

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

3 **Robin J. Hogan¹ and Marco Matricardi¹**

4 ¹European Centre for Medium-Range Weather Forecasts, Reading, UK.

5 **Key Points:**

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

Corresponding author: Robin J. Hogan, r.j.hogan@ecmwf.int

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 parts of the spectrum. In the near-infrared, the large spectral variation in cloud absorption is represented by partitioning the parts of the spectrum where gases are optically thin into three or more 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 near-infrared 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. 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 are 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 convection scheme (Manabe and Wetherald, 1967). The correlated k -distribution (CKD) method (Goody et al., 1989; Lacis and 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 water (NIR) 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 gas-optics scheme. Along with several other modeling centers, the European Centre for Medium-Range Weather Forecasts (ECMWF) 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 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, one per gas. While these files can be read by ECMWF’s ‘ecRad’ radiation scheme ([Hogan and 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 and 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 demonstrated in the atmosphere 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 at different heights 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. \(2014\)](#). 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 problem 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 atmospheric profiles, and in section 4 the models are evaluated in cloudy skies.

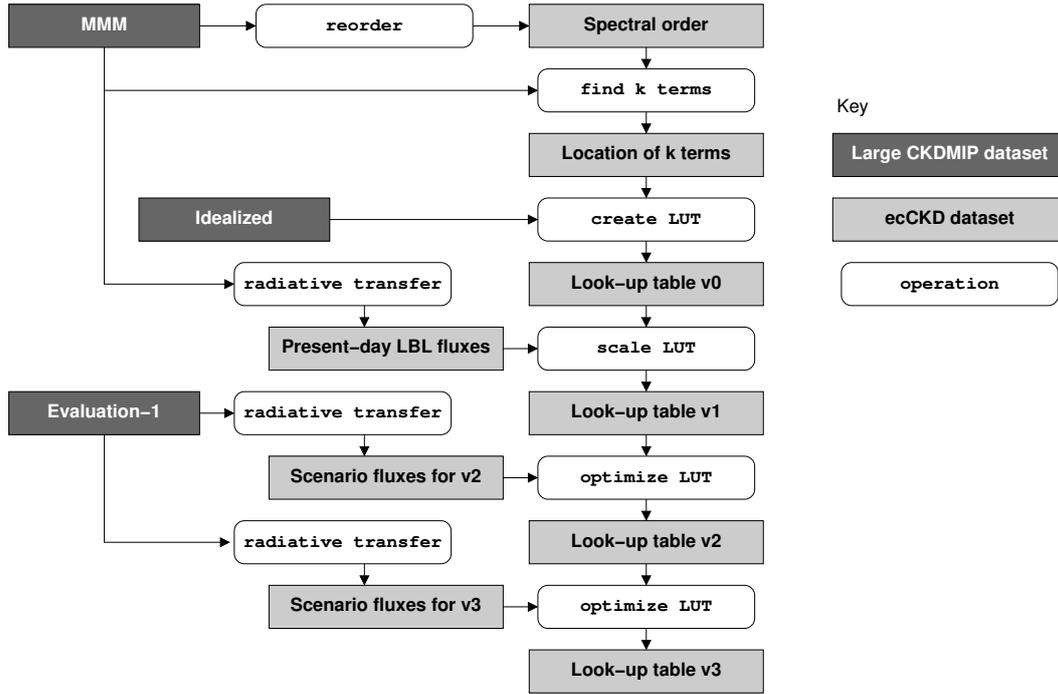


Figure 1. Flowchart illustrating the steps of the ecCKD method (white boxes), which make use of the large CKDMIP datasets of [Hogan and Matricardi \(2020\)](#) and store intermediate information in smaller netCDF files (light grey boxes), ultimately producing a look-up table (LUT) file for use in a radiation scheme.

2 Method

2.1 Overview

The ecCKD tool performs a sequence of tasks shown in Fig. 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 and Matricardi, 2020](#)), performing LBL radiation calculations on them as needed. The datasets 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 spectrum is 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 look-up tables (LUTs) is created for each gas 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 in order 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 dataset described by [Hogan and Matricardi \(2020\)](#), appropriate for the terrestrial atmosphere. The absorbing gases considered are H₂O, O₃, CO₂, CH₄, N₂O, O₂, N₂, CFC-11 and CFC-12. Except for

133 H₂O, the molar absorption of all these gases can be considered independent of their concen-
 134 tration, so a very wide range of climate scenarios can be considered by simply scaling the
 135 absorptions. According to Meinshausen et al. (2017), this choice of gases represents 94.5%
 136 of anthropogenic greenhouse warming in terms of longwave radiative forcing since 1750, and
 137 a further 38 more minor gases (representing the remaining 5.5% of radiative forcing) can be
 138 adequately represented by using an increased ‘equivalent’ concentration of CFC-11. This ap-
 139 proach is common in the various phases of the Coupled Model Intercomparison Project given
 140 that most or all climate-model radiation schemes are unable to represent all the 43 well-mixed
 141 greenhouse gases listed by Meinshausen et al. (2017). However, it should be noted that there
 142 is no reason in principle why other gases could not be added to ecCKD if LBL calculations
 143 were performed to provide additional absorption spectra in the appropriate format.

144 As described by Hogan and Matricardi (2020), the spectral resolution of the dataset is
 145 variable, being finest in the strong absorption bands of CO₂ because at mesospheric altitudes
 146 the CO₂ lines are only Doppler broadened and thus become very narrow. This results in a total
 147 of 7 211 999 spectral points in the longwave and 3 126 494 in the shortwave. The total volume
 148 of the dataset is around 1 TB.

149 2.3 Gas representation

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

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

153 where ψ_j is the mole fraction of gas j . The four terms on the right-hand-side represent the four
 154 different ways that gases can be represented. The first is the *background* term, a 2D look-up
 155 table representing the combined contribution from all gases with a constant, pre-defined mole
 156 fraction. For a model intended for climate simulations, this would typically include only O₂
 157 and N₂, but for a CKD model targeting present-day NWP we could include the contribution
 158 from well-mixed greenhouse gases. The fact that the LUT includes a dependence on pres-
 159 sure means that a pressure dependence of these gases can be represented. The second term
 160 represents gases 1 to l , whose absorption varies *linearly* with mole fraction; in this case k_i^j
 161 is the molar absorption coefficient of gas j , i.e. the absorption cross-section per mole of the
 162 gas. The third term represents gases $l + 1$ to n whose absorption varies *nonlinearly* with con-
 163 centration; in this case a 3D LUT is used for molar absorption coefficient, with an additional
 164 dependence on the mole fraction of the gas in question. In the terrestrial atmosphere only H₂O
 165 is in this category, and the representation here allows the contribution from the water vapour
 166 continuum (both self and foreign) to be treated completely with no need to separate the line
 167 and continuum contributions to the absorption, as is done by many existing CKD models.
 168 The fourth term has what we refer to as a *relative-linear* dependence of absorption on mole
 169 fraction: a ‘reference’ mole fraction, ψ_j^{ref} , is defined for the gas, typically the mean surface
 170 present-day concentration. The absorption by present-day concentrations of the gas is then
 171 folded into the background term, while the fourth term represents the additional absorption
 172 due to perturbations (which may be negative) of concentrations from ψ_j^{ref} .

173 It is up to the user which of the four representations to use for each gas, and the choice
 174 depends particularly on what range of greenhouse-gas concentrations will need to be simulated
 175 by the target CKD model. The example models generated in this paper are intended to simulate
 176 the climate scenarios proposed by Hogan and Matricardi (2020). In both the shortwave and
 177 longwave, we represent O₂ and N₂ absorption by the background term, CO₂ and O₃ absorption
 178 as linear terms, H₂O as a nonlinear term, and CH₄ and N₂O as relative-linear terms. In the

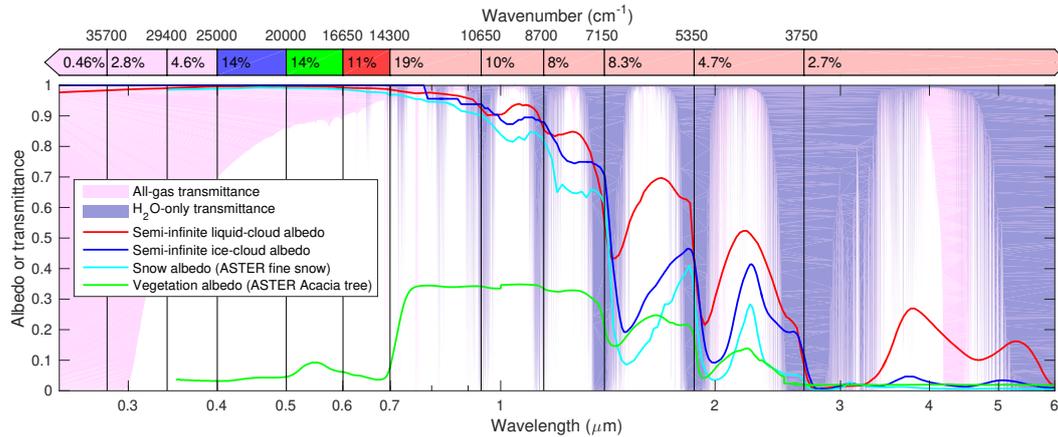


Figure 2. Spectral features of surface and atmosphere informing shortwave band selection. The shading indicates the zenith transmittance of the CKDMIP median atmosphere due to both all gases (including Rayleigh scattering) and water vapor only. The red and dark blue lines depict the albedo of semi-infinite 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 dataset. The bar at the top shows the fraction of incoming solar energy in a number of candidate bands.

179 longwave, CFC-11 and CFC-12 are represented by linear terms, while in the shortwave they
 180 are neglected (see Table 4 of Hogan and Matricardi, 2020).

181 2.4 Band selection

182 The selection of bands is entirely specified by the user, while the partitioning of each
 183 band into k -terms is automated. Many CKD models select band boundaries in order to mini-
 184 mize the number of absorbing gas species in a band, to ensure that the assumption of random
 185 overlap of gas absorption is valid, or to cap the error due to assuming the Planck function to
 186 be constant across each band. Since ecCKD is not subject to these limitations (any number of
 187 gases can be handled with arbitrary overlap, and the Planck function is computed exactly for
 188 each k -term), the choice of bands is driven by (1) the need to represent spectral variations of
 189 the properties of clouds, aerosols and the surface, and (2) the needs of downstream users for
 190 irradiances in specific bands.

191 In the longwave, the radiative effect of clouds is dominated by a fairly narrow range
 192 of wavelengths in the infrared atmospheric window. We therefore hypothesize that adequate
 193 accuracy can be achieved without the use of bands, i.e. treating the entire longwave spectrum
 194 as a single band. This is verified *a posteriori* in section 4, provided that the optical properties
 195 of clouds are computed separately per k -term.

196 In the shortwave there are multiple concerns that suggest the need for band boundaries,
 197 as illustrated in Fig. 2. From the point of view of gases, the important line absorption occurs
 198 almost entirely at wavelengths longer than 0.625 μm (wavenumbers less than 16 000 cm^{-1}). At
 199 shorter wavelengths, continuum absorption and Rayleigh scattering dominate, and since these
 200 tend to vary monotonically with wavelength, the k -terms selected automatically by ecCKD in
 201 this range tend to span contiguous ranges of wavelength and therefore behave exactly as bands.
 202 From the perspective of Numerical Weather Prediction (NWP) and reanalysis applications,
 203 there are several specific bands that are useful for downstream products, and it is preferable
 204 to specify them manually rather than leave ecCKD to place k -terms at arbitrary locations in

205 wavelength. Photosynthetically Active Radiation (PAR) is a common product consisting of the
 206 surface downwelling irradiance in the range 0.4–0.7 μm . For generating color imagery (similar
 207 to that produced by Lopez, 2020), this can be further split at 0.5 and 0.6 μm . The ecCKD tool
 208 works in wavenumber rather than wavelength, with bands specified to the nearest 50 cm^{-1} ,
 209 so in the example shortwave CKD models generated in this paper we use wavenumbers of
 210 14 300, 16 650, 20 000 and 25 000 cm^{-1} to bound the ‘red’, ‘green’ and ‘blue’ (RGB) bands.
 211 The solar energy in each of these is shown in Fig. 2. For more specialist applications additional
 212 bands could be considered, such as finer visible bands for modeling of marine biology (e.g.
 213 Ciavatta et al., 2014), or finer ultraviolet bands for computing photolysis rates or UV index.

214 In the NIR part of the shortwave spectrum, the optimum CKD model produced by ec-
 215 CKD for clear skies over surfaces with spectrally constant albedo would treat the entire region
 216 (wavelengths longer than around 0.7 μm) in a single band, and indeed this is the FSCK ap-
 217 proach taken by Pawlak et al. (2004). However, as shown in Fig. 2, there are very large
 218 variations in both cloud and surface albedo in this region, with sharp changes tending to occur
 219 in the regions of largest gas absorption. This suggests that to capture cloud and surface heating
 220 rates, band boundaries should be located between some of the NIR windows, particularly at
 221 1.4 μm (around 7150 cm^{-1}). To reconcile these two conflicting requirements, ecCKD offers
 222 the capability to use ‘sub-bands’, which we demonstrate in section 2.7: only parts of the NIR
 223 spectrum that are optically thin enough for clouds and the surface to be important are parti-
 224 tioned into sub-bands, while the wavenumbers corresponding to optically thick parts of the
 225 spectrum are treated as a single band.

226 Thus, the CKD models generated in this paper to demonstrate the capabilities of ecCKD
 227 use the FSCK approach in the longwave. In the shortwave we use the ‘RGB’ band structure: a
 228 single ultraviolet band, red, green and blue bands, and a single NIR band. The latter employs
 229 sub-bands to allow for spectrally varying cloud and surface properties.

230 2.5 Reordering the spectrum

231 The longwave and shortwave parts of the spectrum are considered separately. The first
 232 task shown in Fig. 1 is to reorder the spectrum of each gas in order of increasing absorption
 233 within each band. We seek a unique mapping independent of height, and therefore aim to sort
 234 the high-resolution CKDMIP spectra in terms of the approximate height of the peak cooling in
 235 the longwave and height of the peak heating in the shortwave, using the median CKDMIP at-
 236 mosphere with present-day (2020) greenhouse gas concentrations. In the longwave we follow
 237 a method very similar to that proposed by Hogan (2010): for each gas a LBL radiative transfer
 238 calculation is performed with all other gas concentrations set to zero and an idealized profile
 239 of temperature increasing linearly with the logarithm of pressure from -100°C at 0.01 hPa to
 240 $+15^\circ\text{C}$ at 1000 hPa. Sorting is in order of the height of peak cooling rate. This method fails
 241 for low column optical depths, τ , where heating rate peaks at the surface, so when $\tau < 0.5$ we
 242 sort by τ instead. In the shortwave a simpler approach is taken: the spectra are ordered by the
 243 height at which the optical depth from top-of-atmosphere (TOA) reaches 0.25, which is the
 244 height at which direct radiation for a solar zenith angle of 60° will have fallen to around 60%
 245 of the TOA value.

246 The result of the reordering is written into a spectral-order file for each gas, containing
 247 the integer rank r of each wavenumber point. Suppose the entire spectrum contains N discrete
 248 wavenumbers indexed 1 to N , and a particular band corresponds to wavenumbers indexed m
 249 to n . The ranks r_m to r_n will consist of the integers m to n but reordered. In the following
 250 sections we follow previous authors and introduce a coordinate variable for the reordered
 251 spectrum, g . In ecCKD, this simply maps the integer ranks for the bands to the range 0–1,
 252 i.e. an element of the reordered spectrum with integer rank r in a particular band would have
 253 $g(r) = (r - r_m)/(r_n - r_m)$.

254

2.6 Partitioning g space for individual gases

255

256

257

258

259

260

261

The next step shown in Fig. 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).

262

263

264

265

266

267

268

269

270

271

272

273

274

275

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. The appendix 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:

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

276

277

278

279

280

281

where H_j^{CKD} and H_j^{LBL} are the heating rates predicted by the ecCKD and line-by-line 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. In order to weight the stratosphere and troposphere on an approximately equal basis, the vertical profile is weighted by the square-root of pressure, i.e. 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.

282

283

284

285

286

287

288

289

290

291

292

293

294

295

296

297

Since this function is called multiple times by the algorithm described in the appendix, 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 use the median present-day profile from the CKDMIP ‘MMM’ dataset. 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 dataset, which contains the minimum values in any of the 25 000 profiles of the original Eresmaa and McNally (2014) dataset. The minimum concentrations of the well-mixed greenhouse gases 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 well-mixed greenhouse gases. For simulation of past and future climate, we use the ‘Glacial Maximum’ values proposed by Hogan and Matricardi (2020), i.e. the minimum greenhouse gas concentrations found in the last million years.

298

299

300

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 (2) are computed simply by summing

301 the LBL spectral irradiances from the wavenumbers corresponding to the requested range of
 302 g . A single zenith angle is used in each hemisphere (equivalent to the two-stream method
 303 but without scattering). The ‘CKD’ terms in (2) are computed by first averaging the optical
 304 depths of the target gas across the wavenumbers corresponding to the requested range of g ,
 305 but retaining the full spectral resolution for the other gases. This way (2) quantifies the error
 306 purely associated with approximating the target gas. The user can select the method used to
 307 average the optical depths of the target gas, the default being a linear average of the layer
 308 transmittances weighted by the Planck function at the temperature of the layer.

309 In the shortwave, the radiative transfer calculations are limited to the direct (unscattered)
 310 solar beam, which contains almost all of the sensitivity to gas absorption, and reduces com-
 311 putational cost. Thus, the upwelling terms in (2) are omitted and the heating rates consider
 312 only heating by the direct beam. Optical depths are averaged weighting by the TOA incom-
 313 ing solar spectral irradiance, which following Hogan and Matricardi (2020) is taken to be the
 314 1986–2018 average of the Coddington et al. (2016) climate data record.

315 Shortwave partitioning includes the option to use sub-bands in the NIR, as introduced in
 316 section 2.4. Since water vapor dominates in the NIR, this can be achieved by dividing water
 317 vapor alone into sub-bands. Each of the high-resolution spectral points in the NIR band has
 318 both a wavenumber and a g value indicating the water vapor absorption strength in the NIR
 319 region. The spectral points with $g < g_{\text{crit}}$, where g_{crit} is some user-specified critical value,
 320 are deemed to be optically thin enough that sub-bands are needed, so these points are further
 321 grouped according to their wavenumber into user-specified sub-bands, although within the
 322 sub-bands the ordering by g is preserved. For $g \geq g_{\text{crit}}$, water vapor is optically thick enough
 323 that no grouping by wavenumber is needed. This is illustrated graphically at the end of the
 324 next section.

325 2.7 Partitioning g space for multiple gases

326 After each of the gases have had their reordered spectra partitioned into intervals in g
 327 space, they are combined to obtain a final set of k -terms. This is achieved using the ‘hypercube
 328 partition method’ of Hogan (2010): for m active gases in a particular band we consider an
 329 m -dimensional unit hypercube where dimension j represents the g space for gas j . Each
 330 wavenumber point in the band lies at a particular location in this m -dimensional space, and
 331 an individual k -term represents a subregion of the full hypercube such that the wavenumber
 332 points that lie within the subregion for a particular k -term are then treated together in a single
 333 quasi-monochromatic radiative transfer calculation.

334 Hogan (2010) described an optimal way to partition the hypercube into k -terms, and
 335 illustrated this in both the $m = 2$ case (two gases; his Fig. 4) and $m = 3$ case (his Fig. 5).
 336 The first k -term consists of the intersection of the first g interval for each gas, i.e. the weakest
 337 absorption. The remaining k -terms each correspond to one of the remaining g intervals for
 338 one of the gases, thereby ‘specializing’ in that gas. Thus, if the number of intervals required
 339 to partition the g space for gas j is n_j , then the number of k -terms required for m gases is
 340 $n_{\text{total}} = 1 + \sum_j^m (n_j - 1)$. The method supports arbitrary spectral overlap of the spectra of
 341 individual gases, and is an improvement on the approach of Ritter and Geleyn (1992) and
 342 Edwards (1996), which requires $n_{\text{total}} = 1 + \sum_j^m n_j$ terms and makes the assumption that the
 343 spectra of individual gases are randomly overlapped.

344 Figure 3 illustrates the resulting partitioning of the spectrum for the two longwave and
 345 two shortwave example CKD models. Vectorization and computational efficiency favors n_{total}
 346 being a power of two, so we have chosen the error tolerances to obtain 16 or 32 k -terms.
 347 Consider first Fig. 3c, showing a model in which the entire longwave spectrum has been treated
 348 as a single band and partitioned into a total of 32 k -terms, in approximate order of increasing
 349 mean optical depth. All k -terms except the first are dominated by the contribution from one
 350 particular gas, indicated on the right of the panel. The ‘full-spectrum’ nature of this CKD
 351 model is clear from widely separated parts of the spectrum being represented by single quasi-

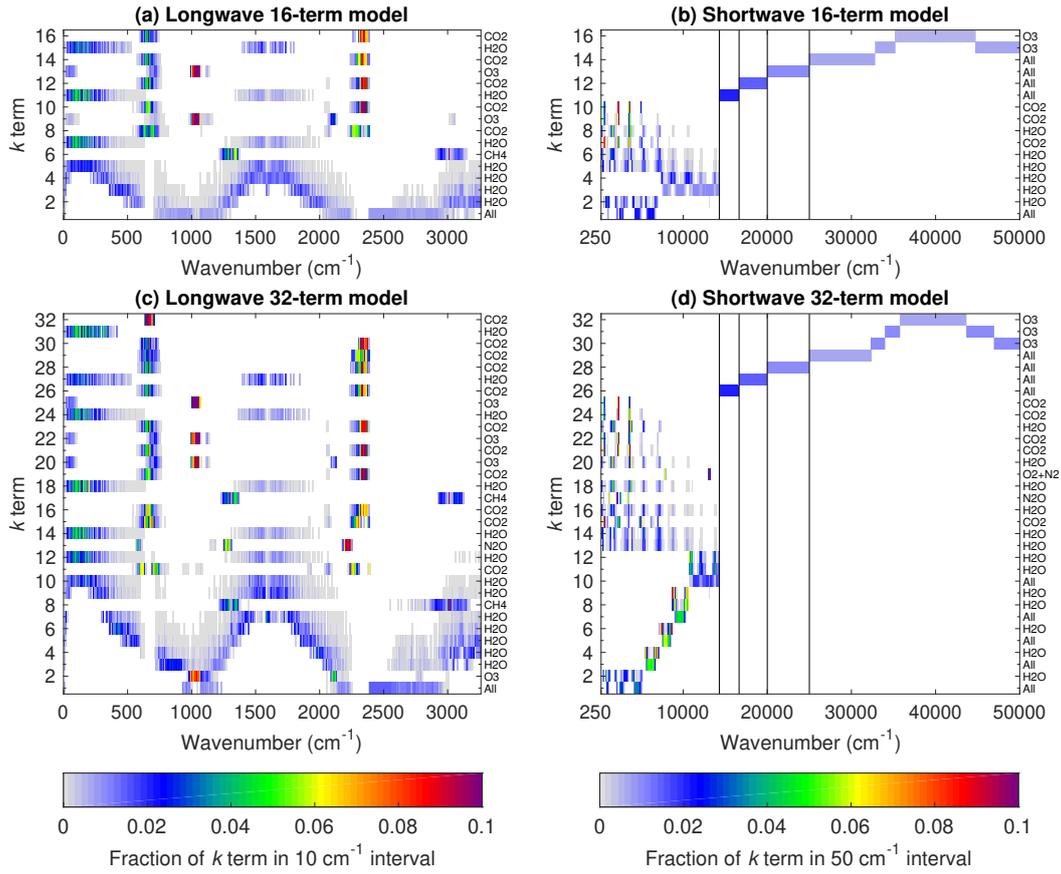


Figure 3. The contribution of each part of the spectrum to each k -term in the four CKD models considered in the text. The fractions sum to unity along each row. The main gas represented by each k term is indicated on the right-hand-side of each panel, although in principle all gases can contribute to 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.

352 monochromatic k -terms; for example, term 28 represents strong CO_2 absorption from both the
 353 $4.3 \mu\text{m}$ and the $15 \mu\text{m}$ CO_2 bands.

354 The 32-term shortwave model in Fig. 3d is very different. The vertical black lines de-
 355 delineate the user-specified ‘RGB’ bands explained in section 2.4. The partitioning algorithm
 356 finds that only one k -term is required for each of the red, green and blue bands, and only four
 357 for the entire ultraviolet. Terms 30 and 31 each combine the effects of similar levels of O_3
 358 absorption on opposite sides of the Hartley band, which peaks at around $40\,000 \text{ cm}^{-1}$. In the
 359 NIR, we use sub-bands for water vapor as explained in section 2.6, with a value of $g_{\text{crit}} = 0.7$
 360 and additional wavenumber partitions at 5350 , 7150 , 8700 and 10650 cm^{-1} (bounding most
 361 of the NIR windows in Fig. 2). In Fig. 3d we see that k -terms 1–12 represent the intersection
 362 of regions of relatively weak water-vapor absorption ($g < g_{\text{crit}}$) with weak absorption in all
 363 the other gases; these use five sub-bands in order to resolve spectral variations in clouds and
 364 the surface. The k -terms 13–25 represent regions of strong water vapor absorption ($g \geq g_{\text{crit}}$)
 365 or strong absorption by one of the other gases. The 16-term shortwave model in Fig. 3b takes
 366 the same approach but with reduced k -terms via the use of $g_{\text{crit}} = 0.65$ and only one additional
 367 wavenumber partition at 7150 cm^{-1} .

368

2.8 Creating initial look-up table

369

370

371

372

373

374

375

376

377

378

379

380

This step creates a first estimate of the LUTs in (1) using the ‘Idealized’ CKDMIP dataset, and indeed we use the same points (described in section 3.3 of Hogan and Matricardi, 2020): 53 logarithmically spaced points in pressure with 10 points per decade; 6 points in temperature, 20 K apart; and 12 logarithmically spaced points in water vapor mixing ratio with 2 points per decade. As shown in Fig. 1, the creation of the LUT involves reading in the location of the k -terms, i.e. 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 dataset 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 final step in the creation of the LUT is to convert from layer optical depth to molar absorption coefficient as used in (1).

381

382

383

384

385

386

387

388

389

390

391

The LUT file contains additional variables that are added at this point and remain unchanged by the subsequent steps shown in Fig. 1. The fraction of the spectrum contributing to each k -term (i.e. the information shown in Fig. 3) 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.

392

2.9 Scaling shortwave look-up table entries

393

394

395

396

397

398

399

400

401

402

403

404

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 in order 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.

405

406

407

408

409

410

411

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 .

412

413

414

415

416

417

418

In the ‘scale LUT’ step in Fig. 1 we take this approach using the median present-day profile from the CKDMIP ‘MMM’ dataset 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

419 be modified by the same proportion, and that the correction factor does not vary with the other
 420 dimensions of the LUT (temperature and water-vapor mixing ratio). Nonetheless, shortwave
 421 radiative transfer calculations using the v1 LUT are significantly more accurate than v0, and
 422 there is still the opportunity (described in the following section) for a global optimization of
 423 all the coefficients in the LUT.

424 2.10 Optimizing look-up table entries

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

435 We improve the basic method of Hogan (2010) in a number of ways. As shown in Fig.
 436 1, the optimization may be performed in several steps, as we have found that better perfor-
 437 mance is achieved if the major gases are optimized first, with minor gases being optimized
 438 individually in subsequent steps. We define the state vector \mathbf{x} to contain all the variables to be
 439 optimized in one of these steps, specifically the natural logarithm of all the non-zero entries in
 440 the look-up for each gas being optimized. With 53 points in pressure, 6 in temperature, 12 in
 441 water vapor concentration and (for example) 32 k -terms, this leads to 122 112 state variables
 442 for water vapor and 10 176 for each other gas. The cost function to be minimized is given by

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

443 The first term ensures the stability of the minimization by penalizing the squared differences
 444 between the state vector and the *a priori* LUT elements \mathbf{x}_a , i.e. those from the previous step
 445 in Fig. 1. The error covariance matrix \mathbf{B} provides a complete description of the weighting of
 446 this term, with its diagonal elements containing the square of the user-specified root-mean-
 447 squared (RMS) error \mathbf{x}_a . We use a value of 8 in this paper, allowing the natural logarithm of
 448 the LUT elements to stray significantly from their prior values in the optimization, although
 449 in practice the RMS difference between the elements of \mathbf{x} before and after optimization for
 450 an individual gas is around 0.25. The off-diagonal elements of \mathbf{B} specify error covariances
 451 between LUT values, and have the important effect of spreading information provided by the
 452 training profiles into adjacent parts of the LUT. We model the error correlation coefficient of
 453 adjacent LUT coefficients along the pressure, temperature and water-vapor-concentration axes
 454 as ρ , and coefficients n steps apart along these axes as ρ^n . No correlation is assumed between
 455 k -terms or gases. Even though \mathbf{B} is large, its inverse is very sparse and the first term in (3) is
 456 efficient to compute. Empirically we find that $\rho = 0.8$ provides the best results. The second
 457 term in (3) expresses the sum of the squared errors in heating rates and irradiances at TOA
 458 and the surface over p atmospheric profiles, where E_j has the same form as in (2) but using
 459 broadband heating rates and irradiances.

460 The cost function is minimized using the quasi-Newton ‘L-BFGS’ algorithm of Liu
 461 and Nocedal (1989), which requires the vector of gradients $\partial J/\mathbf{x}$ to be computed. This is
 462 achieved by coding the entire algorithm in C++ using the combined automatic differentiation,
 463 array and optimization library ‘Adept’ (Hogan, 2014), version 2.1 of which also includes an
 464 implementation of the L-BFGS algorithm. An optimization step typically takes several tens
 465 of minutes to complete.

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

475 In the case of CKD models targeting climate applications (such as those depicted in
 476 Fig. 3), the first step optimizes the coefficients of H₂O, O₃, CO₂ and the background term in
 477 (1). The background term represents not only O₂ and N₂, but also the present-day ‘reference’
 478 concentrations of CH₄ and N₂O, with these gases all treated at this stage as having a constant
 479 mixing ratio with pressure. The training data consist of LBL calculations performed on the
 480 50 ‘Evaluation-1’ CKDMIP profiles for six CKDMIP CO₂ scenarios (surface concentrations
 481 from 180 to 2240 ppmv), i.e. a total of 300 profiles.

482 In the second step the coefficients of CH₄ are optimized using LBL calculations on the
 483 CKDMIP scenarios in which CH₄ is perturbed from 350 to 3500 ppbv. To avoid the CH₄
 484 coefficients being tuned to correct for remaining errors in the previous step, we train on the
 485 difference in heating-rate profiles and irradiances between perturbed and present-day CH₄
 486 calculations, which is equivalent to bias-correcting the present-day CKD calculations from
 487 the previous step. This ensures the CH₄ coefficients are optimized to give the most accurate
 488 radiative forcing when perturbed from present-day concentrations. The third step takes exactly
 489 the same approach but optimizes the N₂O coefficients training on the Evaluation-1 profiles in
 490 which concentrations are perturbed in the range 190–540 ppbv. In the shortwave this yields
 491 the ‘final’ LUT ready for use in a radiation scheme. In the longwave we perform one further
 492 step to optimize the coefficients of CFC-11 and CFC-12.

493 3 Clear-sky evaluation

494 In this section we evaluate the gas-optics models generated in the previous section in
 495 clear skies. The CKDMIP Evaluation-2 dataset is used, which consists of LBL calculations
 496 on 50 independent profiles, including ones with extremes of temperature, ozone and humidity.
 497 Surface longwave emissivity and shortwave albedo are spectrally constant at 1.0 and 0.15, re-
 498 spectively, the latter being the approximate mean albedo of the Earth’s surface. This approach
 499 is the same as that of Hogan and Matricardi (2020) to evaluate the RRTMG gas-optics model,
 500 except that they used the CKDMIP Evaluation-1 dataset. Naturally, when the ecCKD models
 501 are evaluated against the Evaluation-1 dataset (not shown) they perform better because these
 502 profiles were used in the training.

503 Figure 4 evaluates the performance of the FSCK-16 and FSCK-32 models, for present-
 504 day greenhouse gas concentrations. Even though relatively few k -terms are used compared to
 505 other CKD models, the errors are small; the root-mean-squared (RMS) error in heating rates
 506 from the surface to the upper stratosphere (4 hPa) is only 0.15 and 0.1 K d⁻¹ for the 16- and
 507 32-term models, respectively, approximately doubling in the mesosphere.

508 Figure 5 evaluates the instantaneous radiative forcing associated with perturbing the
 509 five main greenhouse-gas concentrations from their present-day values. Note that the CFC-
 510 11 concentrations here correspond to artificially increased values to approximately represent
 511 38 further greenhouse gases (Meinshausen et al., 2017). In large part, both models capture
 512 the forcing associated with large perturbations to concentrations, including up to eight times
 513 preindustrial concentrations of CO₂. The 16-term model performs slightly worse in some sce-
 514 narios, tending to underestimate the magnitude of the surface forcing associated with reducing
 515 CO₂ concentrations to glacial-maximum values of 180 ppmv, as well as struggling with the

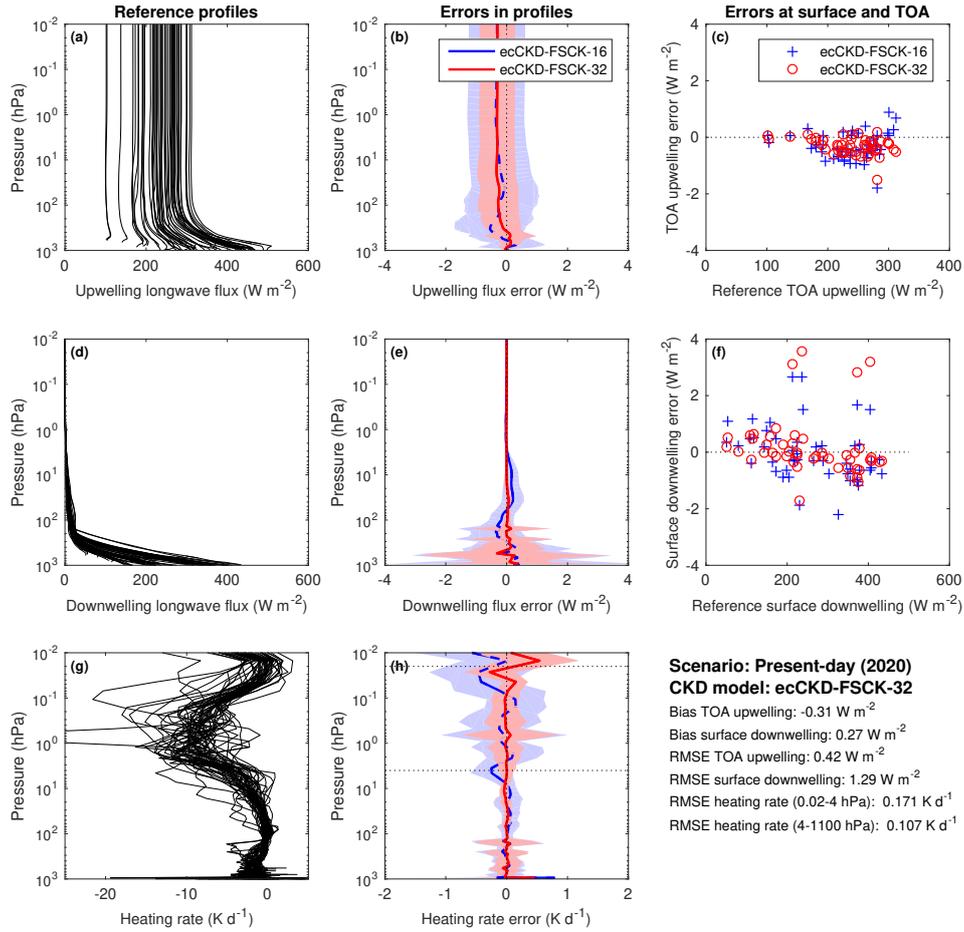


Figure 4. Evaluation of clear-sky longwave irradiances and heating rates from the two ecCKD models for the 50 independent profiles of the CKDMIP Evaluation-2 dataset with present-day concentrations of the well-mixed greenhouse gases. Panels a, d and g show quantities from the reference LBL calculations, while panels b, e and 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 TOA and downwelling surface irradiances. The statistics of the comparison are summarized in the lower right, including the root-mean-squared error (RMSE) in heating rate (weighted by the cube-root of pressure) in two ranges of pressure indicated by the horizontal dotted lines in panel h.

516 extreme CH₄ concentrations. As can be seen in Figs. 3a and 3c, the improvement of the 32-
 517 term model for CH₄ can be attributed to its use of two CH₄-specific *k* terms, rather than just
 518 one.

519 Figure 6 presents the corresponding present-day evaluation of irradiances and heating
 520 rates for the RGB-16 and RGB-32 shortwave models. Again, the errors are modest given
 521 the small number of *k*-terms, with the RMS error in heating rates from the surface to 4 hPa
 522 being 0.1 and 0.06 K d⁻¹ for the 16- and 32-term models, respectively. The much larger
 523 mid-mesosphere heating-rate error for the 16-term model is associated with its poorer repre-
 524 sentation of the 4.3 μm CO₂ band; Figs. 3b and 3 show that it used only three CO₂-specific
 525 *k*-terms, compared to five for the 32-band model. This also explains the difference in how well
 526 the two models capture the shortwave CO₂ forcing shown in Fig. 7. This figure also indicates

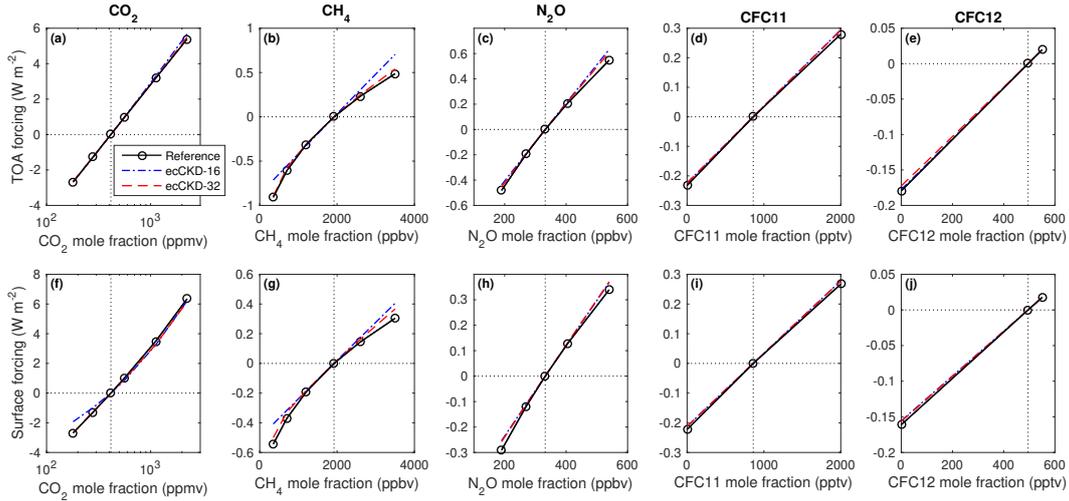


Figure 5. Comparison of reference LBL and ecCKD calculations of the instantaneous longwave clear-sky radiative forcing from perturbing each of the five well-mixed greenhouse gases from their present-day (2020) values at (a–e) top of atmosphere and (f–j) the surface, averaged over the 50 profiles of the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

Table 1. Summary of the properties of the two atmospheric profiles used to evaluate the representation of liquid and ice clouds in section 4. Both are taken from the CKDMIP dataset. The size distribution of the liquid cloud was modeled as a gamma distribution with a shape parameter of 2.

Profile	Liquid cloud	Ice cloud
Cloud effective radius	10 μm	30 μm
Cloud pressure range	726.6–907.1 hPa	184.5–457.3 hPa
Optical properties	Mie theory	Baum et al. (2014)
Origin	Evaluation-2 profile 29	Evaluation-1 profile 28
Location	31.4°S, 3.5°W	38.9°N, 25.6°W
Date and time	19 March 2014, 18 UTC	11 March 2014, 00 UTC
Surface pressure	1017.4 hPa	1021.6 hPa
Surface temperature	22.2°C	14.7°C

527 that the CH_4 forcing in the two models is similar; in fact neither model uses CH_4 -specific
 528 k -terms, but rather includes the optical-depth contribution of CH_4 in all the other k -terms. The
 529 16-term model also has no N_2O -specific k -terms and Fig. 7 shows that this leads to it tending
 530 to overestimate the N_2O forcing by around a factor of two (although the magnitude of the
 531 shortwave forcing of this gas is only a tenth of the longwave). The 32-term model introduces
 532 a single N_2O -specific k -term and is able to achieve a much greater accuracy.

533 4 Cloudy-sky evaluation

534 As discussed in section 2.4, clouds, aerosols and the surface can exhibit significant spectral
 535 variations in optical properties. One of the features of ecCKD is that each k -term has a
 536 unique mapping to specific parts of the spectrum (see Fig. 3), and this mapping is available to
 537 downstream applications (such as the ecRad radiation scheme) so that optical properties can
 538 be averaged separately for each k term. Nonetheless, the use of the full-spectrum correlated- k

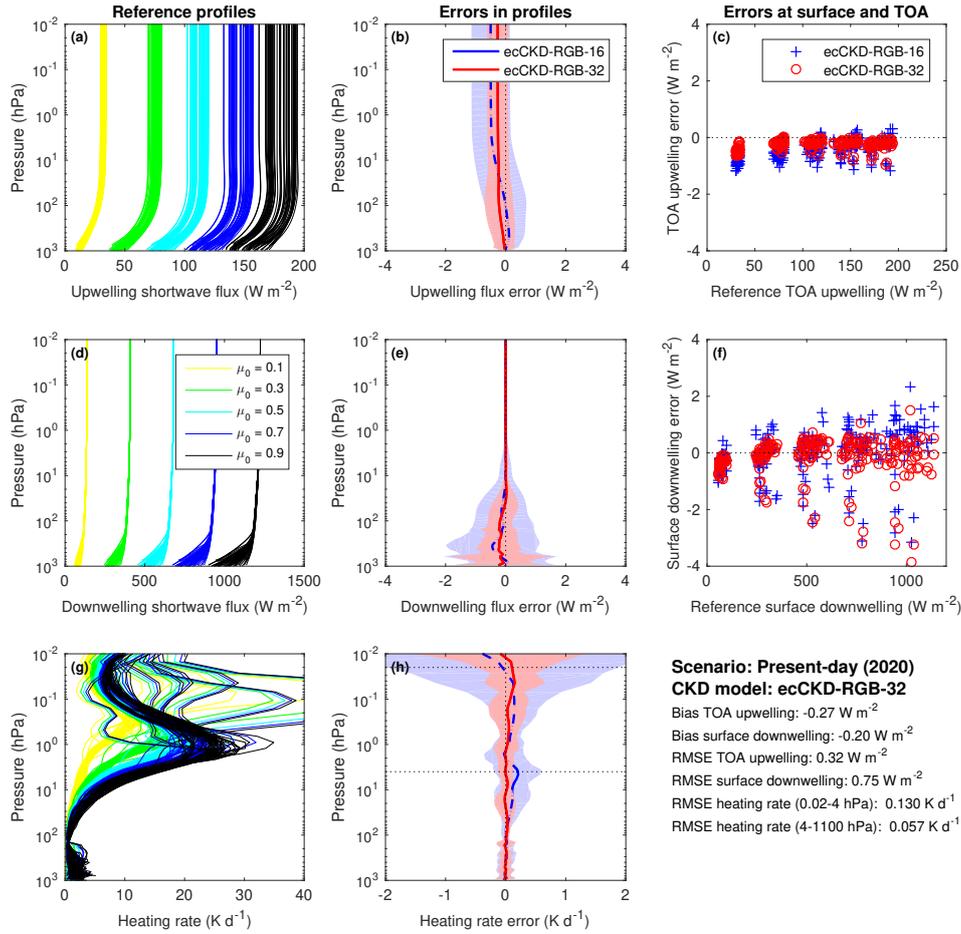


Figure 6. Similar to Fig. 4 but for the shortwave. The reference LBL calculations in panels a, d and g are for all 50 CKDMIP Evaluation-2 profiles at five values of the cosine of the solar zenith angle, μ_0 (0.1, 0.3, 0.5, 0.7 and 0.9). The subsequent evaluation considers all 250 combinations. The five clusters of points in panels c and f correspond to the five values of μ_0 .

539 approach means that individual k -terms can represent widely separated points in the spectrum.
 540 In this section we test the impact on the accuracy of calculations of the radiative effect of liquid
 541 and ice clouds, using two real-world profiles from the CKDMIP dataset, summarized in Table
 542 1. In each case, 26 LBL calculations have been performed for water paths ranging from 10^{-4}
 543 to 10 kg m^{-2} , plus an additional clear-sky calculation, with a vertically constant cloud mixing
 544 ratio between two pressure bounds. The optical properties of liquid clouds are computed
 545 using Mie theory at 396 wavenumbers from 5 to $50\,000 \text{ cm}^{-1}$, while the ice properties are
 546 taken from the Baum et al. (2014) ‘General Habit Mixture’ available at 445 wavenumbers be-
 547 tween 101 and $50\,251 \text{ cm}^{-1}$. When used in LBL calculations, the mass-extinction coefficient,
 548 single scattering albedo and asymmetry factor are interpolated linearly in wavenumber space,
 549 but clamped when used at wavenumbers outside the range provided. The radiative transfer
 550 calculations use a no-scattering solver in the longwave, and the two-stream method with a
 551 solar zenith angle of 60° in the shortwave, both from the CKDMIP software package. The
 552 plane-parallel approximation is adopted, i.e. clouds are taken to be horizontally uniform with
 553 a cloud fraction of unity.

554 The equivalent ecCKD calculations use a version of the ecRad offline radiative transfer
 555 package that supports ecCKD gas-optics models, and a radiative transfer solver equivalent to

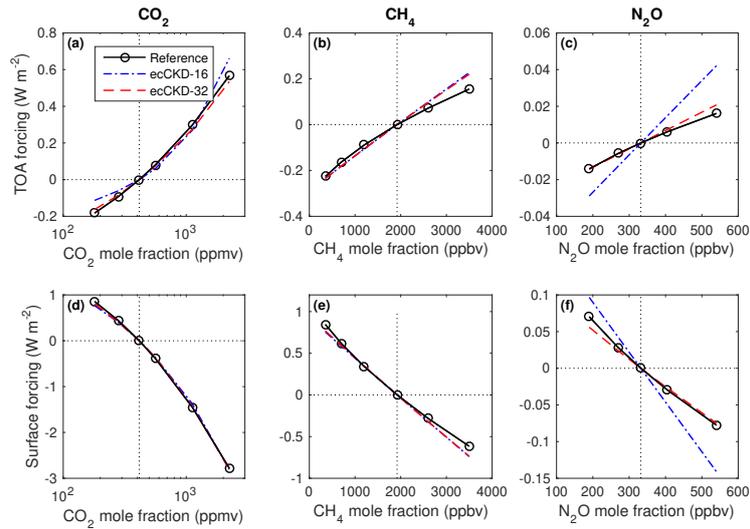


Figure 7. As Fig. 5 but for the instantaneous shortwave radiative forcing by CO_2 , CH_4 and N_2O . The results for the five solar zenith angles have been averaged, so the values shown here represent a daytime average.

556 that used for the LBL calculations. Cloud optical properties are computed by averaging the
 557 same Mie and Baum et al. (2014) data according to the parts of the spectrum corresponding to
 558 each individual k -term (as shown in Fig. 3). Following the approach of Edwards and Slingo
 559 (1996), we consider both ‘thin’ and ‘thick’ spectral averaging. The former is appropriate in the
 560 optically thin limit and simply involves averaging of the mass-extinction coefficient, the mass-
 561 absorption coefficient, and averaging asymmetry factor weighted by scattering coefficient.
 562 The latter is more appropriate in the optically thick limit and is intended to provide the exact
 563 cloud albedo in the limit of infinite optical depth (although in practice it is not exact in this
 564 limit except in the absence of gas absorption). For a little extra accuracy, we apply delta-
 565 Eddington scaling (Joseph et al., 1976) before performing the spectral average. A further
 566 weighting is used in the averaging to approximately represent the energy at each wavenumber;
 567 in the longwave we use the Planck function at a representative atmospheric temperature of
 568 0°C and in the shortwave at an effective solar temperature of 5777 K. In the longwave, no
 569 benefit was found from using a different reference temperature for liquid and ice clouds. For
 570 comparison we also show the results of the RRTMG gas-optics model using the same cloud
 571 optical properties, but since no information is available on the exact wavenumbers used for
 572 each of its k -terms, the optical properties are averaged to its 16 longwave and 14 shortwave
 573 bands.

574 Figure 8a depicts the LBL calculations of ‘true’ longwave cloud radiative effect at TOA
 575 and the surface for the liquid-cloud profile versus water path, with the error in these quantities
 576 shown in Figs. 8b and 8c. The errors for all models are generally less than 1 W m^{-2} . The
 577 RRTMG model performs best, although the errors associated with the two ecCKD models are
 578 still small, being up to around 5% for the FSCK-16 model and 2% for the FSCK-32 model.
 579 Naturally, the RRTMG model with its 16 longwave bands is able to capture the spectral vari-
 580 ation of cloud optical properties, but it is nonetheless surprising how well the FSCK models
 581 perform when they consider the entire longwave spectrum in a single band. This is because (as
 582 revealed by the LBL calculations shown in Fig. S1 of the Supporting Information) over 92%
 583 of the radiative effect of this cloud at the surface and TOA is in the $8\text{--}13 \mu\text{m}$ ($769\text{--}1250 \text{ cm}^{-1}$)
 584 longwave atmospheric window, within which the variation of cloud optical properties is quite
 585 modest, and certainly much less than in the NIR. Outside the longwave atmospheric window,
 586 clouds make a much weaker contribution to longwave cloud radiative effect either because of

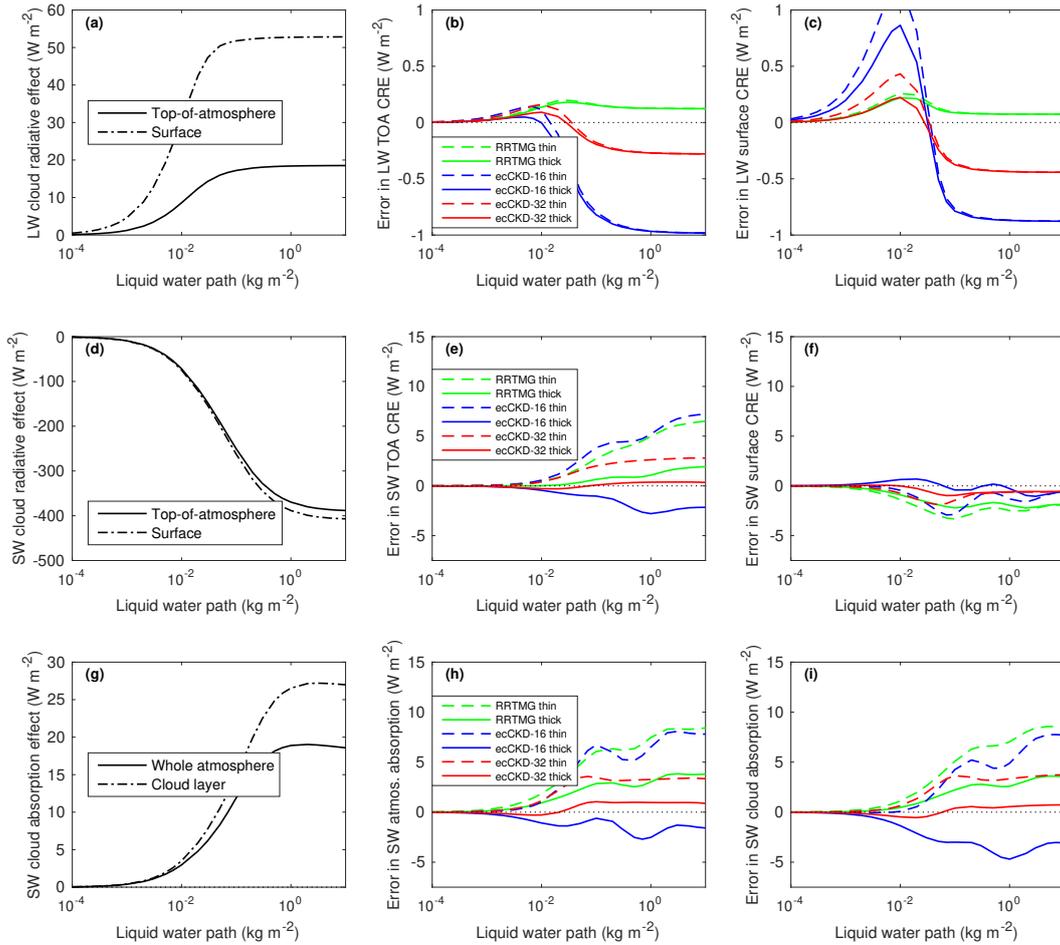


Figure 8. Error in cloud radiative effect (CRE) due to spectral discretization of RRTMG and the various ecCKD models, for a low-level liquid cloud with effective radius $10 \mu\text{m}$ in Profile 29 of the CKDMIP Evaluation-2 dataset. Panel a depicts the longwave cloud radiative effect (i.e. the change to net irradiance due to cloud) at top-of-atmosphere (TOA) and the surface as a function of LWP, for the LBL calculations. Panels b and c depict the error in these quantities for the various 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’, i.e. 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.

587 the much stronger gas absorption or the much weaker Planck function. The top row of Fig.
 588 9 shows the equivalent evaluation for the ice cloud profile, where the errors for all longwave
 589 gas-optics models are even less, both in an absolute and a relative sense, due to ice particles
 590 having less variation in their optical properties than liquid droplets across the longwave spec-
 591 trum (see Fig. S2 in the Supporting Information). Overall, these results demonstrate that the
 592 longwave FSCK method is a viable approach for use in weather and climate models.

593 The middle row of Fig. 8 depicts the equivalent evaluation but in the shortwave where
 594 the magnitude of the cloud radiative effect is much larger. The best performing model is
 595 clearly RGB-32 using ‘thick’ averaging, with errors of only around 0.4 W m^{-1} (0.1%) for any
 596 value of liquid water path. This provides *a posteriori* justification for the use of five NIR sub-
 597 bands in Fig. 3d, bounded at the points shown in Fig. 2 where cloud optical properties tend to
 598 change most rapidly. The RGB-16 model incurs a larger error due to its employing only two

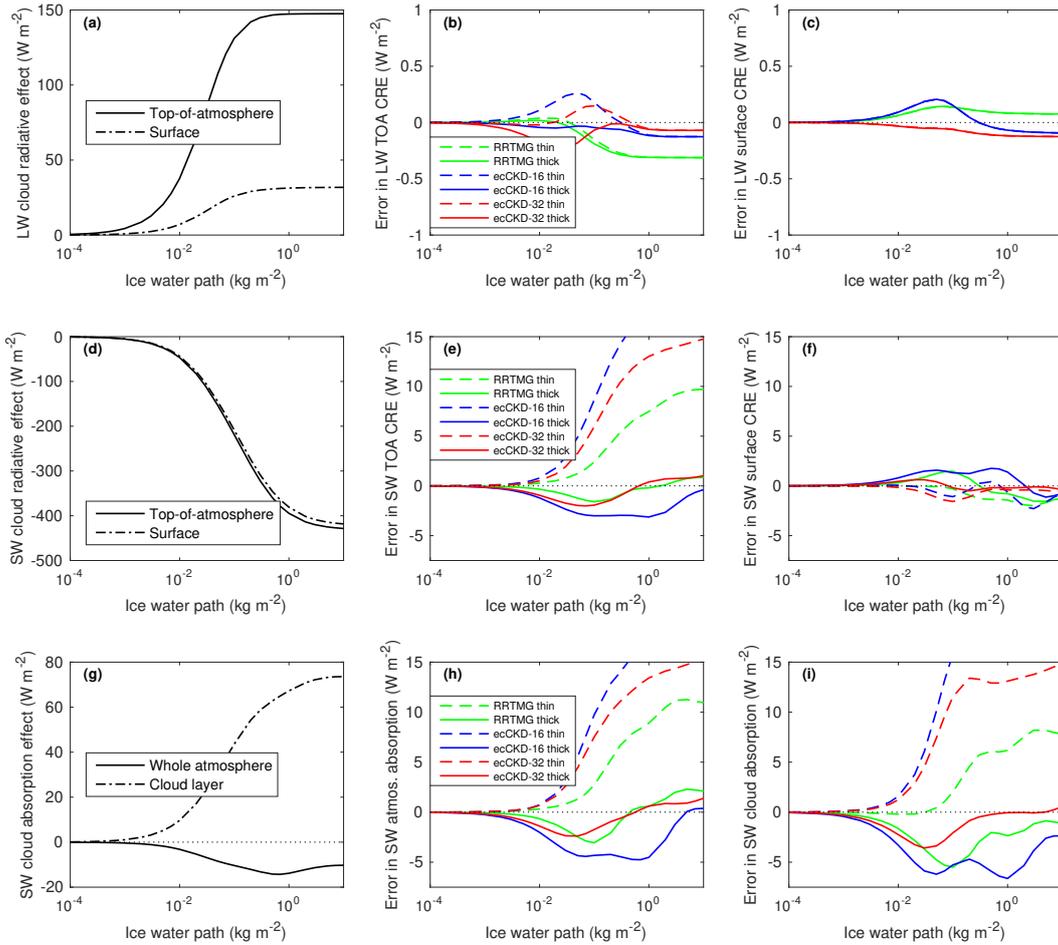


Figure 9. As Fig. 9 but for an ice cloud with an effective radius of $30 \mu\text{m}$ in Profile 28 of the CKDMIP Evaluation-1 dataset.

599 sub-bands. The result for ice clouds in the middle row of Fig. 9 are broadly similar. It is clear
 600 from Figs. 8e and 9e that the most accurate calculations are achieved using thick rather than
 601 thin spectral averaging, for any of the gas-optics models. This conclusion is different from
 602 that of [Edwards and Slingo \(1996\)](#), who advocated thick averaging for liquid clouds and thin
 603 for ice clouds.

604 The bottom rows of Figs. 8 and 9 consider the effect of the cloud on shortwave absorption,
 605 both by the entire atmosphere and by the cloud layer alone. Again, the RGB-32 model
 606 with thick averaging performs best. One of the interesting features of Fig. 9g is that the ef-
 607 fect of the ice cloud is to increase shortwave absorption in the cloud layer itself, as would be
 608 expected, but to reduce absorption overall by reflecting sunlight that would otherwise have
 609 been absorbed by gases lower in the atmosphere. Thus, the sign of the impact of the cloud on
 610 whole-atmosphere absorption is dependent on two competing effects, and while the absolute
 611 magnitude of the errors shown in whole-atmosphere and cloud-layer absorption (Figs. 9h and
 612 9i) are similar, the relative error in the latter is much larger; indeed, the RGB-16 model with
 613 thin averaging predicts that the effect of the ice cloud on atmospheric absorption is to increase
 614 rather than to decrease it.

5 Conclusions

In this paper, we have introduced a free software tool ‘ecCKD’ for generating fast correlated- k -distribution (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 look-up table coefficients for each gas in order to minimize errors against hundreds of training profiles (extending the approach of [Hogan, 2010](#)). In the shortwave, the introduction of ‘sub-bands’ enables the full-spectrum correlated- k (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, i.e. 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, 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-hour 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 0.5 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 similar accuracy to RRTMG but using only 25 rather than 78 terms in the NIR. We also find that ‘thick’ averaging of cloud spectral properties ([Edwards and Slingo, 1996](#)) provides the best agreement with LBL calculations for both liquid and ice clouds.

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 and Bozzo, 2018](#)). Additionally, the use of a simple look-up table to compute optical depths (Eq. 1) makes it straightforward to incorporate the CKD models into different types of radiation scheme, including explicit 3D solvers (e.g. [Jakub and Mayer, 2016](#)).

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:

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

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

662 where the error $E(g_{i-1}, g_i)$ is a non-differentiable user-supplied function, and the second con-
 663 dition states that the fractional range of errors, $F = [\max(E) - \min(E)]/\bar{E}$, should be no larger
 664 than the user-supplied tolerance F_{tol} , typically 0.02. In our case, $E(g_{i-1}, g_i)$ is given by (2) and
 665 involves LBL calculations with a computational cost proportional to the width of the interval
 666 $g_i - g_{i-1}$. Therefore, a good partitioning algorithm should not require an excessive number of
 667 calculations of E , especially ones for wide g intervals. We are not aware of an off-the-shelf
 668 algorithm for performing this partitioning, so this appendix describes our solution to the prob-
 669 lem. While it is not likely to be the fastest possible algorithm, it almost always converges to a
 670 solution that satisfies the conditions above.

671 The first task is to find the number of intervals required, n , in order that condition (A1)
 672 is satisfied. This is achieved by partitioning g space starting at the lower end such that for each
 673 interval (except possibly the last), $0.95E_{\text{tol}} \leq E \leq E_{\text{tol}}$. We start with a test value of $g_1 = 0.75$
 674 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
 675 result is less than $0.95E_{\text{tol}}$ then a new lower bound for g_1 has been found, and the next test value
 676 is selected by extrapolating (but not beyond $g_1 = 1$) from the old and new lower bounds to
 677 where we would expect $E(0, g_1) = E_{\text{tol}}$ assuming a linear variation of E with g_1 . On the other
 678 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
 679 found by linearly interpolating between the lower and upper bounds on g_1 . The new test value
 680 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
 681 $g_1 = 1$ and $E \leq E_{\text{tol}}$. If there is remaining g space to partition then the process is repeated
 682 to compute g_2 and so on, until the process returns $g_n = 1$. We now know how many intervals
 683 are required, and have candidate values for g_i , but usually the error associated with the final
 684 interval, $E(g_{n-1}, 1)$, is significantly less than all the other errors.

685 The second task is to find the interior boundaries of the intervals ($g_1, g_2 \cdots g_{n-1}$) such
 686 that condition (A2) is satisfied. Note that the outer boundaries are already fixed at $g_0 = 0$ and
 687 $g_n = 1$. In the simple case of $n = 2$ we have only a single value to find, g_1 ; it is straightforward
 688 to progressively refine this value until (A2) is satisfied. In the more general $n > 2$ case, we
 689 use the candidate values of g_i from the first task above to compute the cumulative error as

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

690 A new set of candidate values g'_i is found by attempting to repartition the total error, $C(1)$,
 691 evenly amongst the n intervals. This is achieved by using linear interpolation into the C func-
 692 tion to compute the g'_i values such that $C(g'_i) = iC(1)/n$. The errors are recomputed and the
 693 process is repeated until (A2) is satisfied. If at any iteration the fractional range F increases
 694 then the iteration is not successful and instead a ‘shuffle’ step is performed. This consists
 695 of looping through adjacent pairs of intervals and adjusting the g point between them until
 696 their errors agree to within 2%. Thus, for intervals 1 and 2 we adjust g_1 until $E(g_0, g_1)$ and
 697 $E(g_1, g_2)$ satisfy (A2), then do the same for intervals 2 and 3 and so on up to intervals $n - 1$
 698 and n , followed by a pass back down to intervals 1 and 2. This is usually enough that subse-
 699 quent partitioning iterations using (A3) lead to a reduction of F . Any further shuffle operations
 700 proceed in the opposite direction through g space as the previous one.

701 Acknowledgments

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

704 Data Availability Statement

705 The ecCKD software is available from <https://github.com/ecmwf-ifs/ecckd>. It makes
 706 use of data from the CKDMIP project available via the links at <http://confluence.ecmwf.int/>

707 [display/CKDMIP](#). The netCDF files defining the 16- and 32-term gas-optics models described
708 in this paper are available at the ecCKD web site <https://confluence.ecmwf.int/x/XwU0Dw>.

709 References

- 710 Barker, H. W., Qu, Z., Dhanraj, V., and Cole, J. N. S. (2021) Partial validation of a lossy
711 compression approach to computing radiative transfer in cloud system-resolving models.
712 *Q. J. R. Meteorol. Soc.*, *147*, 363–381, <https://doi.org/10.1002/qj.3922>.
- 713 Baum, B. A., Yang, P., Heymsfield, A. J., Bansemmer, A., Merrelli, A., Schmitt, C., & Wang,
714 C. (2014). Ice cloud bulk single-scattering property models with the full phase matrix at
715 wavelengths from 0.2 to 100 μm . *J. Quant. Spectrosc. Radiat. Transfer*, *146*, 123–139,
716 <https://doi.org/10.1016/j.jqsrt.2014.02.029>.
- 717 Bennartz, R., & Fischer, J. (2000). A modified k -distribution approach applied to narrow
718 band water vapour and oxygen absorption estimates in the near infrared. *J. Quant. Spec-*
719 *trosc. Radiat. Transfer*, *66*, 539–553, [https://doi.org/10.1016/S0022-4073\(99\)00184-3](https://doi.org/10.1016/S0022-4073(99)00184-3).
- 720 Bucholtz, A. (1995). Rayleigh-scattering calculations for the terrestrial atmosphere. *App.*
721 *Opt*, *34*, 2765–2773, <https://doi.org/10.1364/AO.34.002765>.
- 722 Ciavatta, S., Torres, R., Martinez-Vicente, V., Smyth, T., Dall’Olmo, G., Polimene,
723 L., Icarus Allen, J. (2014). Assimilation of remotely-sensed optical properties to
724 improve marine biochemistry modelling. *Progress in Oceanography*, *127*, 74–95,
725 <https://doi.org/10.1016/j.pocean.2014.06.002>.
- 726 Coddington, O., Lean, J. L., Pilewskie, P., Snow, M., & Lindholm, D. (2016). A solar ir-
727 radiance climate data record. *Bull. Am. Meteorol. Soc.*, *97*, 1265–1282, [https://doi.org/](https://doi.org/10.1175/BAMS-D-14-00265.1)
728 [10.1175/BAMS-D-14-00265.1](https://doi.org/10.1175/BAMS-D-14-00265.1).
- 729 DeAngelis, A., Qu, X., Zelinka, M., & Hall, A. (2015). An observational radiative con-
730 straint on hydrologic cycle intensification. *Nature*, *528*, 249–253, [https://doi.org/10.1038/](https://doi.org/10.1038/nature15770)
731 [nature15770](https://doi.org/10.1038/nature15770).
- 732 Doppler, L., Preusker, R., Bennartz, R., & Fischer, J. (2013). k -bin and k -IR: k -distribution
733 methods without correlation approximation for non-fixed instrument response function
734 and extension to the thermal infrared—Applications to satellite remote sensing. *J. Quant.*
735 *Spectrosc. Radiat. Transfer*, *133*, 382–395, <https://doi.org/10.1016/j.jqsrt.2013.09.001>.
- 736 Edwards, J. M. (1996). Efficient calculation of infrared fluxes and cooling rates us-
737 ing the two-stream equations. *J. Atmos. Sci.*, *53*, 1921–1932, [https://doi.org/10.1175/](https://doi.org/10.1175/1520-0469(1996)053<1921:ECOIFA>2.0.CO;2)
738 [1520-0469\(1996\)053<1921:ECOIFA>2.0.CO;2](https://doi.org/10.1175/1520-0469(1996)053<1921:ECOIFA>2.0.CO;2).
- 739 Edwards, J. M., & Slingo, A. (1996). Studies with a flexible new radiation code: 1. Choos-
740 ing a configuration for a large-scale model. *Q. J. R. Meteorol. Soc.*, *122*, 689–719.
741 <https://doi.org/10.1002/qj.49712253107>.
- 742 Eresmaa, R., & McNally, A. P. (2014). Diverse profile datasets from the ECMWF 137-
743 level short-range forecasts. NWP-SAF Document NWPSAF_EC_TR_017, available from
744 <https://www.nwpsaf.eu/site/software/atmospheric-profile-data/>.
- 745 Fu, Q., & Liou, K. N. (1992). On the correlated k -distribution method for radiative transfer
746 in nonhomogeneous atmospheres. *J. Atmos. Sci.*, *49*, 2139–2156, [https://doi.org/10.1175/](https://doi.org/10.1175/1520-0469(1992)049<2139:OTCDMF>2.0.CO;2)
747 [1520-0469\(1992\)049<2139:OTCDMF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1992)049<2139:OTCDMF>2.0.CO;2).
- 748 Fu, Q., Liou, K. N., Cribb, M. C., Charlock, T. P., & Grossman, A. (1997). Multiple scatter-
749 ing parameterization in thermal infrared radiative transfer. *J. Atmos. Sci.*, *54*, 2799–2812,
750 [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).
- 751 Goody, R., West, R., Chen, L., & Crisp, D. (1989). The correlated- k method for radiation
752 calculations in nonhomogeneous atmospheres. *J. Quant. Spectrosc. Radiat. Transfer*, *42*,
753 539–550, [https://doi.org/10.1016/0022-4073\(89\)90044-7](https://doi.org/10.1016/0022-4073(89)90044-7).
- 754 Hogan, R. J. (2010). The full-spectrum correlated- k method for longwave atmospheric ra-
755 diation using an effective Planck function. *J. Atmos. Sci.*, *67*, 2086–2100, [https://doi.org/](https://doi.org/10.1175/2010JAS3202.1)
756 [10.1175/2010JAS3202.1](https://doi.org/10.1175/2010JAS3202.1).
- 757 Hogan, R. J. (2014). Fast reverse-mode automatic differentiation using expression tem-
758 plates in C++. *ACM Trans. Mathematical Softw.*, *40*, 26:1–26:16. <https://doi.org/10.1145/>

- 2560359.
- 759
760 Hogan, R. J., & Bozzo, A. (2018). A flexible and efficient radiation scheme for the
761 ECMWF model. *J. Adv. Modeling Earth Sys.*, *10*, 1990–2008, [https://doi.org/10.1029/](https://doi.org/10.1029/2018MS001364)
762 [2018MS001364](https://doi.org/10.1029/2018MS001364).
- 763 Hogan, R. J., & M. Matricardi (2020). Evaluating and improving the treatment of gases in
764 radiation schemes: the Correlated K-Distribution Model Intercomparison Project (CKD-
765 MIP). *Geosci. Model Dev.*, *13*, 6501–6521, <https://doi.org/10.5194/gmd-13-6501-2020>.
- 766 Hogan, R. J., Ahlgrimm, M., Balsamo, G., Beljaars, A. C. M., Berrisford, P., Bozzo, A., Di
767 Giuseppe, F., Forbes, R. M., Haiden, T., Lang, S., Mayer, M., Polichtchouk, I., Sandu, I.,
768 Vitart F., & Wedi, N. (2017). Radiation in numerical weather prediction. ECMWF Tech.
769 Memo. No. 816, 48 pp, <http://doi.org/10.21957/2bd5dkj8x>.
- 770 Hogan, R. J., Schäfer, S. A. K., Klinger, C., Chiu, J.-C., & Mayer, B. (2016). Representing
771 3D cloud-radiation effects in two-stream schemes: 2. Matrix formulation and broadband
772 evaluation. *J. Geophys. Res.*, *121*, 8583–8599, <https://doi.org/10.1002/2016JD024875>.
- 773 Jakub, F., & Mayer, B. (2016) 3-D radiative transfer in large-eddy simulations – experiences
774 coupling the TenStream solver to the UCLA-LES. *Geosci. Model Dev.*, *9*, 1413–1422,
775 <http://doi.org/10.5194/gmd-9-1413-2016>.
- 776 Joseph, J. H., Wiscombe, W. J., & Weinman, J. A. (1976) The delta-Eddington approxi-
777 mation for radiative flux-transfer. *J. Atmos. Sci.*, *33*, 2452–2459, [http://doi.org/10.1175/](http://doi.org/10.1175/1520-0469(1976)033<2452:TDEAFR>2.0.CO;2)
778 [1520-0469\(1976\)033<2452:TDEAFR>2.0.CO;2](http://doi.org/10.1175/1520-0469(1976)033<2452:TDEAFR>2.0.CO;2).
- 779 Kato, S., Ackerman, T. P., Mather, J. H., & Clothiaux, E. E. (1999). The *k*-distribution
780 method and correlated-*k* approximation for a shortwave radiative transfer model. *J. Quant.*
781 *Spectrosc. Radiat. Trans.*, *62*, 109–121, [https://doi.org/10.1016/S0022-4073\(98\)00075-2](https://doi.org/10.1016/S0022-4073(98)00075-2).
- 782 Lacis, A., & Oinas, V. (1991). A description of the correlated k-distribution method
783 for modeling nongray gaseous absorption, thermal emission, and multiple scattering
784 in vertically inhomogeneous atmospheres. *J. Geophys. Res.*, *96*, 9027-9063,
785 <https://doi.org/10.1029/90JD01945>.
- 786 Lopez, P. (2020). Forecasting the past: views of Earth from the Moon and beyond. *Bull. Am.*
787 *Meteorol. Soc.*, *101*, E1190–E1200, <https://doi.org/10.1175/BAMS-D-19-0254.1>.
- 788 Lu, P., Zhang, H., & Li, J. (2011). Correlated k-distribution treatment of cloud optical prop-
789 erties and related radiative impact. *J. Atmos. Sci.*, *68*, 2671–2688, [https://doi.org/10.1175/](https://doi.org/10.1175/JAS-D-10-05001.1)
790 [JAS-D-10-05001.1](https://doi.org/10.1175/JAS-D-10-05001.1).
- 791 Manabe, S., & Wetherald, R. T. (1967). Thermal equilibrium of the atmosphere with a given
792 distribution of relative humidity. *J. Atmos. Sci.*, *24*, 241–259, [https://doi.org/10.1175/](https://doi.org/10.1175/1520-0469(1967)024<0241:TEOTAW>2.0.CO;2)
793 [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).
- 794 Meinshausen, M., Vogel, E., Nauels, A., Lorbacher, K., Meinshausen, N., Etheridge, D. M.,
795 Fraser, P. J., Montzka, S. A., Rayner, P. J., Trudinger, C. M., Krummel, P. B., Beyerle,
796 U., Canadell, J. G., Daniel, J. S., Enting, I. G., Law, R. M., Lunder, C. R., O’Doherty, S.,
797 Prinn, R. G., Reimann, S., Rubino, M., Velders, G. J. M., Vollmer, M. K., Wang, R. H.
798 J., & Weiss, R. (2017). Historical greenhouse gas concentrations for climate modelling
799 (CMIP6). *Geosci. Model Dev.*, *10*, 2057–2116, <https://doi.org/10.3929/ethz-b-000191830>.
- 800 Mlawer, E. J., Taubman, S. J., Brown, P. D., Iacono, M. J., & Clough, S. A. (1997). Radi-
801 ative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model
802 for the longwave. *J. Geophys. Res. Atmos.*, *102*, 16 663–16 682, [https://doi.org/10.1029/](https://doi.org/10.1029/97JD00237)
803 [97JD00237](https://doi.org/10.1029/97JD00237).
- 804 Modest, M. F., & Zhang, H. (2002). The full-spectrum correlated-*k* distribution for ther-
805 mal radiation from molecular gas-particulate mixtures. *J. Heat Transfer*, *124*, 30–38,
806 <http://doi.org/10.1115/1.1418697>.
- 807 Liu, D. C., & Nocedal, J. (1989). On the limited memory method for large scale optimiza-
808 tion. *Math. Programming B*, *45*, 503–528, <https://doi.org/10.1007/BF01589116>.
- 809 Pawlak, D. T., Clothiaux, E. E., Modest, M. F., & Cole, J. N. S. (2004). Full-spectrum
810 correlated-*k* distribution for shortwave atmospheric radiative transfer. *J. Atmos. Sci.*,
811 *61*, 2588-2601, <https://doi.org/10.1175/JAS3285.1>.

- 812 Ritter, B. & Geleyn, J. F. (1992). A comprehensive radiation scheme for numerical weather
813 prediction models with potential applications in climate simulations. *Mon. Weath. Rev.*,
814 *120*, 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).
- 815 Ukkonen, P., Pincus, R., Hogan, R. J., Nielsen, K. P., and Kaas, E. (2020). Accelerating
816 radiation computations for dynamical models with targeted machine learning and code
817 optimization. *J. Adv. Modeling Earth Sys.*, *12*, e2020MS002226. [https://doi.org/10.1029/](https://doi.org/10.1029/2020MS002226)
818 [2020MS002226](https://doi.org/10.1029/2020MS002226).
- 819 Zhang, H., Nakajima, T., Shi, G., Suzuki, T., & Imasu, R. (2003). An optimal approach to
820 overlapping bands with correlated k distribution method and its application to radiative
821 calculations. *J. Geophys. Res.*, *108*, 4641, <https://doi.org/10.1029/2002JD003358>.