Author’s response

Eradiate: An Accurate and Flexible Radiative Transfer Model for Earth Observation and Atmospheric Science

Dear Editor,

Please find attached the revised version of our manuscript “Eradiate: An Accurate and Flexible Radiative Transfer Model for Earth Observation and Atmospheric Science” submitted for publication to your journal Geoscientific Model Development.

We thank you and the reviewers for your helpful comments and suggestions. We implemented as many of the suggestions we received as possible, and we attach the comments we made to describe them to this letter.

We hope our responses and the changes we made to the manuscript address the concerns of our reviewers.

Sincerely,

Vincent Leroy

Response to RC1

Comment URL: https://doi.org/10.5194/egusphere-2025-4861-RC1

We would like to thank Reviewer 1 for their helpful comments and suggestions. Please find below our response to all points (reviewer recommendations and comments in italics, and changes made to the manuscripts within quotes):

- **More explicitly position Eradiate with respect to established RT solvers such as MYSTIC, highlighting specific physical and numerical distinctions.**
  We understand and share the Reviewer's desire to see the differences with other RTMs highlighted; however, when faced with the task of summarizing these differences, we found that it would result inevitably in a biased and incomplete comparison, as we could select the features and models we want, to highlight the qualities of ours. For instance, why limit ourselves to a comparison with MYSTIC, and omit DART, LESS, htrdr, SMART-G, 3DMCPOL, etc.? We therefore decided not to include such a comparison in this description paper, and prefer to leave it to a benchmarking paper. We however hope that the description we make helps users understand Eradiate's major features at the time of writing (1D atmosphere modelling, 3D surface modelling, CKD molecular absorption parametrization, polarization support, graphics-style path tracing, deep Python integration).
  That said, a major difference Eradiate has with respect to MYSTIC is that it is fully available to the public without restrictions.
- **Elaborate on the practical implications of CG techniques for scaling scene complexity in EO applications compared to traditional RT/MC solvers.**
  We thank the Reviewer for bringing up this topic. We added in our introduction a discussion of some topics that can bring obvious and immediate benefits to EO applications. The list is however potentially endless and could be the subject of a paper in itself, so we tried to remain concise.

  "Influence on radiative transfer modelling for remote sensing. *Among the methods and techniques widely adopted in the computer graphics community that are particularly interesting in the context of 3D radiative transfer for EO, we can cite:*

  – hierarchical data structures, such as bounding volume hierarchies (BVH) or k-d trees (Pharr et al., 2023, Chapter 7), which significantly reduce the cost of scene geometry traversal and are required to handle scenes with high geometric complexity;
  – null-collision methods (Galtier et al., 2013; Novák et al., 2018; Miller et al., 2019), a family of grid-free unbiased sampling methods for participating media, which decouple medium sampling from the underlying data structures, allowing for flexible volume data storage and lookup;
  – volumetric hierarchical data structures, such as VBD (Museth, 2013, 2021), useful to store efficiently information about complex heterogeneous participating media such as clouds.

Many more examples, ranging from Monte Carlo sampling techniques to software architecture, can motivate a convergence between the two scientific communities.

Physically based rendering has had an influence on radiative transfer modelling [...]"

- **Include performance-oriented benchmarks, particularly regarding convergence rates and computational efficiency.**
  This is connected to recommendation #1: designing meaningful benchmarks is a challenging task that goes beyond the scope of a description paper. Comparison to other models only makes sense in a constrained framework, for comparable features, with similar hardware. And comparison between different algorithms or implementations within Eradiate is not yet relevant since this is the first major version that is released. We therefore prefer to avoid venturing into a benchmarking exercise which would have been too shallow to provide real insight about Eradiate's performance.
- **Discuss the potential role of differentiable rendering within the Eradiate framework.**
  We indeed left this topic aside. A full discussion of how to apply Eradiate to inverse problem is premature at this point due to practical limitations that effectively prevent from using the differentiable backends of Mitsuba; but our intent to is make this possible in the future. We therefore added a paragraph to the outlook section that shows that we are aware of this.

  "*A topic that was not explored so far by Eradiate is differentiable rendering. Mitsuba 3, and its underlying computational core Dr.Jit, have built-in infrastructure to differentiate simulations automatically, allowing to output gradients of a primal Monte Carlo estimator with respect to any variable in the scene (Jakob et al., 2022b). The simulation can then be included in an optimization loop to solve inverse problems. Inverse 3D radiative transfer problems in EO are extremely complex and still require a large amount of research for differentiable rendering techniques to be effective in that context; Eradiate has potential to contribute to that effort.*

  Eradiate is free software […]"

Response to RC2

Comment URL: https://doi.org/10.5194/egusphere-2025-4861-RC2

We would like to thank Reviewer 2 for their helpful comments and suggestions. Please find below our response to all points (reviewer recommendations and comments in italics, and changes made to the manuscripts within quotes):

- **However, what is missing is a (detailed) discussion of the results. The results are presented in a very brief manner, and there is almost no discussion and explanation of the results, which is needed to understand their implications and the capabilities of the model.**
  Section 8 shows examples of educational interest that do not bring significantly new scientific results by themselves: detailed discussions require much more detailed studies and exceed the scope of this description paper. For that reason, we decided to keep the discussions brief. We however acknowledge the desire of the Reviewer to see the presented examples exploited to illustrate the capabilities of the model: we added to Section 8 comments explaining the results in more detail and emphasizing the utility of the demonstrated capabilities where relevant.
  - “[...] extracted from the HAMSTER database (Roccetti et al., 2024).

    Figure 10 presents the results. The top panel shows the solar irradiance spectrum measured by TSIS-1 (Coddington et al., 2021), which was used as input for the simulations. The bottom panel displays the TOA radiance spectra for a viewing geometry defined by a zenith angle of 60◦ and an azimuth angle of 75°. The corresponding solar zenith and azimuth angles are 30° and 160°, respectively. The blue line represents the TOA radiance for the US Standard Atmosphere (Anderson et al., 1986). The red line, which shows lower radiance values, includes an additional aerosol layer between 0 and 3 km altitude with an optical thickness of 2, superimposed on the molecular atmosphere. The aerosol optical properties correspond to the “continental” aerosol type in the 6S radiative transfer model (Kotchenova et al., 2006). The simulated spectrum clearly reveals the oxygen A and B absorption bands centred near 760 nm and 686 nm, respectively. Absorption features around 580 nm and 950 nm are associated with water vapour. Eradiate also provides the variance of the simulated radiance, from which the standard deviation can be derived. The shaded region indicates two standard deviations. This uncertainty is clearly non-negligible in the simulation that includes aerosols, whereas it is very small and not visible in the molecular-only case. The standard deviation scales with the square root of the number of samples, which was set to 100 in these calculations.”
  - “*Fig. 13 shows that, even in a clear-sky scenario (i.e. without aerosols), the three quantities differ significantly. This can have significant impact: for instance, drones are now often used to measure the BOA hemispherical conical reflectance factor (HCRF), which is close to the BOA HDRF, and is identified to the TOC BRF. From a radiometric point of view, this is not correct; however, if*

one can tolerate a certain level of inaccuracy, quantified by Schunke et al. (2023), the BOA HCRF can be used as a TOC BRF."

- "*The highly directional features of the scene shown in Fig. 15 result in strong anisotropy in the canopy's reflective pattern, first at the TOC, and also at the TOA, showing a typical example where accounting for 3D surface features makes a significant difference when simulating satellite observations.*"
- "*Topography is part of the complex 3D surface features that can cause deviations from the usual flat surface assumption mentioned in Section 1, e.g. when simulating satellite images above bright desert PICS, as illustrated by Govaerts (2015).*"
- "*In these scenes, a 100 km-sized area is divided into 100 m tiles populated with 3D geometry. Tiles are categorized based on a land cover map, and each type is assigned a set of 3D shapes and a background reflection model. These scenes exploit two important features:*

  – Instancing: The 3D shapes associated to a land cover type are loaded only once, then clone at every location where the land cover type is detected. In practice, this means that the geometry corresponding to the land cover type is loaded only once, which drastically reduces the memory footprint of the scene.
  – Land cover-based material dispatching (see also Section 4.1.3): The background BRDFs of all land cover types are loaded into an index, which is then mapped spatially by a texture encoding the material with an 8-bit integer. This compact representation allows to map many reflection models with a high spatial resolution at a very low memory cost.

  The output (see Fig. 16) shows an example of a coastal area near the city of Dakar, with both an in situ render and a synthetic satellite image. A similar setup is used to create Fig. 5.

  This setup showcases the potential of Eradiate to simulate large scenes sourced from remote sensing data, with sensors at satellite level, and resolving metre-sized geometric features in a 100 km-sized area. The size of the scene and number of objects are such that they can be handled only thanks to features offered by the radiometric kernel, which demonstrates some of the many benefits brought by repurposing rendering software."

- **Lines 49-50: "The world is intrinsically three-dimensional, but most simulation tools, for performance reasons, currently represent the scene using a plane-parallel geometry...". Please provide examples of such simulation tools.**
  Here, by "simulation tools", we meant "radiative transfer models". This was clarified, and we cross-referenced a list of examples given later in the paper.

  "[…] The world is intrinsically three-dimensional, but *most radiative transfer models (see Sect. 2.1 for examples)*, for performance reasons, [...]"

- **Lines 184-192: Please add that you assume that the scatterers are totally randomly oriented, as Eq. 3 makes sense only under this assumption. If this is not assumed, then adjust the equation accordingly. Furthermore, explain the R matrices and the relationship of θ_{xx} to the incoming and outgoing directions.**
  We made it clear that this expression of the VRTE only applies to media with randomly oriented particles, and that Eradiate does not support oriented particles. For the expressions of the rotation matrices, we referred interested readers to reference publications, as we consider that such level of detail in modelling goes a bit beyond the scope of this paper.

  "[…] and θin (respectively θout) is the angle between the pre-scattering propagation and scattering (respectively scattering and post-scattering) reference frames. *For further detail on the VRTE, we refer interested readers to reference publications (Mishchenko et al., 1994; Emde et al., 2010; Cornet et al., 2010). It should be stressed that Eq. 3 is valid only for randomly oriented scattering particles: Eradiate does not handle oriented scattering particles.*"
- **Lines 205-208: Please give references and briefly explain the methods.**
  We provided brief descriptions of the variance reduction techniques but also referred readers to provided references for fundamentals of path tracing, both from an atmospheric RT and graphics perspective. We thought going into the details of these numerical methods, which are nowadays state of the art in computer graphics and fairly popular in atmospheric RT, was beyond the scope of this description paper.

  "Eradiate uses backward MCRT methods to sample the solution of the radiative transfer equation (see Fig. 1). *Such methods perform a random walk from the sensor to connect it with light emitters in the scene (Mayer, 2009; Pharr et al., 2023). The algorithms used in Eradiate include numerical techniques that are commonly found in rendering system and / or atmospheric RTMs (Veach, 1997):*

  – local estimate, also known as next event estimation or direct illumination sampling, which samples the illumination at each node of the path constructed by the random walk rather than waiting to hit an emitter by chance, reducing variance significantly;
  – Russian roulette termination, which avoids constructing unnecessarily long paths by triggering path termination with a certain probability (and associated correction to avoid bias);
  – multiple importance sampling, used e.g. to combine light- and reflectance-driven samples, also reducing variance.

  In plane-parallel geometry, [...]"
- **Lines 226-230: Please briefly explain the different BSDF models and their purpose (e.g., which are used for vegetation, which for soil, which include**

polarization, etc.). A table might be helpful here.
We added the requested information in the form of a table (Table 1) as suggested.

- **Lines 237-238. The reference to Fig. 3 seems to be in the wrong place. It should be placed after "vegetated canopies" as there are no urban environments in Fig. 3.**
  We applied this suggestion.

- **Lines 269-271: Must the database be predefined by the user? This paragraph is a bit confusing.**
  This paragraph and the previous were rephrased to clearly distinguish between the molecular absorption database (Eradiate), and HITRAN's absorption line database (external, upstream), and stress that users can bring their own data. Additional detail about shipped absorption databases was also added in the following paragraphs.

  "Molecular absorption. *The thermophysical profile is used to query a molecular absorption database and build an absorption coefficient vertical profile. At the time of writing, Eradiate ships with 6 molecular absorption databases created from the HITRAN 2020 spectroscopic database (Gordon et al., 2022) using the RADIS (Pannier and Laux, 2019) software package. Additional cross-section data, not found in HITRAN's absorption line database (e.g. ozone photodissociation data published by Gorshelev et al. (2014); Serdyuchenko et al. (2014)) are also incorporated in the shipped cross-section databases. Users can use their own molecular absorption data, provided that it complies with the documented data format. Table 2 lists all shipped databases.* [...]

  In monochromatic mode, the user inputs the sequence of wavelengths for which the RTE is solved. In that mode, *the atmospheric absorption database (either brought by the user or selected from the list of built-in monochromatic databases mentioned in Table 2)* should have a spectral resolution that is fine enough to resolve the complex line structure of the molecular absorption spectrum."

- **Lines 271-276: Please explain the CKD mode in more detail. From your description, it is difficult to understand how the CKD mode works. Maybe a sketch could help to explain the CKD mode.**
  This paragraph was extended with a brief description of the method, the associated terminology, and references to comprehensive sources introducing the method to non-specialists.

  "In CKD mode, *the spectral dimension is handled using the CKD method. Eradiate mostly follows the terminology of Hogan and Matricardi (2020). The spectrum is divided into a a set of spectral intervals, known as bins. This method alleviates the need to cover very densely the spectral dimension to resolve the high-frequency spectral features of molecular absorption spectra by reordering absorption coefficient values (denoted k) in each bin against a pseudo-frequency (denoted g) that represents the cumulative probability of the absorption coefficient over that bin.*

The resulting function k(g) is monotonically increasing and smooth, which allows to compute spectral integrals using a low-order quadrature rule (typically Gaussian) with good accuracy. The quadrature points are called g-points. In practice, this reduces by several orders of magnitude the number of spectral samples to process to achieve satisfactory accuracy. For further detail on this method, see, e.g., Lacis and Oinas (1991). In the CKD mode, the spectral discretization is driven by the selected absorption coefficient database. The bins that will be processed during the simulation [...]"

- **Line 281: "110 cm-1". The unit seems to be wrong.**
  The units are correct: this database has constant spacing in the wavenumber space. We updated the manuscript to clarify this with a complete list of all shipped absorption databases (see also point 6).

- **Lines 288-292: Is it correct that Eradiate does not provide tools to calculate the scattering properties using the T-matrix method or Mie theory? If yes, please rephrase the paragraph to make this clear. If not, please explain how the scattering properties can be calculated in Eradiate.**
  We updated the manuscript to clarify that Eradiate does not ship components to compute single-scattering properties and provided an additional reference to libRadtran's mie tool, which produces compatible output.

  "*No tool is included to derive single-scattering properties, but Eradiate can import data in the NetCDF format supported by the libRadtran RTM: Thus, users can provide customized input generated, e.g., using the MOPSMAP online tool (Gasteiger and Wiegner, 2018), or libRadtran's mie tool (Emde et al., 2016).*"

- **Lines 307-310: When you write that the sun is located at an infinite distance, the angular size of the sun (your emitter) is always infinitesimally small, but in the first sentence you write that the second illumination model is characterized by the angular size of the emitter. Either I am missing something or something in the text is not clear. Either way, this paragraph needs to be revised for clarity. Furthermore, a simple sketch could be helpful to explain the two illumination models.**
  While unintuitive, this is possible. We clarified that point and illustrated the difference between the two illumination models with a figure.

  "Illumination model. *Eradiate provides two illumination models that describe radiance boundary conditions at the outer boundary of the computational domain, giving the value of the radiance field as a function of direction. The first one, a perfectly directional model, describes the radiance distribution as a delta Dirac distribution. The second one is more realistic and accounts for the fact that solar illumination, at the top of the atmosphere, is not perfectly directional, and implements in practice an environment map, also known as infinite area light in computer graphics (Pharr et al., 2023). Both illumination models are parametrized by the local solar zenith and azimuth angles and the selected irradiance spectrum. Fig. 6 illustrate the conceptual difference between the two models.*"

- **Lines 325-326: Can the reference surface be set to any direction? If yes, please explicitly state this. If not, rephrase the sentence accordingly.**
  This sentence has been rephrased for improved clarity.

  “A special variant of the distant radiancemeter *records the directional flux leaving a target location defined by an arbitrarily positioned and oriented flat surface.* It can, in practice, be used to compute the directional or total flux reflected by a surface.”

- **Lines 435-438: The sentence starting with "ROMC submissions..." is confusing. It is difficult to understand what the actual differences are and what the purpose of each mode is. Please rephrase the sentence to make it clear.**
  This sentence was rephrased to explain the purpose of each ROMC mode.

  “ROMC submissions can be done in two modes:

  – The DEBUG mode allows repeated submissions on many scenes and reveals to the user a reference against which they can compare and adjust. This mode is designed to help RTM developers debug their code to reach the community consensus.
  – The VALIDATE mode allows only one submission on a restricted set of randomly chosen scenes different from those available in DEBUG mode, and exposes no “ground truth” for comparison before submission is completed. This mode is designed as a blind exercise to allow RTM developers and users to get a traceable proof of the performance of their model.”

- **Lines 439-449: Please explain briefly the different scenarios and measurements.**
  We added a brief description of the ROMC scenarios and directed interested readers to the RAMI / ROMC websites and publications. We also added information on scenes and measurements to Fig. 9.

  “It provides a framework for autonomous canopy RTM benchmarking based on scenarios developed during the first three phases of the Radiative transfer Model Intercomparison (RAMI) (Pinty et al., 2001, 2004; Widlowski et al., 2007). *ROMC scenarios are based on abstract canopies with diverse structures, e.g. homogeneous vegetated covers with different leaf angle distributions, clusters of spherical leaf clouds, or more realistic setups similar to forests. Canopy definitions are general, leaving to users and their RTM the choice of the representation: We chose to represent leaves as disks, but similar performance could be achieved by representing them using triangulated meshes. Readers interested in a complete list and description of the ROMC scenarios are referred to the original RAMI-3 publication (Widlowski et al., 2007) and the RAMI website.*

  ROMC users can submit simulation results for various measurements [...]”

- **Lines 468-474: Please explain/discuss the results that are shown in Fig. 8 in more detail. Saying that it is comparable to Fig. 8 of Emde et al. (2015) is not sufficient. Furthermore, please provide some more details about the scenario, e.g., what is the geometry of this scenario, where is the observation location, what aerosol was used, and apart from ozone, were there any other absorbers, etc.**
  We added details about the scenario and additional comments on the results.

  "[...] intercomparison of six participating independently developed radiative transfer codes.

  Fig. 9 shows the Eradiate results for case B3. The simulation uses the US Standard Atmosphere (Anderson et al., 1986) together with a typical aerosol profile (Shettle, 1984), with a vertical aerosol optical thickness of 0.2. Aerosol optical properties are provided on the IPRT website5 and correspond to prolate spheroids with an aspect ratio of 3. The particle size distribution has a mode radius of 390 nm. The complex refractive index is 1.52 − 0.01i and the mass density is 2.6 g cm−3. These properties are representative of the coarse mode of mineral dust aerosols (Hess et al., 1998). Simulations are performed at a wavelength of 350 nm in the ultraviolet spectral region, where Rayleigh scattering is strong and absorption is dominated by ozone. Radiances are calculated for sensors located at the surface, at 1 km altitude, and at the top of the atmosphere (TOA). The figure can be directly compared with Fig. 8 in Emde et al. (2015). The upper panels show the radiation field for a sensor at TOA viewing downward, while the lower panels show the radiation field for a surface sensor viewing upward. The solar zenith angle is 30° and the solar azimuth angle is 0°. The solar position is indicated by a yellow star in the polar plots. The subplots display the Stokes vector components: the intensity I, the linear polarization components Q and U, and the circular polarization component V . In this case, V is three to four orders of magnitude smaller than the linear polarization components. At the surface, the intensity I exhibits a pronounced forward-scattering peak in the solar direction. The degree of polarization is largely controlled by Rayleigh scattering and reaches its maximum at scattering angles close to 90°.

  Table 5 shows [...]"

- **Line 498: What does "PICS" stand for? Please explain the acronym.**
  The acronym was defined.

- **Line 507: What is the *panellus* correlated-k distribution? Please explain.**
  The table listing the different CKD databases was referenced.

  "[…] we employ the correlated-k distribution database *panellus (1 nm resolution, see Table 2)* to generate spectra [...]"

- **Line 547: I do not understand the purpose of this sentence. Please rephrase.**
  This sentence was modified to improve clarity (see point 2, "Minor issues").

- **Lines 551-553: Please explain what desert aerosol model was used and add some details about the model, e.g., which size distribution, etc.**
  This aerosol model is sourced from the RAMI4ATM benchmark, and it is thoroughly documented on the RAMI4ATM website. We refrained from fully describing the microphysical properties used to generate the single-scattering properties of this dataset for two reasons: (1) we believe that it does not improve significantly the impact of this section, and (2) as other RAMI phases, RAMI4ATM should be reported in a relatively short term in a dedicated publication which will provide these details.

  "[…] The setup includes the US Standard atmosphere and a desert aerosol model (govaerts_2021_desert) *designed to support the RAMI4ATM benchmarking exercise*. Irradiances are evaluated [...]."

- **Lines 647-648: The sentence starting with "Various..." is difficult to read. Please rephrase.**
  This sentence was rewritten with less jargonic language.

  "Eradiate's user interface relies entirely on Python programming. *Various features are implemented to make configuration code less error-prone, briefer and more enjoyable to write:* [...]"

- **Lines 652-653: The sentence starting with "In order ..." is difficult to read. Please rephrase.**
  This sentence we rewritten for improved clarity.

  "[…] with domain-specific units. *To reduce input verbosity, values provided without units are automatically converted to configurable default units.*"

- **Fig. 1: The figure could be misleading as you depict horizontally distributed clouds, but in the text you explicitly state that your atmosphere is 1D. Please clarify this in the text and/or adjust the figure accordingly.**
  This figure only conveys the general idea of path tracing. We updated the caption and clarified that the atmosphere in Eradiate v1.0 is 1D.

  "Eradiate uses a path tracing algorithm which builds light paths by performing random walks from the sensor, *bouncing at surfaces and in the 1D atmosphere*."

- **Table 2: The font size is very small; please increase the font size to make it more readable. Furthermore, the caption is confusing. Please explain what the different measurement abbreviations stand for and what the different scenarios are, or refer to the section where they are explained.**
  (See also point 13: "Lines 439-449") Unfortunately, ROMC authors and maintainers explicitly forbid the modification of these plots and strongly discourage digitization to create custom designs: we are not allowed to modify the plots. We added annotations and footnotes to the table to clarify scenarios and measurements.

- **Fig. 8: Please increase the font size. Furthermore, it seems that the units are missing. What does "B3, eradiate, case 0" in the title of the plot mean?**
  The figure was updated for improved readability, and the title was replaced with a detailed description of the corresponding benchmarking case (see also point 14: "Lines 468-474").
- **Fig. 10: The unit seems to be missing, or the caption is insufficient. Slightly increase the font size.**
  We updated the figure.
- **Fig. 11: What does "VAA" stand for? Please explain the acronym.**
  The acronym was added to the list of acronyms.
- **Fig. 12: It seems that the y-axis label is missing. Please explicitly state in the caption which of the three scenarios includes an atmosphere and which does not. The labels of the different lines are inconsistent with the text (lines 541-546). Please adjust.**
  We updated the figure.