

A tool for generating fast k -distribution gas-optics models for weather and climate applications

Article

Published Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 3.0

open access

Hogan, R. J. ORCID: <https://orcid.org/0000-0002-3180-5157>
and Matricardi, M. ORCID: <https://orcid.org/0000-0001-7514-9473> (2022) A tool for generating fast k -distribution gas-optics models for weather and climate applications. *Journal of Advances in Modeling Earth Systems*, 14 (10). e2022MS003033. ISSN 1942-2466 doi: <https://doi.org/10.1029/2022MS003033> Available at <https://centaur.reading.ac.uk/110544/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1029/2022MS003033>

Publisher: American Geophysical Union

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

A Tool for Generating Fast k -Distribution Gas-Optics Models for Weather and Climate Applications

 Robin J. Hogan¹  and Marco Matricardi¹ 
¹European Centre for Medium-Range Weather Forecasts, Reading, UK

Key Points:

- We describe a free software tool for generating accurate gas optics models for radiation schemes and test with line-by-line calculations
- Extra efficiency is achieved via the use of the full-spectrum correlated- k method in the thermal- and near-infrared (NIR)
- The spectral properties of clouds are treated accurately via the use of sub-bands in the NIR and several other techniques

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

 R. J. Hogan,
r.j.hogan@ecmwf.int

Citation:

 Hogan, R. J., & Matricardi, M. (2022). A tool for generating fast k -distribution gas-optics models for weather and climate applications. *Journal of Advances in Modeling Earth Systems*, 14, e2022MS003033. <https://doi.org/10.1029/2022MS003033>

 Received 6 FEB 2022
Accepted 29 JUL 2022

Author Contributions:

Conceptualization: Robin J. Hogan
Data curation: Robin J. Hogan, Marco Matricardi
Formal analysis: Robin J. Hogan
Investigation: Robin J. Hogan
Methodology: Robin J. Hogan
Resources: Robin J. Hogan, Marco Matricardi
Software: Robin J. Hogan
Validation: Robin J. Hogan

© 2022 The Authors. Journal of Advances in Modeling Earth Systems published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Abstract One of the most important components of an atmospheric radiation scheme is its treatment of gas optical properties, which determines not only the accuracy of its radiative forcing calculations fundamental to climate prediction, but also its computational cost. This paper describes a free software tool “ecCKD” for generating fast gas-optics models by optimally dividing the spectrum into pseudo-monochromatic spectral intervals (known as k -terms) according to a user-specified error tolerance and the range of greenhouse-gas concentrations that needs to be simulated. The models generated use the correlated k -distribution method in user-specified bands, but can also generate accurate “full-spectrum correlated- k ” models that operate on the entire longwave or near-infrared (NIR) parts of the spectrum. In the NIR, the large spectral variation in cloud absorption is represented by partitioning the parts of the spectrum where gases are optically thin into 2–6 sub-bands, while allowing k -terms for the optically thicker parts of the spectrum (where clouds and surface reflectance are less important) to span the entire NIR spectrum. Candidate models using only 16 and 32 k -terms in each of the shortwave and longwave are evaluated against line-by-line calculations on clear and cloudy profiles. The 32-term models are able to accurately capture the radiative forcing of varying greenhouse gases including CO₂ concentrations spanning a factor of 12, and heating rates at pressures down to 1 Pa.

Plain Language Summary A crucial component of atmospheric computer models used to make climate projections and weather forecasts is the “gas optics scheme,” which represents the interaction of sunlight and infrared radiation with greenhouse gases. This paper describes a free software tool “ecCKD” that uses a number of novel techniques to generate new gas optics schemes that are computationally faster than most existing schemes while still being very accurate. For example, the schemes are able to simulate variations in carbon dioxide concentration spanning a factor of 12 and methane concentration spanning a factor of 10. Users of ecCKD can generate schemes that are optimized for specific applications, such as short-term weather forecasting or simulating past climates. A special focus has been placed on the near-infrared part of the solar spectrum to ensure that the schemes work well when computing the interactions of sunlight simultaneously with gases and clouds, important to ensure that the impact of clouds on weather and climate is well simulated.

1. Introduction

Perhaps the most fundamental part of a climate model is the gas-optics module of its radiation scheme; in fact, one of the most influential (and indeed Nobel-prize-winning) studies of the climatic impact of increased greenhouse gases used a climate model consisting of little more than a radiation scheme coupled to a convective-adjustment scheme (Manabe & Wetherald, 1967). The correlated k -distribution (CKD) method (Goody et al., 1989; Lacis & Oinas, 1991) has emerged as the leading technique for treating the radiative effects of gases that is fast enough to use in 3D weather and climate models. Many models still use older alternatives; DeAngelis et al. (2015) reported a large spread in the magnitude of the near-infrared (NIR) water vapor feedback amongst 14 climate models, but those using the CKD method were found to be much the most accurate compared to benchmark line-by-line (LBL) radiation calculations.

CKD models are very time consuming to develop from scratch, so when writing a new radiation scheme one must usually incorporate one of the small number of off-the-shelf models, even though it may not be optimized for ones particular application. The perceived high computational cost of radiation schemes has spurred numerous ideas to accelerate them, such as replacement of the radiation scheme (or only its gas-optics component) by a neural network (e.g., Ukkonen et al., 2020), or sub-sampling of model columns (Barker et al., 2021). The computational cost of a radiation scheme scales with the number of pseudo-monochromatic calculations (hereafter referred to as “ k -terms” or spectral intervals) required to represent the entire spectrum, which is determined by the

Visualization: Robin J. Hogan
Writing – original draft: Robin J. Hogan
Writing – review & editing: Robin J. Hogan, Marco Matricardi

gas-optics scheme. Along with several other modeling centers, the European Centre for Medium-Range Weather Forecasts (ECMWFs) uses the “RRTMG” (Rapid Radiative Transfer Model for General Circulation Models; Mlawer et al., 1997) CKD gas-optics model, which employs 112 terms in the shortwave and 140 in the longwave. However, there is a significant variation; Hogan et al. (2017) reported that the number of k -terms in seven global weather forecast models spanned more than a factor of 3.7. This leads us to ask: what is the minimum number of k -terms that a CKD model needs while still being sufficiently accurate for weather and climate applications?

This paper describes the ECMWF free-software tool “ecCKD,” which allows users to generate CKD models tailored for their own applications. We aim to adopt some of the best features of CKD models reported in the literature. For example, Edwards and Slingo (1996) described a radiation scheme that is *flexible* in the sense that the spectral discretization and gas optical properties are configured at run-time by a spectral file. Their use of the concept of “equivalent extinction” put additional constraints on the calculations required by the downstream solver such that the spectral file was, in practice, only compatible with the Edwards and Slingo (1996) radiation scheme. We use a self-describing netCDF file that consists of little more than a set of look-up tables (LUTs), one per gas. While these files can be read by ECMWF’s “ecRad” radiation scheme (Hogan & Bozzo, 2018), it would in principle be straightforward to interface them to other radiation schemes.

The CKD method works by grouping non-contiguous parts of the spectrum, where the gaseous optical properties are similar, into a single pseudo-monochromatic spectral interval. One of the most interesting developments to the original CKD technique is the full-spectrum correlated- k (FSCK) approach (e.g., Modest & Zhang, 2002), based on the insight that for clear-sky radiative transfer, the parts of the spectrum grouped together in one interval can be very separated in wavelength, thus dispensing with the need for bands. Since fewer bands leads to fewer intervals overall, this is a powerful way to reduce computational cost. The FSCK idea has been prototyped for atmospheric applications in the shortwave (Pawlak et al., 2004) and longwave (Hogan, 2010). The ecCKD tool is capable of producing CKD models using either traditional bands or treating the whole spectrum in a single band. It can also produce hybrid models to address the challenge posed by the large spectral variation in surface albedo and cloud optical properties in the NIR.

The classical CKD method involves reordering the absorption spectra separately for different atmospheric conditions (pressure, temperature, and H_2O concentration) and assuming perfect rank correlation between these spectra, implicitly allowing radiation to change wavelength as it traverses the atmosphere. We prefer each spectral interval of a CKD scheme to correspond to a unique set of wavelengths, independent of height, an approach taken by, for example, Bennartz and Fischer (2000), Hogan (2010), and Doppler et al. (2013). This has sometimes been referred to as the *uncorrelated* k -distribution method, although we prefer to consider this as a variant of the CKD method since its accuracy still relies on the high correlation of absorption spectra at each height, even though perfect rank correlation is not assumed. By reporting the spectral mapping in the spectral file, the optical properties of clouds and aerosols can be averaged accurately to each spectral interval. Furthermore, by allowing optical properties to be specified in individual spectral intervals, rather than only in bands, we overcome the challenge identified by Lu et al. (2011) that the optical properties of clouds and water vapor are correlated within a band.

The paper is organized as follows. The steps of the method are described in Section 2, and are illustrated via the production of two candidate CKD models in the shortwave and two in the longwave. Section 3 evaluates these models using independent LBL calculations for 50 independent clear-sky atmospheric profiles, and in Section 4 the models are evaluated in cloudy skies.

2. Method

2.1. Overview

The ecCKD tool performs a sequence of tasks shown in Figure 1, each of which is configurable by the user. Rather than computing atmospheric absorption spectra from scratch, ecCKD takes as input pre-computed absorption spectra from the Correlated K-Distribution Model Intercomparison Project (CKDMIP; Hogan & Matricardi, 2020), performing LBL radiation calculations on them as needed. The data sets are described in Section 2.2, and the representation of gases in Section 2.3. The user specifies the bands to be used (Section 2.4), within which the spectra are reordered separately for each gas (Section 2.5). Sections 2.6 and 2.7 describe how each band is partitioned into spectral intervals, also known as k -terms or g -points, according to a user-specified error tolerance. Section 2.8 then describes how an initial set of gas-absorption LUTs is created for each gas

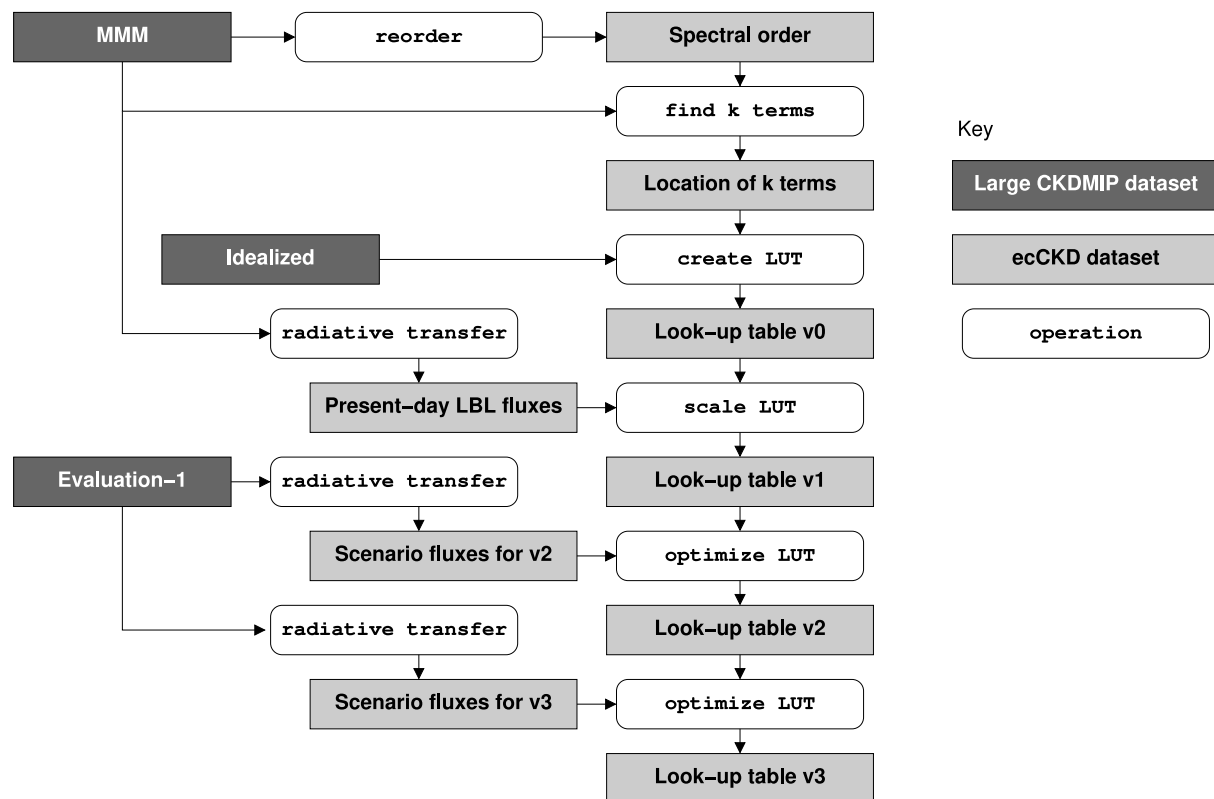


Figure 1. Flowchart illustrating the steps of the ecCKD method (white boxes), which make use of the large Correlated K-Distribution Model Intercomparison Project (CKDMIP) data sets of Hogan and Matricardi (2020) and store intermediate information in smaller netCDF files (light gray boxes), ultimately producing a look-up table file for use in a radiation scheme. “MMM” refers to the CKDMIP data set containing the median, minimum and maximum profiles from a much larger database.

and each k -term, which constitutes a functioning but possibly inaccurate CKD model in the form of a file that can be used to configure the gas optical properties in a radiation scheme. The subsequent steps then refine these LUTs; in the shortwave the absorptions are scaled so as to produce an exact profile of direct irradiances for each k -term for one particular representative atmosphere (Section 2.9). In both the longwave and shortwave, a number of optimization steps are performed to refine the LUT coefficients to minimize the errors in irradiances and heating rates for a set of training profiles (Section 2.10).

2.2. Data

The ecCKD tool makes use of the CKDMIP LBL spectral absorption data set described by Hogan and Matricardi (2020), appropriate for the terrestrial atmosphere. The gases considered are H_2O , O_3 , O_2 , N_2 , and the well-mixed greenhouse gases (WMGHGs) CO_2 , CH_4 , N_2O , CFC-11, and CFC-12. Except for H_2O , the molar absorption of all these gases can be considered independent of their concentration, so a very wide range of climate scenarios can be considered by simply scaling the absorptions. According to Meinshausen et al. (2017), these five WMGHGs represent 94.5% of anthropogenic greenhouse warming in terms of longwave radiative forcing since 1750, and a further 38 more minor gases (representing the remaining 5.5% of radiative forcing) can be adequately represented by using an increased “equivalent” concentration of CFC-11. This approach is common in the various phases of the Coupled Model Intercomparison Project given that most or all climate-model radiation schemes are unable to represent all 43 WMGHGs listed by Meinshausen et al. (2017). However, there is no reason why other gases could not be added to ecCKD if LBL calculations were performed to provide additional absorption spectra in the appropriate format.

As described by Hogan and Matricardi (2020), the spectral resolution of the data set is variable, being finest in the strong absorption bands of CO_2 because at mesospheric altitudes the CO_2 lines are only Doppler broadened and thus become very narrow. This results in a total of 7,211,999 spectral points in the longwave and 3,126,494

in the shortwave. The resolution of the data set in terms of pressure, temperature and H₂O mole fraction is given in Section 2.8, and the total volume is around 1 TB.

2.3. Gas Representation

Ultimately, a CKD model produced by ecCKD computes molar absorption coefficient k_i (the absorption cross section per mole of *all* gases) in spectral interval i at pressure p and temperature T as the sum of the contribution from m gases, as follows:

$$k_i(p, T, \psi_1 \dots \psi_m) = k_i^0(p, T) + \sum_{j=1}^l k_i^j(p, T) \psi_j + \sum_{j=l+1}^n k_i^j(p, T, \psi_j) \psi_j + \sum_{j=n+1}^m k_i^j(p, T) \times (\psi_j - \psi_j^{\text{ref}}), \quad (1)$$

where ψ_j is the mole fraction of gas j . The four terms on the right-hand-side represent the four different ways that gases can be represented. The first is the *background* term, a 2D LUT representing the combined contribution from all gases with a constant, pre-defined mole fraction. For a model intended for climate simulations, this would typically include only O₂ and N₂, but for a CKD model targeting present-day Numerical Weather Prediction (NWP) we could include the contribution from WMGHGs. The fact that the LUT includes a dependence on pressure means that a pressure dependence of the concentration of these gases can be represented (e.g., as shown in Figure 2 of Hogan & Matricardi, 2020). The second term represents gases 1 to l , whose absorption varies *linearly* with mole fraction; in this case k_i^j is the molar absorption coefficient of gas j , that is, the absorption cross-section per mole of the gas. The third term represents gases $l + 1$ to n whose absorption varies *nonlinearly* with concentration; in this case a 3D LUT is used for molar absorption coefficient, with an additional dependence on the mole fraction of the gas in question. In the terrestrial atmosphere only H₂O is in this category, and the representation here allows the contribution from the water vapor continuum (both self and foreign) to be treated completely with no need to separate the line and continuum contributions to the absorption, as is done by many existing CKD models.

The fourth term in Equation 1 has what we refer to as a *relative-linear* dependence of absorption on mole fraction: a “reference” mole fraction, ψ_j^{ref} , is defined for the gas, typically the mean surface present-day concentration. The absorption by present-day concentrations of the gas is then folded into the background term, while the fourth term represents the additional absorption due to perturbations (which may be negative) of concentrations from ψ_j^{ref} . This approach is useful for some minor greenhouse gases where one k -term may be approximating a large range of absorptions, resulting in the transmittances behaving as if the dependence of k_i^j on ψ_j is not perfectly linear. We have not found it necessary to use a full nonlinear treatment for these gases (as for H₂O), but the relative-linear term can be thought of as a linearization around the present-day concentration.

It is up to the user which of the four representations to use for each gas, and the choice depends particularly on what range of greenhouse-gas concentrations will need to be simulated by the target CKD model. The example models generated in this paper are intended to simulate the climate scenarios proposed by Hogan and Matricardi (2020). In both the shortwave and longwave, we represent O₂ and N₂ absorption by the background term, CO₂ and O₃ absorption as linear terms, H₂O as a nonlinear term, and CH₄ and N₂O as relative-linear terms. In the longwave, CFC-11 and CFC-12 are represented by linear terms, while in the shortwave they are neglected (see Table 4 of Hogan & Matricardi, 2020). If ecCKD were to be applied to an extraterrestrial atmosphere then the choice of gases and how to represent them would need to be reconsidered, and if processes such as collision-induced absorption were important then in principle an additional term could be added to Equation 1 representing absorption dependent on the concentration of two different gases.

2.4. Band Selection

The selection of bands is entirely specified by the user, while the partitioning of each band into k -terms is automated. Many CKD models select band boundaries to minimize the number of absorbing gas species in a band, to ensure that the assumption of random overlap of gas absorption is valid, or to cap the error due to assuming the Planck function to be constant across each band. Since ecCKD is not subject to these limitations (any number of gases can be handled with arbitrary overlap, and the Planck function is computed exactly for each k -term),

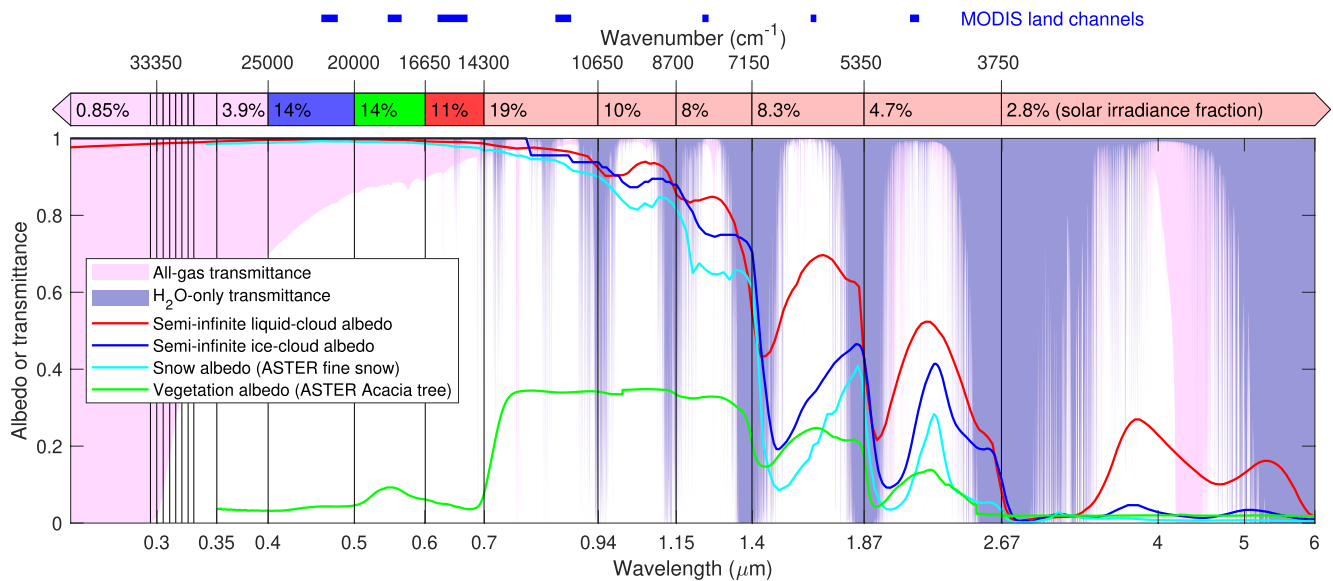


Figure 2. Spectral features of surface and atmosphere informing shortwave band selection. The shading indicates the zenith transmittance of the Correlated K-Distribution Model Intercomparison Project “median” atmosphere due to both all gases (including Rayleigh scattering) and H₂O only. The red and dark blue lines depict the albedo of semi-infinite liquid and ice clouds with effective radii of 10 and 30 μm, respectively, the latter assuming optical properties for the Baum et al. (2014) “General Habit Mixture.” The cyan and green lines depict the albedo of fine snow and vegetation (Acacia tree) from the ASTER data set. The vertical lines delimit the 19 bands of the “Window” band structure discussed in the text, the bar at the top showing the fraction of incoming solar energy in most of them.

the choice of bands is driven by (a) the need to represent spectral variations of the properties of clouds, aerosols and the surface, and (b) the needs of downstream users for irradiances in specific bands. The models generated are given names of the form ecCKD-*B-N*, where *B* is the name of the band structure and *N* is the total number of *k*-terms.

In the longwave, the radiative effect of clouds is dominated by a fairly narrow range of wavelengths in the infrared atmospheric window. We therefore hypothesize that adequate accuracy can be achieved without the use of bands, that is, treating the entire longwave spectrum as a single band (the “FSCK” band structure), provided that the optical properties of clouds are computed separately per *k*-term. This is verified a posteriori in Section 4, where we also evaluate an ecCKD model generated with the “Narrow” 13-band structure suggested by Hogan and Matricardi (2020).

In the shortwave there are multiple concerns that suggest the need for bands, as illustrated in Figure 2. We first define a “Window” band structure with 19 bands (delimited by the vertical black lines) suitable for reference calculations in an atmospheric model, and then describe a simpler structure that exploits more of the efficiencies possible with ecCKD. From the point of view of gases, the important line absorption occurs almost entirely at wavelengths longer than 0.625 μm (wavenumbers less than 16,000 cm⁻¹). At shorter wavelengths, continuum absorption and Rayleigh scattering dominate, and since these tend to vary monotonically with wavelength, the *k*-terms selected automatically by ecCKD in this range tend to span contiguous ranges of wavelength and therefore behave exactly as bands. From the perspective of NWP and reanalysis applications, there are several specific bands that are useful for downstream products, and it is preferable to specify them manually rather than leave ecCKD to place *k*-terms at arbitrary locations in wavelength. Photosynthetically Active Radiation is a common product consisting of the surface downwelling irradiance in the range 0.4–0.7 μm. To generate real-color imagery (similar to that produced by Lopez, 2020), this is further split at 0.5 and 0.6 μm to define red, green and blue bands. The ecCKD tool works in wavenumber rather than wavelength, with shortwave bands specified to the nearest 50 cm⁻¹, so the exact boundaries are at wavenumbers of 14,300, 16,650, 20,000, and 25,000 cm⁻¹. The solar energy in each of these is shown in Figure 2. The “Window” band structure also uses 7 fine ultraviolet (UV) bands of width 500 cm⁻¹ (around 5 nm) for online calculation of UV index. For more specialist applications, additional bands could be considered such as finer visible bands for modeling of marine biology (e.g., Ciavatta et al., 2014) or additional UV bands for computing photolysis rates. The NIR part of the shortwave spectrum consists of a sequence of windows in the gas absorption within which the albedo of clouds and the surface tends

to step down when moving to longer wavelengths. Since ecCKD imposes no constraints on the number of active gases in each band, we align the “Window” band structure to the location of the windows to best resolve the very large variation in cloud and surface albedo.

For faster radiation calculations, which is the primary focus of this paper, we define the simpler “RGB” band structure consisting of the red, green and blue bands as above, but merging the UV and NIR each into single bands. Treating the entire NIR (wavelengths longer than around $0.7 \mu\text{m}$) in a single band is essentially the FSKC approach taken by Pawlak et al. (2004), but it would clearly be a poor approximation to treat clouds and surface properties as spectrally constant in this region. Therefore, ecCKD offers the capability to use “sub-bands,” which we demonstrate in Section 2.7 and evaluate in cloudy skies in Section 4: only parts of the NIR gas-absorption spectrum that are optically thin enough for clouds and the surface to be important are partitioned into sub-bands, while the wavenumbers corresponding to optically thick parts of the spectrum are treated as a single band.

2.5. Reordering the Spectrum

The longwave and shortwave parts of the spectrum are considered separately. The first task shown in Figure 1 is to reorder the spectrum of each gas in order of increasing absorption within each band. We seek a unique mapping independent of height, and therefore aim to sort the high-resolution CKDMIP spectra in terms of the approximate height of the peak cooling in the longwave and peak heating in the shortwave. We use the “median” atmosphere from the CKDMIP “MMM” data set, which is described by Hogan and Matricardi (2020) and contains the profiles of the median, minimum and maximum temperature, H_2O and O_3 from the 25,000-profile data set of Eresmaa and McNally (2014). This is combined with present-day (2020) greenhouse gas concentrations. In the longwave we follow a method very similar to that proposed by Hogan (2010): for each gas a LBL radiative transfer calculation is performed with all other gas concentrations set to zero and an idealized profile of temperature increasing linearly with the logarithm of pressure from -100°C at 0.01 hPa to $+15^\circ\text{C}$ at 1,000 hPa; this ensures that the height of the peak cooling varies monotonically with the strength of the absorption, which is not guaranteed with a more realistic temperature profile. Sorting is in order of the height of peak cooling rate. This method fails for low column optical depths, τ , where cooling rate peaks at the surface, so when $\tau < 0.5$ we sort by τ instead. In the shortwave a simpler approach is taken: the spectra are ordered by the height at which the optical depth from top-of-atmosphere (TOA) reaches 0.25, which is the height at which direct radiation for a solar zenith angle of 60° will have fallen to around 60% of the TOA value. The accuracy of the final CKD model is fairly insensitive to the exact value of threshold optical depth; we find that changing it to 0.5 has a negligible effect on the results. Figure S1 in Supporting Information S1 depicts NIR H_2O molecular absorption versus cumulative probability at different pressure levels, and illustrates that, despite imperfect rank correlation of the absorption spectra between levels, the procedure above ensures that at a particular pressure the absorption coefficients are most accurately reordered for the wavenumbers that contribute most to the solar heating at that pressure.

The result of the reordering is written into a spectral-order file for each gas, containing the integer rank r of each wavenumber point. Suppose the entire spectrum contains N discrete wavenumbers indexed 1 to N , and a particular band corresponds to wavenumbers indexed m to n . The ranks r_m to r_n will consist of the integers m to n but reordered. In the following sections, we follow previous authors and introduce a coordinate variable for the reordered spectrum, g . In ecCKD, this simply maps the integer ranks for the bands to the range 0–1, that is, an element of the reordered spectrum with integer rank r in a particular band would have $g(r) = (r - r_m)/(r_n - r_m)$.

2.6. Partitioning g Space for Individual Gases

The next step shown in Figure 1 is to partition each band into k -terms, each corresponding to a fixed set of wavenumber points. In a radiation scheme, each k -term would be treated by an independent quasi-monochromatic radiative transfer calculation, so generally more terms correspond to a more accurate but more computationally costly scheme. We use a refined version of the algorithm described by Hogan (2010): first the spectrum is partitioned separately for each gas in each band (described in this section), then the partitions for each gas are merged taking account of the spectral overlap of gases (described in Section 2.7).

For each gas and band we need to find the number of intervals, n , into which g space (defined in the previous section) is to be divided, and their boundaries $g_0, g_1 \dots g_n$, with the upper and lower bounds already defined as $g_0 = 0$ and $g_n = 1$. Most previous papers require the user to specify n and define the boundaries according to a

fixed mathematical rule such as Gaussian Quadrature (e.g., Kato et al., 1999), but this is not adaptive to the spectra of individual gases. In ecCKD, the user provides a single error tolerance, and the tool attempts to partition g space such that the error associated with each g interval is approximately equal to this tolerance. The smaller the tolerance, the more g intervals will be needed and a greater overall accuracy should be achieved. Appendix A describes a general algorithm for partitioning g space given a function $E(g_{i-1}, g_i)$ that returns the error associated with treating the wavenumbers corresponding to the range g_{i-1} to g_i by a single quasi-monochromatic radiation calculation. This function is similar to a cost function in estimation theory, and following Hogan (2010) is formulated as the mean squared error in heating rate over l layers, but with an additional term (weighted by f) penalizing errors in surface and TOA irradiances:

$$E(g_{i-1}, g_i) = \sum_{j=1}^l w_j (H_j^{\text{CKD}} - H_j^{\text{LBL}})^2 + f \left[\left(F_{\uparrow\text{TOA}}^{\text{CKD}} - F_{\uparrow\text{TOA}}^{\text{LBL}} \right)^2 + \left(F_{\downarrow\text{surf}}^{\text{CKD}} - F_{\downarrow\text{surf}}^{\text{LBL}} \right)^2 \right], \quad (2)$$

where H_j^{CKD} and H_j^{LBL} are the heating rates predicted by the ecCKD and LBL models in layer j , while $F_{\uparrow\text{TOA}}$ and $F_{\downarrow\text{surf}}$ are the upwelling irradiance at TOA and the downwelling irradiance at the surface, respectively. To weight the stratosphere and troposphere on an approximately equal basis, we follow Hogan (2010) and weight the vertical profile by the square-root of pressure, that is, the weight term is given by $w_j = (p_{j+1/2}^{1/2} - p_{j-1/2}^{1/2}) / p_{l+1/2}^{1/2}$, where $p_{j+1/2}$ is the pressure at the interface between layers j and $j + 1$, and $p_{l+1/2}$ is the surface pressure.

Since this function is called multiple times by the algorithm described in Appendix A, it cannot be too computationally costly, so partitioning is performed using only a single profile. For temperature and the concentration of the target gas, we again use the “median” present-day CKDMIP atmosphere. The partitioning needs to account for the presence of other gases, which can dominate in some parts of the spectrum. This tends to reduce the error associated with representing the target gas and therefore reduces the number of g intervals required, but depends on the concentration of these other gases. We use the minimum concentration of these gases that the CKD model is intended to simulate. For H_2O and O_3 we use the “minimum” concentration profile from the CKDMIP MMM data set. The minimum concentrations of the WMGHGs depends on what application the CKD model is to be used for. For a CKD model to be used solely in NWP, it is appropriate to simply use present-day concentrations for the WMGHGs. For simulation of past and future climate, we use the “Glacial Maximum” values proposed by Hogan and Matricardi (2020), that is, the minimum concentrations found in the last million years.

In the longwave, the LBL radiative transfer calculation is performed with the present-day concentration of the target gas and the “minimum” concentrations of all other gases, and when the function $E(g_{i-1}, g_i)$ is called, the LBL terms in Equation 2 are computed simply by summing the LBL spectral irradiances from the wavenumbers corresponding to the requested range of g . A single zenith angle is used in each hemisphere, equivalent to the two-stream method but without scattering. The “CKD” terms in Equation 2 are computed by first averaging the optical depths of the target gas across the wavenumbers corresponding to the requested range of g , but retaining the full spectral resolution for the other gases. This way Equation 2 quantifies the error purely associated with approximating the target gas. The user can select the method used to average the optical depths of the target gas, the default being a linear average of the layer transmittances weighted by the Planck function at the temperature of the layer.

In the shortwave, the radiative transfer calculations are limited to the direct (unscattered) solar beam, which contains almost all of the sensitivity to gas absorption, and reduces computational cost. Thus, the upwelling terms in Equation 2 are omitted and the heating rates consider only heating by the direct beam. Optical depths are averaged weighting by the TOA incoming solar spectral irradiance, which following Hogan and Matricardi (2020) is taken to be the 1986–2018 average of the Coddington et al. (2016) climate data record.

Shortwave partitioning includes the option to use sub-bands in the NIR, as introduced in Section 2.4. Since H_2O dominates in the NIR, this can be achieved by dividing H_2O alone into sub-bands. Each of the high-resolution spectral points in the NIR band has both a wavenumber and a g value indicating the H_2O absorption strength in the NIR region. The spectral points with $g < g_{\text{crit}}$, where g_{crit} is some user-specified critical value, are deemed to be optically thin enough that sub-bands are needed, so these points are further grouped according to their wavenumber into user-specified sub-bands, although within the sub-bands the ordering by g is preserved. For $g \geq g_{\text{crit}}$,

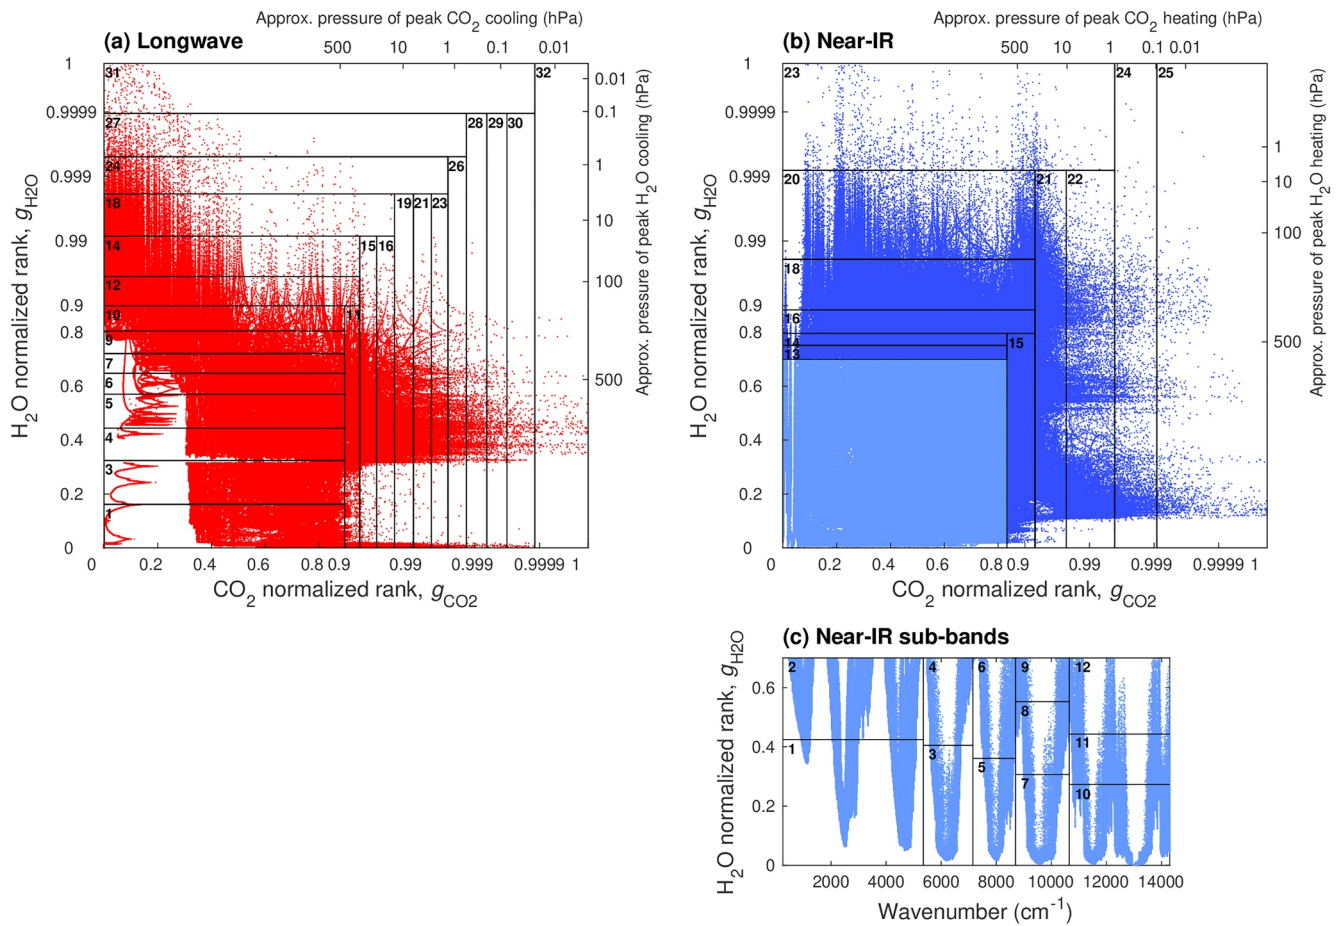


Figure 3. Scatterplot of the entire (a) longwave and (b) near-infrared spectra, where each point corresponds to a wavenumber in the high-resolution line-by-line data set, and the axes show the normalized rank (g) of the CO_2 and H_2O absorption for that wavenumber as defined in Section 2.5. The scales are linear for $g < 0.9$ and > 0.9999 , and logarithmic in $1 - g$ in the range $0.9 < g < 0.9999$. The equivalent pressures of peak heating and cooling for atmospheres containing only one gas are shown to the top and right of each axis. The numbered rectangles indicate the k -term into which the wavenumbers are grouped for the ecCKD-FSCK-32 longwave model and the near-infrared band of the ecCKD-RGB-32 shortwave model, where missing numbers correspond to k -terms specializing in gases other than CO_2 or H_2O . Panel (c) shows how points with weak gas absorption in the lower-left rectangle in panel (b) are allocated to 12 k -terms, grouped into five sub-bands to resolve spectral variations in cloud, aerosol, and surface properties.

H_2O is optically thick enough that no grouping by wavenumber is needed. This is illustrated graphically at the end of the next section.

2.7. Partitioning g Space for Multiple Gases

After each of the gases have had their reordered spectra partitioned into intervals in g space, they are combined to obtain a final set of k -terms. This is achieved using the “hypercube partition method” of Hogan (2010): for m active gases in a particular band we consider an m -dimensional unit hypercube where dimension j represents the g space for gas j . Figure 3a provides a visualization of two of these dimensions (corresponding to H_2O and CO_2) for the entire longwave spectrum, where each red dot represents an individual wavenumber. To generate an FSCK model, our task is to divide this space up into subregions (rectangles in the 2D case), each representing a k -term, such that the wavenumber points that lie within the subregion are treated together in a single quasi-monochromatic radiative transfer calculation. Consider the case where the error tolerance chosen in Section 2.6 leads to the following numbers of g intervals for each gas: $n_{\text{H}_2\text{O}} = 14$, $n_{\text{CO}_2} = 12$, $n_{\text{O}_3} = 5$, $n_{\text{CH}_4} = 3$, and $n_{\text{N}_2\text{O}} = 2$. If we defined a k -term as the intersection of a single g interval from each gas, the number of k terms required would be the product of the number of g intervals: 5,040 in this case, far too many for a weather or climate model. Hogan (2010) described an automated procedure to optimally partition the hypercube, which

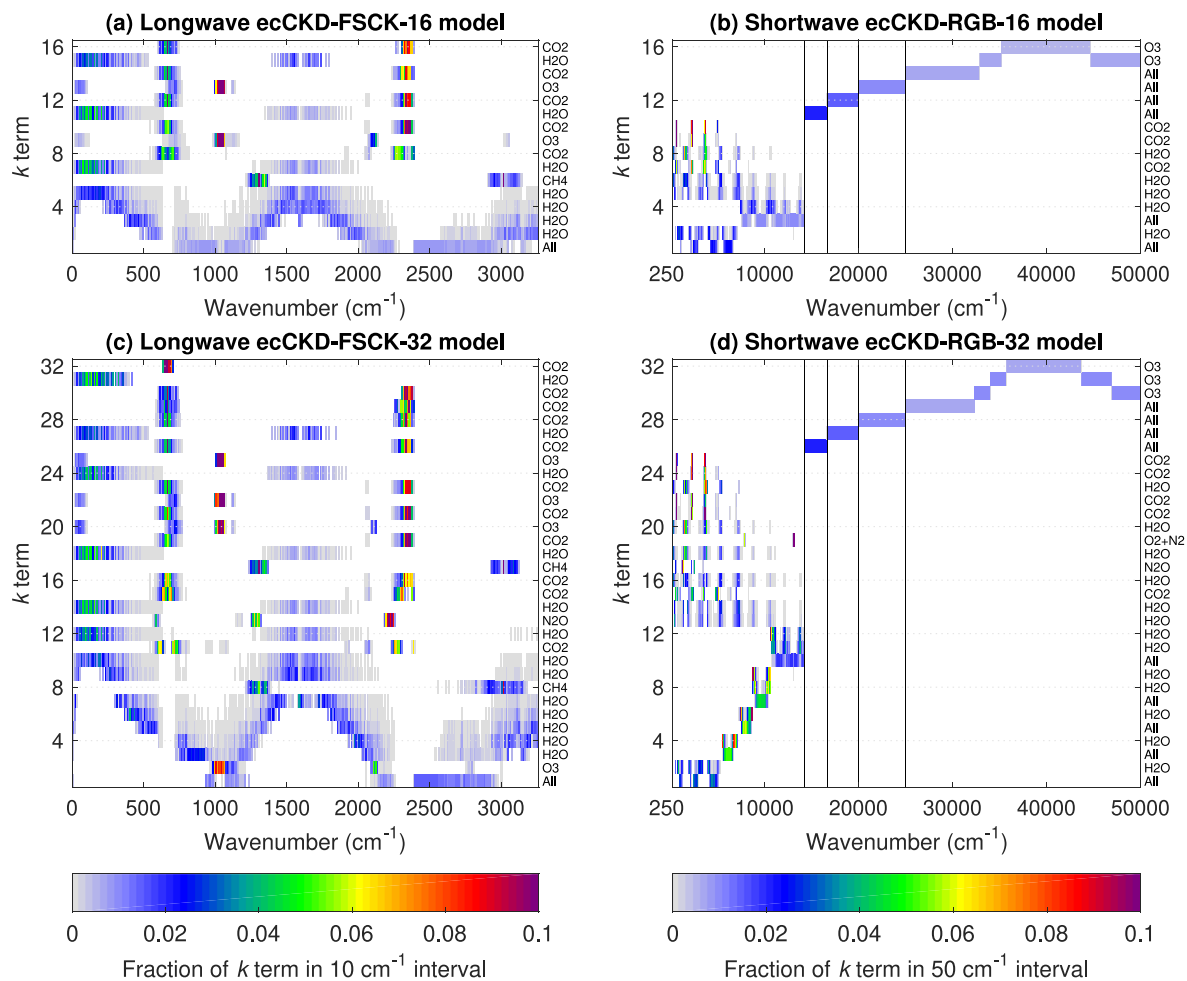


Figure 4. The contribution of each part of the spectrum to each k -term in four of the correlated k -distribution models considered in the text. The fractions sum to unity along each row. The main gas represented by each k -term is indicated on the right-hand-side of each panel, although in principle all gases can contribute to the optical depth for all k -terms. The first (least optically thick) k -term in each band or sub-band is marked “All” since its boundaries are determined by all gases.

recognizes that usually the absorption of one gas dominates over all the others. For example, the red points toward the right of Figure 3a represent wavenumbers for which CO_2 absorption is much stronger than H_2O , and therefore there is no need to resolve variations in H_2O absorption. In this algorithm, the first k -term (numbered 1 in Figure 3a) consists of the intersection of the first g interval for each gas, that is, the weakest absorption. The remaining k -terms are assigned in order of the approximate pressure level of their peak heating or cooling, and each correspond to one of the remaining g intervals for one of the gases, thereby “specializing” in that gas. For example, terms 11, 15, and 16 specialize in CO_2 , although it should be stressed that the (weaker) contribution to the absorption from other gases is still included via Equation 1. Note that the missing numbers in Figure 3a represent terms specializing in other gases that exist outside the plane depicted in this 2D slice; for example, term 2 is for O_3 and term 8 is for CH_4 . If the number of intervals required to partition g space for gas j is n_j , then the number of k -terms required for m gases is $n_{\text{total}} = 1 + \sum_j^m (n_j - 1)$, which is 32 in the case in Figure 3a. The method supports arbitrary spectral overlap of the spectra of individual gases, and is an improvement on the approach of Ritter and Geleyn (1992), which requires $n_{\text{total}} = 1 + \sum_j^m n_j$ terms and makes the assumption that the spectra of individual gases are randomly overlapped. The reader is referred to Hogan (2010) for a more detailed description of the hypercube partition method, and a visualization of the partitioning in the 3D case.

Figure 4c illustrates the contribution of each part of the longwave spectrum to the k -terms of the ecCKD-FSCK-32 model described above, with the gas it is specializing in shown on the right. The “full-spectrum” nature of this CKD model is clear from widely separated parts of the spectrum being represented by single quasi-monochromatic k -terms; for example, term 28 represents strong CO_2 absorption from both the 4.3 μm and the 15 μm CO_2 bands.

Vectorization and computational efficiency favor n_{total} being a power of two, so we have chosen the error tolerances to obtain 16 or 32 k -terms.

The 32-term shortwave model in Figure 4d is very different. The vertical black lines delineate the user-specified “RGB” bands explained in Section 2.4. The partitioning algorithm finds that only one k -term is required for each of the red, green and blue bands, and only four for the entire UV. Terms 30 and 31 each combine the effects of similar levels of O_3 absorption on opposite sides of the Hartley band, which peaks at around $40,000 \text{ cm}^{-1}$. The partitioning of the NIR band is visualized in Figure 3b, and as explained in Section 2.6, the intersection of $0 < g_{\text{H}_2\text{O}} < g_{\text{crit}}$ with the weakest-absorbing g interval for all the other gases (light blue in the figure) is treated separately, where a user-specified value of $g_{\text{crit}} = 0.7$ has been chosen in this case. As shown in Figures 3c and 4d, these wavenumbers are grouped into sub-bands at the additional wavenumber partitions of $5,350$, $7,150$, $8,700$ and $10,650 \text{ cm}^{-1}$ (bounding most of the NIR windows in Figure 2). Each sub-band is then partitioned into g intervals using the algorithm in the previous section, resulting in k -terms 1–12 that are able both to resolve spectral variations in clouds and the surface, and to represent variable H_2O absorption. The k -terms 13–25 represent regions of strong H_2O absorption ($g \geq g_{\text{crit}}$) or strong absorption by one of the other gases. The ecCKD-RGB-16 shortwave model in Figure 4b takes the same approach but with reduced k -terms via the use of $g_{\text{crit}} = 0.65$ and only one additional wavenumber partition at $7,150 \text{ cm}^{-1}$.

2.8. Creating Initial Look-Up Table

This step creates a first estimate of the LUTs in Equation 1 using the “Idealized” CKDMIP data set, and indeed we use the same points (described in Section 3.3 of Hogan & Matricardi, 2020): 53 logarithmically spaced points in pressure from 0.007 to $1,100 \text{ hPa}$ with 10 points per decade; 6 points in temperature, 20 K apart; and 12 logarithmically spaced points in H_2O mixing ratio with 2 points per decade. As shown in Figure 1, the creation of the LUT involves reading in the location of the k -terms, that is, a file containing the indices of the wavenumber points in the high-resolution spectrum that contribute to each term. The layer optical depths in the Idealized data set for the relevant wavenumber points are averaged to each k -term separately for each gas, weighting by the local Planck function in the longwave and the solar spectral irradiance in the shortwave. As in Section 2.6, the default averaging method is linear in layer transmittance for a zenith angle of 60° . The minimum and maximum values from the relevant wavenumber points are also stored and used to bound the possible values in the optimization step described in Section 2.10. The final step in the creation of the LUT is to convert from layer optical depth to molar absorption coefficient as used in Equation 1.

The LUT file contains additional variables that are added at this point and remain unchanged by the subsequent steps shown in Figure 1. The fraction of the spectrum contributing to each k -term (i.e., the information shown in Figure 4) is provided to enable subsequent averaging of cloud, aerosol and surface properties to k -terms, with a resolution of 10 cm^{-1} in the longwave and 50 cm^{-1} in the shortwave. In the longwave we provide the Planck function for each k -term as a LUT versus temperature between 120 and 350 K at 1 K intervals, computed by simply integrating the Planck function over the wavenumber points contributing to each k -term. In the shortwave we provide the solar spectral irradiance for each k -term. Also provided is the Rayleigh molar scattering coefficient, computed for each wavenumber using the Bucholtz (1995) formula and averaged across the parts of the spectrum contributing to each shortwave k -term weighted by the solar spectral irradiance.

2.9. Scaling Shortwave Look-Up Table Entries

At this point, the LUT entries have been computed only from consideration of the spectroscopy at the local pressure level, and do not necessarily perform well in radiative transfer traversing multiple levels. One of the reasons for this in the shortwave is that each k -term represents the average of a range of absorption strengths that are highly correlated in the vertical. As the solar beam traverses the atmosphere, radiation in the more optically thick parts of the spectrum is attenuated more rapidly. This means that lower in the atmosphere the optically thick parts are less important and the effective average molar absorption coefficient for the k -term should be lower than the one computed weighting by the TOA solar spectral irradiance, as in Section 2.8. A similar effect occurs in the longwave, so in both parts of the spectrum non-local effects need to be considered to derive the optimum LUT entries. This section describes the first part of this refinement in the shortwave, while Section 2.10 describes a subsequent more general optimization performed in both the shortwave and longwave.

For a single atmospheric profile of temperature, pressure and gas concentrations, it is possible to derive a profile of layer optical depths for each k -term that reproduces the LBL profile of direct-beam shortwave irradiance exactly, for a particular value of the cosine of the solar zenith angle μ_0 . If we define $F_{j+1/2}$ as the LBL direct irradiance at the interface between layers j and $j + 1$ (counting down from TOA) integrated over the parts of the spectrum corresponding to an individual k -term, then the Beer-Lambert law states that $F_{j+1/2} = F_{j-1/2} \exp(-\tau_j/\mu_0)$. This can be inverted to obtain τ_j , the effective optical depth of layer j .

In the “scale LUT” step in Figure 1 we take this approach using the “median” present-day profile from the CKDMIP “MMM” data set and $\mu_0 = 1/2$, yielding a profile of τ_j values for each k -term. The same values are computed using the ecCKD v0 LUT, and the ratio of the LBL and ecCKD optical depths is calculated to provide a correction factor that varies with pressure and k -term. The correction factor is then interpolated to the pressure grid of the LUT and all the molar absorption coefficients in the file are multiplied by it, producing v1 of the LUT. We have made some significant assumptions here: that the absorption of all gases should be modified by the same proportion, and that the correction factor does not vary with the other dimensions of the LUT (temperature and H₂O mixing ratio). Nonetheless, shortwave radiative transfer calculations using the v1 LUT are significantly more accurate than v0, and there is still the opportunity (described in the following section) for a global optimization of all the coefficients in the LUT.

2.10. Optimizing Look-Up Table Entries

The final task is to optimize the coefficients of the LUTs to minimize the errors in predicted irradiance and heating-rate profiles in a set of training profiles in a least squares sense. The need for this step in terms of non-local dependencies was explained at the start of Section 2.9, but it also tunes the coefficients to mitigate any errors caused by simplifications in the formulation of the CKD model. For example, Equation 1 assumes that the average optical depths from each individual gas in a k -term can be simply summed, whereas Zhang et al. (2003) argued that a more complex formulation was necessary (including much narrower bands) to treat non-random spectral overlap in the parts of the spectrum contributing to a k -term. We find that this complexity and additional cost is unnecessary if the coefficients can be optimized as described in this section.

We improve the basic method of Hogan (2010) in a number of ways. As shown in Figure 1, the optimization may be performed in several steps, as we have found that better performance is achieved if the major gases are optimized first, with minor gases being optimized individually in subsequent steps. We define the state vector \mathbf{x} to contain all the variables to be optimized in one of these steps, specifically the natural logarithm of all the non-zero entries in the look-up for each gas being optimized. With 53 points in pressure, 6 in temperature, 12 in H₂O concentration and (e.g.,) 32 k -terms, this leads to 122,112 state variables for H₂O and 10,176 for each other gas. The cost function to be minimized is given by

$$J = (\mathbf{x} - \mathbf{x}_a)^T \mathbf{B}^{-1} (\mathbf{x} - \mathbf{x}_a) + \sum_{j=1}^p E_j. \quad (3)$$

The first term ensures the stability of the minimization by penalizing the squared differences between the state vector and the a priori LUT elements \mathbf{x}_a , that is, those from the previous step in Figure 1. The error covariance matrix \mathbf{B} provides a complete description of the weighting of this term, with its diagonal elements containing σ_a^2 , the square of the user-specified root-mean-squared (RMS) error in \mathbf{x}_a . We find the best results for $\sigma_a = 8$ in the longwave and $\sigma_a = 2$ in the shortwave, allowing the natural logarithm of the LUT elements to stray significantly from their prior values in the optimization, although in practice the RMS difference between the elements of \mathbf{x} before and after optimization for an individual gas is around 0.25. The off-diagonal elements of \mathbf{B} specify error covariances between LUT values, and have the important effect of spreading information provided by the training profiles into adjacent parts of the LUT. We model the error correlation coefficient of adjacent LUT coefficients along the pressure, temperature and H₂O-concentration axes as ρ , and coefficients n steps apart along these axes as ρ^n . No correlation is assumed between k -terms or gases. Even though \mathbf{B} is large, its inverse is very sparse and the first term in Equation 3 is efficient to compute. Empirically we find that $\rho = 0.8$ provides the best results.

The second term in Equation 3 expresses the sum of the squared errors in heating rates and irradiances at TOA and the surface over p atmospheric profiles, where E_j has the same form as in Equation 2 but rather than penalize errors for individual g intervals, we penalize errors in both broadband values and (when FSCK is not being used)

in individual bands, with the weighting between the two under user control. In practice, broadband irradiances are improved by some compensation of errors between bands, but without the irradiance in individual bands being noticeably compromised.

The cost function is minimized using the quasi-Newton “L-BFGS” algorithm of Liu and Nocedal (1989), which requires the vector of gradients $\partial J/\mathbf{x}$ to be computed. This is achieved by coding the entire algorithm in C++ using the combined automatic differentiation, array and optimization library “Adept” (Hogan, 2014), version 2.1 of which also includes an implementation of the L-BFGS algorithm. We use a bounded version of L-BFGS, constraining the individual absorption coefficients in \mathbf{x} to lie between the minimum and maximum possible values computed in Section 2.8. An optimization step typically takes several tens of minutes to complete.

Naturally, for a CKD model to be used in climate projections we wish to calculate not only the most accurate profiles of irradiances and heating rates, but also the radiative forcing associated with perturbations to both major and minor greenhouse gases. Unfortunately, if all gases are optimized simultaneously, the scheme tends to adjust minor gases to try to offset errors in major gases. This problem can be overcome via several separate optimization steps as shown in Figure 1, first for the major gases and then the minor. We acknowledge that the multi-step approach is somewhat “ad hoc” and there is scope to improve it in future versions of the software, but as will be shown in Section 3, it does produce models that can accurately compute radiative forcing.

In the case of CKD models targeting climate applications (such as those depicted in Figure 4), the first step optimizes the coefficients of H_2O , O_3 , CO_2 , and the background term in Equation 1. The background term represents not only O_2 and N_2 , but also the present-day “reference” concentrations of CH_4 and N_2O , with these gases all treated at this stage as having a constant mixing ratio with pressure. The training data consist of LBL calculations performed on the 50 “Evaluation-1” CKDMIP profiles (covering a wide range of temperature, H_2O and O_3 concentrations), each of which is used with six CKDMIP CO_2 scenarios (surface concentrations from 180 to 2,240 ppmv), that is, a total of 300 profiles. In practice, the effectiveness of the optimization is limited by how well the training profiles span parameter space, and with only 50 base profiles, the error covariance matrix is key for spreading information. This is why steps to improve the initial LUT, such as the scaling described in Section 2.9 are important despite the optimization afterward.

In the second step the coefficients of CH_4 are optimized using LBL calculations on the CKDMIP scenarios in which CH_4 is perturbed from 350 to 3,500 ppbv. To avoid the CH_4 coefficients being tuned to correct for remaining errors in the previous step, we train on the difference in heating-rate profiles and irradiances between perturbed and present-day CH_4 calculations, which is equivalent to bias-correcting the present-day CKD calculations from the previous step. This ensures the CH_4 coefficients are optimized to give the most accurate radiative forcing when perturbed from present-day concentrations. The third step takes exactly the same approach but optimizes the N_2O coefficients training on the Evaluation-1 profiles in which concentrations are perturbed in the range 190–540 ppbv. In the shortwave this yields the “final” LUT ready for use in a radiation scheme. In the longwave we perform one further step to optimize the coefficients of CFC-11 and CFC-12.

3. Clear-Sky Evaluation

In this section, we evaluate the gas-optics models generated in the previous section in clear skies. The CKDMIP Evaluation-2 data set is used, which consists of LBL calculations on 50 independent profiles, including ones with extremes of temperature, ozone and humidity. Surface longwave emissivity and shortwave albedo are spectrally constant at 1.0 and 0.15, respectively, the latter being the approximate mean albedo of the Earth's surface. This approach is the same as that of Hogan and Matricardi (2020) to evaluate the RRTMG gas-optics model, except that they used the CKDMIP Evaluation-1 data set.

Figure 5 evaluates the performance of the longwave ecCKD-FSCK-16 and -32 models, for present-day greenhouse gas concentrations. Even though relatively few k -terms are used compared to other CKD models, the errors are small; the RMS error in heating rates from the surface to the upper stratosphere (4 hPa) is only 0.15 and 0.11 K d^{-1} for the 16- and 32-term models, respectively, approximately doubling in the mesosphere. Figure S2 in Supporting Information S1 evaluates these models before the optimization step described in Section 2.10, and the much larger errors highlight the importance of the optimization.

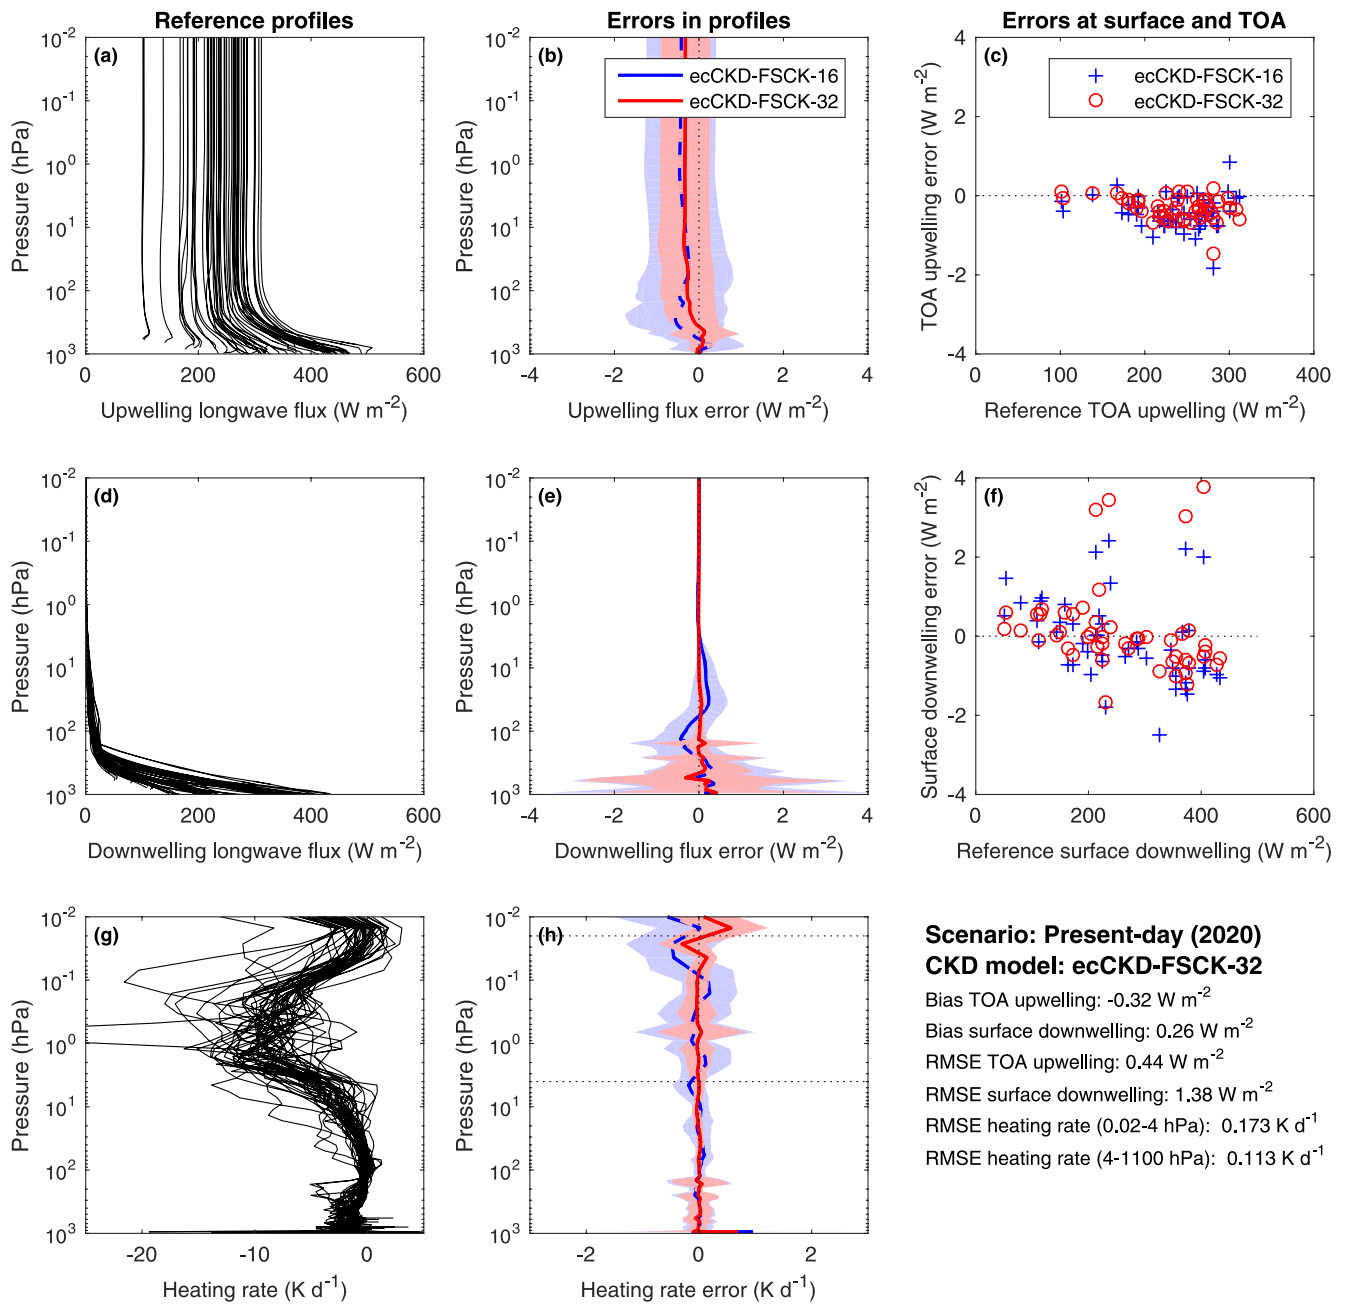


Figure 5. Evaluation of clear-sky longwave irradiances and heating rates from two ecCKD models for the 50 independent profiles of the Correlated K-Distribution Model Intercomparison Project Evaluation-2 data set with present-day concentrations of the well-mixed greenhouse gases. Panels (a, d, g) show quantities from the reference line-by-line calculations, while panels (b, e, h) show the corresponding biases in the ecCKD calculations using an identical radiative transfer solver with four angles per hemisphere. The shaded regions encompass 95% of the errors (estimated as 1.96 multiplied by the standard deviation of the error). Panels (c and f) depict instantaneous errors in upwelling top-of-atmosphere and downwelling surface irradiances. The statistics of the comparison are summarized in the lower right, including the root-mean-squared error in heating rate (weighted by the cube-root of pressure) in two ranges of pressure indicated by the horizontal dotted lines in panel (h).

Figure 6 evaluates the instantaneous radiative forcing associated with perturbing the five main greenhouse-gas concentrations from their present-day values. Note that the CFC-11 concentrations here correspond to artificially increased values to approximately represent 38 further greenhouse gases (Meinshausen et al., 2017). In large part, both models capture the forcing associated with large perturbations to concentrations, including up to eight times preindustrial concentrations of CO_2 . The 16-term model performs slightly worse in some scenarios, tending to underestimate the magnitude of the surface forcing associated with reducing CO_2 concentrations to

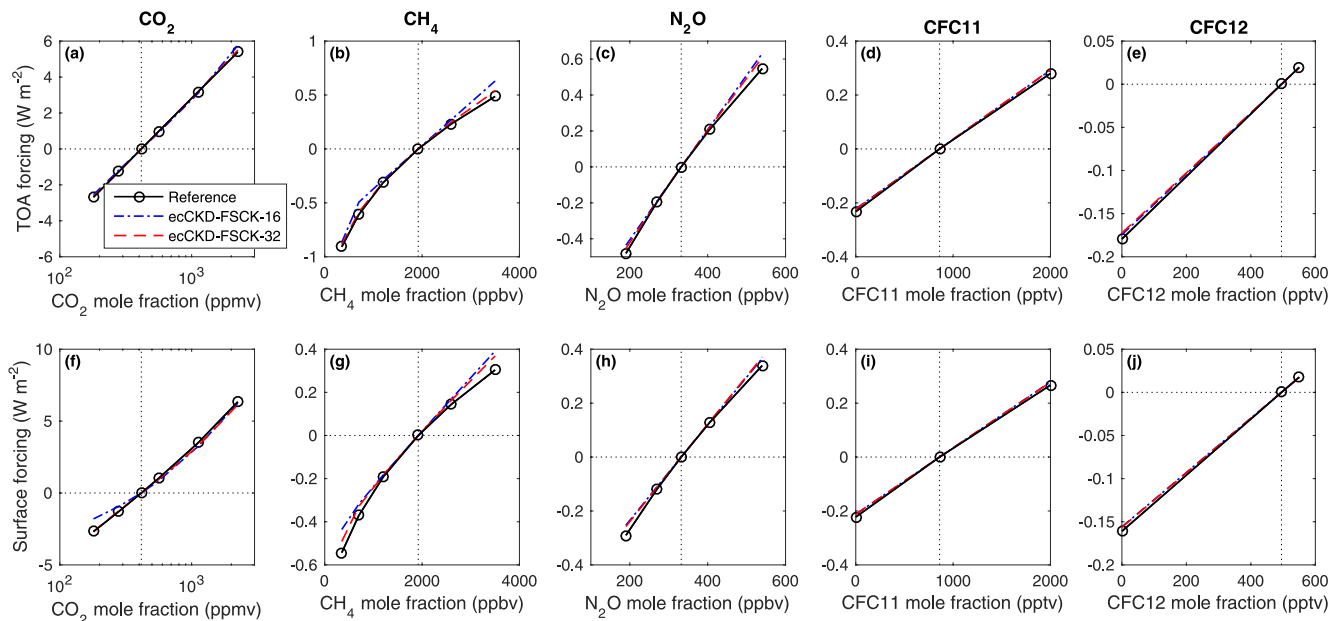


Figure 6. Comparison of reference line-by-line and ecCKD calculations of the instantaneous longwave clear-sky radiative forcing from perturbing each of the five well-mixed greenhouse gases from their present-day (2020) values at (a–e) top of atmosphere and (f–j) the surface, averaged over the 50 profiles of the Correlated K-Distribution Model Intercomparison Project Evaluation-2 data set. The black circles correspond to scenarios 5–22 proposed by Hogan and Matricardi (2020).

glacial-maximum values of 180 ppmv, as well as struggling with the extreme CH_4 concentrations. As can be seen in Figures 4a and 4c, the improvement of the 32-term model for CH_4 can be attributed to its use of two CH_4 -specific k -terms, rather than just one.

Figure 7 presents the corresponding present-day evaluation of irradiances and heating rates for the shortwave ecCKD-RGB-16 and -32 models. Again, the errors are modest given the small number of k -terms, with the RMS error in heating rates from the surface to 4 hPa being 0.1 and 0.06 K d^{-1} for the 16- and 32-term models, respectively. The much larger mid-mesosphere heating-rate error for the 16-term model is associated with its poorer representation of the 4.3 μm CO_2 band; Figures 4b and 4d show that it used only three CO_2 -specific k -terms, compared to five for the 32-band model. This also explains the difference in how well the two models capture the shortwave CO_2 forcing shown in Figure 8. This figure also indicates that the CH_4 forcing in the two models is similar; in fact neither model uses CH_4 -specific k -terms, but rather includes the optical-depth contribution of CH_4 in all the other k -terms. The 16-term model also has no N_2O -specific k -terms and Figure 8 shows that this leads to it tending to overestimate the N_2O forcing by around a factor of two (although the magnitude of the shortwave forcing of this gas is only a tenth of the longwave). The 32-term model introduces a single N_2O -specific k -term and is able to achieve a much greater accuracy. The Supporting Information provides an evaluation of earlier versions of these models just after being initially created (Figure S3 in Supporting Information S1) and after the scaling step described in Section 2.9 (Figure S4 in Supporting Information S1), highlighting the importance of both the scaling and optimizations steps. Figure S5 in Supporting Information S1 evaluates the performance of individual bands of the final ecCKD-RGB-32 model, confirming that the weighting of broadband irradiance in the optimization does not compromise the accuracy of individual bands.

To explore the trade-off between efficiency and accuracy, Figure 9 depicts the biases and RMS errors in TOA and surface irradiances, as well as heating rates, for CKD models generated with between 8 and 64 k -terms. Naturally the errors tend to decrease with more terms, although beyond around 32 terms the improvement is only very modest. Similar behavior was reported by Hogan (2010) for atmospheres containing single gases, which he hypothesized to be due to imperfect rank correlation of the spectra at different heights. This implies we have hit the fundamental limit of the correlated- k method, at least for the FSCK and RGB band structures. It is also noticeable how much larger the errors in surface irradiances (both bias and RMS error) are when evaluating against independent data rather than against the training data used for the optimization step. This suggests the training

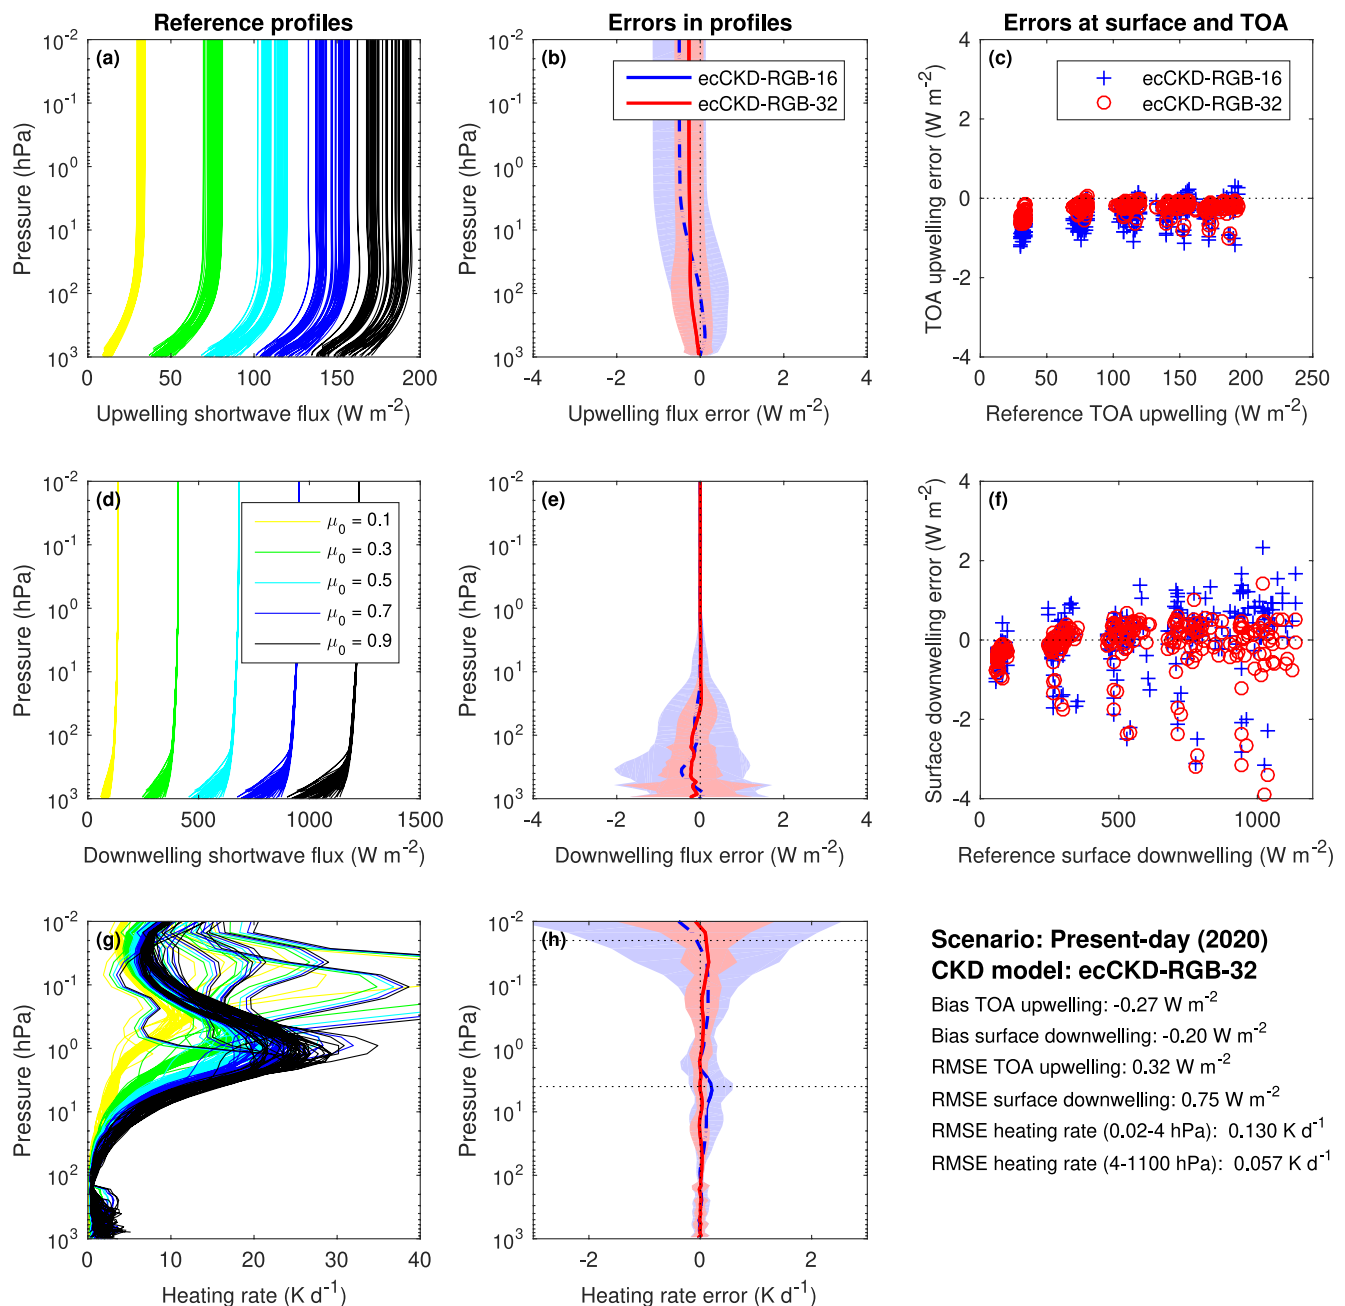


Figure 7. Similar to Figure 5 but for the shortwave. The reference line-by-line calculations in panels (a, d, g) are for all 50 Correlated K-Distribution Model Intercomparison Project Evaluation-2 profiles at five values of the cosine of the solar zenith angle, μ_0 (0.1, 0.3, 0.5, 0.7, and 0.9). The subsequent evaluation considers all 250 combinations. The five clusters of points in panels (c and f) correspond to the five values of μ_0 .

data set is not large enough to tightly constrain all corners of the LUT, and for this reason when generating CKD models to use in the ECMWF model we train on both CKDMIP data sets (“Evaluation-1” and “Evaluation-2”).

4. Cloudy-Sky Evaluation

As discussed in Section 2.4, clouds, aerosols and the surface can exhibit significant spectral variations in optical properties. One of the features of ecCKD is that each k -term has a unique mapping to specific parts of the spectrum (see Figure 4), and this mapping is available to downstream applications (such as the ecRad radiation scheme) so that optical properties can be averaged separately for each k -term. Nonetheless, the use of the

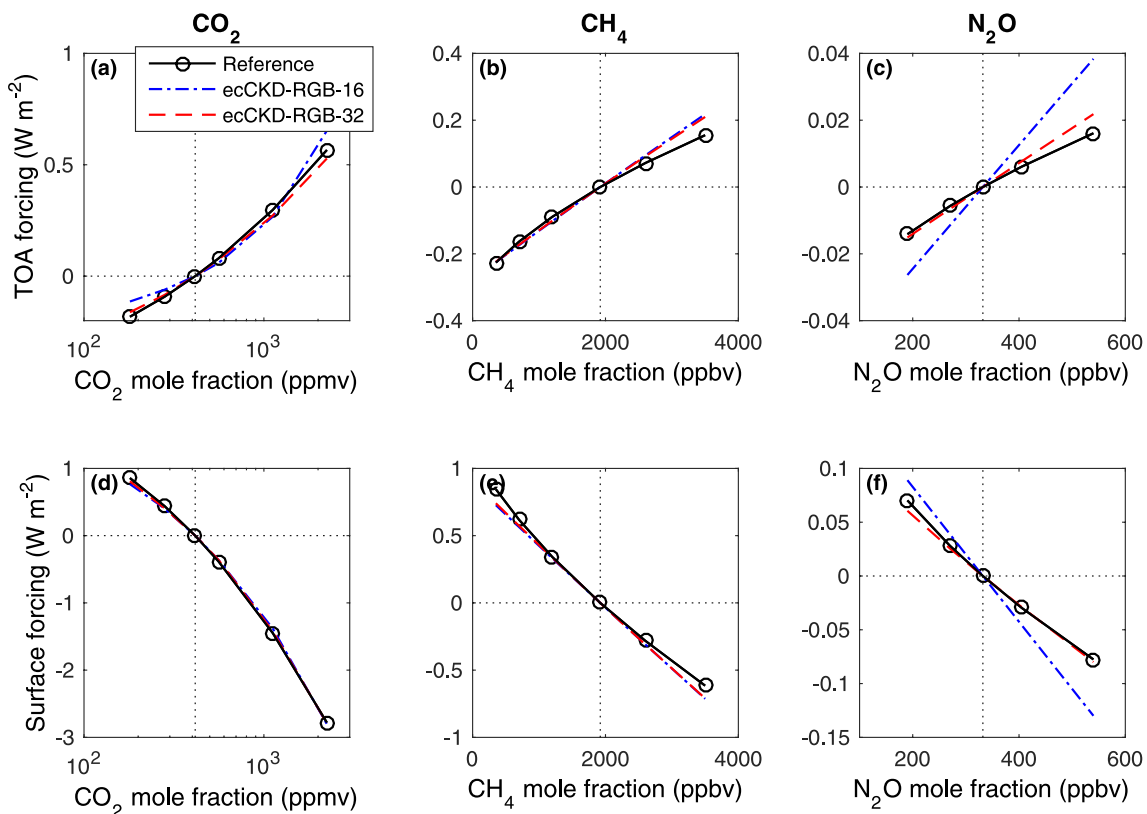


Figure 8. Same as Figure 6, but for the instantaneous shortwave radiative forcing by CO_2 , CH_4 , and N_2O . The results for the five solar zenith angles have been averaged, so the values shown here represent a daytime average.

FSCK approach means that individual k -terms can represent widely separated points in the spectrum. In this section, we test the impact on the accuracy of calculations of the radiative effect of liquid and ice clouds, using two real-world profiles from the CKDMIP data set, summarized in Table 1. Each cloudy layer of the original profile has been divided into 10 and the relative humidity increased to 100%. Twenty-six LBL calculations have been performed on each profile, for water paths ranging from 10^{-4} to 10 kg m^{-2} (plus an additional clear-sky calculation), with a vertically constant cloud mixing ratio between two pressure bounds. The optical properties of liquid clouds are computed using Mie theory at 396 wavenumbers from 5 to $50,000$ cm^{-1} , while the ice properties are taken from the Baum et al. (2014) “General Habit Mixture” available at 445 wavenumbers between 101 and $50,251$ cm^{-1} . When used in LBL calculations, the mass-extinction coefficient, single scattering albedo and asymmetry factor are interpolated linearly in wavenumber space, but clamped when used at wavenumbers outside the range provided. The radiative transfer calculations use a no-scattering solver in the longwave, and the two-stream method with a solar zenith angle of 60° in the shortwave, both from the CKDMIP software package. The plane-parallel approximation is adopted, that is, clouds are taken to be horizontally uniform with a cloud fraction of unity.

The equivalent ecCKD calculations use a version of the ecRad offline radiative transfer package that supports ecCKD gas-optics models, and a radiative transfer solver equivalent to that used for the LBL calculations. Cloud optical properties are computed by averaging the same Mie and Baum et al. (2014) data according to the parts of the spectrum corresponding to each individual k -term (as shown in Figure 4). Following the approach of Edwards and Slingo (1996), we consider both “thin” and “thick” spectral averaging. The former is appropriate in the optically thin limit and simply involves averaging of the mass-extinction coefficient, the mass-absorption coefficient, and averaging asymmetry factor weighted by scattering coefficient. The latter is more appropriate in the optically thick limit and is intended to provide the exact cloud albedo in the limit of infinite optical depth (although in practice it is not exact in this limit except in the absence of gas absorption). For a little extra accuracy, we apply delta-Eddington scaling (Joseph et al., 1976) before performing the spectral average. A further

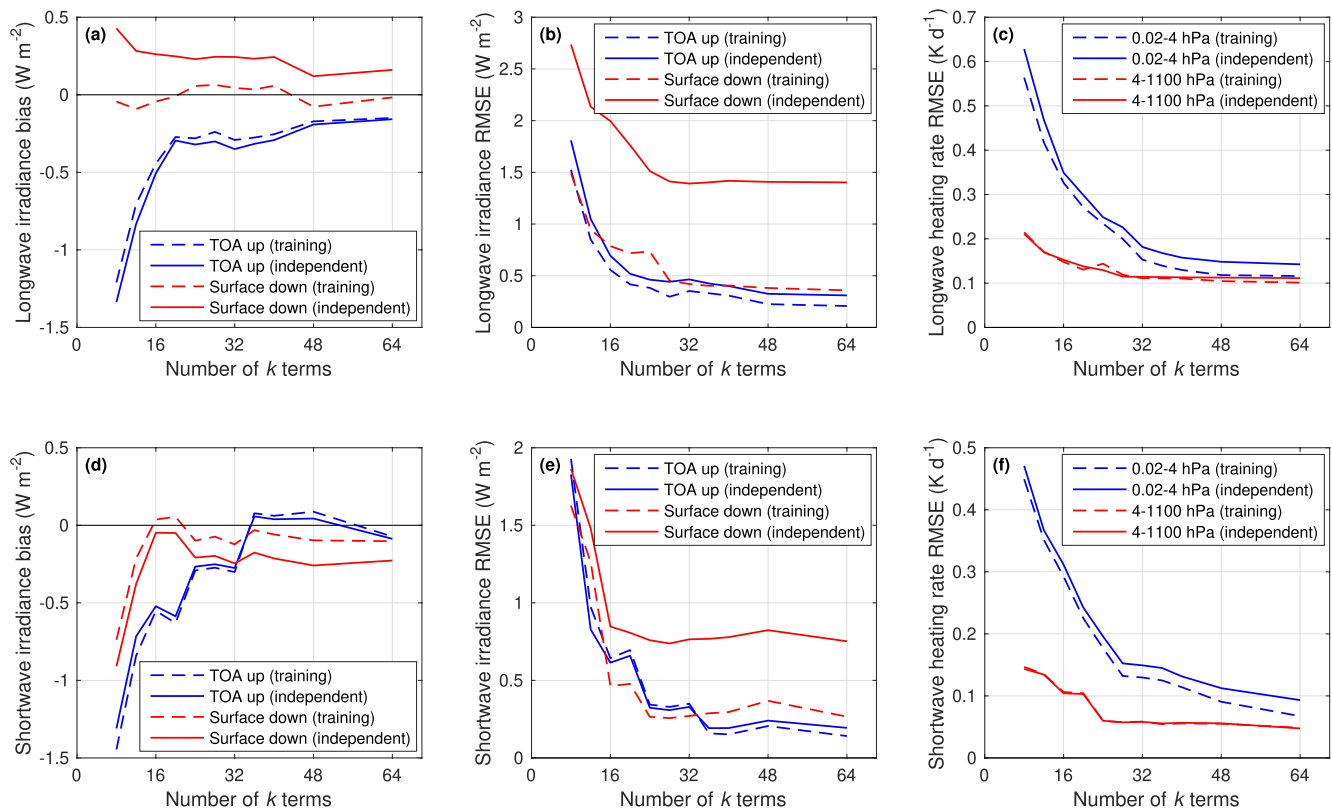


Figure 9. Various metrics of the accuracy of ecCKD models as a function of the number of k -terms for (top row) longwave full-spectrum correlated- k models and (bottom row) shortwave models, as evaluated using (dashed lines) 900 profiles used as part of the training (i.e., the 50 Correlated K-Distribution Model Intercomparison Project [CKDMIP] “Evaluation-1” profiles with greenhouse gases perturbed in 18 CKDMIP scenarios), and (solid lines) an independent set of profiles (i.e., the 50 CKDMIP “Evaluation-2” profiles with the same greenhouse gas scenarios). “RMSE” denotes root-mean-squared error. The shortwave models use the “RGB” band structure except for the 8-term model which uses only two bands on either side of $16,000\text{ cm}^{-1}$. The number of sub-bands used in the near-infrared band of the RGB models has been chosen to be commensurate with the overall accuracy of the scheme; thus the 12-term model does not use sub-bands, the 16-term model uses two, the 20- and 24-term models use three, those with 28–40 terms use five and the 48- and 64-term models use six.

weighting is used in the averaging to approximately represent the energy at each wavenumber; in the longwave we use the Planck function at a representative atmospheric temperature of 0°C and in the shortwave at an effective solar temperature of 5777 K . In the longwave, no benefit was found from using a different reference temperature for liquid and ice clouds. Of primary interest is the accuracy of the fast FSCK and RGB models with 16 and 32 k -terms, which we compare to results from the longwave ecCKD-Narrow-64 and shortwave ecCKD-Window-64 models, which use much narrower bands (see Section 2.4 and Figure S6 in Supporting Information S1). Also shown will be results from the RRTMG gas-optics model (140 k -terms in the longwave and 112 in the shortwave) using the same cloud optical properties, but since no information is available on the exact wavenumbers used for each of its k -terms, the optical properties are averaged to its 16 longwave and 14 shortwave bands.

Figure 10a depicts the LBL calculations of “true” longwave cloud radiative effect at TOA and the surface for the liquid-cloud profile versus water path, with the error in these quantities for various gas-optics models shown in Figures 10b and 10c. The errors for all models are less than 2 W m^{-2} . The ecCKD-Narrow-64 model performs best, although the errors associated with the two FSCK models are still small, being up to around 5% for ecCKD-FSCK-16% and 2% for ecCKD-FSCK-32. Naturally, the use of

Table 1
Summary of the Properties of the Two Atmospheric Profiles Used to Evaluate the Representation of Liquid and Ice Clouds in Section 4

Profile	Liquid cloud	Ice cloud
Cloud effective radius	10 μm	30 μm
Cloud pressure range	726.6–907.1 hPa	184.5–404.6 hPa
Optical properties	Mie theory	Baum et al. (2014)
Origin	Evaluation-2 profile 29	Evaluation-1 profile 28
Location	31.4°S and 3.5°W	38.9°N and 25.6°W
Date and time	19 March 2014, 18 UTC	11 March 2014, 00 UTC
Surface pressure	1,017.4 hPa	1,021.6 hPa
Surface temperature	22.2°C	14.7°C

Note. Both are taken from the Correlated K-Distribution Model Intercomparison Project data set. The size distribution of the liquid cloud was modeled as a gamma distribution with a shape parameter of 2.

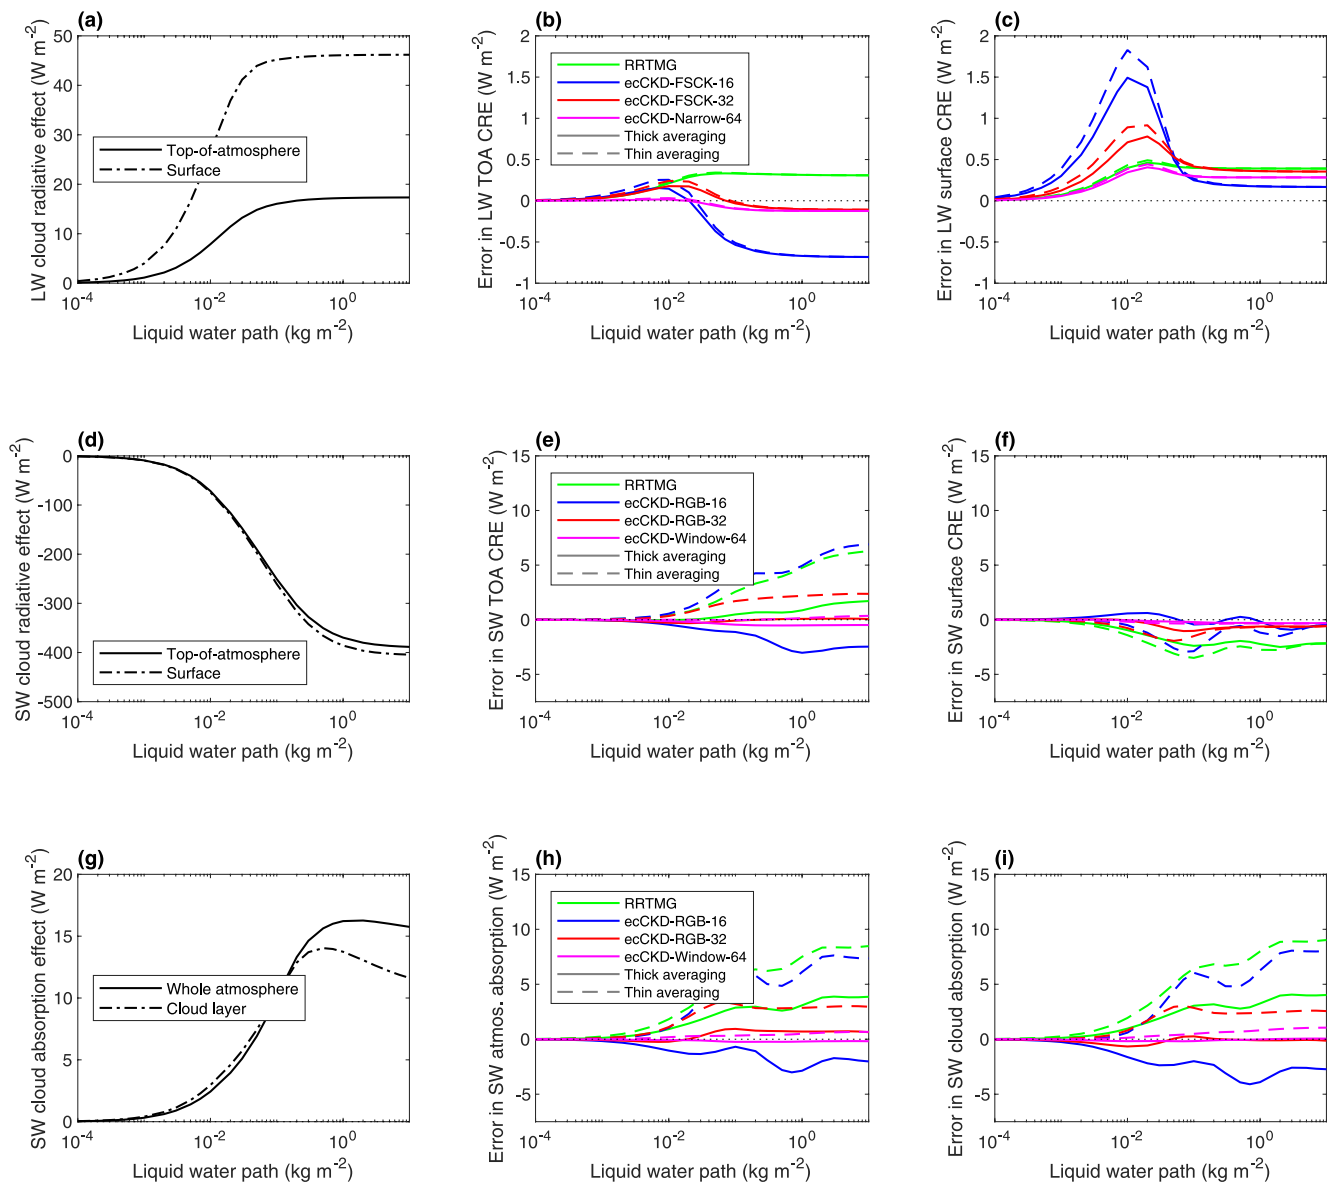


Figure 10. Error in cloud radiative effect (CRE) due to spectral discretization of Rapid Radiative Transfer Model for General Circulation Model (RRTMG) and the various ecCKD models, for a low-level liquid cloud with effective radius $10\ \mu\text{m}$ in Profile 29 of the Correlated k -Distribution Model Intercomparison Project Evaluation-2 data set. Panel (a) depicts the longwave cloud radiative effect (i.e., the change to net irradiance due to cloud) at top-of-atmosphere and the surface as a function of LWP, for the line-by-line calculations. Panels (b) and (c) depict the error in these quantities for the various correlated k -distribution (CKD) models, and for thin and thick spectral averaging. Panels (d)–(f) show the same but in the shortwave with a solar zenith angle of 60° . Panel (g) shows the “cloud absorption effect,” that is, the absorption by the entire atmosphere, and by the cloud layer alone, minus the corresponding clear-sky absorptions. Panels (h) and (i) show the error in these quantities for the CKD models.

narrow bands enables the spectral variation of cloud optical properties to be represented, but it is nonetheless surprising how well the FSCK models perform when they consider the entire longwave spectrum in a single band. This is because (as revealed by the LBL calculations shown in Figure S7 in Supporting Information S1) over 92% of the radiative effect of this cloud at the surface and TOA is in the $8\text{--}13\ \mu\text{m}$ ($769\text{--}1,250\ \text{cm}^{-1}$) longwave atmospheric window, within which the variation of cloud optical properties is quite modest, and certainly much less than in the NIR. Outside the longwave atmospheric window, clouds make a much weaker contribution to longwave cloud radiative effect either because of the much stronger gas absorption or the much weaker Planck function. The top row of Figure 11 shows the equivalent evaluation for the ice cloud profile, where the errors for all longwave gas-optics models are even less, both in an absolute and a relative sense, due to ice particles

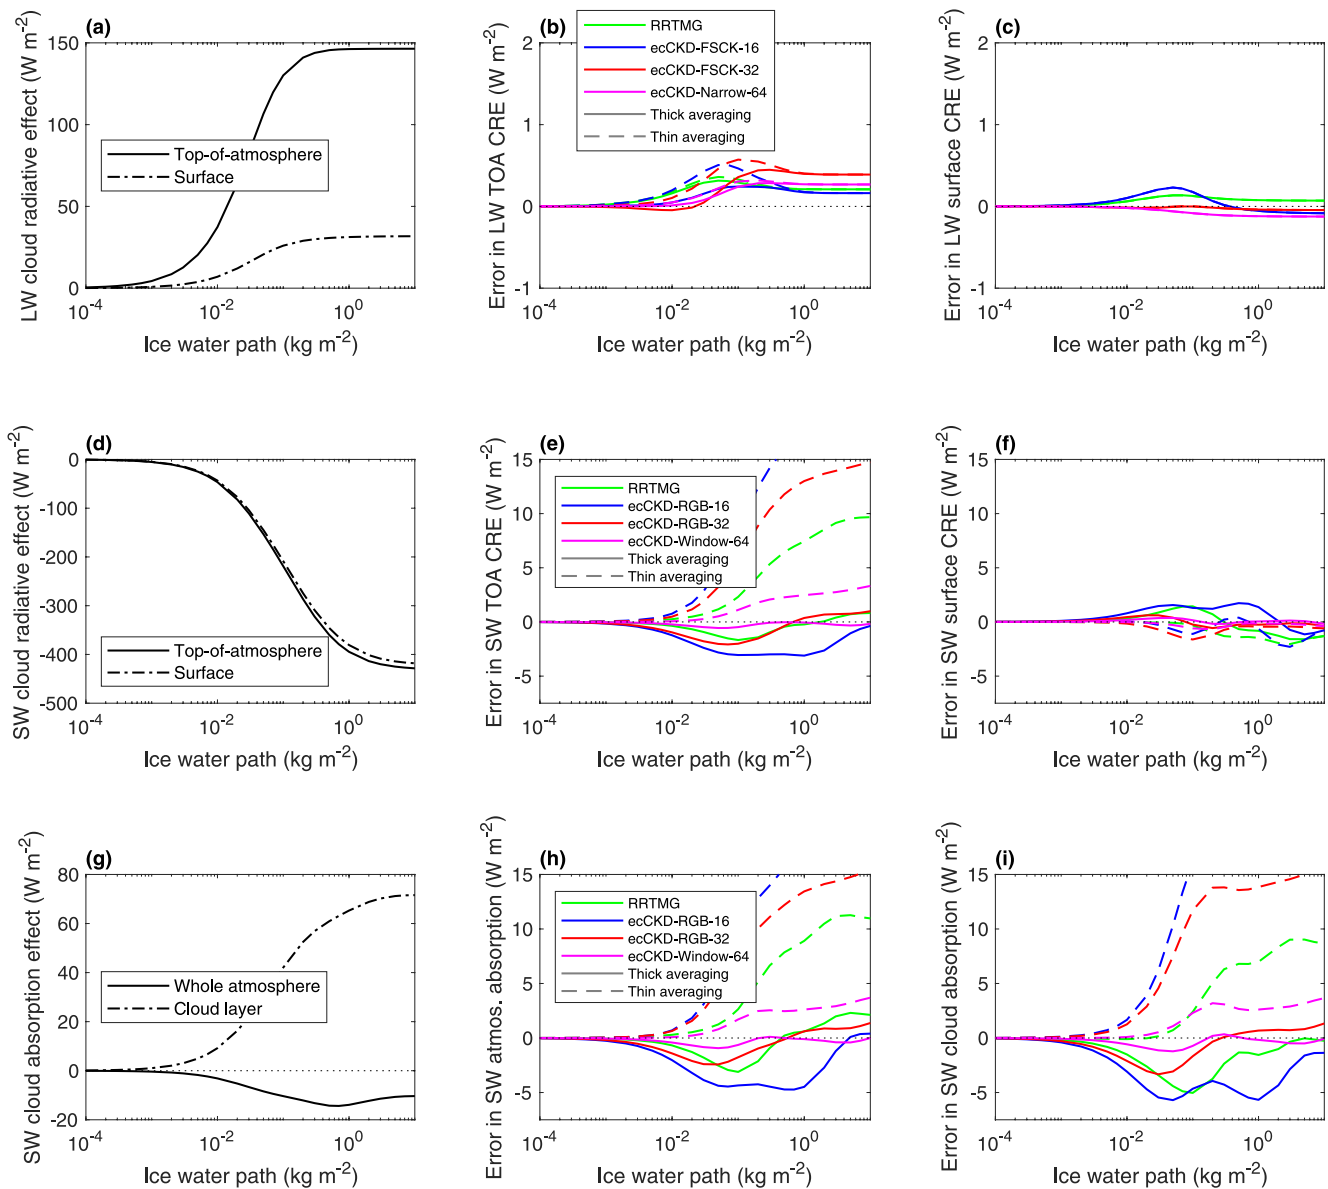


Figure 11. Same as Figure 10, but for an ice cloud with an effective radius of 30 μm in Profile 28 of the Correlated K-Distribution Model Intercomparison Project Evaluation-1 data set.

having less variation in their optical properties than liquid droplets across the longwave spectrum (see Figure S8 in Supporting Information S1). Figures S9 and S10 in Supporting Information S1 show good performance in longwave heating profiles for all models. Overall, these results indicate that the longwave FSCK method is a viable approach for use in weather and climate models, although further work would be required to confirm this result in very dry atmospheres where windows open in the far infrared, and where spectral variations in surface emissivity may also become important.

The middle row of Figure 10 depicts the equivalent evaluation but in the shortwave where the magnitude of the radiative effect of low clouds is much larger. The best performing models are clearly ecCKD-RGB-32 and ecCKD-Window-64 using “thick” averaging, with errors of no more than 1 W m^{-2} (0.4%) for any value of liquid water path. This provides a posteriori justification for the use of five NIR sub-bands in Figure 4d, bounded at the points shown in Figure 2 where cloud optical properties tend to change most rapidly. The ecCKD-RGB-16 model incurs a larger error due to its employing only two sub-bands. The result for ice clouds in the middle row of Figure 11 show the most accuracy for ecCKD-Window-64 and slightly less for ecCKD-RGB-32 and RRTMG.

Figures 10e and 11e suggest that for all gas-optics models the most accurate calculations are achieved using thick rather than thin spectral averaging, except for ice clouds with IWC less than around 0.03 kg m^{-2} where thin averaging is slightly more accurate.

The bottom rows of Figures 10 and 11 consider the effect of the cloud on shortwave absorption, both by the entire atmosphere and by the cloud layer alone. Again, the ecCKD-Window-64 and ecCKD-RGB-32 models with thick averaging performs best, although Figures S9 and S10 in Supporting Information S1 show that the latter is poorer at simulating the vertical profile of shortwave absorption. An interesting feature of Figure 11g is that the effect of the ice cloud is to increase shortwave absorption in the cloud layer itself, as would be expected, but to reduce absorption overall by reflecting sunlight that would otherwise have been absorbed by gases lower in the atmosphere. Thus, the sign of the impact of the cloud on whole-atmosphere absorption is dependent on two competing effects, and while the absolute magnitude of the errors shown in whole-atmosphere and cloud-layer absorption (Figures 11h and 11i) are similar, the relative error in the latter is much larger; indeed, the ecCKD-RGB-16 model with thin averaging predicts that the effect of the ice cloud on atmospheric absorption is to increase rather than to decrease it.

5. Conclusions

In this paper, we have introduced a free software tool “ecCKD” for generating fast CKD gas-optics models for use in the radiation schemes of atmospheric models. The CKD models generated are both accurate and efficient, needing considerably fewer k -terms than most others in the literature. This is achieved via the use of algorithms to optimally partition the k distribution for each gas, and to optimize the LUT coefficients for each gas to minimize errors against hundreds of training profiles (extending the approach of Hogan, 2010). In the shortwave, the introduction of “sub-bands” enables the FSCK approach to treat the entire NIR as a single band, while still enabling the large spectral differences in cloud and surface albedo to be resolved.

The tool has been demonstrated by generating and testing CKD models with only 16 and 32 k -terms in each of the shortwave and longwave, that is, nearly a factor of 8 and 4 times fewer, respectively, than the total number used operationally at ECMWF. When evaluated against independent data, the 32-term models are shown to be very accurate in clear skies, with RMS heating-rate errors of less than 0.18 K d^{-1} from the troposphere to the mid-mesosphere. The radiative forcing of the main anthropogenic greenhouse gases is captured accurately, including CO_2 varying over a factor of 12 and CH_4 over a factor of 10. The 32-term models have been found to perform well when run online in the ECMWF forecast model, to be explored in a future paper. The 16-term models are naturally somewhat less accurate, but would be suitable for short forecasts such as 12-hr forecasts performed repeatedly in a data-assimilation cycle in which efficiency is paramount.

We have used LBL calculations for profiles containing liquid and ice clouds with a large range of water contents to verify the accuracy of the FSCK approach in cloudy situations. In the longwave, provided that cloud properties are averaged over each k -term rather than per band, errors in irradiances calculated using the 32-term model are less than 0.7 W m^{-2} , an important demonstration of the viability of the longwave FSCK approach for cloudy terrestrial atmospheres. In the shortwave, the use of sub-bands in the NIR gives the 32-term ecCKD model comparable accuracy to RRTMG but using only 25 rather than 78 terms in the NIR. A 64-term shortwave ecCKD model with an explicit band for each NIR window and a total of 48 terms in the NIR is found to be considerably more accurate than either RRTMG or the 32-term ecCKD model in cloudy skies.

The tool described in this paper offers a number of opportunities for users of radiation schemes. Principally, it allows optimized CKD models to be generated for specific applications, from present-day NWP to palaeoclimate simulations of periods when atmospheric composition was very different. Moreover, the fact that the CKD models generated tend to be faster while of similar accuracy to existing models frees up computer time to improve the accuracy of other parts of the radiation scheme, such as the use of more than two streams (e.g., Fu et al., 1997), inclusion of 3D effects (e.g., Hogan et al., 2016), and calling the scheme more frequently in time and space (e.g., Hogan & Bozzo, 2018). Additionally, the use of a simple LUT to compute optical depths (Equation 1) makes it straightforward to incorporate the CKD models into different types of radiation scheme, including explicit 3D solvers (e.g., Jakub & Mayer, 2016). It would also be possible to add the capability for ecCKD to generate CKD models suitable for satellite data assimilation by simply replacing Equation 2 by a cost function that penalizes only errors in TOA radiances.

Appendix A: Equipartition Algorithm

Section 2.6 outlined the partitioning of a reordered spectrum into intervals such that each interval was associated with around the same mean-squared error in a radiation calculation, and less than a user-specified tolerance E_{tol} . The 1D space to be partitioned is denoted g and ranges from 0 to 1. We seek the boundaries of n intervals, denoted $g_0, g_1 \dots g_n$, such that the following two conditions are satisfied:

$$E(g_{i-1}, g_i) \leq E_{\text{tol}} \quad \text{for all } i; \quad (\text{A1})$$

$$F \leq F_{\text{tol}}, \quad (\text{A2})$$

where the error $E(g_{i-1}, g_i)$ is a non-differentiable user-supplied function, and the second condition states that the fractional range of errors, $F = [\max(E) - \min(E)]/\bar{E}$, should be no larger than the user-supplied tolerance F_{tol} , typically 0.02. In our case, $E(g_{i-1}, g_i)$ is given by Equation 2 and involves LBL calculations with a computational cost proportional to the width of the interval $g_i - g_{i-1}$. Therefore, a good partitioning algorithm should not require an excessive number of calculations of E , especially ones for wide g intervals. We are not aware of an off-the-shelf algorithm for performing this partitioning, so this appendix describes our solution to the problem. While it is not likely to be the fastest possible algorithm, it almost always converges to a solution that satisfies the conditions above.

The first task is to find the number of intervals required, n , in order that Equation A1 is satisfied. This is achieved by partitioning g space starting at the lower end such that for each interval (except possibly the last), $0.95E_{\text{tol}} \leq E \leq E_{\text{tol}}$. We start with a test value of $g_1 = 0.75$ and compute $E(0, g_1)$, noting that there is a lower bound on g_1 of 0 where $E(0, 0) = 0$. If the result is less than $0.95E_{\text{tol}}$ then a new lower bound for g_1 has been found, and the next test value is selected by extrapolating (but not beyond $g_1 = 1$) from the old and new lower bounds to where we would expect $E(0, g_1) = E_{\text{tol}}$ assuming a linear variation of E with g_1 . On the other hand, if $E(0, g_1) > E_{\text{tol}}$ then an upper bound for g_1 has been found, and the next test value is found by linearly interpolating between the lower and upper bounds on g_1 . The new test value is used to compute $E(0, g_1)$ and the process is repeated until either $0.95E_{\text{tol}} \leq E \leq E_{\text{tol}}$, or $g_1 = 1$ and $E \leq E_{\text{tol}}$. If there is remaining g space to partition then the process is repeated to compute g_2 and so on, until the process returns $g_n = 1$. We now know how many intervals are required, and have candidate values for g_i , but usually the error associated with the final interval, $E(g_{n-1}, 1)$, is significantly less than all the other errors.

The second task is to find the interior boundaries of the intervals ($g_1, g_2 \dots g_{n-1}$) such that Equation A2 is satisfied. Note that the outer boundaries are already fixed at $g_0 = 0$ and $g_n = 1$. In the simple case of $n = 2$ we have only a single value to find, g_1 ; it is straightforward to progressively refine this value until Equation A2 is satisfied. In the more general $n > 2$ case, we use the candidate values of g_i from the first task above to compute the cumulative error as

$$C(g_i) = \sum_{j=1}^i E(g_{j-1}, g_j). \quad (\text{A3})$$

A new set of candidate values g'_i is found by attempting to repartition the total error, $C(1)$, evenly amongst the n intervals. This is achieved by using linear interpolation into the C function to compute the g'_i values such that $C(g'_i) = iC(1)/n$. The errors are recomputed and the process is repeated until Equation A2 is satisfied. If at any iteration the fractional range F increases then the iteration is not successful and instead a “shuffle” step is performed. This consists of looping through adjacent pairs of intervals and adjusting the g point between them until their errors agree to within 2%. Thus, for intervals 1 and 2 we adjust g_1 until $E(g_0, g_1)$ and $E(g_1, g_2)$ satisfy Equation A2, then do the same for intervals 2 and 3 and so on up to intervals $n - 1$ and n , followed by a pass back down to intervals 1 and 2. This is usually enough that subsequent partitioning iterations using Equations A3 lead to a reduction of F . Any further shuffle operations proceed in the opposite direction through g space as the previous one.

Data Availability Statement

The ecCKD software is available from <https://doi.org/10.5281/zenodo.6670050> and <https://github.com/ecmwf-ifs/ecckd>. It makes use of data from the CKDMIP project, available via the links at <http://confluence.ecmwf.int/display/CKDMIP>. The netCDF files defining the 16-, 32-, and 64-term gas-optics models described in this paper are available at the ecCKD web site <https://confluence.ecmwf.int/x/XwUODw>.

Acknowledgments

The authors thank Steve English and Richard Forbes for constructive comments on the original manuscript. This work has benefited from valuable discussions with Robert Pincus.

References

- Barker, H. W., Qu, Z., Dhanraj, V., & Cole, J. N. S. (2021). Partial validation of a lossy compression approach to computing radiative transfer in cloud system-resolving models. *Quarterly Journal of the Royal Meteorological Society*, *147*(734), 363–381. <https://doi.org/10.1002/qj.3922>
- Baum, B. A., Yang, P., Heymsfield, A. J., Bansemmer, A., Merrelli, A., Schmitt, C., & Wang, C. (2014). Ice cloud bulk single-scattering property models with the full phase matrix at wavelengths from 0.2 to 100 μm . *Journal of Quantitative Spectroscopy & Radiative Transfer*, *146*, 123–139. <https://doi.org/10.1016/j.jqsrt.2014.02.029>
- Bennartz, R., & Fischer, J. (2000). A modified k -distribution approach applied to narrow band water vapour and oxygen absorption estimates in the near infrared. *Journal of Quantitative Spectroscopy & Radiative Transfer*, *66*(6), 539–553. [https://doi.org/10.1016/S0022-4073\(99\)00184-3](https://doi.org/10.1016/S0022-4073(99)00184-3)
- Bucholtz, A. (1995). Rayleigh-scattering calculations for the terrestrial atmosphere. *Applied Optics*, *34*(15), 2765–2773. <https://doi.org/10.1364/AO.34.002765>
- Ciavatta, S., Torres, R., Martinez-Vicente, V., Smyth, T., Dall'Olmo, G., Polimene, L., & Icarus Allen, J. (2014). Assimilation of remotely-sensed optical properties to improve marine biochemistry modelling. *Progress in Oceanography*, *127*, 74–95. <https://doi.org/10.1016/j.pocean.2014.06.002>
- Coddington, O., Lean, J. L., Pilewskie, P., Snow, M., & Lindholm, D. (2016). A solar irradiance climate data record. *Bulletin of the American Meteorological Society*, *97*(7), 1265–1282. <https://doi.org/10.1175/BAMS-D-14-00265.1>
- DeAngelis, A., Qu, X., Zelinka, M., & Hall, A. (2015). An observational radiative constraint on hydrologic cycle intensification. *Nature*, *528*(7581), 249–253. <https://doi.org/10.1038/nature15770>
- Doppler, L., Preusker, R., Bennartz, R., & Fischer, J. (2013). k -bin and k -IR: K -distribution methods without correlation approximation for non-fixed instrument response function and extension to the thermal infrared—Applications to satellite remote sensing. *Journal of Quantitative Spectroscopy & Radiative Transfer*, *133*, 382–395. <https://doi.org/10.1016/j.jqsrt.2013.09.001>
- Edwards, J. M., & Slingo, A. (1996). Studies with a flexible new radiation code: 1. Choosing a configuration for a large-scale model. *Quarterly Journal of the Royal Meteorological Society*, *122*(531), 689–719. <https://doi.org/10.1002/qj.49712253107>
- Eresmaa, R., & McNally, A. P. (2014). Diverse profile datasets from the ECMWF 137-level short-range forecasts (NWP-SAF Document NWPSAF_EC_TR_017). Retrieved from <https://www.nwpsaf.eu/site/software/atmospheric-profile-data/>
- Fu, Q., Liou, K. N., Cribb, M. C., Charlock, T. P., & Grossman, A. (1997). Multiple scattering parameterization in thermal infrared radiative transfer. *Journal of the Atmospheric Sciences*, *54*(24), 2799–2812. [https://doi.org/10.1175/1520-0469\(1997\)054<2799:MSPITI>2.0.CO;2](https://doi.org/10.1175/1520-0469(1997)054<2799:MSPITI>2.0.CO;2)
- Goody, R., West, R., Chen, L., & Crisp, D. (1989). The correlated- k method for radiation calculations in nonhomogeneous atmospheres. *Journal of Quantitative Spectroscopy & Radiative Transfer*, *42*(6), 539–550. [https://doi.org/10.1016/0022-4073\(89\)90044-7](https://doi.org/10.1016/0022-4073(89)90044-7)
- Hogan, R. J. (2010). The full-spectrum correlated- k method for longwave atmospheric radiation using an effective Planck function. *Journal of the Atmospheric Sciences*, *67*(6), 2086–2100. <https://doi.org/10.1175/2010JAS3202.1>
- Hogan, R. J. (2014). Fast reverse-mode automatic differentiation using expression templates in C++. *ACM Transactions on Mathematical Software*, *40*(4), 1–16. <https://doi.org/10.1145/2560359>
- Hogan, R. J. (2022). The ecCKD software package for generating gas optics models for atmospheric radiation schemes (v1.2). Zenodo. <https://doi.org/10.5281/zenodo.6670050>
- Hogan, R. J., Ahlgrimm, M., Balsamo, G., Beljaars, A. C. M., Berrisford, P., Bozzo, A., et al. (2017). *Radiation in numerical weather prediction* (Technical Memorandum No. 816, pp. 48). ECMWF. <https://doi.org/10.21957/2bd5dkj8x>
- Hogan, R. J., & Bozzo, A. (2018). A flexible and efficient radiation scheme for the ECMWF model. *Journal of Advances in Modeling Earth Systems*, *10*(8), 1990–2008. <https://doi.org/10.1029/2018MS001364>
- Hogan, R. J., & Matricardi, M. (2020). Evaluating and improving the treatment of gases in radiation schemes: The Correlated K -Distribution Model Intercomparison Project (CKDMIP). *Geoscientific Model Development*, *13*(12), 6501–6521. <https://doi.org/10.5194/gmd-13-6501-2020>
- Hogan, R. J., Schäfer, S. A. K., Klinger, C., Chiu, J.-C., & Mayer, B. (2016). Representing 3D cloud-radiation effects in two-stream schemes: 2. Matrix formulation and broadband evaluation. *Journal of Geophysical Research: Atmospheres*, *121*(14), 8583–8599. <https://doi.org/10.1002/2016JD024875>
- Jakub, F., & Mayer, B. (2016). 3-D radiative transfer in large-eddy simulations – Experiences coupling the TenStream solver to the UCLA-LES. *Geoscientific Model Development*, *9*(4), 1413–1422. <https://doi.org/10.5194/gmd-9-1413-2016>
- Joseph, J. H., Wiscombe, W. J., & Weinman, J. A. (1976). The delta-Eddington approximation for radiative flux-transfer. *Journal of the Atmospheric Sciences*, *33*(12), 2452–2459. [https://doi.org/10.1175/1520-0469\(1976\)033<2452:TDEAFR>2.0.CO;2](https://doi.org/10.1175/1520-0469(1976)033<2452:TDEAFR>2.0.CO;2)
- Kato, S., Ackerman, T. P., Mather, J. H., & Clothiaux, E. E. (1999). The k -distribution method and correlated- k approximation for a short-wave radiative transfer model. *Journal of Quantitative Spectroscopy and Radiative Transfer*, *62*(1), 109–121. [https://doi.org/10.1016/S0022-4073\(98\)00075-2](https://doi.org/10.1016/S0022-4073(98)00075-2)
- Lacis, A., & Oinas, V. (1991). A description of the correlated k -distribution method for modeling nongray gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres. *Journal of Geophysical Research*, *96*(D5), 9027–9063. <https://doi.org/10.1029/90JD01945>
- Liu, D. C., & Nocedal, J. (1989). On the limited memory method for large scale optimization. *Mathematical Programming B*, *45*(1–3), 503–528. <https://doi.org/10.1007/BF01589116>
- Lopez, P. (2020). Forecasting the past: Views of Earth from the Moon and beyond. *Bulletin of the American Meteorological Society*, *101*(11), E1190–E1200. <https://doi.org/10.1175/BAMS-D-19-0254.1>
- Lu, P., Zhang, H., & Li, J. (2011). Correlated k -distribution treatment of cloud optical properties and related radiative impact. *Journal of the Atmospheric Sciences*, *68*(11), 2671–2688. <https://doi.org/10.1175/JAS-D-10-05001.1>
- Manabe, S., & Wetherald, R. T. (1967). Thermal equilibrium of the atmosphere with a given distribution of relative humidity. *Journal of the Atmospheric Sciences*, *24*(3), 241–259. [https://doi.org/10.1175/1520-0469\(1967\)024<0241:TEOTAW>2.0.CO;2](https://doi.org/10.1175/1520-0469(1967)024<0241:TEOTAW>2.0.CO;2)

- Meinshausen, M., Vogel, E., Nauels, A., Lorbacher, K., Meinshausen, N., Etheridge, D. M., et al. (2017). Historical greenhouse gas concentrations for climate modelling (CMIP6). *Geoscientific Model Development*, *10*(5), 2057–2116. <https://doi.org/10.3929/ethz-b-000191830>
- Mlawer, E. J., Taubman, S. J., Brown, P. D., Iacono, M. J., & Clough, S. A. (1997). Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave. *Journal of Geophysical Research*, *102*(D14), 16663–16682. <https://doi.org/10.1029/97JD00237>
- Modest, M. F., & Zhang, H. (2002). The full-spectrum correlated-*k* distribution for thermal radiation from molecular gas-particulate mixtures. *Journal of Heat Transfer*, *124*(1), 30–38. <https://doi.org/10.1115/1.1418697>
- Pawlak, D. T., Clothiaux, E. E., Modest, M. F., & Cole, J. N. S. (2004). Full-spectrum correlated-*k* distribution for shortwave atmospheric radiative transfer. *Journal of the Atmospheric Sciences*, *61*(21), 2588–2601. <https://doi.org/10.1175/JAS3285.1>
- Ritter, B., & Geleyn, J. F. (1992). A comprehensive radiation scheme for numerical weather prediction models with potential applications in climate simulations. *Monthly Weather Review*, *120*(2), 303–325. [https://doi.org/10.1175/1520-0493\(1992\)120<0303:ACRSFN>2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120<0303:ACRSFN>2.0.CO;2)
- Ukkonen, P., Pincus, R., Hogan, R. J., Nielsen, K. P., & Kaas, E. (2020). Accelerating radiation computations for dynamical models with targeted machine learning and code optimization. *Journal of Advances in Modeling Earth Systems*, *12*, e2020MS002226. <https://doi.org/10.1029/2020MS002226>
- Zhang, H., Nakajima, T., Shi, G., Suzuki, T., & Imasu, R. (2003). An optimal approach to overlapping bands with correlated k distribution method and its application to radiative calculations. *Journal of Geophysical Research*, *108*(D20), 4641. <https://doi.org/10.1029/2002JD003358>