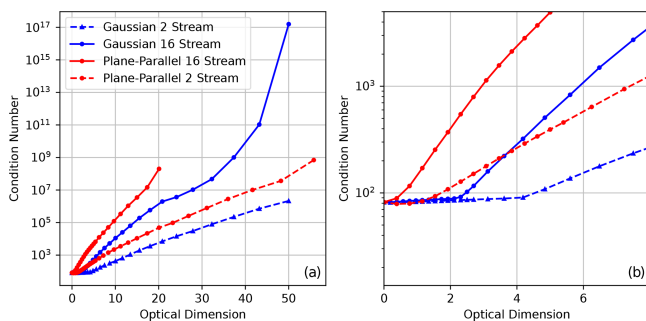

Figure 16. Similar to Fig. 15 but using base states that are plane-parallel, horizontally homogeneous extinction fields. Note the difference in color scale from Figs. 10 and 15. Panels (a), (b), and (c) have isotropic phase functions. Panels (d), (e), and (f) have Mie phase functions, and panels (g), (h), and (i) have Mie phase functions with a reduced angular resolution (see text for details). Panels (a), (d), and (g) have a black surface. Panels (b), (e), and (h) have Lambertian surface albedos of 0.2, while panels (c), (f), and (i) have Lambertian surface albedos of 0.7.

Figure 17. The dependence of the condition number of the approximate Jacobian on the optical dimension which is the vertical optical thickness of plane-parallel homogeneous (red lines) and the optical radius of 3D Gaussian (blue lines) extinction fields. Panel (b) is a magnified version of panel (a), highlighting the different behaviors at small optical depths. The plane-parallel 16 stream curve stops at maximum optical thickness 20, since, beyond this, it becomes truly ill-posed (condition number tending to infinity).

sibly multi-layered combinations of them, along with their aerosol environment. Further studies quantifying the effectiveness of tomographic retrievals in these different cloud regimes are warranted. In Part 2 of this study, we focus on cumuliform cloud types to test the effectiveness of the method in filling the observational gap in operational cloud property retrievals which do not perform well in situations with broken cumulus.

It remains to be seen how effective tomography will be for very thick, cumuliform clouds and for moderately thick stratiform clouds, all of which are strongly ill-conditioned. It also remains to be seen what regularization schemes or prior information is required to improve retrievals in these conditions. There are also methods that have been proposed to mitigate the ill-conditioning of the inverse problem through the use of tailor-made preconditioning schemes (Niu et al., 2010; Tian et al., 2010).

7 Summary

In this study, we have introduced and validated an algorithm for retrieving the 3D volumetric properties of clouds using multi-angle, multi-pixel radiances and 3D radiative transfer. The retrieval utilizes an iterative, optimization-based solution to the generalized least squares problem to find a best-fitting state vector parameterizing the atmosphere. The iterative retrieval is made computationally tractable through the use of an approximate Jacobian calculation introduced by Levis et al. (2015, 2017, 2020) that has been extended to accommodate open and periodic horizontal boundary conditions and an improved treatment of non-black surfaces. We implemented this retrieval in a new software package, AT3D, which we have made publicly available.

We presented the basic physical principles of inverse radiative transfer from a linearized perspective. We identified that the iterative retrieval will tend to ill-posedness as the optical depth of the medium increases. This is due to the increasing smoothing effects of multiple scattering, which are ill-conditioned under inversion. This ill-conditioning of the inversion is also highly sensitive to the numerical treatment of the forward-scattering peak for highly peaked phase functions, such as cloud droplets at solar wavelengths. In the SHDOM solver used in the retrieval algorithm, this manifests as a sensitivity to both the phase function and also the angular resolution used in the solver. When forward-scattering peaks are strong and angular resolutions are low, the ill-conditioning is mitigated.

Our linear analysis of the cloud tomography problem also indicates that the fastest reductions in the cost function will occur by modifying regions of the cloud optically close to the Sun and to the sensors, where the magnitude of the elements of the Jacobian matrix is the largest. As the cloud becomes optically thick, the magnitude of the elements of the Jacobian matrix becomes exponentially larger in the regions closest to the Sun and sensors than those furthest away. This may cause a retrieval using local optimization, such as AT3D described here, to converge to a local minimum when the target medium is optically thick, if no other constraints are employed in the retrieval – other than the multi-angle radiances.

We presented the derivation of the approximate Jacobian as an approximation to an adjoint radiative transfer problem and evaluated its accuracy. Errors in the elements of the approximate Jacobian matrix which contain derivatives of radiances with respect to the 3D volume extinction coefficient increase from 2 % to 12 % for media with cloud-like single-scattering properties over surfaces with Lambertian albedos less than 0.2, as the maximum optical depths of the medium increase from 0.2 and 100. When the albedo of the Lambertian surface is 0.7, then the errors are larger, reaching 20 % for media with maximum optical depths of 100. Errors are smaller for media with phase functions with strong forward-scattering peaks, especially when a low angular resolution is utilized in the SHDOM solver.

The elements of the approximate Jacobian matrix that contain the derivatives of radiances with respect to the volume extinction coefficients of the plane-parallel media that provide the open horizontal boundary conditions to the 3D radiative transfer problem are very inaccurate. Errors in these elements of the approximate Jacobian matrix exceed 50 %, unless the plane-parallel media are optically thin (~ 0.1). The larger errors in the elements of the Jacobian matrix that correspond to the horizontal boundary conditions compared to those that correspond to the volumetric optical properties of the medium are due to the accuracy of the approximate calculation of the Jacobian matrix, which has lower accuracy in lower dimensional (2D and 1D) transport problems. These results indicate that AT3D will likely only be useful for jointly retrieving volumetric optical properties and horizontal boundary conditions when the horizontal boundary conditions correspond to optically thin plane-parallel media like the clear atmosphere.

The approximate Jacobian captures the key information content in the reference Jacobians and has a similar dependence of ill-conditioning on the optical depth of the medium as the reference. Our numerical tests using the approximate Jacobian also indicate that the retrieval problem becomes much more ill-posed, as measured by the condition number, when the target medium forms an infinite slab geometry compared to when it forms a finite geometry. This difference is due to the inability of the sensors to distinguish between rearrangements in the extinction field at the bottom of an optically thick plane-parallel layer and indicates that tomographic retrievals will be most beneficial in optically thinner stratiform clouds or broken fields of cumulus.

We therefore judge that the approximate Jacobian, and by extension the retrieval method currently available in AT3D for tomographic retrievals, is most suitable for retrievals in thin, cirriform clouds and trade cumulus over oceanic surfaces and their adjacent aerosols. The successful application of the retrieval to a broader variety of clouds and surface types is also possible but will likely require the incorporation of additional constraints. In Part 2 of this study, we will examine the implications of the ill-conditioning and errors in the approximate Jacobian in idealized tomographic retrievals of simulated clouds. In Part 2, we focus on the retrieval of the 3D volume extinction coefficient using monochromatic

radiance to address these fundamental concerns. The AT3D software is already able to examine a much wider variety of problems and can be used to explore microphysical retrievals using polarimetric (Levis et al., 2020) or multi-spectral measurements. We hope that, in making this software package publicly available, we will encourage the development of this retrieval method and other next-generation remote sensing retrievals that utilize 3D radiative transfer modeling.

Appendix A: Glossary

Table A1. Mathematical terms introduced in Sect. 3.

Variables	Description	Reference
$\boldsymbol{a}$	State vector	Eq. (1)
$\widetilde{\boldsymbol{a}}$	Retrieved state vector	Eq. (2)
$\boldsymbol{a}_0$	Initial value of state vector	Fig. 1
$\Delta\boldsymbol{a}_m$	Update to state vector at mth iteration of optimization	Eq. (3)
a_{tol}	Absolute error tolerance when testing for agreement between floats	Eq. (C1)
$\mathbf{B}(\boldsymbol{r})$	Polarized volume emission source	Near Eq. (13)
$B_\lambda(T)$	Planck blackbody radiance function	Near Eq. (13)
$\boldsymbol{b}$	A vector of output from AT3D on which to test for accuracy	Eq. (C1)
$\boldsymbol{c}$	A vector of reference output used to evaluate AT3D	Eq. (C1)
D	Spatial domain of RT	Near Eq. (7)
∂D_n	The boundary of a spatial domain	Eq. (11)
D_1	Spatial domain of 3D RT	Eq. (7)
$D_2 - D_9$	Spatial domain of 2D and 1D RT (evens 1D; odds 2D)	Eq. (8)
$D_n \times \mathbb{S}^2$	Internal set of radiative transfer (space and direction)	Fig. 3
δ	Dirac delta distribution	Eq. (17)
$\boldsymbol{F}(\boldsymbol{a})$	Forward model	Eq. (1)
$F_{i,n}(\boldsymbol{x})$	The contribution of the nth RT domain to the ith component of the forward model	Eq. (42)
$\boldsymbol{F}_0$	The flux along the solar beam	Eq. (17)
$f(\boldsymbol{r}, \boldsymbol{\Omega})$	The volume source of the RTE	Eq. (13)
$f_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})$	The volume source of the RTE for diffuse radiance	Eq. (26)
$\hat{f}_{\text{d}}(\boldsymbol{r}, \boldsymbol{\Omega})$	The effective volume source of the RTE for diffuse radiance	Eq. (33)
$g(\boldsymbol{r}, \boldsymbol{\Omega})$	The boundary source of the RTE	Eq. (15)
$g^{\text{BOT}}(\boldsymbol{r}, \boldsymbol{\Omega})$	Boundary source at domain bottom (thermal emission)	Eq. (27)
$g^{\text{TOP}}(\boldsymbol{r}, \boldsymbol{\Omega})$	Boundary source at domain top (isotropic emission)	Eq. (27)
$g_{\odot}(\boldsymbol{r}, \boldsymbol{\Omega})$	Boundary source at domain top due to collimated solar illumination	Eq. (17)
$g^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})$	Boundary source on domain sides	Eq. (18)
$g_{\odot}^{\text{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})$	Boundary source on domain sides due to direct, unscattered solar illumination	Eq. (22)

Table A1. Continued.

Variables	Description	Reference
$\boldsymbol{g}_{\mathrm{d}}^{\mathrm{SIDE}}(\boldsymbol{r}, \boldsymbol{\Omega})$	Boundary source on domain sides not due to direct, unscattered solar illumination	Eq. (22)
$\boldsymbol{g}_{\mathrm{d}}$	The boundary source of the RTE for diffuse radiance	Eq. (27)
$\hat{\boldsymbol{g}}_{\mathrm{d}}(\boldsymbol{r}, \boldsymbol{\Omega})$	The effective boundary source of the RTE for diffuse radiance	Eq. (34)
$\boldsymbol{\Gamma}_n^-$	Incoming radiation at the boundary of the nth RT domain	Eq. (11)
$\boldsymbol{\Gamma}_n^+$	Outgoing radiation at the boundary of the nth RT domain	Eq. (11)
$\boldsymbol{\Gamma}_n^{\mathrm{S}}$	Incoming radiation at the horizontal sides of the nth RT domain	Eq. (37)
$H_L(x)$	A Heaviside function with transition at length L	Eq. (38)
i	Index of forward model	Eq. (5)
$\boldsymbol{I}_n(\boldsymbol{r}, \boldsymbol{\Omega}) = [I, Q, U, V]^{\mathrm{T}}$	Stokes vector (polarized radiance field) of the nth domain	Eq. (12)
$\boldsymbol{I}_\odot(\boldsymbol{r}, \boldsymbol{\Omega})$	Direct, unscattered solar radiance	Eq. (21)
$\boldsymbol{I}_{\mathrm{d}}(\boldsymbol{r}, \boldsymbol{\Omega})$	Diffuse radiance; all radiance that is not direct, unscattered solar radiance	Eq. (21)
j	Index of state vector elements	Eq. (5)
k	Index of sub-pixel rays in a sensor response function	Eq. (32)
$\mathbf{K}$	Jacobian matrix containing the derivatives of the forward model with respect to elements of the state vector	Eq. (4)
$\kappa(\mathbf{K})$	Condition number (of Jacobian matrix)	Eq. (6)
$\boldsymbol{l}$	Lower bounds on state vector	Eq. (2)
l	Path length	Eq. (24)
l_e	Path length over a subgrid interval	Eq. (D2)
L_x	Length of 3D RT domain in x direction	Eq. (7)
L_y	Length of 3D RT domain in y direction	Eq. (7)
L_z	Length of 3D RT domain in z direction	Eq. (7)
$\mathcal{L}[\cdot]$	The transport operator of the RTE	Eq. (12)
m	Iteration number in L-BFGS-B optimization	Eq. (2)
M	The number of past L-BFGS-B iterations over which gradient information is stored to approximate the Hessian of the cost function	Near Eq. (3)
n	Domain index of RT domains	Eq. (11)
$\boldsymbol{n}_n(\boldsymbol{r})$	Unit vector outer normal to the boundary of the nth domain	Eq. (11)
$\boldsymbol{O}_i$	The polarization analyzer of the sensor response function of the ith element of the forward model	Eq. (32)
$\boldsymbol{\Omega}$	Direction of propagation	Eq. (11)
$\boldsymbol{\Omega}_{\mathrm{sun}}$	Direction of propagation of solar illumination	Eq. (17)
$\boldsymbol{\Omega}_{ik}$	Propagation direction of radiance that is sampled by the kth sub-pixel ray of the sensor response function for the ith element of the forward model	Eq. (32)
$\omega(\boldsymbol{r})$	Single-scattering albedo	Eq. (12)
$\boldsymbol{P}_i$	Sensor response function of the ith measurement	Eq. (32)
$\boldsymbol{p}_{ik,n}$	The volume response function of the kth sub-pixel ray of the sensor response function for the ith element of the forward model for the nth RT domain	Eq. (39)
$\boldsymbol{q}_{ik,n}$	The attenuated boundary response function of the kth sub-pixel ray of the sensor response function for the ith element of the forward model for the nth RT domain	Eq. (40)

Table A1. Continued.

Variables	Description	Reference
$\boldsymbol{r} = (x, y, z)^{\mathrm{T}}$	Position vector	Eq. (7)
$\boldsymbol{r}_{ik}$	Position of the kth sub-pixel ray for the sensor response function of the ith measurement	Eq. (32)
r_{tol}	Relative error tolerance when testing for agreement between floats	Eq. (C1)
$R(\boldsymbol{a})$	Regularization term in cost function	Eq. (1)
$\mathcal{R}$	Reflection operator part of the boundary condition of the RTE	Eq. (15)
$\mathbf{R}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}')$	Polarized bidirectional reflectance distribution function	Eq. (16)
$\mathbf{S}_{\epsilon}$	Error co-variance matrix	Eq. (1)
$\mathbb{S}^2$	Unit sphere in 3D which is the domain of directions of RT	Eq. (11)
$\mathrm{d}S_{\boldsymbol{\Omega}}$	An element of surface parameterized by the propagation direction	Eq. (29)
$\mathrm{d}V_{\boldsymbol{r}}$	An element of volume parameterized by the position vector	Eq. (29)
$\mathcal{S}[\cdot]$	Volume streaming operator	Eq. (35)
$\sigma(\boldsymbol{r})$	Volume extinction coefficient	Eq. (12)
s_{largest}	Largest singular value of Jacobian matrix	Eq. (6)
s_{smallest}	Smallest singular value of Jacobian matrix	Eq. (6)
$\mathcal{T}[\cdot]$	Streaming operator	Eq. (23)
$\mathrm{T}(\boldsymbol{r}, \boldsymbol{r}')$	Transmission between two positions	Eq. (24)
$\mathcal{U}$	Solution operator of the RTE	Eq. (28)
$\boldsymbol{u}$	Upper bounds on state vector	Eq. (2)
$\boldsymbol{v}$	Test field	Eq. (29)
$\boldsymbol{w}$	Test field	Eq. (29)
w_k	The weights of the sub-pixel rays forming the sensor response function	Eq. (32)
χ^2	Cost function	Eq. (1)
$\boldsymbol{y}$	Measurement vector	Eq. (1)
$\mathbf{Z}(\boldsymbol{r}, \boldsymbol{\Omega}, \boldsymbol{\Omega}')$	Scattering-phase matrix	Eq. (12)
$\oplus$	Set union operator	Near Eq. (12)

Appendix B: Modifications to the SHDOM solver

The SHDOM solver in AT3D is derived from the original implementation of SHDOM in Fortran 77 and Fortran 90 (Evans, 1998). We have made some modifications to the solver and optical property schemes to ensure the differentiability of radiances with respect to all optical properties and microphysical properties. The SHDOM solver uses two different grids, namely one to represent the optical properties, known as the property grid, and one to solve the RT problem. The property grid is regular. The RT grid is based on a regular grid, but with optional local grid refinement that makes the grid irregular, and is referred to as the adaptive grid. Specific choices must be made to both prepare the optical properties on the property grid and to interpolate them onto the RT grid.

The original optical property generation scheme (PROPGEN) prepared the optical properties through external mixing of different participating particle species. This means that the total volume extinction coefficient and single-scatter albedo are calculated by summing the volume extinction coefficients and the volume scattering coefficients. The calculation of the effective phase functions of the mixture would require calculating the weighted mean of the phase functions, where the weights are the fractional contribution of each species to the total volume scattering coefficient. This scheme would typically produce a unique phase function for

each point, which would cause a substantial memory burden in the SHDOM solution, especially if polarization is considered. This is also an issue when the adaptive grid is used in the SHDOM solution, as the number of grid points can grow, with each one requiring a new unique phase function.

To alleviate this, the original SHDOM PROPGEN program limited the number of unique phase function mixtures by only adding a new phase function if a tolerance on the phase function accuracy is not met by any of the phase functions within the current set. This smaller set of unique phase functions is then stored in a lookup table (LUT), with a corresponding pointer at each grid point. This scheme is non-differentiable, as a thresholding operation is used to select the phase functions.

When a new adaptive grid point requires new optical properties, the extinction and single-scattering albedo at each point can be calculated with linear interpolation. To avoid new unique phase functions, a less accurate nearest-neighbor interpolation is used. Specifically, the new adaptive grid point inherits the phase function of a property grid point if either (a) the two are co-located or (b) the property grid point is the one with the largest scattering coefficient amongst the property grid points surrounding the new adaptive grid point. This scheme is also not differentiable.

We have replaced the PROPGEN program and modified the SHDOM solver used in AT3D itself to accommodate a new scheme for representing phase functions that is differentiable. Specifically, we employ an online mixing of phase functions, when required during the SHDOM solution procedure, which is during the convolution of the phase function and the radiance field and referred to as the SOURCE function computation. We note that the delta-M scaling occurs on the RT grid, consistent with the original SHDOM implementation. This operation is performed by the subroutine COMPUTE_SOURCE in the shdomsub1.f file. In the online mixing, we store a limited LUT for each scattering species. We may adopt a nearest-neighbor or linear interpolation from this LUT to define the phase function of each species at each grid point. We then calculate the total phase function of all species at each grid point by performing a weighted average over the phase function of each species. This latter part of the procedure is the exact part of the phase function calculation. The approximation can occur, depending on the choice of phase functions within the LUT and the choice of interpolation rule. When the number of unique phase functions is small, each entry in the LUT can be the species phase function. On the other hand, when the number of unique phase functions is large, we approximate the dependence of the phase function on the particle's microphysical properties using linear interpolation. This approximation occurs for each individual particle species, but the mixing is still exact. Both the mixing and linear interpolation from the LUT are differentiable. The downside of the interpolation from the LUT is that some error will be induced compared to using the exact Mie calculations for each phase function.

The final result of this is that we have replaced the memory burden of storing phase functions with the computational burden or performing weighted sums of phase functions. The computational burden is relatively low because the number of phase function expansion coefficients that need to be mixed for the SHDOM solution is relatively small. To be precise, it is small compared to the total number of expansion coefficients needed to store a unique phase function. It is also small compared to the number of spherical harmonics in the SOURCE function computation. As such, the computational cost of the SOURCE function computation and the solver itself is relatively invariant, as shown in Table B1. Given that most simulations will have, at most, cloud liquid, cloud ice, and one or two types of aerosol, the computational expense of performing these simulations is only around 10 % larger for typical angular accuracies. This is relatively large because the SOURCE computation is actually called three times during each solution procedure to reduce the memory expense of the SHDOM solver. Multi-species SHDOM solutions are much more expensive for lower angular accuracies, but the solutions are extremely fast in these cases, resulting in minimal changes in wall time.

Appendix C: SHDOM solver verification

We have made substantial modifications to the implementation of the SHDOM solver in AT3D. These include the modifications to the representation of optical properties in the SOURCE function calculation, described above in Appendix B, and also through the Python wrapping of the SHDOM solution procedure. Due to these changes, it is critical to perform a thorough verification of the new software to ensure that bugs have not been introduced and that the efficacy of the translation of the algorithm into code remains intact (Kanewala and Bieman, 2014). As such, we have performed a verification of the model to ensure consistency of the implementation of SHDOM in AT3D with the original SHDOM implementation and with several analytic benchmarks. The comparisons against the original SHDOM implementation were the most informative, as they provided strict bounds on the behavior of the model in complex cases, thereby allowing us to diagnose several errors, including one in the original SHDOM implementation. The details of all of the tests in AT3D, including the solver verifications described here, can be found in the AT3D test folder, including input files for the original SHDOM code.

In our quantitative verification of the AT3D SHDOM solver, we must compare arrays of floating-point numbers. Given our changes to the optical property and SOURCE function calculations in the solver, it is not possible to test for bit-perfect reproduction of the original SHDOM solver. The comparison of outputs is also difficult, due to the need to verify the solution when the adaptive grid is used. The adaptive grid scheme utilizes a thresholding operation on a splitting

Table B1. The ratio of the computational expense of the SHDOM solution procedure with the new SOURCE function computation for a different number of particle species against the equivalent exact mixture for a varying angular resolution of the SHDOM solver. The media used in the SHDOM solutions are generated from uniform distributions of the effective particle radius, extinction, and single-scattering albedo. The medium has 832 radiative transfer grid points. Half of the particle species in each medium have cloud-sized effective particle radii (10 to 30 µm), and the other half have aerosol-sized effective particle radii (0.3–0.5 µm). The timing ratios are averaged over 10 stochastic realizations for each configuration, with the standard deviations of the ratios shown in parentheses.

Number of spherical harmonics	Number of particle species			
	2	4	6	8
4	1.52 (0.07)	2.16 (0.13)	2.68 (0.22)	3.42 (0.36)
16	1.24 (0.15)	1.57 (0.08)	1.94 (0.15)	2.26 (0.18)
64	1.15 (0.11)	1.39 (0.16)	1.53 (0.16)	1.66 (0.09)
256	1.06 (0.06)	1.11 (0.06)	1.18 (0.05)	1.28 (0.08)
1024	1.02 (0.1)	1.05 (0.07)	1.07 (0.07)	1.19 (0.06)

criterion, which is a floating-point number, to decide when to refine the grid. This thresholding is very sensitive to even small changes in the inputs or numerical operations at the level of the numerical precision of the single precision floats. The changes in the solution in the RT solution (e.g., fluxes and radiances), due to one additional refinement of the grid, can be much larger (i.e., several percent) than changes in the inputs. As such, the adaptive grid scheme can amplify small differences in inputs into much larger differences in the outputs and is numerically unstable. Care must be taken to manage this issue when performing the comparisons with original SHDOM. To test for the presence of significant differences between a vector of output from the AT3D solver, $\boldsymbol{b}$, and a vector of reference output (e.g., from the original SHDOM solver), $\boldsymbol{c}$, we test whether the absolute differences between the N elements of the vectors exceed the sum of a prescribed relative error, r_{tol}, and absolute error, a_{tol}. Specifically, we require

$$|b_i - c_i| \leq a_{\text{tol}} + r_{\text{tol}} |c_i| \ \forall \, i \in 1, \ldots, N \qquad \text{(C1)}$$

to hold for every element. For all tests, r_{tol} is set to 10^{-5}, while a_{tol} is set according to the expected truncation error, which varies depending on the nature of the reference. When comparing between AT3D and SHDOM radiances, we use output written to file with five significant figures (standard SHDOM output), and we require $a_{\text{tol}} = 10^{-5}$, as most outputs are of the order of unity. We also use a custom, untruncated output of the SOURCE function from SHDOM. SHDOM uses single-precision floats, and so we expect agreement to within seven significant figures. As such, we judge the test to be successful in cases with an untruncated output if the agreement is to within $a_{\text{tol}} < 10^{-6}$, though we test as strictly as we are able by reporting the smallest value of a_{tol} for which each test passes. For example, we verified that the modified SOURCE function computation in AT3D, due to the new treatment of the phase matrices, passes the consistency test with $a_{\text{tol}} = 10^{-8}$.

We proceed with our description of the verification process in the order of increasing complexity. We begin with the analytic benchmarks, some comparisons of AT3D against SHDOM in plane-parallel atmospheres, and then a comparison against SHDOM in a complex 3D case.

C1 Analytic and numerical benchmarks for 1D RT

We compare the AT3D solver against analytic benchmarks to ensure the absolute accuracy of the solver in simple situations (Jones and Di Girolamo, 2018). The first such situation is a non-scattering, homogeneously absorbing atmosphere with a Lambertian surface, where radiances are verified with $a_{\text{tol}} = 10^{-8}$. We next consider an isothermal atmosphere with surface reflection to verify the thermal RTE solution. The nadir radiance is verified against an analytical solution to 0.005 % for a Lambertian surface with an albedo of 0.5, a surface temperature of 300 K, and an atmospheric temperature of 288 K, with 50 optical paths ranging from 0.03 to 15. The accuracy of this result depends on the angular resolution for calculating the reflected radiation at the surface. In total, 128 zenith discrete ordinate bins and 256 azimuthal discrete ordinate bins are used. The analytic solution in this case is given by Eq. (C9) in Jones and Di Girolamo (2018). We also verified a case with a non-scattering, homogeneously absorbing, and isothermal atmosphere with solar surface reflection and emission, which is a combination of two previous cases, to the same level of accuracy.

We compare RT solutions from AT3D and SHDOM for molecular (Rayleigh) scattering in plane-parallel atmospheres over all of the surface BRDFs available in SHDOM and AT3D with varying parameters (e.g., surface wind speed) to verify that the surface BRDFs, the Rayleigh scatter calculations, and 1D RT are consistent between the two solvers. The details of the comparison can be found in the AT3D code. We compare hemispheric downwelling fluxes, hemispheric upwelling fluxes, and direct fluxes at the surface, as well as top-of-atmosphere (TOA) radiances for 20 angles

spaced equally in cosine of zenith over the upwelling hemisphere. This is done through comparison with SHDOM output written to file; i.e., it is verified with $a_\text{tol} = 10^{-5}$.

C2 Intercomparison against SHDOM

We use a 3D radiative transfer setup with three Stokes components to compare the RTE solution between SHDOM and AT3D for a cloud distributed with the AT3D code. This cloud was utilized in Levis et al. (2015, 2020). Both AT3D and SHDOM use optical properties read from a file prepared by SHDOM's PROPGEN program to avoid amplification of errors due to the adaptive grid. In brief, the simulation uses 8 zenith discrete ordinate bins, 16 azimuthal discrete ordinate bins, a splitting accuracy of 0.1, a spherical harmonic accuracy of 0.01, and a solution accuracy of 10^{-4}. The final output of the SHDOM solver is the spherical harmonic expansion of the angular convolution of the radiance field with the phase function, denoted as the polarized SOURCE vector, on which we perform the consistency test. When the adaptive grid matches between AT3D and the SHDOM reference, as in our example cloud, the solutions are consistent with $a_\text{tol} = 5 \times 10^{-7}$.

Larger differences in the SOURCE vector may occur when comparing SHDOM solutions across different machines and compilers when the adaptive grid is used because if the adaptive grid splitting occurs differently, then the SOURCE vectors will not have a one-to-one correspondence. The input files and scripts to reproduce the SHDOM benchmarks used in the test are distributed with AT3D, along with the static SOURCE output from the original SHDOM. The static output may not be consistent with other machines and may need to be regenerated. Given the good agreement between the solvers for the SOURCE vector when the adaptive grid splitting is consistent, we judge that the AT3D implementation of the SHDOM solver is in good agreement with the original SHDOM implementation.

We also test the radiance calculation in the complex situation of 3D RT with the adaptive grid. We find that, even in a simplified situation where the original SOURCE function computation is used, there is substantial disagreement in the calculated radiances, despite the good agreement of the SOURCE function computations. We identify discrepancies in the radiances that can, in rare cases, reach up to 1.5 % error in the intensity and, as such, are clearly noticeable, even when using the truncated output of SHDOM written to file as a reference. The root mean square errors are much smaller, being less than 0.004 %, and are therefore are effectively undetectable in model intercomparisons. Given the good agreement of the SOURCE vectors, the differences in the radiances between AT3D and the original SHDOM were traced to differences in the subroutines used to calculate the integrals for the formal solution of the RTE (Eq. 36). The original SHDOM implementation uses the CALCULATE_RADIANCE subroutine, which is used in the original SHDOM radiance output mode. AT3D's radiance calculation builds on the VISUALIZE_RADIANCE subroutine, which is used in the original SHDOM visualization output mode. The visualization output mode is more flexible in accommodating a unique viewing angle for each calculated radiance.

We do not have an absolute benchmark for the radiance calculation in such a complex situation with the adaptive grid that is sufficiently precise. Instead, we determined the cause of the discrepancy through code analysis and judged that the AT3D implementation is more physically correct, based on the following considerations. The discrepancy only occurs when the adaptive grid is used and is due to differing implementations of the interpolation of the SOURCE extinction product onto the characteristic of the radiance when moving between cells with different resolution in the adaptive grid. In SHDOM, the SOURCE extinction product at the face of the most recently exited cell is always used as the SOURCE extinction product at the entry point of the next cell. However, if going to a higher-resolution cell, then the SOURCE extinction product at the entry point can be estimated with higher accuracy, using the higher-resolution grid points in the new cell. This latter procedure is used in AT3D, which we judge to be more accurate. We note that the significant discrepancies are limited to rare cases. The errors will be largest when there is a large SOURCE function difference between the SOURCE function at the higher- and lower-resolution grid points. Therefore, the errors will also be minimized when a low value of splitting accuracy is used. Our example used a relatively large splitting accuracy of 0.1, and even then, the bug correction produced a very low root mean square change in the radiance field (0.004 %). As such, the correction of this bug will not substantially affect holistic benchmarks of radiances in 3D RT, such as reciprocity (e.g., Di Girolamo, 1999, 2002) or previous model intercomparisons, as intermodel differences in radiances are much larger.

Appendix D: Implementation of the approximate Jacobian matrix calculation

Here, we describe how the inner products used to calculate the entries in the Jacobian matrix (Eq. 58) are numerically evaluated. These details are relevant to the performance of AT3D, as there are several strategies that can be employed when performing forward-adjoint-based linearization of a model, depending on whether adjoints are formed from continuous, discrete, or numerical algorithms (Klose and Hielscher, 2002). First, we consider the evaluation of the integral involving the volume streaming operator, which is the second term in Eq. (58). This integral is simply a line integral, as it only needs to be evaluated at the specific position and angle ($\boldsymbol{r}_{ik}\boldsymbol{\Omega}_{ik}$) corresponding to a quadrature point in the sensor's response function.

$$S\left[\left.\Delta\hat{\boldsymbol{f}}_j\right|_{D_n\times\mathbf{S}^2}\right](\boldsymbol{r}_{ik},\boldsymbol{\Omega}_{ik})$$

$$=\int_0^l \Delta\boldsymbol{f}_j\left(\boldsymbol{r}_{ik}-l'\boldsymbol{\Omega}_{ik},\boldsymbol{\Omega}_{ik}\right)$$

$$\exp\left(-\int_0^{l'}\sigma\left(\boldsymbol{r}_{ik}-l''\boldsymbol{\Omega}_{ik}\right)\mathrm{d}l''\right)\mathrm{d}l' \tag{D1}$$

This integral is approximated by dividing it into E subintervals, with optical paths smaller than a specified tolerance (default 0.1) and with the endpoints l_e and l_{e+1}.

$$S\left[\left.\Delta\hat{\boldsymbol{f}}_j\right|_{D_n\times\mathbf{S}^2}\right](\boldsymbol{r}_{ik},\boldsymbol{\Omega}_{ik})$$

$$=\sum_e \exp\left(-\int_0^{l_e}\sigma\left(l'\right)\mathrm{d}l'\right)\int_{l_e}^{l_{e+1}}\Delta\boldsymbol{f}_j\left(l'\right)$$

$$\exp\left(-\int_{l_e}^{l'}\sigma\left(l''\right)\mathrm{d}l''\right)\mathrm{d}l' \tag{D2}$$

We use a formula for the integration over each subinterval that assumes that $\Delta\boldsymbol{f}_j$ and the extinction σ vary linearly along the path. This formula is accurate to the second order in the subinterval path length.

$$\int_{l_e}^{l_{e+1}}\Delta\boldsymbol{f}_j\left(l'\right)\exp\left(-\int_{l_e}^{l'}\sigma\left(l''\right)\mathrm{d}l''\right)\mathrm{d}l'$$

$$\approx\left(1-e^{-\frac{1}{2}(l_{e+1}-l_e)(\sigma(l_e)+\sigma(l_{e+1}))}\right)\frac{2}{\sigma\left(l_e\right)+\sigma\left(l_{e+1}\right)}$$

$$\left\{\frac{1}{2}\left(\Delta\boldsymbol{f}_j\left(l_e\right)+\Delta\boldsymbol{f}_j\left(l_{e+1}\right)\right)+\frac{l_{e+1}-l_e}{12}\right.$$

$$\left(\sigma\left(l_{e+1}\right)\Delta\boldsymbol{f}_j\left(l_n\right)-\sigma\left(l_e\right)\Delta\boldsymbol{f}_j\left(l_{n+1}\right)\right)$$

$$\left.\left(1-\frac{l_{e+1}-l_e}{20}\left[\sigma\left(l_e\right)-\sigma\left(l_{e+1}\right)\right]\right)\right\} \tag{D3}$$

This formula is approximate but consistent with the discretization of the forward model, with one exception. Namely, the first term, $\Delta\boldsymbol{f}_j$, involves the radiance, which does not actually vary linearly, since the SOURCE extinction product varies linearly in the forward model. To evaluate Eq. (D3), all that is needed is to evaluate $\Delta\boldsymbol{f}_j$ and the extinction at the end points of each subinterval. We adopt a trilinear interpolation rule from the RT grid points for the volume extinction coefficient and all terms in $\Delta\boldsymbol{f}_j$, apart from the radiance. For the radiance, we directly evaluate it at the endpoints of each subinterval by evaluating the forward model.

The remaining complication is the evaluation of $\Delta\boldsymbol{f}_j$ in the required direction. If the delta-M approximation and TMS correction are not employed, then all angularly varying quantities, including the single-scattering terms in $\Delta\boldsymbol{f}_j$, are represented in a basis of (real) generalized spherical harmonics as in the forward model. On the other hand, if the delta-M approximation and TMS correction are employed, then the single-scattering terms in $\Delta\boldsymbol{f}_j$ are computed using a LUT of phase functions dependent on scattering angle, which is again consistent with the forward model.

Second, we consider the boundary integral in the radiance derivative calculation, which is the first term in Eq. (58):

$$\mathcal{T}\left[\Delta\hat{\boldsymbol{g}}_{n,j}^0\left(\boldsymbol{r},\boldsymbol{\Omega}\right)|_{\boldsymbol{\Gamma}_n^-}\right](\boldsymbol{r}_{ik},\boldsymbol{\Omega}_{ik})=\mathcal{T}\left[\mathcal{R}\left[\left.\frac{\partial\boldsymbol{I}_\odot}{\partial a_j}\right|_{\boldsymbol{\Gamma}_+^n}\right]\right.$$

$$+\frac{\partial\mathcal{R}}{\partial a_j}\left[\boldsymbol{I}_{n,\odot}+\boldsymbol{I}_{n,\mathrm{d}}\right]+\frac{\partial\boldsymbol{g}_n^{\mathrm{TOP}}\left(\boldsymbol{r},\boldsymbol{\Omega}\right)}{\partial a_j}$$

$$\left.+\frac{\partial\boldsymbol{g}_n^{\mathrm{BOT}}\left(\boldsymbol{r},\boldsymbol{\Omega}\right)}{\partial a_j}+\frac{\partial\boldsymbol{g}_n^{\mathrm{SIDE},0}\left(\boldsymbol{r},\boldsymbol{\Omega}\right)}{\partial a_j}\right]. \tag{D4}$$

Currently, in AT3D, only the first term is relevant, as all other terms are neglected. Note that the horizontal boundary derivative term, $\frac{\partial\boldsymbol{g}_n^{\mathrm{SIDE},0}(\boldsymbol{r},\boldsymbol{\Omega})}{\partial a_j}$, is made up of volume and boundary terms in other auxiliary domains so that the first term is the only true boundary term that is considered. The direct radiance derivative is calculated at every RT grid point, including on the boundary. The reflection operator $\mathcal{R}$ in Eq. (16) acts on these grid points. The result is then bilinearly interpolated along the boundary to the position where it needs to be evaluated, which is the intersection of the line of sight of the sensor with the boundary ($\boldsymbol{r}_{ik}-l\boldsymbol{\Omega}_{ik}\boldsymbol{\Omega}_{ik}$).

Appendix E: Verification of the approximate Jacobian matrix calculation

It is critical to perform a proper verification of the approximate calculation of the Jacobian matrix before we can draw scientific conclusions about the behavior of the proposed algorithm with numerical experiments (Kanewala and Bieman, 2014). To verify the approximate Jacobian matrix, we use a combination of benchmarks calculated from simple scenarios with analytic solutions and finite-differencing benchmarks. All of the analytic benchmarks come from consideration of the non-scattering thermal emission or surface reflectance cases.

If we have an emissive surface and a homogeneously absorbing isothermal atmosphere, then we can easily calculate the derivative of the measured radiance with respect to the homogeneous extinction in terms of the black body radiances of the atmosphere B_A and the surface B_S, as follows:

$$\left.\frac{\partial I}{\partial\sigma}\right|_{\sigma_0}=\frac{h}{\mu}e^{\sigma_0 h/\mu}\left(B_\mathrm{A}-B_\mathrm{S}\right), \tag{E1}$$

where h is the geometric thickness of the layer, and μ is the cosine of the viewing zenith angle. This test can be used to

determine correct discrete implementation of the derivatives consistent with the formal, integral solution of the RTE used to calculate the radiances in SHDOM. We compare this analytic benchmark with the approximate Jacobian and a derivative estimated from a two-point central difference. In this simple scenario, with $h = 0.4$ and $\mu = 1.0$, the relative error in the approximate Jacobian increases abruptly from 10^{-6} to 5×10^{-3}, when $\sigma_0 = 10.0$. This is due to the saturation of the error, as many subintervals in the integration reach an optical path length of 0.1. Errors are substantially reduced if a smaller optical path length is used. The error rapidly balloons in optically thicker scenarios ($\sigma_0 = 100$) due to a rounding error affecting the signal, as the signal is attenuated below the precision of single-precision floats. The error in the two-point central difference also grows as σ_0 increases but is also dependent on the step size. If a step size is chosen as the maximum of 1 % or 0.01, then the error is maintained below 6 % for $\sigma_0 < 10$, where the absolute error also peaks. This indicates that, despite the use of the approximate integration of the subintervals, the approximate Jacobian tends to be a bit more accurate than the finite differencing with respect to the analytic calculation for a case where the approximate Jacobian should be exact.

In the remaining tests, we can only utilize a finite-differencing reference, as we invoke more complex, randomly generated media (white noise) to test the correct application of the chain rule to the interpolation rules used for the optical properties. First, we test the calculation of the direct-beam radiance derivatives, which are verified with $a_{\text{tol}} = 10^{-5}$ in Eq. (C1) for both periodic and open BCs and for reflective Lambertian surfaces to ensure that the surface boundary term is correctly calculated. Second, derivatives of measured radiances with respect to the extinction, asymmetry parameter, and single-scatter albedo are verified in the optically thin limit and achieve, respectively, $a_{\text{tol}} = 9.2 \times 10^{-6}$, $a_{\text{tol}} = 1.3 \times 10^{-6}$, $a_{\text{tol}} = 5.3 \times 10^{-7}$, both with and without delta-M scaling.

Appendix F: Calculating the Jacobian matrix by finite differencing

Here, we describe our procedure for calculating the Jacobian matrix using finite differencing, to produce the results in Sect. 4, and the reference results, against which the approximate Jacobian calculation is compared in Sect. 5. The step size for computing the derivatives is chosen to loosely balance the truncation error and rounding error. For finite differencing, the solution accuracy and step size must be chosen so that computational noise does not dominate the derivative calculation. We choose the solution accuracy as 10^{-6}. We cannot perform a perfect optimization of the step size, as the optimal step size for each element of the Jacobian matrix varies with the measurement, and not just the state, so it would require a separate derivative calculation for each ele-

ment of the Jacobian matrix. Instead, we note, from our verification of the approximate derivatives, that the dominant error contribution appears to be due to rounding error. As such, we have chosen a relatively large step size that is still small enough so that it does not unduly compromise accuracy for measurements that are optically close to the sensor and Sun, which have larger, higher-order derivatives (Box, 2002) and therefore a larger truncation error.

To perform the numerical differentiation using the central difference scheme, we must evaluate the forward model twice for each element of the state vector for a total of 2201 times for each base state, which results in a total of almost 240 000 3D RTE solutions for all base states. Though they are still each relatively small, we still accelerate these solutions to reduce the computational expense by noting that the forward- and backward-perturbed RTE solutions will be very close to the base state RTE solution. As such, we initialize the SHDOM solution of the perturbed RTE solutions with the RTE solution of the base state. This method is inspired by the acceleration method for multi-spectral SHDOM solutions described in Evans (1998). It rapidly accelerates the convergence of the perturbed RTE solutions, especially in the optically thickest isotropically scattering cases, where a reduction of ~ 100 in the number of SHDOM solution iterations is achieved. This finite-differencing procedure described above is more computationally efficient than solving the tangent-linear model in Eq. (46) directly using SHDOM. The tangent-linear model tends to take more iterations to converge than a forward problem (e.g., twice as much) and cannot be accelerated, as the tangent-linear problem is very different from the base state RTE problem. The slower convergence of a tangent-linear model is possibly due to the inappropriateness of the initial radiance field used in the SHDOM solution iterations, though we did not investigate this. The adaptive grid is used in these RTE solutions with a splitting accuracy of 0.03. The adaptive grid from the base state RTE solution is used in the forward- and backward-perturbation solutions without further grid splitting. This eliminates the adaptive grid as a source of computational noise in the finite differencing.

Appendix G: Computational cost of the approximate linearization

We now compare the computational cost of the approximate linearization to that of the radiance calculation, specifically the line integrations described in Appendix D, with the forward model. The timing of the SHDOM solutions is therefore not considered, though it is documented elsewhere (Evans, 1998; Pincus and Evans, 2009). In particular, we compare the timing of the subroutine LEVISAPPROX_GRADIENT (in shdomsub4.f) that evaluates both the radiances and the approximate Jacobian to the subroutine RENDER (in shdom-

sub4.f) that only evaluates radiances. As such, we always expect the ratio of their timing to exceed unity.

The evaluation of the approximate gradient is expected to be much more computationally expensive than REN-DER because of the evaluation of the direct-beam derivatives $\frac{\partial I_\odot(r, \Omega)}{\partial a_j}$ in Δf_j, which requires its own traversal of the grid along the solar direction from each grid point. In our case, the implementation pre-computes the pointers to each grid point along these solar paths and their relative contributions to the direct-beam derivative. This pre-processing is not included in our timings reported here, as they contributed negligibly to the relative timing. In LEVISAPPROX_GRADIENT, operations must still be performed on each property grid point along the solar direction. This results in an overall quadratic scaling of LEVISAPPROX_GRADIENT with the number of property grid points. Without this calculation, we expect a linear scaling with the number of adaptive RT grid points for LEVISAPPROX_GRADIENT along with RENDER. To confirm this, we compute timings both with and without this calculation, referred to as "with exact single scatter" and "without exact single scatter", as this is the calculation that ensures accuracy of the approximate Jacobian in the single-scattering limit.

Our timing is based on a homogeneous, plane-parallel cloud. We compute the approximate Jacobian and radiances for a single pixel viewing at nadir repeated 10 000 times to ensure minimal influence of overhead on the calculated time per ray. The Sun is also at zenith. The cloud has a varying number of grid points in the vertical, which controls computational cost, as the line integrals are vertical, and angular resolution. Each cell is optically thin, so that there is one subinterval per cell, and there is no adaptive grid splitting. This also affects the relative timing of the approximate Jacobian to the radiance calculation, as the quadratic cost of the exact single-scatter equation scales with the property grid (not the RT grid). As such, we are examining a worst-case scenario for the relative timing of the approximate Jacobian.

The computational cost as a function of the number of spherical harmonics varies with the number of RT grid points, which can also vary relative to the property grid, based on the optical depth and grid splitting of the medium. Note that we compute only extinction derivatives, which do not require computation of the additional angular integrals in Δf_j beyond what is already computed for the radiance calculation. As such, the scaling of the relative timing with the number of spherical harmonics (NSHs) is likely underestimated compared to derivatives with respect to the phase function. In Fig. G1, we compare the timing ratio of LEVIS-APPROX_GRADIENT to RENDER for each ray. Computations are performed on a 2.3 GHz Intel Core i5, and timings are computed using the Python function time.process_time. The larger timing ratios for lower numbers of spherical harmonics shown in Fig. G1 are due to the larger number of interpolation calculations required per subinterval in the ap-
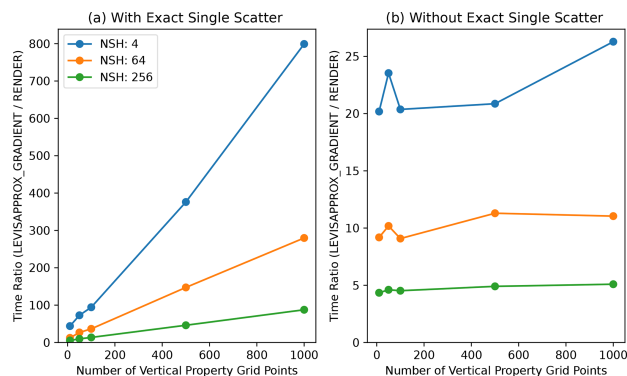
Figure G1. Relative timings of the approximate Jacobian and radiance calculation (LEVISAPPROX_GRADIENT) and the radiance calculation only (RENDER) as a function of the number of grid points for different numbers of spherical harmonics (NSHs). Panel **(a)** shows the relative timing of the default approximate Jacobian calculation, which includes a term ensuring exact single-scatter derivatives. Panel **(b)** shows the relative timing without that term. See the text for details.

proximate Jacobian calculation than in the radiance calculation. The number of spherical harmonic calculations per grid point are common to both subroutines. As expected, the relative timing is independent of the grid size when the exact single-scatter calculation is not performed.

References

Ahn, E., Huang, Y., Siems, S. T., and Manton, M. J.: A Comparison of Cloud Microphysical Properties Derived From MODIS and CALIPSO With In Situ Measurements Over the Wintertime Southern Ocean, J. Geophys. Res.-Atmos., 123, 11120–11140, https://doi.org/10.1029/2018JD028535, 2018.

Alexandrov, M. D., Emde, C., Van Diedenhoven, B., and Cairns, B.: Application of Radon Transform to Multi-Angle Measurements Made by the Research Scanning Polarimeter: A New Approach to Cloud Tomography. Part I: Theory and Tests on Simulated Data, Frontiers in Remote Sensing, 2, 791130, https://doi.org/10.3389/frsen.2021.791130, 2021.

Arridge, S. R. and Schotland, J. C.: Optical tomography: forward and inverse problems, Inverse Problems, 25, 123010, https://doi.org/10.1088/0266-5611/25/12/123010, 2009.

Bal, G.: Inverse transport theory and applications, Inverse Problems, 25, 053001, https://doi.org/10.1088/0266-5611/25/5/053001, 2009.

Bal, G. and Jollivet, A.: Stability estimates in stationary inverse transport, Inverse Probl. Imag., 2, 427–454, https://doi.org/10.3934/ipi.2008.2.427, 2008.

Bellouin, N., Quaas, J., Gryspeerdt, E., Kinne, S., Stier, P., Watson-Parris, D., Boucher, O., Carslaw, K. S., Christensen, M., Daniau, A.-L., Dufresne, J.-L., Feingold, G., Fiedler, S., Forster, P., Gettelman, A., Haywood, J. M., Lohmann, U., Malavelle, F., Mauritsen, T., McCoy, D. T., Myhre, G., Mülmenstädt, J., Neubauer, D., Possner, A., Rugenstein, M., Sato, Y., Schulz, M., Schwartz, S. E., Sourdeval, O., Storelvmo, T., Toll, V., Winker, D., and Stevens, B.: Bounding Global Aerosol Radiative Forcing of Climate Change, Rev. Geophys., 58, e2019RG000660, https://doi.org/10.1029/2019RG000660, 2020.

Bodas-Salcedo, A., Andrews, T., Karmalkar, A. V., and Ringer, M. A.: Cloud liquid water path and radiative feedbacks over the Southern Ocean, Geophys. Res. Lett., 43, 10938–10946, https://doi.org/10.1002/2016GL070770, 2016.

Box, M. A.: Radiative perturbation theory: a review, Environ. Modell. Softw., 17, 95–106, https://doi.org/10.1016/S1364-8152(01)00056-1, 2002.

Buras, R. and Mayer, B.: Efficient unbiased variance reduction techniques for Monte Carlo simulations of radiative transfer in cloudy atmospheres: The solution, J. Quant. Spectrosc. Ra., 112, 434–447, https://doi.org/10.1016/j.jqsrt.2010.10.005, 2011.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C.: A Limited Memory Algorithm for Bound Constrained Optimization, SIAM J. Sci. Comput., 16, 1190–1208, https://doi.org/10.1137/0916069, 1995.

Cahalan, R. F., Oreopoulos, L., Marshak, A., Evans, K. F., Davis, A. B., Pincus, R., Yetzer, K. H., Mayer, B., Davies, R., Ackerman, T. P., Barker, H. W., Clothiaux, E. E., Ellingson, R. G., Garay, M. J., Kassianov, E., Kinne, S., Macke, A., O'hirok, W., Partain, P. T., Prigarin, S. M., Rublev, A. N., Stephens, G. L., Szczap, F., Takara, E. E., Várnai, T., Wen, G., and Zhuravleva, T. B.: THE I3RC: Bringing Together the Most Advanced Radiative Transfer Tools for Cloudy Atmospheres, B. Am. Meteorol. Soc., 86, 1275–1294, https://doi.org/10.1175/BAMS-86-9-1275, 2005.

Chance, B., Cooper, C. E., Delpy, D. T., Reynolds, E. O. R., Arridge, S. R., and Schweiger, M.: Image reconstruction in optical tomography, Philos. T. Roy. Soc. B, 352, 717–726, https://doi.org/10.1098/rstb.1997.0054, 1997.

Chazette, P., Totems, J., Baron, A., Flamant, C., and Bony, S.: Trade-wind clouds and aerosols characterized by airborne horizontal lidar measurements during the EUREC[4]A field campaign, Earth Syst. Sci. Data, 12, 2919–2936, https://doi.org/10.5194/essd-12-2919-2020, 2020.

Chen, K., Li, Q., and Wang, L.: Stability of stationary inverse transport equation in diffusion scaling, Inverse Problems, 34, 025004, https://doi.org/10.1088/1361-6420/aa990c, 2018.

Cornet, C. and Davies, R.: 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, D04202, https://doi.org/10.1029/2007JD008921, 2008.

Culver, J. P., Ntziachristos, V., Holboke, M. J., and Yodh, A. G.: Optimization of optode arrangements for diffuse optical tomography: A singular-value analysis, Opt. Lett., 26, 701–703, https://doi.org/10.1364/OL.26.000701, 2001.

Czerninski, I. and Schechner, Y. Y.: Accelerating Inverse Rendering By Using a GPU and Reuse of Light Paths, arXiv [preprint], https://doi.org/10.48550/arXiv.2110.00085, 30 September 2021.

Davies, R.: The Effect of Finite Geometry on the Three-Dimensional Transfer of Solar Irradiance in Clouds, J. Atmos. Sci., 35, 1712–1725, https://doi.org/10.1175/1520-0469(1978)035<1712:TEOFGO>2.0.CO;2, 1978.

Davis, A., Marshak, A., Cahalan, R., and Wiscombe, W.: The Landsat Scale Break in Stratocumulus as a Three-Dimensional Radiative Transfer Effect: Implications for Cloud Remote Sensing, J. Atmos. Sci., 54, 241–260, https://doi.org/10.1175/1520-0469(1997)054<0241:TLSBIS>2.0.CO;2, 1997.

Davis, A. B. and Marshak, A.: Multiple Scattering in Clouds: Insights from Three-Dimensional Diffusion/P_1 Theory, Nucl. Sci. Eng., 137, 251–280, https://doi.org/10.13182/NSE01-A2190, 2001.

Davis, A. B. and Marshak, A.: Photon propagation in heterogeneous optical media with spatial correlations: enhanced mean-free-paths and wider-than-exponential free-path distributions, J. Quant. Spectrosc. Ra., 84, 3–34, https://doi.org/10.1016/S0022-4073(03)00114-6, 2004.

Davis, A. B., Forster, L., Diner, D. J., and Mayer, B.: Toward Cloud Tomography from Space using MISR and MODIS: The Physics of Image Formation for Opaque Convective Clouds, arXiv [preprint], https://doi.org/10.48550/arXiv.2011.14537, 27 July 2021.

Deutschmann, T., Beirle, S., Frieß, U., Grzegorski, M., Kern, C., Kritten, L., Platt, U., Prados-Román, C., Puķīte, J., Wagner, T., Werner, B., and Pfeilsticker, K.: The Monte Carlo atmospheric radiative transfer model McArtim: Introduction and validation of Jacobians and 3D features, J. Quant. Spectrosc. Ra., 112, 1119–1137, https://doi.org/10.1016/j.jqsrt.2010.12.009, 2011.

Di Girolamo, L., Liang, L., and Platnick, S.: A global view of one-dimensional solar radiative transfer through oceanic water clouds, Geophys. Res. Lett., 37, L18809, https://doi.org/10.1029/2010GL044094, 2010.

Diner, D. J., Xu, F., Garay, M. J., Martonchik, J. V., Rheingans, B. E., Geier, S., Davis, A., Hancock, B. R., Jovanovic, V. M., Bull, M. A., Capraro, K., Chipman, R. A., and McClain, S. C.: The Airborne Multiangle SpectroPolarimetric Imager (AirMSPI): a new tool for aerosol and cloud remote sensing, Atmos. Meas. Tech., 6, 2007–2025, https://doi.org/10.5194/amt-6-2007-2013, 2013.

Doicu, A. and Efremenko, D. S.: Linearizations of the Spherical Harmonic Discrete Ordinate Method (SHDOM), Atmosphere, 10, 292, https://doi.org/10.3390/atmos10060292, 2019.

Doicu, A., Mishchenko, M. I., and Trautmann, T.: Electromagnetic scattering by discrete random media illuminated by a Gaussian beam II: Solution of the radiative transfer equation, J. Quant. Spectrosc. Ra., 256, 107297, https://doi.org/10.1016/j.jqsrt.2020.107297, 2020.

Doicu, A., Doicu, A., Efremenko, D., and Trautmann, T.: Cloud tomographic retrieval algorithms. I: Surrogate minimization method, J. Quant. Spectrosc. Ra., 277, 107954, https://doi.org/10.1016/j.jqsrt.2021.107954, 2022a.

Doicu, A., Doicu, A., Efremenko, D., and Trautmann, T.: Cloud tomographic retrieval algorithms. II: Adjoint method, J. Quant. Spectrosc. Ra., 108177, https://doi.org/10.1016/j.jqsrt.2022.108177, 2022b.

Dubovik, O., Holben, B. N., Lapyonok, T., Sinyuk, A., Mishchenko, M. I., Yang, P., and Slutsker, I.: Non-spherical aerosol retrieval method employing light scattering by spheroids, Geophys. Res. Lett., 29, 54-1–54-4, https://doi.org/10.1029/2001GL014506, 2002.

Dubovik, O., Herman, M., Holdak, A., Lapyonok, T., Tanré, D., Deuzé, J. L., Ducos, F., Sinyuk, A., and Lopatin, A.: Statistically optimized inversion algorithm for enhanced retrieval of aerosol properties from spectral multi-angle polarimetric satellite observations, Atmos. Meas. Tech., 4, 975–1018, https://doi.org/10.5194/amt-4-975-2011, 2011.

Dwight, R. P. and Brezillon, J.: Effect of Approximations of the Discrete Adjoint on Gradient-Based Optimization, AIAA J., 44, 3022–3031, https://doi.org/10.2514/1.21744, 2006.

Dzambo, A. M., L'Ecuyer, T., Sinclair, K., van Diedenhoven, B., Gupta, S., McFarquhar, G., O'Brien, J. R., Cairns, B., Wasilewski, A. P., and Alexandrov, M.: Joint cloud water path and rainwater path retrievals from airborne ORACLES observations, Atmos. Chem. Phys., 21, 5513–5532, https://doi.org/10.5194/acp-21-5513-2021, 2021.

Efremenko, D. S., Loyola, D. G., Doicu, A., and Spurr, R. J. D.: Multi-core-CPU and GPU-accelerated radiative transfer models based on the discrete ordinate method, Comput. Phys. Commun., 185, 3079–3089, https://doi.org/10.1016/j.cpc.2014.07.018, 2014.

Emde, C., Barlakas, V., Cornet, C., Evans, F., Korkin, S., Ota, Y., Labonnote, L. C., Lyapustin, A., Macke, A., Mayer, B., and Wendisch, M.: IPRT polarized radiative transfer model intercomparison project – Phase A, J. Quant. Spectrosc. Ra., 164, 8–36, https://doi.org/10.1016/j.jqsrt.2015.05.007, 2015.

Emde, C., Buras-Schnell, R., Kylling, A., Mayer, B., Gasteiger, J., Hamann, U., Kylling, J., Richter, B., Pause, C., Dowling, T., and Bugliaro, L.: The libRadtran software package for radiative transfer calculations (version 2.0.1), Geosci. Model Dev., 9, 1647–1672, https://doi.org/10.5194/gmd-9-1647-2016, 2016.

Emde, C., Barlakas, V., Cornet, C., Evans, F., Wang, Z., Labonotte, L. C., Macke, A., Mayer, B., and Wendisch, M.: IPRT polarized radiative transfer model intercomparison project – Three-dimensional test cases (phase B), J. Quant. Spectrosc. Ra., 209, 19–44, https://doi.org/10.1016/j.jqsrt.2018.01.024, 2018.

Endo, S., Fridlind, A. M., Lin, W., Vogelmann, A. M., Toto, T., Ackerman, A. S., McFarquhar, G. M., Jackson, R. C., Jonsson, H. H., and Liu, Y.: RACORO continental boundary layer cloud investigations: 2. Large-eddy simulations of cumulus clouds and evaluation with in situ and ground-based observations, J. Geophys. Res.-Atmos., 120, 5993–6014, https://doi.org/10.1002/2014JD022525, 2015.

Eppstein, M. J., Fedele, F., Laible, J., Zhang, C., Godavarty, A., and Sevick-Muraca, E. M.: A comparison of exact and approximate adjoint sensitivities in fluorescence tomography, IEEE T. Med. Imaging, 22, 1215–1223, https://doi.org/10.1109/TMI.2003.818165, 2003.

Evans, K. F.: The Spherical Harmonics Discrete Ordinate Method for Three-Dimensional Atmospheric Radiative Transfer, J. Atmos. Sci., 55, 429–446, https://doi.org/10.1175/1520-0469(1998)055<0429:TSHDOM>2.0.CO;2, 1998.

Evans, K. F.: Spherical Harmonics Discrete Ordinate Method, University of Colorado, Boulder [code], https://nit.coloradolinux.com/~evans/shdom/shdom.tar.gz, last access: 23 March 2023.

Ewald, F., Groß, S., Wirth, M., Delanoë, J., Fox, S., and Mayer, B.: Why we need radar, lidar, and solar radiance observations to constrain ice cloud microphysics, Atmos. Meas. Tech., 14, 5029–5047, https://doi.org/10.5194/amt-14-5029-2021, 2021.

Eytan, E., Khain, A., Pinsky, M., Altaratz, O., Shpund, J., and Koren, I.: Shallow Cumulus Properties as Captured by Adiabatic Fraction in High-Resolution LES Simulations, J. Atmos. Sci., 79, 409–428, https://doi.org/10.1175/JAS-D-21-0201.1, 2022.

Fielding, M. D., Chiu, J. C., Hogan, R. J., and Feingold, G.: A novel ensemble method for retrieving properties of warm cloud in 3-D using ground-based scanning radar and zenith radiances, J. Geophys. Res.-Atmos., 119, 10912–10930, https://doi.org/10.1002/2014JD021742, 2014.

Forster, L., Davis, A. B., Diner, D. J., and Mayer, B.: Toward Cloud Tomography from Space Using MISR and MODIS: Locating the "Veiled Core" in Opaque Convective Clouds, J. Atmos. Sci., 78, 155–166, https://doi.org/10.1175/JAS-D-19-0262.1, 2020.

Fu, D., Di Girolamo, L., Liang, L., and Zhao, G.: Regional Biases in MODIS Marine Liquid Water Cloud Drop Effective Radius Deduced Through Fusion With MISR, J. Geophys. Res.-Atmos., 124, 13182–13196, https://doi.org/10.1029/2019JD031063, 2019.

Gao, M., Franz, B. A., Knobelspiesse, K., Zhai, P.-W., Martins, V., Burton, S., Cairns, B., Ferrare, R., Gales, J., Hasekamp, O., Hu, Y., Ibrahim, A., McBride, B., Puthukkudy, A., Werdell, P. J., and Xu, X.: Efficient multi-angle polarimetric inversion of aerosols and ocean color powered by a deep neural network forward model, Atmos. Meas. Tech., 14, 4083–4110, https://doi.org/10.5194/amt-14-4083-2021, 2021.

Garay, M. J., Davis, A. B., and Diner, D. J.: Tomographic reconstruction of an aerosol plume using passive multiangle observations from the MISR satellite instrument, Geophys. Res. Lett., 43, 12590–12596, https://doi.org/10.1002/2016GL071479, 2016.

Gerber, H., Jensen, J. B., Davis, A. B., Marshak, A., and Wiscombe, W. J.: Spectral Density of Cloud Liquid Water Content at High Frequencies, J. Atmos. Sci., 58, 497–503, https://doi.org/10.1175/1520-0469(2001)058<0497:SDOCLW>2.0.CO;2, 2001.

Girolamo, L. D.: Reciprocity principle applicable to reflected radiance measurements and the searchlight problem, Appl. Optics, 38, 3196–3198, https://doi.org/10.1364/AO.38.003196, 1999.

Girolamo, L. D.: Reciprocity principle for radiative transfer models that use periodic boundary conditions, J. Quant. Spectrosc. Ra., 73, 23–27, https://doi.org/10.1016/S0022-4073(01)00166-2, 2002.

Gordon, I. E., Rothman, L. S., Hargreaves, R. J., Hashemi, R., Karlovets, E. V., Skinner, F. M., Conway, E. K., Hill, C., Kochanov, R. V., Tan, Y., Wcisło, P., Finenko, A. A., Nelson, K., Bernath, P. F., Birk, M., Boudon, V., Campargue, A., Chance, K. V., Coustenis, A., Drouin, B. J., Flaud, J.-M., Gamache, R. R., Hodges, J. T., Jacquemart, D., Mlawer, E. J., Nikitin, A. V., Perevalov, V. I., Rotger, M., Tennyson, J., Toon, G. C., Tran, H., Tyuterev, V. G., Adkins, E. M., Baker, A., Barbe, A., Canè, E., Császár, A. G., Dudaryonok, A., Egorov, O., Fleisher, A. J., Fleurbaey, H., Foltynowicz, A., Furtenbacher, T., Harrison, J. J., Hartmann, J.-M., Horneman, V.-M., Huang, X., Karman, T., Karns, J., Kassi, S., Kleiner, I., Kofman, V., Kwabia-Tchana, F., Lavrentieva, N. N., Lee, T. J., Long, D. A., Lukashevskaya, A. A., Lyulin, O. M., Makhnev, V. Yu., Matt, W., Massie, S. T., Melosso, M., Mikhailenko, S. N., Mondelain, D., Müller, H. S. P., Naumenko, O. V., Perrin, A., Polyansky, O. L., Raddaoui, E., Raston, P. L., Reed, Z. D., Rey, M., Richard, C., Tóbiás, R., Sadiek, I., Schwenke, D. W., Starikova, E., Sung, K., Tamassia, F., Tashkun, S. A., Vander Auwera, J., Vasilenko, I. A., Vigasin, A. A., Villanueva, G. L., Vispoel, B., Wagner, G., Yachmenev, A., and Yurchenko, S. N.: The HITRAN2020 molecular spectroscopic database, J. Quant. Spectrosc. Ra., 277, 107949, https://doi.org/10.1016/j.jqsrt.2021.107949, 2022.

Grosvenor, D. P., Sourdeval, O., Zuidema, P., Ackerman, A., Alexandrov, M. D., Bennartz, R., Boers, R., Cairns, B., Chiu, J. C., Christensen, M., Deneke, H., Diamond, M., Feingold, G., Fridlind, A., Hünerbein, A., Knist, C., Kollias, P., Marshak, A., McCoy, D., Merk, D., Painemal, D., Rausch, J., Rosenfeld, D., Russchenberg, H., Seifert, P., Sinclair, K., Stier, P., van Diedenhoven, B., Wendisch, M., Werner, F., Wood, R., Zhang, Z., and Quaas, J.: Remote Sensing of Droplet Number Concentration in Warm Clouds: A Review of the Current State of Knowledge and Perspectives, Rev. Geophys., 56, 409–453, https://doi.org/10.1029/2017RG000593, 2018.

Guillaume, A., Kahn, B. H., Yue, Q., Fetzer, E. J., Wong, S., Manipon, G. J., Hua, H., and Wilson, B. D.: Horizontal and Vertical Scaling of Cloud Geometry Inferred from *CloudSat* Data, J. Atmos. Sci., 75, 2187–2197, https://doi.org/10.1175/JAS-D-17-0111.1, 2018.

Hack, J. J., Caron, J. M., Danabasoglu, G., Oleson, K. W., Bitz, C., and Truesdale, J. E.: CCSM–CAM3 Climate Simulation Sensitivity to Changes in Horizontal Resolution, J. Climate, 19, 2267–2289, https://doi.org/10.1175/JCLI3764.1, 2006.

Hasekamp, O. P. and Landgraf, J.: Linearization of vector radiative transfer with respect to aerosol properties and its use in satellite remote sensing, J. Geophys. Res.-Atmos., 110, D04203, https://doi.org/10.1029/2004JD005260, 2005.

Hill, P. G., Hogan, R. J., Manners, J., and Petch, J. C.: Parametrizing the horizontal inhomogeneity of ice water content using CloudSat data products, Q. J. Roy. Meteor. Soc., 138, 1784–1793, https://doi.org/10.1002/qj.1893, 2012.

Holodovsky, V., Schechner, Y. Y., Levin, A., Levis, A., and Aides, A.: In-situ multi-view multi-scattering stochastic tomography, in: 2016 IEEE International Conference on Computational Photography (ICCP), Evanston, IL, 13–15 May 2016, IEEE, 1–12, https://doi.org/10.1109/ICCPHOT.2016.7492869, 2016.

Hoyer, S. and Hamman, J.: xarray: N-D labeled Arrays and Datasets in Python, Journal of Open Research Software, 5, 10, https://doi.org/10.5334/jors.148, 2017.

Huang, D., Liu, Y., and Wiscombe, W.: Determination of cloud liquid water distribution using 3D cloud tomography, J. Geophys. Res.-Atmos., 113, D13201, https://doi.org/10.1029/2007JD009133, 2008.

Huang, D., Gasiewski, A., and Wiscombe, W.: Tomographic retrieval of cloud liquid water fields from a single scanning microwave radiometer aboard a moving platform – Part 2: Observation system simulation experiments, Atmos. Chem. Phys., 10, 6699–6709, https://doi.org/10.5194/acp-10-6699-2010, 2010.

Jiang, W., Zhan, Y., Xi, S., Huang, D. D., and Lu, J.: Compressive Sensing-Based 3-D Rain Field Tomographic Reconstruction Using Simulated Satellite Signals, IEEE T. Geosci. Remote, 60, 1–13, https://doi.org/10.1109/TGRS.2021.3063617, 2022.

Jimenez, P. A., Hacker, J. P., Dudhia, J., Haupt, S. E., Ruiz-Arias, J. A., Gueymard, C. A., Thompson, G., Eidhammer, T., and Deng, A.: WRF-Solar: Description and Clear-Sky Assessment of an Augmented NWP Model for Solar Power Prediction, B. Am. Meteorol. Soc., 97, 1249–1264, https://doi.org/10.1175/BAMS-D-14-00279.1, 2016.

Jones, A. L. and Di Girolamo, L.: Design and Verification of a New Monochromatic Thermal Emission Component for the I3RC Community Monte Carlo Model, J. Atmos. Sci., 75, 885–906, https://doi.org/10.1175/JAS-D-17-0251.1, 2018.

Kahn, R. A., Li, W.-H., Moroney, C., Diner, D. J., Martonchik, J. V., and Fishbein, E.: Aerosol source plume physical characteristics from space-based multiangle imaging, J. Geophys. Res.-Atmos., 112, D11205, https://doi.org/10.1029/2006JD007647, 2007.

Kanewala, U. and Bieman, J. M.: Testing scientific software: A systematic literature review, Inform. Software Tech., 56, 1219–1232, https://doi.org/10.1016/j.infsof.2014.05.006, 2014.

Kato, S. and Marshak, A.: Solar zenith and viewing geometry-dependent errors in satellite retrieved cloud optical thickness: Marine stratocumulus case, J. Geophys. Res.-Atmos., 114, D01202, https://doi.org/10.1029/2008JD010579, 2009.

King, N. J. and Vaughan, G.: Using passive remote sensing to retrieve the vertical variation of cloud droplet size in marine stratocumulus: An assessment of information content and the potential for improved retrievals from hyperspectral measurements, J. Geophys. Res.-Atmos., 117, D15206, https://doi.org/10.1029/2012JD017896, 2012.

Klose, A. D. and Hielscher, A. H.: Optical tomography using the time-independent equation of radiative transfer – Part 2: inverse model, J. Quant. Spectrosc. Ra., 72, 715–732, https://doi.org/10.1016/S0022-4073(01)00151-0, 2002.

Langmore, I., Davis, A. B., and Bal, G.: Multipixel Retrieval of Structural and Optical Parameters in a 2-D Scene With a Path-Recycling Monte Carlo Forward Model and a New Bayesian Inference Engine, IEEE T. Geosci. Remote, 51, 2903–2919, https://doi.org/10.1109/TGRS.2012.2217380, 2013.

Lebsock, M. and Su, H.: Application of active spaceborne remote sensing for understanding biases between passive cloud water path retrievals, J. Geophys. Res.-Atmos., 119, 8962–8979, https://doi.org/10.1002/2014JD021568, 2014.

Lee, B., Di Girolamo, L., Zhao, G., and Zhan, Y.: Three-Dimensional Cloud Volume Reconstruction from the Multi-angle Imaging SpectroRadiometer, Remote Sens.-Basel, 10, 1858, https://doi.org/10.3390/rs10111858, 2018.

Lee, E. K. H., Wardenier, J. P., Prinoth, B., Parmentier, V., Grimm, S. L., Baeyens, R., Carone, L., Christie, D., Deitrick, R., Kitzmann, D., Mayne, N., Roman, M., and Thorsbro, B.: 3D Radiative Transfer for Exoplanet Atmospheres. gCMCRT: A GPU-accelerated MCRT Code, Astrophys. J., 929, 180, https://doi.org/10.3847/1538-4357/ac61d6, 2022.

Levis, A., Schechner, Y. Y., Aides, A., and Davis, A. B.: Airborne Three-Dimensional Cloud Tomography, in: 2015 IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 7–13 December 2015, IEEE, 3379–3387, https://doi.org/10.1109/ICCV.2015.386, 2015.

Levis, A., Schechner, Y. Y., and Davis, A. B.: Multiple-Scattering Microphysics Tomography, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017, IEEE, 5797–5806, https://doi.org/10.1109/CVPR.2017.614, 2017.

Levis, A., Schechner, Y. Y., Davis, A. B., and Loveridge, J.: Multi-View Polarimetric Scattering Cloud Tomography and Retrieval of Droplet Size, Remote Sens., 12, 2831, https://doi.org/10.3390/rs12172831, 2020.

Liemert, A. and Kienle, A.: Exact and efficient solution of the radiative transport equation for the semi-infinite medium, Sci. Rep.-UK, 3, 2018, https://doi.org/10.1038/srep02018, 2013.

Loeub, T., Levis, A., Holodovsky, V., and Schechner, Y. Y.: Monotonicity Prior for Cloud Tomography, in: Computer Vision – ECCV 2020, edited by: Vedaldi, A., Bischof, H., Brox, T., and Frahm, J. M., Lecture Notes in Computer Science, Springer, Cham, 12363, 283–299, https://doi.org/10.1007/978-3-030-58523-5_17, 2020.

Loveridge, J., Levis, A., Aides, A., Forster, L., and Holodovsky, V.: Atmospheric Tomography with 3D Radiative Transfer, v4.1.2, Zenodo [code], https://doi.org/10.5281/zenodo.7062466, 2022.

Loveridge, J., Levis, A., Di Girolamo, L., Holodovsky, V., Forster, L., Davis, A. B., and Schechner, Y. Y.: Retrieving 3D distributions of atmospheric particles using Atmospheric Tomography with 3D Radiative Transfer – Part 2: Local optimization, Atmos. Meas. Tech., submitted, 2023a.

Loveridge, J., Levis, A., Holodovsky, V., and Forster, L.: Atmospheric Tomography with 3D Radiative Transfer, GitHub [code], https://github.com/CloudTomography/AT3D, last access: 28 March 2023b.

Marchand, R. and Ackerman, T.: Evaluation of radiometric measurements from the NASA Multiangle Imaging Spectroradiometer (MISR): Two- and three-dimensional radiative transfer modeling of an inhomogeneous stratocumulus cloud deck, J. Geophys. Res.-Atmos., 109, D18208, https://doi.org/10.1029/2004JD004710, 2004.

Marshak, A., Davis, A., Wiscombe, W., and Cahalan, R.: Radiative smoothing in fractal clouds, J. Geophys. Res.-Atmos., 100, 26247–26261, https://doi.org/10.1029/95JD02895, 1995.

Marshak, A., Davis, A., Wiscombe, W., and Cahalan, R.: Scale Invariance in Liquid Water Distributions in Marine Stratocumulus. Part II: Multifractal Properties and Intermittency Issues, J. Atmos. Sci., 54, 1423–1444, https://doi.org/10.1175/1520-0469(1997)054<1423:SIILWD>2.0.CO;2, 1997.

Marshak, A., Davis, A., Cahalan, R. F., and Wiscombe, W.: Nonlocal independent pixel approximation: direct and inverse problems, IEEE T. Geosci. Remote, 36, 192–205, https://doi.org/10.1109/36.655329, 1998a.

Marshak, A., Davis, A., Wiscombe, W., and Cahalan, R.: Radiative effects of sub-mean free path liquid water variability observed in stratiform clouds, J. Geophys. Res.-Atmos., 103, 19557–19567, https://doi.org/10.1029/98JD01728, 1998b.

Marshak, A., Platnick, S., Várnai, T., Wen, G., and Cahalan, R. F.: Impact of three-dimensional radiative effects on satellite retrievals of cloud droplet sizes, J. Geophys. Res.-Atmos., 111, D09207, https://doi.org/10.1029/2005JD006686, 2006.

Martelli, F., Binzoni, T., Pifferi, A., Spinelli, L., Farina, A., and Torricelli, A.: There's plenty of light at the bottom: statistics of photon penetration depth in random media, Sci. Rep.-UK, 6, 27057, https://doi.org/10.1038/srep27057, 2016.

Martin, W., Cairns, B., and Bal, G.: Adjoint methods for adjusting three-dimensional atmosphere and sur-

face properties to fit multi-angle/multi-pixel polarimetric measurements, J. Quant. Spectrosc. Ra., 144, 68–85, https://doi.org/10.1016/j.jqsrt.2014.03.030, 2014.

Martin, W. G. K. and Hasekamp, O. P.: A demonstration of adjoint methods for multi-dimensional remote sensing of the atmosphere and surface, J. Quant. Spectrosc. Ra., 204, 215–231, https://doi.org/10.1016/j.jqsrt.2017.09.031, 2018.

McBride, B. A., Martins, J. V., Barbosa, H. M. J., Birmingham, W., and Remer, L. A.: Spatial distribution of cloud droplet size properties from Airborne Hyper-Angular Rainbow Polarimeter (AirHARP) measurements, Atmos. Meas. Tech., 13, 1777–1796, https://doi.org/10.5194/amt-13-1777-2020, 2020.

McFarlane, S. A., Mather, J. H., and Mlawer, E. J.: ARM's Progress on Improving Atmospheric Broadband Radiative Fluxes and Heating Rates, Meteor. Mon., 57, 20.1–20.24, https://doi.org/10.1175/AMSMONOGRAPHS-D-15-0046.1, 2016.

Morrison, H., van Lier-Walqui, M., Fridlind, A. M., Grabowski, W. W., Harrington, J. Y., Hoose, C., Korolev, A., Kumjian, M. R., Milbrandt, J. A., Pawlowska, H., Posselt, D. J., Prat, O. P., Reimel, K. J., Shima, S.-I., van Diedenhoven, B., and Xue, L.: Confronting the Challenge of Modeling Cloud and Precipitation Microphysics, J. Adv. Model. Earth Sy., 12, e2019MS001689, https://doi.org/10.1029/2019MS001689, 2020.

Muller, J.-P., Mandanayake, A., Moroney, C., Davies, R., Diner, D. J., and Paradise, S.: MISR stereoscopic image matchers: techniques and results, IEEE T. Geosci. Remote, 40, 1547–1559, https://doi.org/10.1109/TGRS.2002.801160, 2002.

Muller, J.-P., Denis, M.-A., Dundas, R. D., Mitchell, K. L., Naud, C., and Mannstein, H.: Stereo cloud-top heights and cloud fraction retrieval from ATSR-2, Int. J. Remote Sens., 28, 1921–1938, https://doi.org/10.1080/01431160601030975, 2007.

Nakajima, T. and Tanaka, M.: Algorithms for radiative intensity calculations in moderately thick atmospheres using a truncation approximation, J. Quant. Spectrosc. Ra., 40, 51–69, https://doi.org/10.1016/0022-4073(88)90031-3, 1988.

Niu, H., Lin, Z.-J., Tian, F., Dhamne, S., and Liu, H.: Comprehensive investigation of three-dimensional diffuse optical tomography with depth compensation algorithm, J. Biomed. Opt., 15, 046005, https://doi.org/10.1117/1.3462986, 2010.

Nocedal, J. and Wright, S.: Fundamentals of Unconstrained Optimization, in: Numerical Optimization, Springer New York, New York, NY, 10–29, https://doi.org/10.1007/978-0-387-40065-5_2, 2006.

Painemal, D., Spangenberg, D., Smith Jr., W. L., Minnis, P., Cairns, B., Moore, R. H., Crosbie, E., Robinson, C., Thornhill, K. L., Winstead, E. L., and Ziemba, L.: Evaluation of satellite retrievals of liquid clouds from the GOES-13 imager and MODIS over the midlatitude North Atlantic during the NAAMES campaign, Atmos. Meas. Tech., 14, 6633–6646, https://doi.org/10.5194/amt-14-6633-2021, 2021.

Peterson, P.: F2PY: a tool for connecting Fortran and Python programs, International Journal of Computational Science and Engineering, 4, 296–305, https://doi.org/10.1504/IJCSE.2009.029165, 2009.

Pincus, R. and Evans, K. F.: 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, https://doi.org/10.1175/2009JAS3137.1, 2009.

Ramon, D., Steinmetz, F., Jolivet, D., Compiègne, M., and Frouin, R.: Modeling polarized radiative transfer in the ocean-atmosphere system with the GPU-accelerated SMART-G Monte Carlo code, J. Quant. Spectrosc. Ra., 222–223, 89–107, https://doi.org/10.1016/j.jqsrt.2018.10.017, 2019.

Raschke, E., Ohmura, A., Rossow, W. B., Carlson, B. E., Zhang, Y.-C., Stubenrauch, C., Kottek, M., and Wild, M.: Cloud effects on the radiation budget based on ISCCP data (1991 to 1995), Int. J. Climatol., 25, 1103–1125, https://doi.org/10.1002/joc.1157, 2005.

Ronen, R., Schechner, Y. Y., and Eytan, E.: 4D Cloud Scattering Tomography, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 10–17 October 2021, IEEE, 5520–5529, https://doi.org/10.1109/ICCV48922.2021.00547, 2021.

Ronen, R., Holodovsky, V., and Schechner, Y. Y.: Variable Imaging Projection Cloud Scattering Tomography, IEEE T. Pattern Anal., 1–12, https://doi.org/10.1109/TPAMI.2022.3195920, 2022.

Saito, M., Yang, P., Hu, Y., Liu, X., Loeb, N., Smith Jr., W. L., and Minnis, P.: An Efficient Method for Microphysical Property Retrievals in Vertically Inhomogeneous Marine Water Clouds Using MODIS-CloudSat Measurements, J. Geophys. Res.-Atmos., 124, 2174–2193, https://doi.org/10.1029/2018JD029659, 2019.

Saito, M., Yang, P., Ding, J., and Liu, X.: A Comprehensive Database of the Optical Properties of Irregular Aerosol Particles for Radiative Transfer Simulations, J. Atmos. Sci., 78, 2089–2111, https://doi.org/10.1175/JAS-D-20-0338.1, 2021.

Schilling, K., Schechner, Y. Y., and Koren, I.: CloudCT - Computed Tomography Of Clouds By A Small Satellite Formation, in: 12th IAA symposium on Small Satellites for Earth Observation, Berlin Germany, 6–10 May 2019, https://www.cloudct.space/_files/ugd/aef793_e2334d0332014022aeb4a5891ff12e92.pdf (last access: 20 March 2023), 2019.

Seethala, C.: Evaluating the state-of-the-art of and errors in 1D satellite cloud liquid water path retrievals with large eddy simulations and realistic radiative transfer models, PhD thesis, University of Hamburg, Hamburg, https://doi.org/10.17617/2.1404650, 2012.

Sherwood, S. C., Bony, S., and Dufresne, J.-L.: Spread in model climate sensitivity traced to atmospheric convective mixing, Nature, 505, 37–42, https://doi.org/10.1038/nature12829, 2014.

Sherwood, S. C., Webb, M. J., Annan, J. D., Armour, K. C., Forster, P. M., Hargreaves, J. C., Hegerl, G., Klein, S. A., Marvel, K. D., Rohling, E. J., Watanabe, M., Andrews, T., Braconnot, P., Bretherton, C. S., Foster, G. L., Hausfather, Z., von der Heydt, A. S., Knutti, R., Mauritsen, T., Norris, J. R., Proistosescu, C., Rugenstein, M., Schmidt, G. A., Tokarska, K. B., and Zelinka, M. D.: An Assessment of Earth's Climate Sensitivity Using Multiple Lines of Evidence, Rev. Geophys., 58, e2019RG000678, https://doi.org/10.1029/2019RG000678, 2020.

Shi, H.-J. M., Xie, Y., Byrd, R., and Nocedal, J.: A Noise-Tolerant Quasi-Newton Algorithm for Unconstrained Optimization, arXiv [preprint], https://doi.org/10.48550/arXiv.2010.04352, 9 September 2021.

Sourdeval, O., Labonnote, L. C., Brogniez, G., Jourdan, O., Pelon, J., and Garnier, A.: A variational approach for retrieving ice cloud properties from infrared measurements: applica-

tion in the context of two IIR validation campaigns, Atmos. Chem. Phys., 13, 8229–8244, https://doi.org/10.5194/acp-13-8229-2013, 2013.

Stephens, G. L. and Kummerow, C. D.: The Remote Sensing of Clouds and Precipitation from Space: A Review, J. Atmos. Sci., 64, 3742–3765, https://doi.org/10.1175/2006JAS2375.1, 2007.

Tian, F., Niu, H., Khadka, S., Lin, Z.-J., and Liu, H.: Algorithmic depth compensation improves quantification and noise suppression in functional diffuse optical tomography, Biomed. Opt. Express, 1, 441–452, https://doi.org/10.1364/BOE.1.000441, 2010.

van Harten, G., Diner, D. J., Daugherty, B. J. S., Rheingans, B. E., Bull, M. A., Seidel, F. C., Chipman, R. A., Cairns, B., Wasilewski, A. P., and Knobelspiesse, K. D.: Calibration and validation of Airborne Multiangle SpectroPolarimetric Imager (AirMSPI) polarization measurements, Appl. Optics, 57, 4499–4513, https://doi.org/10.1364/AO.57.004499, 2018.

Van Weverberg, K., Morcrette, C. J., Petch, J., Klein, S. A., Ma, H.-Y., Zhang, C., Xie, S., Tang, Q., Gustafson Jr., W. I., Qian, Y., Berg, L. K., Liu, Y., Huang, M., Ahlgrimm, M., Forbes, R., Bazile, E., Roehrig, R., Cole, J., Merryfield, W., Lee, W.-S., Cheruy, F., Mellul, L., Wang, Y.-C., Johnson, K., and Thieman, M. M.: CAUSES: Attribution of Surface Radiation Biases in NWP and Climate Models near the U.S. Southern Great Plains, J. Geophys. Res.-Atmos., 123, 3612–3644, https://doi.org/10.1002/2017JD027188, 2018.

Vial, J., Bony, S., Stevens, B., and Vogel, R.: Mechanisms and Model Diversity of Trade-Wind Shallow Cumulus Cloud Feedbacks: A Review, in: Shallow Clouds, Water Vapor, Circulation, and Climate Sensitivity, edited by: Pincus, R., Winker, D., Bony, S., and Stevens, B., Springer International Publishing, Cham, 159–181, https://doi.org/10.1007/978-3-319-77273-8_8, 2018.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors: SciPy 1.0: fundamental algorithms for scientific computing in Python, Nat. Methods, 17, 261–272, https://doi.org/10.1038/s41592-019-0686-2, 2020.

Wang, C., Platnick, S., Zhang, Z., Meyer, K., and Yang, P.: Retrieval of ice cloud properties using an optimal estimation algorithm and MODIS infrared observations: 1. Forward model, error analysis, and information content, J. Geophys. Res.-Atmos., 121, 5809–5826, https://doi.org/10.1002/2015JD024526, 2016.

Wang, Y., Guo, M., Zhao, Y., and Jiang, J.: GPUs-RRTMG_LW: high-efficient and scalable computing for a longwave radiative transfer model on multiple GPUs, J. Supercomput., 77, 4698–4717, https://doi.org/10.1007/s11227-020-03451-3, 2021.

Wang, Z., Cui, S., Yang, J., Gao, H., Liu, C., and Zhang, Z.: A novel hybrid scattering order-dependent variance reduction method for Monte Carlo simulations of radiative transfer in cloudy atmosphere, J. Quant. Spectrosc. Ra., 189, 283–302, https://doi.org/10.1016/j.jqsrt.2016.12.002, 2017.

Wiscombe, W. J.: The Delta–M Method: Rapid Yet Accurate Radiative Flux Calculations for Strongly Asymmetric Phase Functions,

J. Atmos. Sci., 34, 1408–1422, https://doi.org/10.1175/1520-0469(1977)034<1408:TDMRYA>2.0.CO;2, 1977.

Xie, X. and Zhang, M.: Scale-aware parameterization of liquid cloud inhomogeneity and its impact on simulated climate in CESM, J. Geophys. Res.-Atmos., 120, 8359–8371, https://doi.org/10.1002/2015JD023565, 2015.

Xu, F., Diner, D. J., Dubovik, O., and Schechner, Y.: A Correlated Multi-Pixel Inversion Approach for Aerosol Remote Sensing, Remote Sens., 11, 746, https://doi.org/10.3390/rs11070746, 2019.

Yao, R., Intes, X., and Fang, Q.: Direct approach to compute Jacobians for diffuse optical tomography using perturbation Monte Carlo-based photon "replay," Biomed. Opt. Express, 9, 4588–4603, https://doi.org/10.1364/BOE.9.004588, 2018.

Ye, J. C., Webb, K. J., Millane, R. P., and Downar, T. J.: Modified distorted Born iterative method with an approximate Fréchet derivative for optical diffusion tomography, J. Opt. Soc. Am. A, 16, 1814–1826, https://doi.org/10.1364/JOSAA.16.001814, 1999.

Zawada, D. J., Bourassa, A. E., and Degenstein, D. A.: Two-dimensional analytic weighting functions for limb scattering, J. Quant. Spectrosc. Ra., 200, 125–136, https://doi.org/10.1016/j.jqsrt.2017.06.008, 2017.

Zawada, D. J., Rieger, L. A., Bourassa, A. E., and Degenstein, D. A.: Tomographic retrievals of ozone with the OMPS Limb Profiler: algorithm description and preliminary results, Atmos. Meas. Tech., 11, 2375–2393, https://doi.org/10.5194/amt-11-2375-2018, 2018.

Zhang, L. and Zhang, G.: Brief review on learning-based methods for optical tomography, J. Innov. Opt. Heal. Sci., 12, 1930011, https://doi.org/10.1142/S1793545819300118, 2019.

Zhang, Z., Ackerman, A. S., Feingold, G., Platnick, S., Pincus, R., and Xue, H.: Effects of cloud horizontal inhomogeneity and drizzle on remote sensing of cloud droplet effective radius: Case studies based on large-eddy simulations, J. Geophys. Res.-Atmos., 117, D19208, https://doi.org/10.1029/2012JD017655, 2012.

Zhao, G. and Di Girolamo, L.: Statistics on the macrophysical properties of trade wind cumuli over the tropical western Atlantic, J. Geophys. Res.-Atmos., 112, D10204, https://doi.org/10.1029/2006JD007371, 2007.

Zhao, H. and Zhong, Y.: Instability of an Inverse Problem for the Stationary Radiative Transport Near the Diffusion Limit, SIAM J. Math. Anal., 51, 3750–3768, https://doi.org/10.1137/18M1222582, 2019.

Zhou, C., Zelinka, M. D., and Klein, S. A.: Impact of decadal cloud variations on the Earth's energy budget, Nat. Geosci., 9, 871–874, https://doi.org/10.1038/ngeo2828, 2016.

Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J.: Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization, ACM T. Math. Software, 23, 550–560, https://doi.org/10.1145/279232.279236, 1997.

Zinner, T., Mayer, B., and Schröder, M.: Determination of three-dimensional cloud structures from high-resolution radiance data, J. Geophys. Res.-Atmos., 111, D08204, https://doi.org/10.1029/2005JD006062, 2006.