

Balancing accuracy, efficiency, and flexibility in radiation calculations for dynamical models

Robert Pincus^{1,2}, Eli J. Mlawer³, Jennifer S. Delamere⁴

¹Cooperative Institute for Environmental Studies, University of Colorado, Boulder, Colorado, USA

²NOAA/Earth System Research Lab, Physical Sciences Division, Boulder, Colorado, USA

³Atmospheric and Environmental Research, Lexington, Massachusetts, USA

⁴Alpenglow Scientific, Fairbanks, Alaska, USA

Key Points:

- RTE+RRTMGP is a new freely-available toolbox for radiation calculations for dynamical models
- RTE+RRTMGP seeks to balance accuracy, efficiency, and flexibility, defined expansively
- Both code and data continue to evolve to explore different balances among these goals

Abstract

This paper describes the initial implementation of new toolbox that seeks to balance accuracy, efficiency, and flexibility in radiation calculations for dynamical models. The toolbox consists of two related code bases: Radiative Transfer for Energetics (RTE) computes fluxes given a fully-specified radiative transfer problem, and RRTM for GCM applications - Parallel (RRTMGP), which maps a physical description of the gaseous atmosphere into a radiative transfer problem. The toolbox is an implementation of well-established ideas, including the use of a k -distribution to represent the spectral variation of absorption by gases and the use of two-stream, plane-parallel methods for solving the radiative transfer equation. The focus is instead on accuracy, by basing the k -distribution on state-of-the-art spectroscopy, and on the sometimes-conflicting goals of flexibility and efficiency. Flexibility is facilitated by making extensive use of computational objects encompassing code and data, the latter provisioned at run time and potentially tailored to specific problems. The computational objects provide robust access to a set of high-efficiency computational kernels that can be adapted to new computational environments. Accuracy is obtained by careful choice of algorithms and through tuning and validation of the k -distribution against benchmark calculations.

1 Why build another radiation parameterization?

The ultimate energy source for all atmospheric motions is electromagnetic radiation emitted by the sun and by the planet and its atmosphere. The flow of radiative energy through the atmosphere depends strongly on the state of the surface and the atmosphere itself. Essentially any model of the atmospheric motions, therefore, has to represent the flow of radiation through the atmosphere, especially the vertical gradients within the atmosphere that give rise to heating and cooling, and the fluxes of radiation absorbed at the surface. Models aimed at understanding climate must also accurately compute the net energy at the top of the atmosphere.

The representation of radiation is one of the most pure exercises in parameterization in atmospheric models because the solution to fully-specified problems is known to great accuracy. (This can be contrasted with convection parameterizations, for example, for which sensitive dependence on initial conditions make fully deterministic prediction essentially impossible, or cloud microphysics, for which some governing equations are not known.) Accuracy across a wide range of clear-sky conditions can be measured

47 by comparison to benchmark models (Oreopoulos et al., 2012; Pincus et al., 2015) which
48 are themselves known to be in excellent agreement with observations (Mlawer et al., 2000;
49 Turner et al., 2004; Alvarado et al., 2013). Benchmark models also exist for clouds though
50 observational validation is far more challenging.

51 The ideas underlying state-of-the-art radiative transfer parameterizations have been
52 established for decades. Radiation is assumed not to propagate in the horizontal (the
53 Independent Column Approximation), reducing the dimensionality of the radiative trans-
54 fer problem. The complex spectral structure of absorption by gases is treated by group-
55 ing optically-similar spectral regions using either a correlated k -distribution (e.g. Lacis
56 & Oinas, 1991; Fu & Liou, 1992) or, less commonly, by modeling transmission using an
57 exponential sum fit of transmissivities (Wiscombe & Evans, 1977). The optical proper-
58 ties of condensed materials, such as clouds and aerosols, are computed in advance, usu-
59 ally as functions of one or more bulk parameters such as effective radius, and fit to ta-
60 bles or functional forms. The resulting problem is solved using versions of the radiative
61 transfer equation in which the angular dependence has been reduced analytically. Though
62 innovations continue, for example in efforts to treat the impact of three-dimensional trans-
63 port on radiation fields (Schäfer et al., 2016; Hogan et al., 2016), major conceptual ad-
64 vances in the parameterization of radiation are infrequent.

65 The maturity of ideas, the near-universal need for radiation parameterizations, and
66 the substantial effort involved in building an end-to-end parameterization mean that ra-
67 diation codes tend to be developed as complete packages, and that these packages, and
68 especially the interfaces to them, have long lifetimes. The codes used by the UK Met Of-
69 fice have their roots in the work of Edwards & Slingo (1996). In the United States many
70 atmospheric models, including both regional and global models developed at the National
71 Center for Atmospheric Research and the National Weather Service’s Global Forecast
72 System, use RRTMG (Mlawer et al., 1997). These packages are comprehensive, using
73 information about the physical state of the atmosphere to provide values of spectrally-
74 integrated radiative flux.

75 But conceptual maturity and the black-box nature of radiation codes can hide im-
76 portant errors. The accuracy of radiation parameterizations can be judged by compar-
77 ison to reference line-by-line models with high angular resolution; every such compar-
78 ison over the last two-and-a-half decades (e.g. Ellingson et al., 1991; Collins, 2001; Ore-

79 opoulos et al., 2012; Pincus et al., 2015) has identified significant parameterization er-
80 rors in the treatment of gaseous absorption and scattering. These errors partly reflect
81 different efforts to balance computational cost and accuracy, but they also arise because
82 groups may be slow to incorporate new spectroscopic knowledge. Updates to the widely-
83 used HITRAN database (Rothman et al., 2009, 2013) over the last decades, for exam-
84 ple, have tended to increase the amount of solar radiation absorbed by water vapor. Un-
85 derestimating this absorption has important consequences for calculations of hydrologic
86 sensitivity (Fildier & Collins, 2015; DeAngelis et al., 2015). The likelihood of errors in-
87 creases when parameterizations are used to make calculations far outside the range of
88 conditions on which they are trained, for example in calculations on exoplanets (e.g. Yang
89 et al., 2016). Even the highly-elevated concentrations of CO₂ frequently used to estimate
90 climate sensitivity (Gregory, 2004) represent a challenge for some parameterizations (Pin-
91 cus et al., 2015).

92 Complete packages developed for one application may not be easy to adapt to un-
93 foreseen uses. Every existing radiation package of which we are aware assumes a partic-
94 ular orientation in the vertical dimension, requiring the reordering of data when the con-
95 vention in the radiation package differs from that of the host model. Many require sep-
96 arate clear- and all-sky calculations at each invocation where only the latter are needed
97 to advance the host model. None that we’re aware of provide the ability to specify an
98 upper boundary condition, so that their use in models with shallow domains requires the
99 specification of an atmospheric profile for use in the radiation scheme alone, complicat-
100 ing implementation and requiring unnecessary computation. In practice, too, most pack-
101 ages tightly couple two conceptually-different problems: the mapping of atmospheric state
102 to optical properties, and the subsequent calculation of fluxes (i.e. determining the ra-
103 diative transfer problem and determining the solution to a given problem). This tight
104 coupling makes it difficult to explore algorithms in which access to the underlying spec-
105 tral structure is required (e.g Pincus & Stevens, 2009, 2013).

106 Finally, while every process parameterization seeks to minimize computational cost,
107 efficiency is an acute concern for radiation packages because each calculation is so time-
108 consuming that radiation is typically computed less frequently than other processes by
109 factors of 10-20 (see, for example, section 2.1 in Hogan & Bozzo, 2018). Computational
110 efficiency is not a static target, however, because computing platforms changes rapidly
111 even if the underlying algorithms do not. Even today an implementation that is efficient

112 on traditional processors is likely to be poorly structured for specialized but highly-efficient
 113 hardware such as general-purpose Graphical Processing Units (GPUs).

114 This paper describes the initial implementation of new toolbox that seeks to bal-
 115 ance accuracy, efficiency, and flexibility in radiation calculations for dynamical models.
 116 The toolbox consists of two related code bases: Radiative Transfer for Energetics (RTE)
 117 computes fluxes given a fully-specified radiative transfer problem, and RRTM for GCM
 118 applications - Parallel (RRTMGP), which maps a physical description of the aerosol- and
 119 cloud-free atmosphere into a radiative transfer problem. Although every line of RTE+RRTMGP
 120 is new, the code descends from RRTMG (Mlawer et al., 1997; Iacono et al., 2000; Clough
 121 et al., 2005), a parameterization with similar capabilities developed roughly 20 years ago.
 122 It also incorporates many of the lessons learned in the development of PSrad (Pincus &
 123 Stevens, 2013), a re-implementation of RRTMG built to explore an idea that required
 124 extensive refactoring of the original code. Like its predecessors RRTMGP uses a k -distribution
 125 for computing the optical properties and source functions of the gaseous atmosphere based
 126 on profiles of temperature, pressure, and gas concentrations, while RTE computes fluxes
 127 using the Independent Column Approximation in plane-parallel geometry.

128 Below we describe how the design of RTE+RRTMGP balances the sometimes-conflicting
 129 goals of accuracy, efficiency and flexibility, explain how the k -distributions are constructed,
 130 and assess the accuracy of the current model against more detailed calculations.

131 **2 An extensible architecture for flexibility**

132 The calculation of radiative fluxes for dynamical models presents a particular com-
 133 putational challenge among parameterizations. To treat the enormous spectral variabil-
 134 ity of absorption by the many optically-active gases in the atmosphere, a relatively small
 135 amount of state information, i.e. profiles of temperature, pressure, and gas concentra-
 136 tions, must be mapped into optical properties (the parameters need to solve the radia-
 137 tive transfer equation) at a number of spectral quadrature points. Fluxes are computed
 138 independently at each spectral quadrature point. Users, however, are normally interested
 139 only in integrals over the spectrum (or portions of it), so spectrally-resolved fluxes are
 140 summed, greatly reducing the amount of data used by the host model.

141 As a result of this structure the radiation problem has an exceptional opportunity
 142 to exploit fine-grained parallelism. Much of the problem is atomic, meaning that calcu-

143 lations are independent in space and the spectral dimension. Transport calculations are
144 independent in the spectral and horizontal dimensions (the latter as a result of the In-
145 dependent Column Approximation), while spectral reduction is independent in both the
146 horizontal and vertical dimensions. Exploiting this parallelism is key to computational
147 efficiency although the optimal ordering varies across different stages of the computa-
148 tion. RTE and RRTMGP operate on multiple columns at a time to exploit this paral-
149 lelism. The column dimension is inner-most; despite good reasons for having the spec-
150 tral dimension vary fastest (Hogan & Bozzo, 2018) this choice allows user control over
151 vector length and can be easily adapted to different architectures.

152 RTE and RRTMGP are agnostic to the ordering of the vertical axis.

153 **2.1 Designing for robustness**

154 Like the recently-developed ecRad package (Hogan & Bozzo, 2018) RTE+RRTMGP
155 cleanly separates conceptually-distinct aspects of the radiation problem from one another.
156 Each component, including the gas optics and source function calculations, any imple-
157 mentations of aerosol and cloud optics, and methods for computing radiative transfer
158 (transport), can be modified or replaced independently. RTE+RRTMGP is implemented
159 in Fortran 2003. Many components are implemented as Fortran classes that package to-
160 gether code and data. As described below many of the classes are user-extensible to per-
161 mit greater flexibility. The radiative transfer solvers are straightforward functions.

162 The Fortran 2003 classes simplify control and information passing, as described be-
163 low, but basic computational tasks are isolated as kernels, simple procedures with language-
164 neutral interfaces. The computational kernels are implemented in Fortran 90 with C-language
165 bindings including the explicit run-time specification of array sizes. Kernels expect san-
166 itized input and do no error checking, so they can be quite compact and efficient. Sep-
167 arating computational kernels from flow control is also intended to enhance flexibility:
168 it would be possible to build front ends in other languages including Python or C++,
169 using the Fortran class structure or any alternative that suited the problem at hand, and
170 still exploit the efficient Fortran kernels. It would also be possible to replace the default
171 kernels with other implementations. We have explored this possibility in developing pro-
172 totype kernels optimized for GPUs using OpenACC directives.

173 The class structure is also aimed at minimizing the amount of data passed to and
 174 from the radiation calculation, reducing latency and increasing efficiency when radiation
 175 is implemented on dedicated computational resources (e.g. Balaji et al., 2016) and es-
 176 pecially on devices such as GPUs.

177 Other conventions aim to make RTE+RRTMGP more portable across platforms
 178 and environments. The precision of all REAL variables is explicitly set via a Fortran KIND
 179 parameter so that a one character change in a single file can produce single- or double-
 180 precision versions of the code. RTE+RRTMGP uses few thresholds but most are expressed
 181 relative to working precision. Most procedures are implemented as functions returning
 182 character strings; empty strings indicate success while non-empty strings contain error
 183 messages. RTE+RRTMGP does not read or write to files: classes which require data such
 184 as lookup tables at initialization use load functions with flat array arguments so that users
 185 can read and distribute data consistent with their local software environment.

186 **2.2 Specifying and solving the radiative transfer equation: RTE**

187 The components of RTE+RRTMGP communicate through sets of spectrally-dependent
 188 optical properties. Optical properties are described by their spectral discretization: the
 189 number of bands and the spectral limits of each band in units of wavenumber (inverse
 190 centimeters). Each band covers a continuous region of the spectrum but bands need not
 191 be disjoint or contiguous. Anticipating the spectral structure provided by gas optics pa-
 192 rameterizations like RRTMGP, each band may be further sub-divided into g -points. Each
 193 spectral point is treated as a independent pseudo-monochromatic calculation.

194 Optical properties may be specified as sets of numerical values on a column/height/spectral
 195 grid. Optical properties might be only the absorption optical depth τ_a , as is required for
 196 computing radiative transfer in the absence of scattering; the set of extinction optical
 197 depth τ_e , single-scattering albedo ω_0 , and asymmetry parameter g used in two-stream
 198 calculations; and the set of τ_e , ω_0 , and phase function moments p required by four-stream
 199 or other discrete ordinates calculations. (The dependence on two spatial coordinates and
 200 a spectral coordinate is left implicit.)

201 Each possible set of values is represented as a discrete sub-class of the general op-
 202 tical properties class. Using a class structure allows user interaction to be greatly sim-
 203 plified. As one example, sets of optical properties on the same grid can be added together

204 in a single call, with the class structure invoking the correct kernel depending on which
205 two sets of optical properties are provided. Single calls allow optical properties to be delta-
206 scaled (Potter, 1970; Joseph et al., 1976) or checked for erroneous values.

207 Solvers compute radiative fluxes given values of optical properties and appropri-
208 ate boundary conditions and source function values. A shortwave solver requires spec-
209 ifying the (pseudo-)spectrally-dependent collimated beam at the top of the model and
210 albedos for direct and diffuse radiation at the surface in addition to the values of opti-
211 cal properties within the atmosphere. A longwave solver requires the surface emissiv-
212 ity and the values of the Planck source functions from the surface and at each layer and
213 level of the atmosphere.

214 Calculations that account for scattering, the usual standard for shortwave radia-
215 tion and a more accurate option for longwave calculations that include clouds (Costa &
216 Shine, 2006; Kuo et al., 2017), use the two-stream formulation of Meador & Weaver (1980)
217 to compute layer transmittance and reflectance and the adding formulation of Shonk &
218 Hogan (2008) to compute transport (i.e. the fluxes that result from interactions among
219 layers). Two-stream coupling coefficients in the shortwave come from the “practical im-
220 proved flux method” formulations of Zdunkowski et al. (1980); the longwave follows Fu
221 et al. (1997). The accuracy of longwave calculations that neglect scattering may be in-
222 creased through the use of first-order Gaussian quadrature using up to three terms us-
223 ing weights and directions of Clough et al. (1992). Longwave calculations assume that
224 the source function varies linearly with optical depth. At this writing RTE does not yet
225 include four-stream or higher-order methods for radiative transport.

226 The set of optical properties provided determines the solution method: when the
227 solvers are called with the sub-class representing $\{\tau, \omega_0, g\}$ the two-stream/adding solver
228 is invoked; if only τ is provided, a calculation neglecting scattering is performed. Solu-
229 tions are computed for each g -point in the set of optical properties independently, allow-
230 ing RTE to solve problems for any spectral structure.

231 All solvers allow for the specification of incoming diffuse radiation at the top of the
232 domain (this flux is otherwise assumed to be 0). We originally imagined that this capa-
233 bility would be most useful in the simulation of very shallow domain by fine-scale mod-
234 els (e.g. Seifert et al., 2015). Experience implementing RTE+RRTMGP in global mod-

235 els, however, suggests that it may also be a useful alternative to the common practice
236 of adding an extra layer above the model top in radiation calculations.

237 Because radiative fluxes are computed from optical properties there is no explicit
238 treatment of clouds, and particularly of internal cloud variability or its structure in the
239 vertical. Subgrid variability may be accounted for by random sampling in the spectral
240 dimension using the Monte Carlo Independent Column Approximation (Pincus et al.,
241 2003) but treatments that rely on an explicit clear/cloudy partitioning, including the Triple-
242 Clouds algorithm for treating partial cloudiness (Shonk & Hogan, 2008) or the SPAR-
243 TACUS extension for treating the subgrid-scale effects of three-dimensional radiative trans-
244 port (Schäfer et al., 2016; Hogan et al., 2016), are not consistent with this framework.

245 The RTE solvers compute fluxes for each spectral point independently but the full
246 spatial and spectral detail is unlikely to be useful in most contexts. It is, on the other
247 hand, hard to know precisely what users might need. One approach would be to imple-
248 ment an expansive set of output variables, perhaps allowing user control over which are
249 computed, but this can be cumbersome and requires changes to the radiation code to
250 add a new output.

251 RTE takes a conceptually more complicated but practically simpler approach: out-
252 put from RTE solvers is provided through a user-extensible Fortran 2003 class. The class
253 must include storage for the desired results and code to compute or reduce those results
254 from the full profiles of fluxes at each spectral point. In particular the class must imple-
255 ment a reduction function (so-named because it reduces the amount of output) with ar-
256 guments specified by RTE. These arguments include the spectral discretization informa-
257 tion and the vertical ordering, enabling the computation of very specific quantities (dur-
258 ing design we had in mind the calculation of photosynthetically-active radiation at the
259 surface). Examples are provided that compute broadband fluxes (spectrally-integrated
260 up, down, net, and direct if available) and fluxes within each band. Users provide this
261 class in the call to the solver; the solvers, in turn, call the reduction function after spectrally-
262 dependent fluxes are calculated, minimizing the amount of information returned from
263 RTE.

2.3 Computing the optical properties of the gaseous atmosphere: RRT-MGP

RTE provides methods for solving a spectrally-detailed radiative transfer problem; its complement, RRTMGP, determines the parameters of such a radiative transfer problem for the gaseous component of the atmosphere given the physical state and composition. RRTMGP encapsulates the calculation of gas optics, i.e. the calculation of τ_a or $\{\tau_e, \omega_0, g\}$ and the associated source functions given pressure, temperature and gas concentrations within the domain. RRTMGP builds on RTE: the classes representing gas optics and the Planck functions extend the generic representation of optical properties, and the gas optics calculation returns a set of optical property values.

RRTMGP includes a general framework for representing gas optics. One piece of this framework is a class describing the concentrations of gases within the atmosphere. The volume mixing ratio of each gas is provided as a name-value pair, where the name is normally the chemical formula (e.g. ‘ch4’ or ‘h2o’). Values may be provided as scalars, if the gas is well-mixed; as profiles assumed constant in the horizontal; or varying in the horizontal and vertical dimensions.

The second piece of the general framework, an abstract gas optics class, defines a minimal set of interfaces for functions that map atmospheric state to optical properties. Codes written to use this generic interface can seamlessly use any concrete instance of the abstract class. This approach is motivated by the desire to explore hierarchies of detail in the treatment of absorption by gases (Vallis et al., 2018; Tan et al., 2019) without requiring substantial re-coding.

RRTMGP gas optics is a concrete instance of the abstract gas optics class that uses a k -distribution to represent the spectral variation of absorption coefficients. Data and code are entirely distinct in RRTMGP’s gas optics: the class is initialized with data provided in a netCDF file (though RRTMGP does not read the file directly, for reasons explained above). The ability to provide data at run time, available for more than 20 years in the radiation codes used by the UK Met Office (Edwards & Slingo, 1996), provides flexibility, including the provisioning of data with accuracy matched to application needs, as well as a way to incorporate new spectroscopic knowledge as it become available, so that models can stay up-to-date without code changes. The class representing gas concentrations must also be supplied when initializing RRTMGP gas optics, so that the ta-

296 bles of absorption coefficients may be thinned to include only those gases for which con-
 297 centrations are provided, reducing impacts on memory and computation time.

298 **2.4 Mapping concepts to software**

299 Figure 1 illustrates the class structure by which RTE+RRTMGP is organized. The
 300 figure highlights the capabilities described in Sections 2.2 and 2.3. Not shown are ini-
 301 tialization and finalization procedures, procedures for extracting subsets from values de-
 302 fined with a column dimension (available for source functions, optical properties, and gas
 303 concentrations), or the procedures by which the spectral discretization can be set and
 304 queried.

305 Figure 1 emphasizes the distinction between optics, which map atmospheric con-
 306 ditions defined on a spatial grid onto spectrally-dependent values of optical properties
 307 and source functions, and stored sets of these values defined on a spatial and spectral
 308 grid. RRTMGP is a map for the gaseous component of the atmosphere. As we note above,
 309 users must provide analogous maps for condensed species. In most applications users will
 310 initialize these maps (e.g. RRTMGP gas optics, user-provided aerosol and cloud optics)
 311 with data at the beginning of a simulation. Each calculation of radiative fluxes made dur-
 312 ing the course of a simulation uses those maps to determine the optical properties of each
 313 component of the atmosphere, defines a set of problems to be solved (e.g. clear-sky as
 314 the sum of gases and aerosols and all-sky as the sum of clear-sky and clouds), and in-
 315 vokes the solvers on each problem, summarizing results to meet (problem-specific) user
 316 requirements.

317 **3 Accuracy and efficiency**

318 **3.1 Developing a new treatment of absorption by gases**

319 RRTMGP treats absorption by gases using a k -distribution (Ambartsumian, 1936;
 320 Goody et al., 1989; Lacis & Oinas, 1991; Fu & Liou, 1992) in which an integral over fre-
 321 quency ν is replaced by an integral over the variable g defined such that absorption co-
 322 efficient $k(g)$ increases monotonically (and hence much more smoothly); this integral is
 323 further approximated by a discrete sum over G quadrature points using an average ab-
 324 sorption coefficient at each point. The mapping $\mathcal{M}_{\nu \rightarrow g}$ is normally computed for a set
 325 of bands within which absorption is dominated by one or two gases though alternatives

326 are possible (Hogan, 2010). The map varies with the state of the atmosphere, so there
 327 is no inherent mapping between g -points and wavelengths. For RRTMGP the bands are
 328 disjoint, contiguous, and essentially span the set of frequencies of radiation emitted by
 329 the sun or earth.

330 As is described in more detail below, the k -distribution is first generated for a range
 331 of atmospheric conditions following an automated procedure, then tuned by adjusting
 332 these absorption coefficients (and the related “Planck fractions” described below) by hand
 333 so that fluxes and their sensitivity to composition perturbations, computed over a set
 334 of training profiles, are in agreement with line-by-line reference calculations. The Ap-
 335 pendix contains greater detail about the k -distribution and how it is discretized.

336 ***3.1.1 Automated generation of a k -distribution***

337 The version of RRTMGP data described here is based on high-accuracy calcula-
 338 tions with the Line-By-Line Radiative Transfer Model (LBLRTM; Clough et al., 2005),
 339 which has undergone extensive cycles of evaluation with observations and subsequent im-
 340 provement (see, e.g., Mlawer et al., 2012; Alvarado et al., 2013) and agrees with well-
 341 calibrated spectrally resolved radiometric measurements. Results below are based on LBLRTM_v12.8,
 342 line parameter file aer_v.3.6 (itself based to a large extent on the HITRAN 2012 line file
 343 described by Rothman et al. (2013)), and continuum model MT_CKD_3.2. All are avail-
 344 able from <https://rtweb.aer.com>. Shortwave calculations are based on the solar source
 345 function of Lean & DeLand (2012).

346 In the automated step, computations of optical depth are made with LBLRTM for
 347 a set of pressure and temperature values spanning the range of present-day conditions
 348 to define the spectral map. Reference volume mixing ratios $\hat{\chi}_i$ for water vapor and ozone
 349 are based on a large number of profiles from the MERRA-2 reanalysis (Randles et al.,
 350 2017) and vary with temperature, with distinct reference values for pressures greater than
 351 or less than 10000 Pa. Other species use a constant reference value.

352 RRTMGP follows RRTMG in defining bands so that absorption within each band
 353 is dominated by no more than two gases termed that band’s “major species.” Some bands
 354 have no major species. Dry air is used as the second major species in bands in which ab-
 355 sorption is dominated by a single gas, which allows the use of a single method in all bands
 356 to account for the dependence of the k -distribution on the varying abundances of ma-

357 jor species. Computations are made for range of relative abundances $0 \leq \eta \leq 1$ of the
 358 two major species where $\eta \equiv \tilde{\chi}_1/(\tilde{\chi}_1+\tilde{\chi}_2)$ and $\tilde{\chi}_i$ denotes volume mixing ratio χ_i nor-
 359 malized by its reference value $\hat{\chi}_i(p, T)$, with concentrations of all other gases held fixed
 360 at their respective reference values. The total optical depth, including contributions from
 361 major and all minor species, determines the spectral map $\mathcal{M}_{\nu \rightarrow g}(p, T, \eta)$.

362 Given this spectral map, the absorption coefficients for the major species are de-
 363 rived from LBLRTM calculations of absorption optical depth $\tau_a(p, T, \eta)$ in single atmo-
 364 spheric layers containing only the major species in question. Optical depth values are
 365 mapped from frequency ν to g , averaged across a pre-determined number G of g inter-
 366 vals, and converted to absorption coefficients $k(g)$ by dividing by the combined column
 367 amount $W = W_1 + W_2 \times \hat{\chi}_1/\hat{\chi}_2$, where W_i is defined as the layer-integrated molecular
 368 amount (molecules cm^{-2}) of major species i .

369 For longwave bands the same mapping $\mathcal{M}_{\nu \rightarrow g}(p, T, \eta)$ is used to calculate the ‘‘Planck
 370 fraction,’’ defined as the fraction of the band-integrated Planck energy (uniquely deter-
 371 mined by T) associated with each g -point within the band. The solar source function
 372 for each g -point is constant at present; the Appendix describes how these values are ob-
 373 tained.

374 The contributions of other absorbing species are handled with less detail than are
 375 major species. A single representative pressure $p_0(T)$ is chosen for each ‘‘minor species.’’
 376 LBLRTM is used to calculate the spectrally-dependent absorption coefficient of this species
 377 in isolation as a function of temperature. The coefficients are ordered using $\mathcal{M}_{\nu \rightarrow g}(p_0(T), T, \eta)$
 378 and averaged within each of the G intervals. Rayleigh scattering optical depths follow
 379 the same approach.

380 Absorption by both major and minor species is treated separately in the upper and
 381 lower atmosphere (pressures above and below 10000 Pa). Distinct sets of gases are used
 382 in each domain. Some gases are considered below 100 hPa but not above, or vice versa,
 383 depending on the degree to which they influence fluxes.

384 The discretization of the k -distribution available at this writing, including details
 385 about the spectral discretization (‘‘bands’’), the gases considered within each band, and
 386 the density of the tabulated data, are provided in the Appendix. Given this tabulated
 387 information, RRTMGP computes absorption coefficients and Planck fractions for arbi-

388 trary atmospheric conditions by linearly interpolating the tabulated values in $\ln(p)$, T ,
 389 and η . Optical depths are computed by multiplying the interpolated absorption coeffi-
 390 cient by the combined column amount of the layer in question. Interpolation algorithms
 391 are as general as possible so that, for example, the same code is used for contributions
 392 that depend only on absorber abundance and those that also depend on the abundance
 393 of other gases, such as collision-induced absorption and foreign continua. Planck source
 394 functions are determined by multiplying these Planck-fractions by band-integrated Planck
 395 source functions tabulated on a fine temperature grid.

396 ***3.1.2 Testing and tuning the correlated k -distribution***

397 We evaluate the accuracy of the initial k -distribution by computing fluxes for a set
 398 of 42 clear-sky atmospheric profiles (Garand et al., 2001) that span a large range of tem-
 399 perature, moisture, and ozone abundances, and include baseline concentrations of other
 400 gases. Results from RTE+RRTMGP for these training atmospheres are compared to LBLRTM
 401 calculations. We minimize differences due to transport algorithms by using the same set
 402 of quadrature angles in LBLRTM and RTE for longwave problems; in the shortwave we
 403 focus on the direct beam since this depends only on the optical depth. For shortwave
 404 assessments the solar zenith angle is 30 degrees; for longwave calculations the surface emis-
 405 sivity is 1.

406 Fluxes computed across the set of atmospheres using the initial k -distribution are
 407 in substantially better agreement with reference calculations than are fluxes computed
 408 with RRMTG (Figure 2), primarily because RRTMGP is based on the same underly-
 409 ing spectroscopy as the benchmark.

410 We also assess the accuracy of RRTMGP in computing instantaneous radiative forc-
 411 ing, i.e. the change in flux for these 42 profiles due to increases, relative to nominal pre-
 412 industrial concentrations, of factors of two and four for CO_2 and CH_4 and the change
 413 between present-day and pre-industrial concentrations of N_2O and halocarbons. The pri-
 414 mary focus is on $4\times\text{CO}_2$ and $2\times\text{CH}_4$.

415 Accuracy assessments for both flux and forcing guide a hand-tuning of the absorp-
 416 tion coefficients and source functions. This tuning is holistic, considering a wide range
 417 of radiative quantities but focusing primarily on broadband flux and heating rate pro-

418 files and the forcing due to individual gases, especially CO₂ and CH₄. Attention is also
 419 paid to flux and heating rate profiles within each band to minimize compensating errors.

420 In calculations with RTE, the optical properties and source functions provided by
 421 RRTMGP gas optics at each g -point are treated as a set of pseudo-monochromatic cal-
 422 culations. This is equivalent to assuming that the spectral mapping (or “correlation” be-
 423 tween ν and g) is constant through the atmosphere, and is what distinguishes a corre-
 424 lated k -distributions used in vertically inhomogeneous atmosphere from a k -distribution
 425 developed for a single layer. The assumption is an important source of error in corre-
 426 lated k -distributions. In many circumstances the true spectral map varies in the verti-
 427 cal, such that the absorption coefficients for a g -value correspond to different sets of fre-
 428 quencies at different altitudes. As one example, in shortwave bands in which ozone and
 429 water vapor both absorb significantly, absorption in the stratosphere is dominated by
 430 ozone with a very different spectral structure than the absorption by water vapor in the
 431 troposphere, yet absorption due to these two gases will map to the same g -values at dif-
 432 ferent altitudes. In such circumstances the lack of consistency with height of the spec-
 433 tral map $\mathcal{M}_{\nu \rightarrow g}(p, T, \eta)$ (a lack of correlation) degrades model accuracy relative to spectrally-
 434 resolved calculations.

435 The hand-tuning attempts to correct for these and any other errors. Major species
 436 coefficients are adjusted as functions of p and η ; minor species coefficients are tuned as
 437 functions of T . The ad hoc and empirical tuning is similar in spirit to, but substantially
 438 less formal than, the work reported by Sekiguchi & Nakajima (2008), who used an ex-
 439 plicit cost function to determine the spectral discretization and integration rules for their
 440 k -distribution.

441 Tuning modestly improves the accuracy of the k -distribution (compare orange and
 442 green boxes in Fig. 2), decreasing both the bulk of errors and the most extreme errors
 443 in our training atmospheres. Forcing is also improved (see the examples in Figure 3). In
 444 interpreting these results recall that the profiles used here are chosen to explore specific
 445 sources of error rather than being strictly representative of the distribution of conditions
 446 in the Earth’s atmosphere.

3.2 Accuracy: validation and verification

Before comparing results from RTE+RRTMGP against reference calculations we verified RTE against ecRad (Hogan & Bozzo, 2018) by computing broadband fluxes for the training atmospheres with both codes using RRTMGP’s representation of gas optics. Differences in fluxes are within 10^{-8} W/m² for direct and diffuse shortwave fluxes, reflecting the fact that both packages make the same choices even though they are entirely independent implementations. Differences in longwave fluxes are as large as 10^{-2} W/m² due to different formulations of the source function.

The accuracy of fluxes at the atmosphere’s boundaries computed by RTE+RRTMGP in its most commonly-used configuration is shown in Figure 4; RRTMG is shown for comparison. Here longwave fluxes are computed with a single angle and total fluxes (diffuse plus direct for the shortwave) computed for the training atmospheres are compared against reference line-by-line calculations using three angles. Calculations with RRTMG use a diffusivity angle that depends on column-integrated water vapor in some bands to mimic the three-angle calculation (e.g. Fig. 2). The lack of this correction in RRTMGP increases the error in downwelling longwave flux at the surface in some atmospheres. (We are currently developing a similar treatment of diffusivity angle for RRTMGP.) Changes in other fluxes are dominated by revisions to spectroscopy, so that RRTMGP is substantially more accurate than RRTMG.

Figure 5 shows the maximum magnitude of heating rate errors. Pressures greater and less than 10000 Pa are shown separately because radiative heating rates are much larger in the latter than the former.

We assess the out-of-sample accuracy of RTE+RRTMGP using 100 profiles chosen by the Radiative Forcing Model Intercomparison Project (RFMIP) protocol (PinCUS et al., 2016). The profiles were drawn from reanalysis so that the weighted sum of fluxes in the profiles reproduces the change in global-mean, annual-mean top-of-atmosphere present-day to pre-industrial forcing (the change in flux between atmospheres with present-day and pre-industrial concentrations of greenhouse gases). Relative to high angular-resolution line-by-line calculations with LBLRTM, fluxes computed by RTE+RRTMGP are accurate to within 0.4% at the top of the atmosphere and 0.2% at the surface; absorption by the atmosphere is accurate to about 0.4% for the longwave and 0.8% for the shortwave (see Table 1). Pre-industrial to present-day changes at the atmospheres boundaries

Table 1. Error (and reference value) of annual-mean, global-mean instantaneous radiative forcing (W/m^2), for present-day relative to pre-industrial conditions, computed using 100 profiles following the protocol of the Radiative Forcing Model Intercomparison Project. Error is computed relative to reference calculations with high angular and spectral resolution. The columns are chosen to characterize the error in forcing; as one consequence average values for fluxes in the present-day (the first set of columns) is affected by sampling error.

	Present-day fluxes		Pre-industrial to present-day change	
	Longwave	Shortwave	Longwave	Shortwave
Top of atmosphere (up)	0.033 (263.197)	0.165 (47.315)	0.148 (-2.845)	0.007 (-0.058)
Net absorption	-0.749 (-180.696)	-0.610 (72.344)	-0.055 (0.803)	-0.051 (0.522)
Surface (down)	0.725 (315.346)	0.026 (245.553)	-0.095 (2.083)	0.065 (-0.534)

479 are accurate to roughly 5% for longwave change and 12% for the (substantially smaller)
 480 shortwave change.

481 3.3 Efficiency

482 As one measure of efficiency we compare the time taken to compute clear-sky flux
 483 profiles for the 1800 atmospheric conditions (100 profiles for each of 18 perturbations to
 484 atmospheric conditions) used in the RFMIP assessment of accuracy. On a dedicated com-
 485 pute node at the National Energy Research Scientific Computing Center, using current
 486 Intel compilers and processing 8 columns at a time, RRTMGP is slower than RRTMG
 487 by roughly a factor of 2.2 in longwave calculations. RRTMG uses substantially fewer spec-
 488 tral points (140 in the longwave) than does RRTMGP (256); even accounting for this
 489 difference RRTMGP remains about 20% slower than its predecessor. The inefficiency is
 490 mostly due to the calculations of gas optics and Planck sources. It arises partly because
 491 RRTMGP takes a general approach to the calculation of gas optical depths, where RRTMG’s
 492 compute paths (e.g. which gases contributed to absorption in each band) were coded by
 493 hand and so were more easily optimized. We are working to refactor a few closely-related
 494 routines to further increase the computational efficiency. In the shortwave, on the other
 495 hand, RRTMGP is about as twice as fast as RRTMG, or almost 4 times faster per g -
 496 point, owing primarily to easily-vectorized codes. We have noted substantial variation

497 in these ratios across computing platforms, operating systems, and compilers, and cau-
498 tion that real-life applications may be less efficient than these idealized tests.

499 **4 Tools and packages**

500 This paper stresses the principles guiding the development and use of RTE+RRTMGP.
501 This is partly because we expect the underlying software to evolve and partly because
502 the principles – designing parameterizations for flexibility and efficiency from the ground
503 up – may be useful in designing other parameterizations. We have stressed our intent
504 to make RTE+RRTMG as flexible as possible with respect to both the computing en-
505 vironment and the context in which radiative calculations are to be made.

506 One consequence of agnosticism with respect to the host model is that users have
507 substantially more responsibility. This is most obvious in the treatment of clouds and
508 aerosol. The RTE+RRTMGP repository includes examples to compute cloud optics (the
509 map from physical state to optical properties), using a class analogous to the RRTMGP
510 gas optics, and to treat cloud overlap with the Monte Carlo Independent Column Ap-
511 proximation (Pincus et al., 2003), using procedures relying on user-generated random
512 numbers. The examples are narrow by design and are directly useful only if the assump-
513 tions about macro- and micro-physics are consistent with the host model’s. The intent
514 of the examples is to be useful as a starting point from which users may build implemen-
515 tations more self-consistent with the host model’s other formulations. The programs used
516 to compute accuracy for RFMIP in section 3.2, also included in the RTE+RRTMGP repos-
517 itory, show how the RRTMGP gas optics is initialized from data and used to compute
518 the inputs needed for RTE, and how output is extracted from RTE, and play a similar
519 role.

520 Many of the concerns that spurred the development of RTE+RRTMGP have mo-
521 tivated other development efforts. One example is the ecRad code (Hogan & Bozzo, 2018),
522 which was developed contemporaneously. Compared to RTE+RRTMGP, ecRad is more
523 complete (it includes treatments for cloud and aerosol optics and carefully-crafted meth-
524 ods for sub-grid scale sampling of homogenous clouds) and more capable (it includes al-
525 ternatives for treating cloud overlap and a parameterization for three-dimensional trans-
526 port within each column). The ecRad package represents a complete solution suitable
527 for users who want to make precisely the same choices or are willing to adapt the inter-

528 nals of the package to their own needs. RTE+RRTMGP, in contrast, is intended as an
529 extensible tool or platform on which user-specific applications can be built by extension
530 rather than modification.

531 Optics computations - the mapping from model state to a radiative transfer prob-
532 lem - are a form of coupling in which detailed information about both representations
533 is required. From this perspective the role of RTE is to provide a reasonably flexible rep-
534 resentation of the radiative transfer problem and a matched set of methods for solution.
535 The coupling of clouds and aerosols to these problems is left to users because the vari-
536 ety of possible macro- and micro-physical descriptions is enormous while the tools re-
537 quired to make the map, such as codes for computing single-scattering properties using
538 Mie-Lorenz theory, are widely accessible. Computing the optical properties of the gaseous
539 atmosphere, on the other hand, requires a small and easily enumerable set of inputs but
540 relies on tools and expertise that is less broadly distributed among the community. These
541 considerations explain our choice to link RTE+RRTMGP in both the software sense and
542 in this description.

543 This paper reports on the initial implementation of RTE+RRTMGP. In particu-
544 lar the assessments of accuracy in Section 3 use a k -distribution with 16 g -points per band,
545 for a total of 256 in the longwave and 224 in the shortwave. Experience developing the
546 predecessor RRTMG from its parent model suggests that much of the accuracy of the
547 underlying k -distribution can be obtained with substantially fewer spectral points (see
548 also Sekiguchi & Nakajima, 2008), making possible substantial increases in efficiency for
549 modest decreases in accuracy. We also anticipate that accuracy in clearly defined appli-
550 cations such as weather forecasting may be able to achieve the same accuracy with less
551 computational cost by reducing the number of spectral points that provide accuracy in
552 instantaneous radiative forcing. We are currently working to provide several sets of ab-
553 sorption coefficients striking different balances between accuracy and efficiency.

554 **Acknowledgments**

555 Code, data, and user documentation for RTE+RRTMGP are available at [https://](https://github.com/RobertPincus/rte-rrtmgp)
556 github.com/RobertPincus/rte-rrtmgp; this manuscript was produced with data and
557 code archived at doi:(to be provided at acceptance). Scripts and data in this paper are
558 available at <https://github.com/RobertPincus/rte-rrtmgp-paper-figures> (doi to

559 be provided at acceptance). The LBLRTM model and associated data used in the ac-
 560 curacy calculations in Section 3 is available at <http://rtweb.aer.com>. We appreciate
 561 design suggestions from Robin J. Hogan of ECMWF and Brian Eaton of NCAR. Ben-
 562 jamin Hillman of Sandia National Lab has identified a number of bugs in prototype im-
 563 plementations. Andre Wehe helped with design and implementation in the early stages
 564 of the project. K. Franklin Evans and Rick Pernak provided infrastructure for regres-
 565 sion testing, verification, and validation. We are grateful for comments from three anony-
 566 mous reviewers, which helped sharpen our writing.

567 This work utilized the Janus and RMACC Summit supercomputers, both of which
 568 were supported by the National Science Foundation (Janus under award number CNS-
 569 0821794, Summit under awards ACI-1532235 and ACI-1532236) and the University of
 570 Colorado Boulder. Janus supercomputer was a joint effort of the University of Colorado
 571 Boulder, the University of Colorado Denver and the National Center for Atmospheric
 572 Research; Summit a joint effort of the University of Colorado Boulder and Colorado State
 573 University. The development of RTE+RRTMPG has been supported to date by the Of-
 574 fice of Naval Research under grants N00014-13-1-0858 and N00014-17-1-2158, by NASA
 575 under grant NNH13CJ92C, and by the Department of Energy’s Office of Biological and
 576 Environmental Sciences under grants DE-SC0012399 and DE-SC0016593.

577 **A RRTMGP’s k -distribution in detail**

578 Tables A.1 and A.2 show the band structure adopted in the present version of RRT-
 579 MGP. The band values in the longwave differ modestly from those in RRTMG. The or-
 580 dering of shortwave bands is strictly monotonic, abandoning the idiosyncratic ordering
 581 of RRTMG. Both changes imply that any fits e.g. for cloud optical properties made for
 582 RRTMG will need to be revisited before use in RRTMGP.

583 The spectral map $\mathcal{M}_{\nu \rightarrow g}(p, T, \eta)$ is computed at pressures $1 \leq p \leq 109600$ Pa
 584 in increments of $\ln(p) = 0.2$, temperatures $160 \leq T \leq 355$ K in 15 K increments, and
 585 $\eta = 0, 1/8, \dots 1$. When computing η the mixing ratio of the second major gas v_2 is set
 586 to the reference value $\hat{v}_2(p, T)$ and v_1 varies except at $\eta = 1$, where $v_2 = 0$ and $v_1 =$
 587 $\hat{v}_1(p, T)$.

588 Band-integrated values of the Planck function are computed in 1 K increments.

589 The g -point dependence of the solar source function S is determined from the ref-
 590 erence line-by-line calculations for the 42 atmospheres used for validation (Sec. 3.1.2).
 591 For each profile i within this set and within each band b we identify the pressure $\check{p}_{i,b}$ at
 592 which the direct solar beam has been depleted by 10% and compute the map at the cor-
 593 responding values of T and η . Although the Garand et al. (2001) atmospheres span a
 594 wide range of temperatures and gas abundances we find relatively little variation among
 595 the maps $\mathcal{M}_{\nu \rightarrow g}^{i,b}(\check{p}_b, T^i(\check{p}_{i,b}), \eta(\check{p}_{i,b}))$. We therefore compute the average map across the
 596 set of profiles and apply this map to the incident solar radiation to determine $S(g)$ for
 597 all profiles.

598 References

- 599 Alvarado, M. J., Payne, V. H., Mlawer, E. J., Uymin, G., Shephard, M. W., Cady-
 600 Pereira, K. E., . . . Moncet, J. L. (2013). Performance of the Line-By-Line Ra-
 601 diative Transfer Model (LBLRTM) for temperature, water vapor, and trace gas
 602 retrievals: recent updates evaluated with IASI case studies. *Atmos. Chem. Phys.*,
 603 *13*(14), 6687–6711.
- 604 Ambartsumian, V. (1936). The effect of absorption lines on the radiative equilib-
 605 rium of the outer layers of stars. *Publ. Astron. Obser. Leningrad*, *6*, 7–18.
- 606 Balaji, V., Benson, R., Wyman, B., & Held, I. (2016, October). Coarse-grained com-
 607 ponent concurrency in Earth system modeling: parallelizing atmospheric radiative
 608 transfer in the GFDL AM3 model using the Flexible Modeling System coupling
 609 framework. *Geosci. Model Dev.*, *9*(10), 3605–3616.
- 610 Clough, S. A., Iacono, M. J., & Moncet, J.-L. (1992). Line-by-line calculations of at-
 611 mospheric fluxes and cooling rates: Application to water vapor. *J. Geophys. Res.*,
 612 *97*(D14), 15761–15785.
- 613 Clough, S. A., Shephard, M. W., Mlawer, E. J., Delamere, J. S., Iacono, M. J.,
 614 Cady-Pereira, K., . . . Brown, P. D. (2005, March). Atmospheric radiative transfer
 615 modeling: a summary of the AER codes. *J. Quant. Spectrosc. Radiat. Transfer*,
 616 *91*(2), 233–244.
- 617 Collins, W. D. (2001, November). Parameterization of Generalized Cloud Overlap
 618 for Radiative Calculations in General Circulation Models. *J. Atmos. Sci.*, *58*(21),
 619 3224–3242.
- 620 Costa, S. M. S., & Shine, K. P. (2006, April). An estimate of the global impact

Table A.1. Current RRTMGP spectral structure for the longwave. The distinction between major and minor absorbers is explained in Section 3.1. Water vapor foreign and self-continua are also included as minor gases for any bands in which water vapor is a major species.

Band	Wavenumber limits (cm^{-1})	absorbers ($p \geq 10000\text{Pa}$)		absorbers ($p < 10000\text{Pa}$)	
		major	minor	major	minor
1	10 - 250	H ₂ O	N ₂	H ₂ O	N ₂
2	250 - 500	H ₂ O		H ₂ O	
3	500 - 630	H ₂ O, CO ₂	N ₂ O	H ₂ O, CO ₂	N ₂ O
4	630 - 700	H ₂ O, CO ₂		O ₃ , CO ₂	
5	700 - 820	H ₂ O, CO ₂	O ₃ , CCL ₄ , CFC-22	O ₃ , CO ₂	CCL ₄ , CFC-22
6	820 - 980	H ₂ O	CO ₂ , CFC-11, CFC-12, HFC-143 a	—	CFC-11, CFC-12, HFC-143 a
7	980 - 1080	H ₂ O, O ₃	CO ₂	O ₃	CO ₂
8	1080 - 1180	H ₂ O	CO ₂ , O ₃ , N ₂ O, CFC-12, CFC-22, HFC-23 HFC-32, HFC-125, HFC-134 a	O ₃	CO ₂ , CO, CFC-12, CFC-22, HFC-23, HFC-32, HFC-125, HFC-134 a
9	1180 - 1390	H ₂ O, CH ₄	N ₂ O, CF ₄ , HFC-134 a, HFC-143 a	CH ₄	N ₂ O, CF ₄ , HFC-134 a, HFC-143 a
10	1390 - 1480	H ₂ O		H ₂ O	
11	1480 - 1800	H ₂ O	O ₂	H ₂ O	O ₂
12	1800 - 2080	H ₂ O, CO ₂		—	
13	2080 - 2250	H ₂ O, N ₂ O	CO ₂ , CO	—	O ₃
14	2250 - 2390	CO ₂		CO ₂	
15	2390 - 2680	H ₂ O, CO ₂	N ₂ O, N ₂	—	
16	2680 - 3250	H ₂ O, CH ₄		CH ₄	

Table A.2. Current RRTMGP spectral structure for the shortwave. The distinction between major and minor absorbers is explained in Section 3.1. Water vapor foreign and self-continua are also included as minor gases for any bands in which water vapor is a major species.

Band	Wavenumber limits (cm^{-1})	absorbers ($p \geq 10000\text{Pa}$)		absorbers ($p < 10000\text{Pa}$)	
		major	minor	major	minor
1	820 - 2680	H ₂ O, CO ₂	CH ₄ , N ₂ O, N ₂	H ₂ O, CO ₂	CH ₄ , N ₂ O, O ₃
2	2680 - 3250	H ₂ O, CH ₄		CH ₄	
3	3250 - 4000	H ₂ O, CO ₂		H ₂ O, CO ₂	
4	4000 - 4650	H ₂ O, CH ₄		CH ₄	
5	4650 - 5150	H ₂ O, CO ₂		CO ₂	
6	5150 - 6150	H ₂ O	CH ₄	H ₂ O	CH ₄
7	6150 - 7700	H ₂ O	CO ₂	H ₂ O, CO ₂	
8	7700 - 8050	H ₂ O, O ₂		H ₂ O, O ₂	
9	8050 - 12850	H ₂ O	O ₂	H ₂ O	O ₃
10	12850 - 16000	H ₂ O, O ₂	O ₃	H ₂ O, O ₂	O ₃
11	16000 - 22650	H ₂ O	O ₃ , O ₂ , NO ₂	O ₃	O ₂ , NO ₂
12	22650 - 29000	—	NO ₂	—	NO ₂
13	29000 - 38000	O ₃		O ₃	
14	38000 - 50000	O ₃ , O ₂		O ₃ , O ₂	

- 621 of multiple scattering by clouds on outgoing long-wave radiation. *Quart. J. Royal*
 622 *Met. Soc.*, 132(616), 885–895.
- 623 DeAngelis, A. M., Qu, X., Zelinka, M. D., & Hall, A. (2015, December). An observa-
 624 tional radiative constraint on hydrologic cycle intensification. *Nature*, 528(7581),
 625 249–253.
- 626 Edwards, J. M., & Slingo, A. (1996, April). Studies with a flexible new radiation
 627 code. I: Choosing a configuration for a large-scale model. *Quart. J. Royal Met.*
 628 *Soc.*, 122(531), 689–719.
- 629 Ellingson, R. G., Ellis, J., & Fels, S. (1991). The intercomparison of radiation codes
 630 used in climate models: Long wave results. *J. Geophys. Res.*, 96(D5), 8929–8953.
- 631 Fildier, B., & Collins, W. D. (2015, October). Origins of climate model discrep-
 632 ancies in atmospheric shortwave absorption and global precipitation changes. *Geo-*
 633 *phys. Res. Lett.*, 42(20), 8749–8757.
- 634 Fu, Q., & Liou, K. N. (1992). On the correlated k -distribution method for radiative
 635 transfer in nonhomogeneous atmospheres. *J. Atmos. Sci.*, 49(22), 2139–2156.
- 636 Fu, Q., Liou, K. N., Cribb, M. C., Charlock, T. P., & Grossman, A. (1997, Decem-
 637 ber). Multiple Scattering Parameterization in Thermal Infrared Radiative Trans-
 638 fer. *J. Atmos. Sci.*, 54(24), 2799–2812.
- 639 Garand, L., Turner, D. S., Larocque, M., Bates, J., Boukabara, S., Brunel, P., . . .
 640 Woolf, H. (2001, October). Radiance and Jacobian intercomparison of radia-
 641 tive transfer models applied to HIRS and AMSU channels. *J. Geophys. Res.*,
 642 106(D20), 24017–24031.
- 643 Goody, R., West, R., Chen, L., & Crisp, D. (1989, December). The correlated-
 644 k method for radiation calculations in nonhomogeneous atmospheres. *J. Quant.*
 645 *Spectrosc. Radiat. Transfer*, 42(6), 539–550.
- 646 Gregory, J. M. (2004). A new method for diagnosing radiative forcing and climate
 647 sensitivity. *Geophys. Res. Lett.*, 31(3), 147.
- 648 Hogan, R. J. (2010, June). The Full-Spectrum Correlated- k Method for Longwave
 649 Atmospheric Radiative Transfer Using an Effective Planck Function. *J. Atmos.*
 650 *Sci.*, 67(6), 2086–2100.
- 651 Hogan, R. J., & Bozzo, A. (2018, August). A Flexible and Efficient Radiation
 652 Scheme for the ECMWF Model. *J. Adv. Model. Earth Syst.*, 10(8), 1990–2008.
- 653 Hogan, R. J., Schäfer, S. A. K., Klinger, C., Chiu, J. C., & Mayer, B. (2016, July).

- 654 Representing 3-D cloud radiation effects in two-stream schemes: 2. Matrix formu-
 655 lation and broadband evaluation. *J. Geophys. Res.*, *121*(14), 8583–8599.
- 656 Iacono, M. J., Mlawer, E. J., Clough, S. A., & Morcrette, J.-J. (2000, June). Im-
 657 pact of an improved longwave radiation model, RRTM, on the energy budget and
 658 thermodynamic properties of the NCAR community climate model, CCM3. *J.*
 659 *Geophys. Res.*, *105*(D11), 14873–14890.
- 660 Joseph, J. H., Wiscombe, W. J., & Weinman, J. A. (1976, December). The Delta-
 661 Eddington Approximation for Radiative Flux Transfer. *J. Atmos. Sci.*, *33*(12),
 662 2452–2459.
- 663 Kuo, C.-P., Yang, P., Huang, X., Feldman, D., Flanner, M., Kuo, C., & Mlawer,
 664 E. J. (2017, December). Impact of Multiple Scattering on Longwave Radiative
 665 Transfer Involving Clouds. *J. Adv. Model. Earth Syst.*, *9*(8), 3082–3098.
- 666 Lacis, A. A., & Oinas, V. (1991). A description of the correlated k -distribution
 667 method for modeling non-grey gaseous absorption, thermal emission, and multiple
 668 scattering in vertically inhomogeneous atmospheres. *J. Geophys. Res.*, *96*(D5),
 669 9027–9063.
- 670 Lean, J. L., & DeLand, M. T. (2012, April). How Does the Sun’s Spectrum Vary? *J.*
 671 *Climate*, *25*(7), 2555–2560.
- 672 Meador, W. E., & Weaver, W. R. (1980). Two-Stream Approximations to Radiative
 673 Transfer in Planetary Atmospheres: A Unified Description of Existing Methods
 674 and a New Improvement. *J. Atmos. Sci.*, *37*(3), 630–643.
- 675 Mlawer, E. J., Brown, P. D., Clough, S. A., Harrison, L. C., Michalsky, J. J.,
 676 Kiedron, P. W., & Shippert, T. (2000, September). Comparison of spectral
 677 direct and diffuse solar irradiance measurements and calculations for cloud-free
 678 conditions. *Geophys. Res. Lett.*, *27*(17), 2653–2656.
- 679 Mlawer, E. J., Payne, V. H., Moncet, J. L., Delamere, J. S., Alvarado, M. J., & To-
 680 bin, D. C. (2012, April). Development and recent evaluation of the MTCKD
 681 model of continuum absorption. *Phil. Trans. Royal Soc. A*, *370*(1968), 2520–
 682 2556.
- 683 Mlawer, E. J., Taubman, S. J., Brown, P. D., Iacono, M. J., & Clough, S. A. (1997,
 684 July). Radiative transfer for inhomogeneous atmospheres: RRTM, a validated
 685 correlated- k model for the longwave. *J. Geophys. Res.*, *102*(D14), 16663–16682.
- 686 Oreopoulos, L., Mlawer, E., Delamere, J., Shippert, T., Cole, J., Fomin, B., ...

- 687 Rossow, W. B. (2012, March). The Continual Intercomparison of Radiation
688 Codes: Results from Phase I. *J. Geophys. Res.*, *117*, D06118.
- 689 Pincus, R., Barker, H. W., & Morcrette, J.-J. (2003). A fast, flexible, approximate
690 technique for computing radiative transfer in inhomogeneous cloud fields. *J. Geo-*
691 *phys. Res.*, *108*(D13), 4376–n/a.
- 692 Pincus, R., Forster, P. M., & Stevens, B. (2016). The Radiative Forcing Model In-
693 tercomparison Project (RFMIP): experimental protocol for CMIP6. *Geosci. Model*
694 *Dev.*, *9*, 3447–3460.
- 695 Pincus, R., Mlawer, E. J., Oreopoulos, L., Ackerman, A. S., Baek, S., Brath, M., ...
696 Schwarzkopf, D. M. (2015, July). Radiative flux and forcing parameterization
697 error in aerosol-free clear skies. *Geophys. Res. Lett.*, *42*(13), 5485–5492.
- 698 Pincus, R., & Stevens, B. (2009). Monte Carlo Spectral Integration: a consistent ap-
699 proximation for radiative transfer in large eddy simulations. *J. Adv. Model. Earth*
700 *Syst.*, *1*(2), n/a–n/a.
- 701 Pincus, R., & Stevens, B. (2013, May). Paths to accuracy for radiation parameteri-
702 zations in atmospheric models. *J. Adv. Model. Earth Syst.*, *5*(2), 225–233.
- 703 Potter, J. F. (1970, September). The Delta Function Approximation in Radiative
704 Transfer Theory. *J. Atmos. Sci.*, *27*(6), 943–949.
- 705 Randles, C. A., da Silva, A. M., Buchard, V., Colarco, P. R., Darmenov, A., Govin-
706 daraju, R., ... Flynn, C. J. (2017, September). The MERRA-2 Aerosol Reanal-
707 ysis, 1980 Onward. Part I: System Description and Data Assimilation Evaluation.
708 *J. Climate*, *30*(17), 6823–6850.
- 709 Rothman, L. S., Gordon, I. E., Babikov, Y., Barbe, A., Chris Benner, D., Bernath,
710 P. F., ... Wagner, G. (2013, November). The HITRAN2012 molecular spectro-
711 scopic database. *J. Quant. Spectrosc. Radiat. Transfer*, *130*(0), 4–50.
- 712 Rothman, L. S., Gordon, I. E., Barbe, A., Benner, D. C., Bernath, P. F., Birk, M.,
713 ... Vander Auwera, J. (2009, June). The HITRAN 2008 molecular spectroscopic
714 database. *J. Quant. Spectrosc. Radiat. Transfer*, *110*(9–10), 533–572.
- 715 Schäfer, S. A. K., Hogan, R. J., Klinger, C., Chiu, J. C., & Mayer, B. (2016, July).
716 Representing 3-D cloud radiation effects in two-stream schemes: 1. Longwave con-
717 siderations and effective cloud edge length. *J. Geophys. Res.*, *121*(14), 8567–8582.
- 718 Seifert, A., Heus, T., Pincus, R., & Stevens, B. (2015, December). Large-eddy sim-
719 ulation of the transient and near-equilibrium behavior of precipitating shallow

- 720 convection. *J. Adv. Model. Earth Syst.*, 7(4), 1918–1937.
- 721 Sekiguchi, M., & Nakajima, T. (2008, November). A k-distribution-based radiation
722 code and its computational optimization for an atmospheric general circulation
723 model. *J. Quant. Spectrosc. Radiat. Transfer*, 109(17-18), 2779–2793.
- 724 Shonk, J. K. P., & Hogan, R. J. (2008, June). Tripleclouds: An Efficient Method for
725 Representing Horizontal Cloud Inhomogeneity in 1D Radiation Schemes by Using
726 Three Regions at Each Height. *J. Climate*, 21(11), 2352–2370.
- 727 Tan, Z., Lachmy, O., & Shaw, T. A. (2019, March). The sensitivity of the jet stream
728 response to climate change to radiative assumptions. *J. Adv. Model. Earth Syst.*,
729 0(ja), 2018MS001492.
- 730 Turner, D. D., Tobin, D. C., Clough, S. A., Brown, P. D., Ellingson, R. G., Mlawer,
731 E. J., ... Shephard, M. W. (2004, November). The QME AERI LBLRTM: A
732 Closure Experiment for Downwelling High Spectral Resolution Infrared Radiance.
733 *J. Atmos. Sci.*, 61(22), 2657–2675.
- 734 Vallis, G. K., Colyer, G., Geen, R., Gerber, E., Jucker, M., Maher, P., ... Thomson,
735 S. I. (2018). Isca, v1.0: a framework for the global modelling of the atmospheres
736 of Earth and other planets at varying levels of complexity. *Geosci. Model Dev.*,
737 11(3), 843–859.
- 738 Wiscombe, W. J., & Evans, J. W. (1977, August). Exponential-sum fitting of radia-
739 tive transmission functions. *J. Comp. Phys.*, 24(4), 416–444.
- 740 Yang, J., Leconte, J., Wolf, E. T., Goldblatt, C., Feldl, N., Merlis, T., ... Abbot,
741 D. S. (2016). Differences in Water Vapor Radiative Transfer among 1D Models
742 Can Significantly Affect the Inner Edge of the Habitable Zone. *Ap. J.*, 826(2),
743 222.
- 744 Zdunkowski, W. G., Welch, R. M., & Korb, G. J. (1980, September). An investi-
745 gation of the structure of typical two-stream methods for the calculation of solar
746 fluxes and heating rates in clouds. *Beiträge zur Physik Atmosphäre*, 53, 147–166.

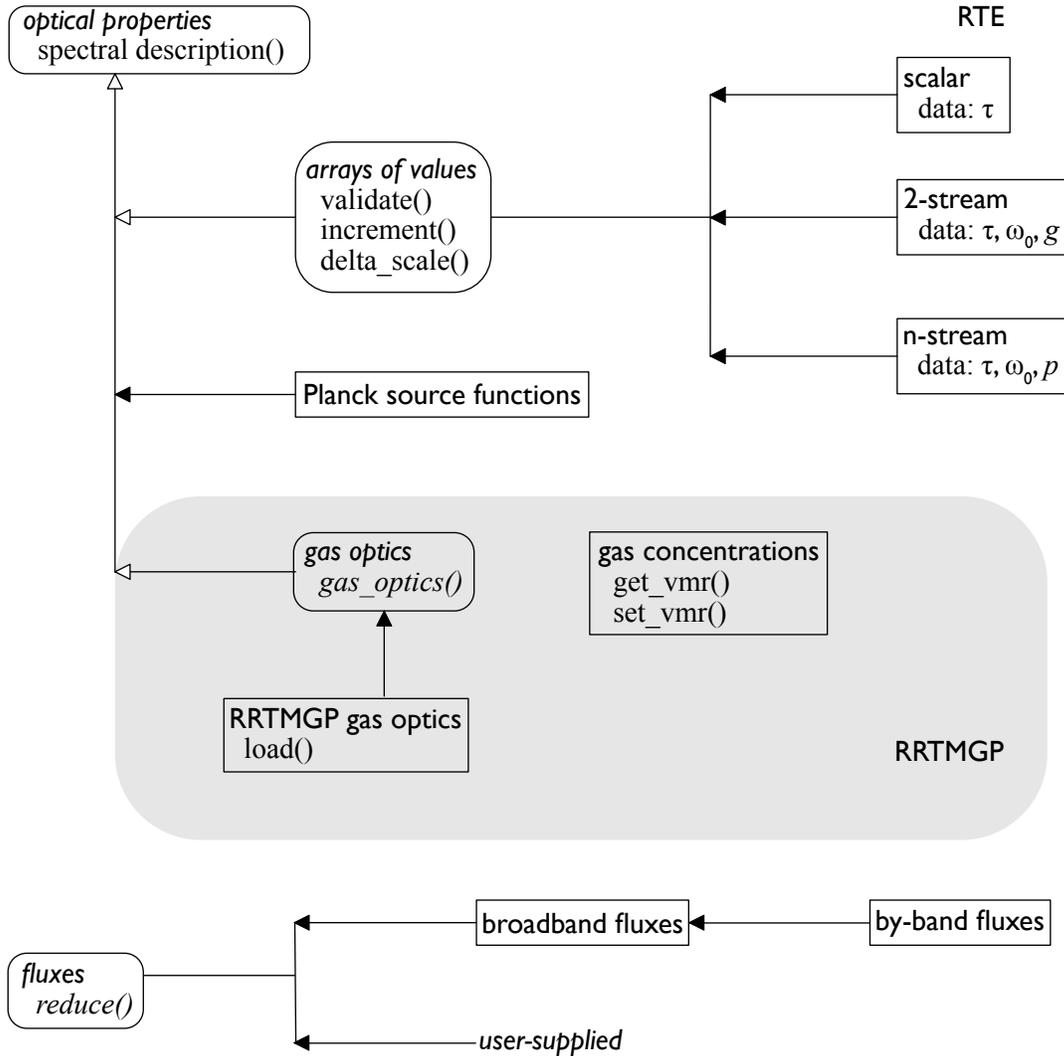


Figure 1. Class organization for RTE+RRTMGP. Class names are in sans serif fonts, data and procedures in serif. Arrows indicate inheritance: classes inherit the data and procedures and/or interfaces provided by their parents. Ovals, open arrowheads, and italicized class names represent abstract classes providing functionality and/or specifying procedures to be provided by descendent classes. Calculations require concrete classes (un-italicized names, rectangles). Solvers are implemented as procedures using these classes as inputs or to compute outputs. The figure illustrates only the most important functionality within each class; most implement more procedures than are shown.

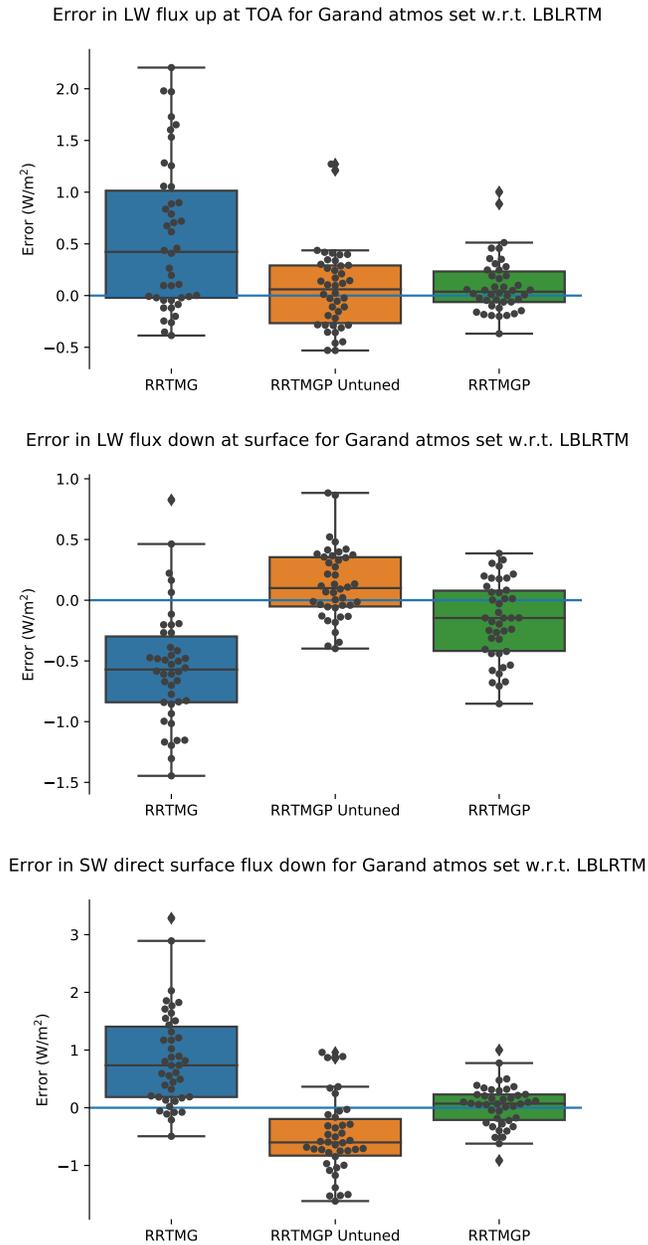


Figure 2. Accuracy of RRTMGP’s new k -distribution, assessed as the difference between fluxes computed with RTE+RRTMGP and those from the reference calculations across the set of training atmospheres. Longwave calculations compare high spectral resolution line-by-line and parameterized calculations using identical transport algorithms while the shortwave comparison focused only on the direct solar beam at the surface and so requires no multiple-scattering calculations.

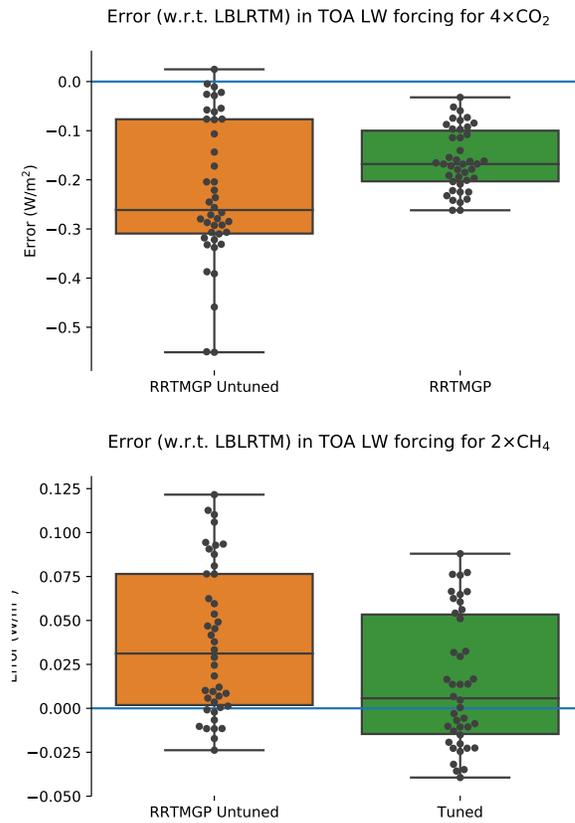


Figure 3. Accuracy of RRTMGP’s new k -distribution for forcing calculations. Shown here are the two primary forcings considered during tuning: impacts on top-of-atmosphere longwave fluxes from concentrations of carbon dioxide quadrupled from pre-industrial concentrations, and doubled methane concentrations. As with fluxes, tuning reduces the largest errors and modestly improves the median error across the training dataset.

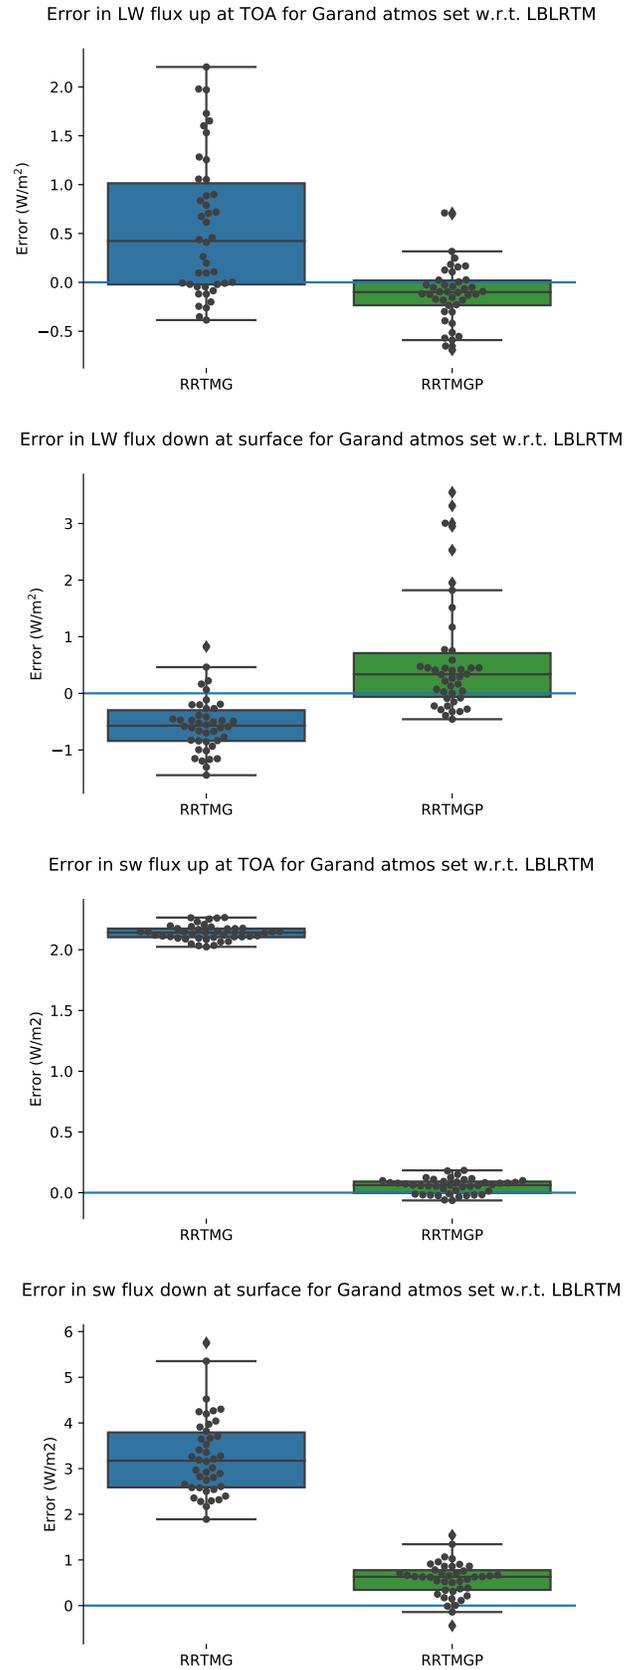


Figure 4. Accuracy of RTE+RRTMGP in producing fluxes at the surface and top-of-

atmosphere as judged against line-by-line calculations on the set of training atmospheres. RTE

uses a single angle calculation (c.f. the three-angle calculation in Fig. 2) for the longwave calcula-

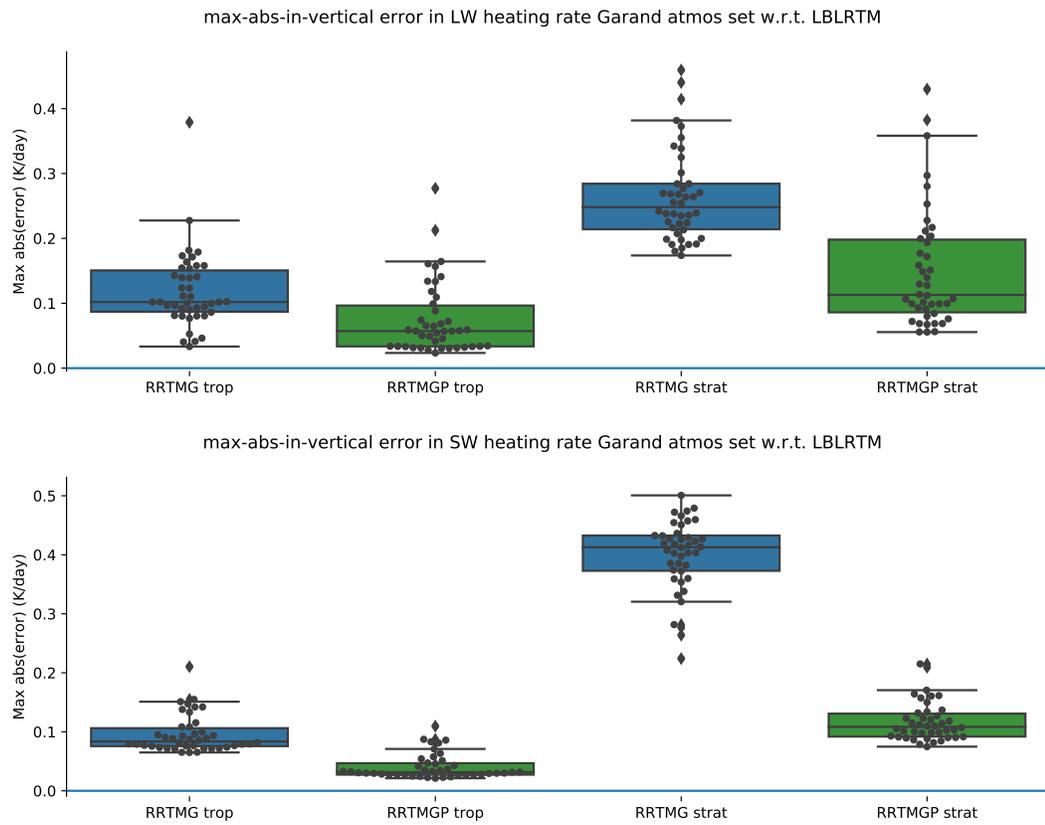


Figure 5. Accuracy of RTE+RRTMGP in producing heating rates.