

# Harnessing machine learning for accurate treatment of overlapping opacity species in GCMs

Aaron David Schneider<sup>1,2</sup>, Paul Mollière<sup>3</sup>, Gilles Louppe<sup>4</sup>, Ludmila Carone<sup>5</sup>, Uffe Gråe Jørgensen<sup>1</sup>, Leen Decin<sup>2</sup> & Christiane Helling<sup>5</sup>

(1) Centre for ExoLife Sciences, Niels Bohr Institute, Øster Voldgade 5, 1350 Copenhagen, Denmark

(2) Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, 3001, Leuven, Belgium

(3) Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

(4) Montefiore Institute, University of Liège, Liège, Belgium

(5) Space Research Institute, Austrian Academy of Sciences, Schmiedlstrasse 6, A-8042 Graz, Austria

November 3, 2023

## ABSTRACT

To understand high precision observations of exoplanets and brown dwarfs, we need detailed and complex general circulation models (GCMs) that incorporate hydrodynamics, chemistry, and radiation. In this study, we specifically examine the coupling between chemistry and radiation in GCMs and compare different methods for mixing opacities of different chemical species in the correlated-k assumption, when equilibrium chemistry cannot be assumed. We propose a fast machine learning method based on DeepSets (DS), which effectively combines individual correlated-k opacities (k-tables). We evaluate the DS method alongside other published methods like adaptive equivalent extinction (AEE) and random overlap with rebinning and resorting (RORR). We integrate these mixing methods into our GCM (expeRT/MITgcm) and assess their accuracy and performance for the example of the hot Jupiter HD 209458 b. Our findings indicate that the DS method is both accurate and efficient for GCM usage, whereas RORR is too slow. Additionally, we observe that the accuracy of AEE depends on its specific implementation and may introduce numerical issues in achieving radiative transfer solution convergence. We then apply the DS mixing method in a simplified chemical disequilibrium situation, where we model the rainout of TiO and VO, and confirm that the rainout of TiO and VO would hinder the formation of a stratosphere. To further expedite the development of consistent disequilibrium chemistry calculations in GCMs, we provide documentation and code for coupling the DS mixing method with correlated-k radiative transfer solvers. The DS method has been extensively tested to be accurate enough for GCMs, however, other methods might be needed for accelerating atmospheric retrievals.

**Key words.** Radiation: dynamics – Radiative transfer – Planets and satellites: atmospheres – Planets and satellites: gaseous planets – Methods: numerical

## 1. Introduction

General circulation models (GCMs) have been applied with wide success to understand the 3D nature of exoplanetary atmospheres (for a review see Showman et al. 2020). These models usually consist of a dynamical core that solves the equations of hydrodynamics, coupled to physical parameterizations with differing complexity that describe the forcing on the temperature and velocity field. A very common physical parameterization in GCMs is that of heating and cooling by (gas) radiative transfer.

With the advent of detailed spectra from medium-resolution space-based telescopes such as the JWST, we will soon have the ability to map the spatial distribution of chemical species in the atmospheres of hot gas giants. Ground-based high-resolution spectroscopy already allows us to measure chemical variations between morning and evening terminators (e.g., Ehrenreich et al. 2020; Kesseli & Snellen 2021; Kesseli et al. 2022). To comprehend these spatial distribution maps, we need 3D numerical models that can couple hydrodynamics, chemistry, and radiative transport. Only one GCM, the unified model (UM), cur-

rently possesses this capability (Drummond et al. 2020; Zamyatina et al. 2023). Lee et al. (2023) employed a faster chemical network, but did not incorporate a chemically consistent radiative transfer solver. However, these studies consistently indicate long runtimes.

One of the challenges of such models is the consistent coupling between chemical abundances and line opacities as used in the radiative transfer (Amundsen et al. 2017). Line opacities of molecular species in the pressure-temperature range of warm and hot gas giants often are a collection of millions of lines, which must be accurately accounted for also in low resolution (fast) radiative transfer. It has therefore been realized early on that rapid calculations would need some simplifications. One of these is the correlated-k method (Goody et al. 1989), which is similar to the method of opacity distribution functions (ODFs) introduced for stellar atmospheres (Gustafsson et al. 1975). The correlated-k method as well as the ODF method converts the wavelength-dependent opacity into distribution functions (called k-tables or ODFs) by sorting opacity values within spectral bins (see Sect. 2). This approach captures the dynamic range of the opacity and allows radiative transfer calculations with a small set of spectral bins (usually 5-30 for GCMs and a few hundred for spectra) and an accuracy of a few percent on bolomet-

Send offprint requests to: A. Schneider,  
e-mail: aaron.schneider@nbi.ku.dk  
or e-mail: aarondavid.schneider@kuleuven.be

ric fluxes (e.g., Amundsen et al. 2014; Leconte 2021; Schneider et al. 2022b).

The main drawback of the correlated-k method is the loss of all wavelength information in a spectral bin, when the opacity is converted to k-tables. Since the wavelength distribution of the opacity is depth-dependent, the loss of wavelength information will at first-order approximation cause erroneous optical depth calculation through the atmosphere, and hence erroneous radiative transfer computation and energy balance. This can be corrected for in statistical ways, as described in detail below using the RORR method (see Sect. 2.1). In the corresponding ODF scheme traditionally used in older stellar atmosphere computations, the problem is the same and was discussed for example in (Saxner & Gustafsson 1984), where it was concluded that the cost in increased computing time as function of the number of individual opacity species made it unfeasible to continue the ODF scheme for cooler stars. Newer stellar models are therefore usually computed based on the opacity sampling scheme, as discussed for example in (Jorgensen 1992; Gustafsson & Jorgensen 1994; Helling & Jorgensen 1998). It is, however, not obvious how one should treat the strong atomic lines correctly in the opacity sampling scheme, and if these are important in exoplanetary atmospheres alongside with the multitude of molecular lines, the correlated-k method may turn out superior, or a new hybrid method may be needed.

In order to still use the correlated-k method, GCMs often use premixed k-tables, which tabulate the k-tables for a given chemical mixture assuming that the gas is in chemical equilibrium and abundances can be constrained as a function of pressure and temperature alone (e.g., Showman et al. 2009; Lee et al. 2021; Schneider et al. 2022b; Deitrick et al. 2022). The k-table at a certain grid point is then recovered by interpolating on the pressure-temperature dependent premixed grid of k-tables. This approach, however, is not chemically correct when the gas is not in chemical equilibrium, which is the case if processes such as chemical kinetics or photochemistry are taken into account in the model.

The UM (Amundsen et al. 2016) is, to our knowledge, currently the only hot Jupiter GCM that can handle k-table mixing on the fly, without the need of premixed tables. Using a 1D radiative transfer code, Amundsen et al. (2017) benchmarked several numerical schemes that can approximate k-table mixing. Treating the k-table mixing on the fly rather than using pre-mixed tables, introduces more freedom in the radiative transfer computation, but for gasses with many opacity sources only if it can be performed sufficiently fast. We therefore extend upon Amundsen et al. (2017) in this paper, by introducing a new machine learning accelerated technique, and by coupling these methods to our GCM (expERT/MITgcm). This paper starts by introducing the correlated-k method and several approximate methods for k-table mixing in Sect. 2. We then discuss the setup of our GCM in Sect. 3, show the benchmarking results in Sect. 4 and a simple disequilibrium application in Sect. 5, where we apply our model in a simple rainout situation, where condensation of TiO and VO are approximated. We finally conclude and discuss the implications of this work in Sect. 6.

## 2. Mixing species

According to the Lambert Beer law, the total transmission  $\mathcal{T}$  of light passing through a homogeneous slab of gas with density  $\rho$  [kg m<sup>-3</sup>] opacity  $\kappa$  [m<sup>2</sup> kg<sup>-1</sup>] and thickness  $d$  [m] in the spectral

window from frequency  $\nu_0$  [Hz] to  $\nu_1$  [Hz] is given by

$$\mathcal{T} = \frac{1}{\nu_1 - \nu_0} \int_{\nu_0}^{\nu_1} \exp(-\kappa(\nu)\rho d) d\nu. \quad (1)$$

In the correlated-k method, the integral of Eq. 1 is solved by substituting  $\nu$ ,

$$\mathcal{T} = \frac{1}{\nu_1 - \nu_0} \int_0^1 \exp(-\kappa(g)\rho d) dg, \quad (2)$$

such that the integration over frequency is substituted by an integration over a new independent variable  $g$ . This independent variable  $g$  represents the cumulative opacity distribution function and is then given by the opacity distribution function  $f$ , such that

$$g(\kappa) = \int_0^\kappa f(\kappa') d\kappa'. \quad (3)$$

The cumulative opacity distribution function  $g$  can be understood as the probability to find an opacity value of less than  $\kappa$  at a specific frequency  $\nu$ . While both Eqs. 1 and 2 are formally identical, since they only differ in the order in which the sum is evaluated, they might differ significantly in the discrete limit, where individual summation points need to represent the value of the opacity for a certain non-zero width  $\Delta\nu$  or  $\Delta g$ . In practice, the correlated-k method divides the total computed frequency range into coarse frequency bins, in which all integrals over frequency are substituted into integrals over  $g$ . The radiative transfer equation for the intensity  $I(\nu, g)$  [W m<sup>-2</sup> sr<sup>-1</sup> Hz<sup>-1</sup>] of the coarse frequency bin between  $\nu_0$  and  $\nu_1$  at the discrete sub bin  $g$  in  $g$ -space then becomes

$$\mathbf{n} \cdot \nabla I(\nu, g) = \rho \kappa(\nu, g) [S(\nu, g) - I(\nu, g)], \quad (4)$$

where  $S(\nu, g)$  [W m<sup>-2</sup> sr<sup>-1</sup> Hz<sup>-1</sup>] is the source function and  $\mathbf{n}$  is the unit vector of the direction in which the intensity is measured. Thus, in each of these coarse frequency bins, the radiative transfer is solved individually for each  $g$  grid value and integrated over  $g$  similarly to Eq. 2 afterward to obtain the intensity  $I$  of the coarse frequency bin from  $\nu_0$  to  $\nu_1$ . In this way, the correlated-k method allows for rapid calculations by requiring less radiative transfer computations for the same level of accuracy.

To obtain the correct k-tables for a mixture of individual species, one would need to sum up the individual contributions of the individual species, weighted with their abundance, and construct the opacity distribution functions of the total opacity. Several methods have been put forward to solve this issue, and we will test some of these in this work. It is important to note that the opacity distribution functions vary as a function of pressure and temperature and mixture in the gas, thus, in order to use the correlated-k method, one requires methods to construct these  $\kappa(\nu, g)$  (k-table) values accurately.

### 2.1. Random overlap with rebinning and resorting (RORR)

Both Saxner & Gustafsson (1984) and Lacis & Oinas (1991) independently introduced a similar method for ODFs and k-tables respectively that treats the mixing of multiple opacity species under the assumption that the distributions of the individual opacity species are not correlated:

$$f_{\text{tot}}(\kappa_1, \dots, \kappa_{N_s}) = f(\kappa_1) \cdot \dots \cdot f(\kappa_{N_s}). \quad (5)$$

In simple terms, this means that, for example, the line cores are randomly distributed and do not systematically occur at the same

frequencies. In the case of correlated-k, this method is called the random overlap with rebinning and resorting (RORR) method. An in depth introduction to the RORR method can be found in Amundsen et al. (2017) and in Mollière et al. (2015), and we will instead just briefly outline its basic function.

The core of the RORR method is the assumption that the opacities of two species are uncorrelated. This will then imply that their transmissions are also uncorrelated (Mollière et al. 2015). From Eq. 1, we can then see that the transmission of both species can be multiplied to get the total transmission. It is then possible to find a k-table of a mixture by convolving their probability distributions (e.g., Mollière et al. 2015; Amundsen et al. 2017). The final result of the convolution calculation can be re-sorted and binned back to the original  $g$  grid for further computations. Repeating this procedure with the combined k-table of the two species and a third species yields the next step. This procedure is then repeated until all opacity species are included in the total k-table (Amundsen et al. (2017) provides a useful visualization of the procedure in their Fig. 1). From the methods outlined in this work, RORR is the slowest but most accurate method. Furthermore, RORR is well benchmarked against line-by-line calculations (e.g., Amundsen et al. 2014; Mollière et al. 2015). It is therefore the method of choice for most correlated-k 1D atmosphere models, such as ATMO (Tremblin et al. 2015), petitRADTRANS (Mollière et al. 2019) or PICASO (Mukherjee et al. 2023).

## 2.2. Premixed k-tables

Assuming equilibrium chemistry, one can create k-tables as lookup tables of pressure and temperature, which can be interpolated on during the radiative transfer calculations. It is important to note that these premixed tables are subject to the exact input to the equilibrium chemistry calculations (e.g., metallicity or C/O ratio) and need to be recomputed if the atmosphere is expected to deviate from these. Premixed tables can be computed in multiple ways. Showman et al. (2009) calculated premixed tables by calculating the distribution functions of the mixture. In Schneider et al. (2022b), we have calculated these k-tables using RORR on the individual k-tables. As already pointed out in Amundsen et al. (2017), the accuracy of this approach is subject to the pressure and temperature resolution of the lookup table, since the equilibrium abundances are expected to vary by many orders of magnitude, much more than the individual opacities themselves. In Schneider et al. (2022b), we have therefore computed the lookup tables, such that they match the pressure grid used in the GCM, removing the need to interpolate in pressure and allowing for a fine grid in temperature (1000 temperature points). While these lookup tables can be very precise, if resolved sufficiently, they come at the cost of flexibility, since they require assumptions on the abundances as a function of temperature and pressure.

## 2.3. Summation

In fact, the easiest way to calculate the mixed k-table from a mixture of different opacity species is to approximate the convolution by a sum:

$$\kappa_{\text{tot}}(\nu, g) = \sum_{i=1}^{N_s} \kappa_i(\nu, g), \quad (6)$$

where  $N_s$  is the number of species and the subscript  $\kappa_i$  is the individual k-table of species  $i$ , weighted with its mass mixing

ratio. While there is no logical justification for this approach, the approach is certainly the fastest method, as it only requires the evaluation of the sum of k-tables. This approach is certainly attractive, it will, however, naturally underestimate  $\kappa_{\text{tot}}$  at small  $g$  and overestimate  $\kappa_{\text{tot}}$  at large  $g$ . This can be best seen in RORR<sup>1</sup>, where the evaluation of the convolution would add  $\kappa_i$  values from larger  $g$  as well as those of smaller  $g$  to the  $\kappa_{\text{tot}}$  values at small  $g$ . This is particularly important, since the small  $g$  values, which encode the small  $\kappa_{\text{tot}}$  values, decide the depth up to which stellar irradiation can be absorbed.

## 2.4. Adaptive equivalent extinction

Adaptive equivalent extinction (AEE) is a variation of equivalent extinction Edwards & Slingo (1996) introduced by Amundsen et al. (2017). The idea is to determine the most important species and then treat all other species as gray within the spectral band. To obtain the major absorber, one first calculates semi-gray opacities for each species. The semi-gray opacity of species  $i$  is calculated as an average of the k-table in a given spectral bin and is weighted with a function  $w$  that depends on the value of  $g$ :

$$\kappa_{\text{av},i}(\nu) = \frac{\int_0^1 \kappa_i(\nu, g)w(\nu, g)dg}{\int_0^1 w(\nu, g)dg}. \quad (7)$$

We show below that the choice of the weighting function  $w$  has a significant impact on the accuracy of the AEE method. A good measure for the importance of the opacity at a given  $g$  value on the accuracy of the radiative transfer calculation is the magnitude of the flux at that value of  $g$  (Amundsen et al. 2017). We thus chose to use the absolute values of the stellar and planetary fluxes through a  $g$  value as a weighting function. However, since we can only know the fluxes, once we have already mixed the k-tables and calculated the radiative transfer, we need to rely on the value of the planetary and stellar flux from a previous radiative transfer calculation, which would be the previous radiative time step in the case of GCMs.

Using these  $\kappa_{\text{av},i}(\nu)$  values, a major absorber is found by vertically integrating the transmission (Eq. 1) from the top of the atmosphere down to an optical depth of one for each species. The first species to reach an optical depth of one is then used as the major absorber in the vertical column. The final total opacity in each spectral bin is then given as

$$\kappa_{\text{tot}}(\nu, g) = \kappa_m(\nu, g) + \sum_{i \neq m}^{N_s} \kappa_{\text{av},i}(\nu), \quad (8)$$

where  $m$  is the species, which has been identified as the major absorber. We note that the UM uses a slightly different and less sophisticated method (called equivalent extinction or EE), where the major absorber is determined locally and without integrating over the atmospheric column. A detailed introduction of the AEE method can be found in Amundsen et al. (2017).

## 2.5. DeepSet approach

The RORR mixing approach has three important attributes. Firstly, the method stays the same, no matter how many species are mixed with each other and in what order. RORR is thus invariant to permutations in the set of opacities that are to be

<sup>1</sup> for a visualization, see Fig. 1 of Amundsen et al. (2017)

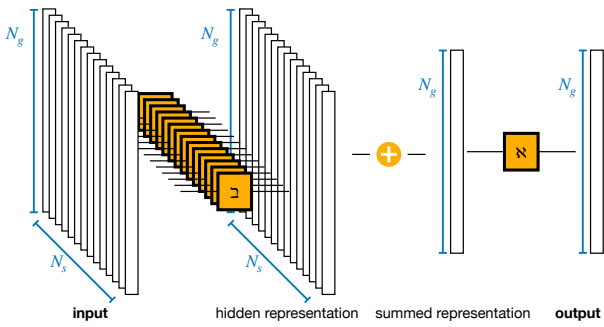


Fig. 1: The neural network used in the DeepSet (DS) approach.

mixed. Secondly, to a first approximation RORR can be approximated by a simple sum as mentioned in Sect. 2.3 and verified below. Lastly, although RORR is computationally expensive compared to the other outlined approaches, in its core for each frequency bin, it only consists of a convolution, a sorting step and an interpolation step, which is repeated  $N_s - 1$ -times. These three attributes greatly constrain a possible emulation of RORR by a machine learning algorithm, since only algorithms with fast inference, versatile input size and permutation invariance can be used.

We tested several architectures, such as a U-net like convolutional network (Ronneberger et al. 2015), gradient boosted regression trees using the abundances as input and the mixed opacities as output using XGboost (Chen & Guestrin 2016). With the U-net we ended up needing too many convolution blocks and with XGboost, we needed structures that were too deep and therefore too memory consuming to get reasonable accuracy. One of the reasons for the poor performance of these methods in our context is that they are too different from a simple sum (see Sect. 2.3). We have therefore settled with a DeepSet approach (Zaheer et al. 2017). A DeepSet for our case of k-table mixing can be written as

$$y = \mathfrak{N} \left( \bigoplus_{i=1}^{N_s} \mathfrak{z}(X_i) \right), \quad (9)$$

where  $y$  is the response of the DeepSet (the mixed k-tables, see below Eq. 13),  $X = \{X_1, \dots, X_{N_s}\}$  is the set of input vectors (the individual k-tables, see below Eq. 12), and  $\mathfrak{N}$  and  $\mathfrak{z}$  are functions.

Simply put, the idea is to use a function  $\mathfrak{z}$  to encode each vector  $X_i$  of a set into a hidden representation  $\mathfrak{z}(X_i)$ . These representations are then summed up<sup>2</sup>, and subsequently decoded by function  $\mathfrak{N}$  to get the output. We illustrate the concept of DeepSets (Zaheer et al. 2017) for the problem of k-table mixing in Fig. 1. We used

$$\mathfrak{z}(X_i) = \max(a_1 \cdot X_i, 0) \quad (10)$$

and

$$\mathfrak{N}(z) = a_2 \cdot z, \quad (11)$$

<sup>2</sup> Note that any permutation invariant operation could be used in this step.

with weights  $a_1$  and  $a_2$ , which are matrices of size  $N_g \times N_g$ . Whereas  $\mathfrak{N}$  is a linear function,  $\mathfrak{z}$  is non-linear due to the inclusion of a rectified linear activation (ReLU). Different, more complex functions could also be used, but we found that the accuracy reached by using these simple functions is already good enough. More complex functions, would thus only result in a slower performance. The weights  $a_1$  and  $a_2$  of the functions  $\mathfrak{N}$  and  $\mathfrak{z}$  are then learned by a neural network.

To train the neural network, we chose to implement the network in Keras (Chollet et al. 2015). The weights are learned by minimizing the mean squared error using the Adam optimizer (Kingma & Ba 2014). We then performed a Bayesian hyperparameter search using hyperopt (Bergstra et al. 2012) to find the best amount of features for the hidden representation and to determine the optimal learning rate of the optimizer. The loss did not improve significantly by the use of more than  $N_g$  features in the hidden representation, and we therefore chose to use  $N_g$  features. Additionally, we found that a learning rate of  $\alpha = 1 \times 10^{-3}$  seems to perform best.

The neural network acts on each frequency bin individually and therefore does not care about the frequency resolution. To generate the training and test data, we computed the mixed k-tables with RORR from  $\approx 8 \times 10^5$  sets of 14 k-tables each. These 14 k-tables were taken from the 11-bin resolution (S1) of the 14 opacity species taken into account in expert/MITgcm Schneider et al. (2022b) (see Sec. 3 for details on the species) and were uniformly randomly weighted with reasonable abundance ranges.<sup>3</sup> The input for the network is not the plain individual k-tables, but instead we scale them with the sum of the k-tables as

$$X_i(\nu, g) = \log \left( \frac{\kappa_i(\nu, g)}{\sum_{j=1}^{N_s} \kappa_j(\nu, g)} \right). \quad (12)$$

Similarly, we scale the targets (e.g., the predictions of the network) by

$$y(\nu, g) = \log \left( \frac{\kappa_{\text{tot}}(\nu, g)}{\sum_{j=1}^{N_s} \kappa_j(\nu, g)} \right). \quad (13)$$

The advantage of this input and output scaling comes three-fold. Foremost, we are interested in minimizing the error of the small values in the k-table, as those are the ones that generate windows in the spectrum, which are very important for the temperature structure, hence taking the logarithm is useful because it pronounces relatively large deviations of small k-table values in the loss. Secondly, we find that the sum of the k-tables is already a reasonable approximation for the mixed k-tables (see Sec. 4). Thirdly, the output scaling nicely captures the positivity of the problem, prohibiting the possibility of negative predictions when reversing the output scaling. It is also important to stress that all species  $\kappa_i(g)$  are fed individually through the same function  $\mathfrak{z}$  (with the same weight) to create a unique non-linear representation for each species.

The main advantage of the DeepSet method is its flexibility, because trained on the individual frequency bins of k-tables, it operates independent of chemical composition or opacity species. It can by construction of the training set operate on any composition and metallicity. Furthermore, we think that it can in principle work on any set of k-tables with shapes that are similar enough to those of the training set. We have tested this by changing the frequency resolution of the training set. Doing so,

<sup>3</sup> We note that the network can, by construction, deal with any number of species without retraining needed

we found that this did not significantly affect the accuracy, when applied to a different frequency resolution than the one being trained on. We therefore think that it would be only necessary to retrain the network, if the discretization of  $g$  values changes. We discuss further numerical considerations of this mixing method in Appendix A.

### 3. Methods

To test the individual mixing methods, we use the 3D GCM `expeRT/MITgcm` (Carone et al. 2020; Schneider et al. 2022b). `expeRT/MITgcm` builds on the dynamical core of the `MITgcm` (Adcroft et al. 1997, 2004), which solves the hydrostatic primitive equations of hydrodynamics on a cubed-sphere grid. In order to accurately account for radiative heating and cooling, `expeRT/MITgcm` solves the radiative transfer using the Feautrier method (Feautrier 1964) with approximate Lambda iteration (Olson et al. 1986) and Ng-acceleration (Ng 1974). The routine that solves the radiative transfer is an adapted version of the flux routine in `petitRADTRANS` (Mollière et al. 2019, 2020). We have incorporated the radiative transfer solver and benchmarked it in `expeRT/MITgcm`. We found in Schneider et al. (2022b), that the combination of 5 frequency bins and 16  $g$  values achieves good enough accuracy for long term convergence studies such as those in Schneider et al. (2022a) and Sainsbury-Martinez et al. (2023). In this work, we use 11 frequency bins, with each 16  $g$  values, which is good enough for comparisons of GCMs to observations. We note that other hot Jupiter GCMs typically use 8  $g$  values, with sometimes a higher frequency resolution of  $\approx 30$  frequency bins (Showman et al. 2009; Amundsen et al. 2016; Lee et al. 2021). The line opacity species used in this work are identical to the ones used in Schneider et al. (2022b) and include  $\text{H}_2\text{O}$  (Polyansky et al. 2018),  $\text{CO}_2$  (Yurchenko et al. 2020),  $\text{CH}_4$  (Yurchenko et al. 2017),  $\text{NH}_3$  (Coles et al. 2019),  $\text{CO}$  (Li et al. 2015),  $\text{H}_2\text{S}$  (Azzam et al. 2016),  $\text{HCN}$  (Barber et al. 2014),  $\text{PH}_3$  (Sousa-Silva et al. 2015),  $\text{TiO}$  (McKemmish et al. 2019),  $\text{VO}$  (McKemmish et al. 2016),  $\text{FeH}$  (Wende et al. 2010),  $\text{Na}$  (Piskunov et al. 1995), and  $\text{K}$  (Piskunov et al. 1995).

We chose to use HD 209458 b as a planet and the setup is identical to the setup in Schneider et al. (2022b), where the models in this work only differ by the method with which opacities are mixed. The different mixing methods and their corresponding labels are laid out in Table 1. To this end, we have implemented each of the above-mentioned mixing methods. `expeRT/MITgcm` can now run in two modes by either mixing k-tables on the fly (utilizing one of the aforementioned methods) or using premixed k-tables. To incorporate on-the-fly mixing in the GCM, we updated our preprocessing toolkit to additionally output a pressure-temperature grid of equilibrium abundances (taken from the `easyCHEM` (Mollière et al. 2017) interface to `petitRADTRANS` (Mollière et al. 2019)), along with a pressure-temperature grid of k-tables for the individual absorbers. In the on-fly mixing mode, abundances and k-tables of each of the considered absorbers are linearly interpolated to the pressure and temperature in the GCM, weighted by their abundance and then mixed by one of the above-mentioned methods.

The weighting in the adaptive equivalent extinction method induces the need for more scattering iterations, because the k-table becomes dependent on the bolometric flux from the previous time-step, effectively inserting a time dependent perturbation into the opacities, because the scattering source function will be subject to these opacity perturbations as well, rendering its guess from the previous time-step less accurate. For performance reasons, we thus found that the weighted adaptive equivalent extinc-

tion method required us to limit the amount of maximum scattering iterations per radiative time step to two instead of 500, which is generally enough for a HD 209458 b like planet with only Rayleigh scattering, since the source function is reused as initial guess in the next radiative time-step (see Schneider et al. 2022b, for an explanation of scattering in `expeRT/MITgcm`) and the source function is thereby naturally iterated on during model convergence. To be consistent in all models, we have thus chosen to generally limit the amount of scattering iterations per radiative time-step to two, if not otherwise stated.

All models have been integrated up to 2000 days with a radiative time-step of 100 s and a dynamical time-step of 25 s, which are typical values for hot Jupiter GCMs (e.g., Showman et al. 2009; Lee et al. 2021; Schneider et al. 2022b). All models use equilibrium chemistry to constrain the abundances. In practice, the code interpolates the abundances on a grid of pressure and temperature. In Section 5, we show a model, where we use the DeepSet mixing of k-tables, but removed all of  $\text{TiO}$  and  $\text{VO}$ , whenever  $\text{TiO}$  and  $\text{VO}$  would reappear in equilibrium chemistry in the gas phase, although it is condensed out further down in the atmosphere. This method is similar to the methods of rainout described elsewhere (e.g., Lodders & Fegley 2002; Marley et al. 2021), but less sophisticated compared to 3D models that include proper chemical transport schemes (e.g., Parmentier et al. 2013; Lee et al. 2023; Drummond et al. 2018). A more detailed description of the algorithm for the detection of rainout is outlined in Appendix C.

### 4. Results

In this work, we compare the mixing methods introduced in Sect. 2 to the slow RORR method. An additional comparison between the premixed method (PRE) used in our GCM, as introduced and used in Schneider et al. (2022b), and the RORR method can be found in Appendix B. Mixing opacities on the fly in a GCM requires a tradeoff between accuracy, performance, and flexibility. When comparing the accuracy of mixing methods, it is important to keep in mind that the low spectral resolution of the k-tables used in GCMs induces an error of a few percent on bolometric fluxes (e.g., Amundsen et al. 2014; Leconte 2021; Schneider et al. 2022b). It is therefore pointless to aim for accuracies of less than one percent, since the overall error will be governed by the chosen spectral resolution. Comparing the individual mixing methods with each other thus needs to consider all of these perspectives. In order to have a fair comparison between the individual methods, we chose to compare all simulations in two aspects: The accuracy on the resulting atmospheric state and the computational time. We include further diagnostics of the accuracy in Appendix D, by looking at fluxes and heating rates.

In order to qualitatively compare the mixing methods, we show the atmospheric state at two pressure layers (0.01 bar and 0.1 bar) together with the zonal wind speed in Fig. 2. Aside from the adaptive equivalent extinction (AEE) model without flux weighting, all models look similar at a first glimpse. The jet strength and day-night temperature contrasts are not significantly affected by these different methods. However, the AEE method seems to produce a significantly higher day-night temperature contrast and a faster jet. These differences are more pronounced at lower pressure, but are still notable at higher pressures of 0.1 bar.

Since most methods produce the same qualitative trend in winds and temperatures, it can be useful to calculate spatial temperature averages and compare those to the RORR models. For

Table 1: Simulations

label	mixing method	reference
RORR	random overlap with rebinning and resorting (RORR)	Sect. 2.1
PRE	premixed k-tables	Sect. 2.2
DS	DeepSet approach	Sect. 2.5
AEE_we	adaptive equivalent extinction with flux weighting	Sect. 2.4
AEE	adaptive equivalent extinction without flux weighting	Sect. 2.4
ADD	sum of all k-tables	Sect. 2.3

**Notes:** Explanation of the mixing methods used in the individual simulations, as labeled in the figures of this work.

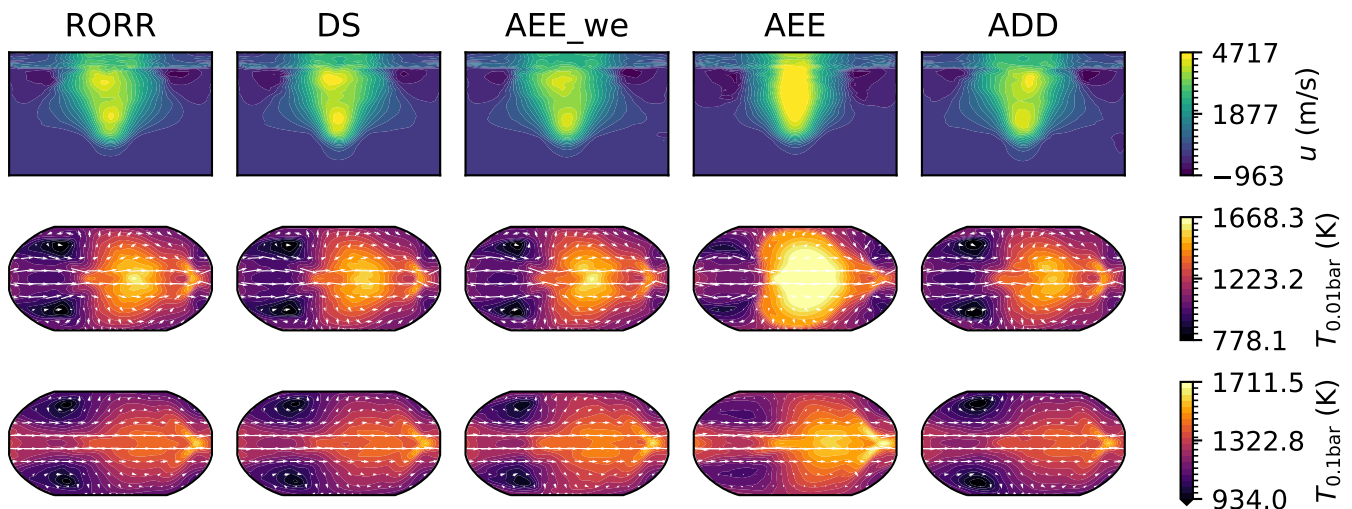


Fig. 2: Zonally averaged eastward winds (first row) and temperature slices at 0.01 bar (second row) and 0.1 bar (third row) for the different models considered. All colors are normalized to the first column (RORR).

visual reasons, we have split these comparisons into two figures, where Fig. 3 compares the DeepSet (DS) and summation (ADD) method to the RORR method and Fig. 4 compares both the weighted and non-weighted adaptive equivalent extinction methods to the RORR method. It was surprising to see that the ADD method performs well, given its simplicity and methodological flaws. However, the temperature is often slightly cooler at higher pressures, which might be related to the underestimation of  $\kappa_{\text{tot}}$  at small  $g$  and the overestimation at large  $g$ . The overestimation of  $\kappa_{\text{tot}}$  at large  $g$  can lead to an enhancement of the absorption of the stellar flux in the upper layers, which cannot penetrate deep enough to cause heating in the deeper layers. These flaws of the ADD method do not seem to persist in the DS mixing, which uses the ADD method in its preprocessing (see Sect. 2.5). Instead, we find that the DS mixing performs very well.

Looking at the adaptive equivalent mixing method with and without weighting, we find that the AEE\_we method is almost as accurate as the DS mixing. The weighting certainly helps to find a good estimate of the major absorber and drastically increases the accuracy of this approach. Looking at the residuals, we see a strong correlation between the error of the weighted and non-weighted method, which points to a general issue of the method instead of an issue with the major absorber. Unlike in the case of simply summing up k-tables, the AEE method tends to not overestimate  $\kappa_{\text{tot}}$  at large  $g$  but instead to underestimate it. This might be explained by the minor absorbers, which flatten a k-table by offsetting  $\kappa_{\text{tot}}$  at the small  $g$  and decreasing the impact of the high  $g$  values. Similar to the overestimation of  $\kappa_{\text{tot}}$  at small

Table 2: Runtime

label	time [h]	relative to PRE
PRE	3.91	1.00
ADD	7.20	1.84
AEE	7.44	1.90
AEE_we	13.37	3.42
DS	10.88	2.78
RORR	25.12	6.43

**Notes:** Runtime of the GCM needed for the first 100 days of the simulation. The node used to run the model includes 2x Intel(R) Xeon(R) Gold 6248R CPU @ 3.00GHz, and we utilized all 48 cores.

$g$  values, the underestimation of  $\kappa_{\text{tot}}$  at large  $g$  values also shifts the location at which irradiation is absorbed, in this case into the opposite direction by leading to less absorption in the uppermost layers.

We thus conclude that both the ADD method and AEE method, with and without weighting, introduce systematic noise to  $\kappa_{\text{tot}}$  at both small and large  $g$  values. This noise can amplify the errors of the AEE and ADD methods (see Appendix D). In contrast, the DS method exhibits no systematic error but instead uniformly distributed random noise. The overall error in temperature estimation is thereby not significantly effected. Therefore, we do not recommend using the ADD method or unweighted AEE method and instead recommend the use of the DS method.



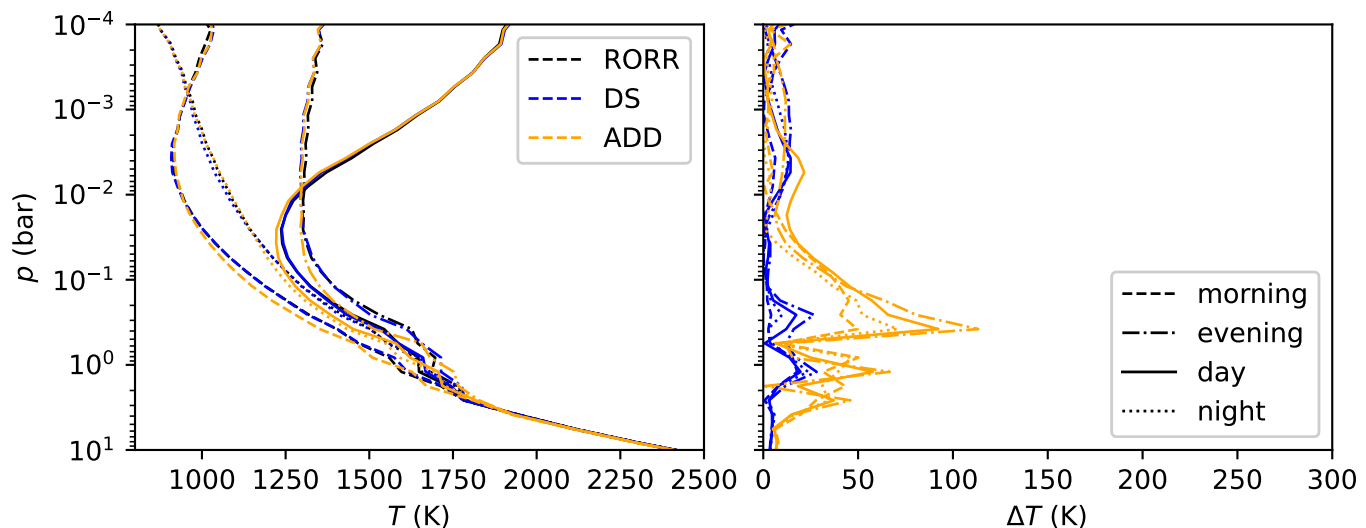


Fig. 3: Temperature profiles for different parts of the atmosphere of the RORR, DS and ADD simulations from Fig. 2. The different colors represent different simulations, whereas the different line-types represent different parts of the atmosphere. The left panel shows the temperature profile and the right panel shows the absolute difference between the temperature profile to the temperature profile of the RORR simulation.

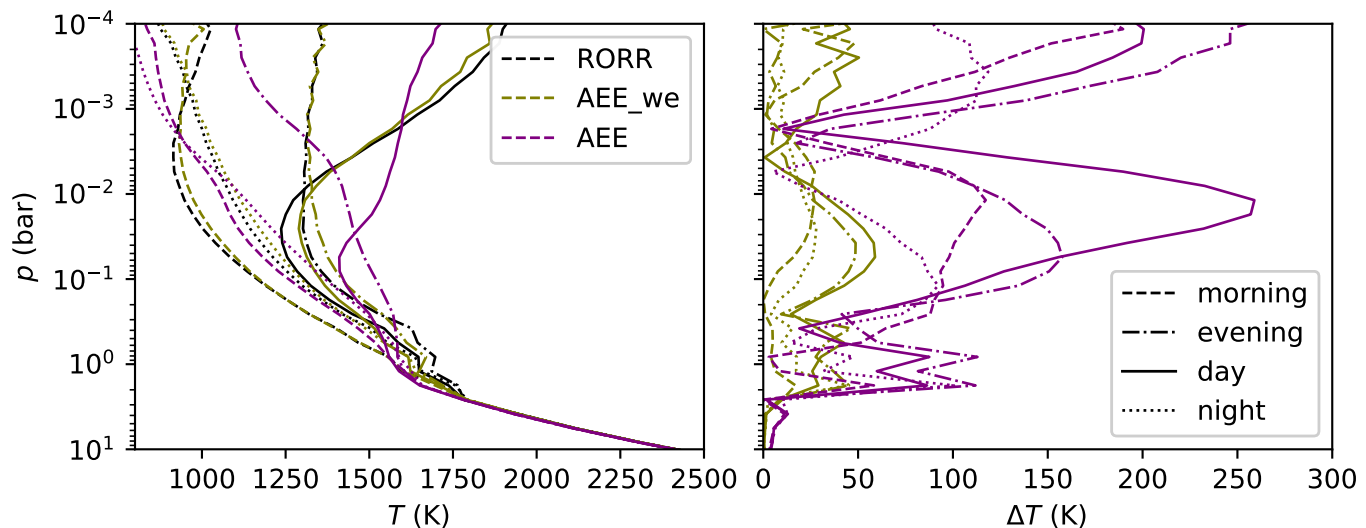


Fig. 4: Temperature profiles for different parts of the atmosphere of the RORR, AEE and weighted AEE (AEE\_we) simulations from Fig. 2. The panels and their meaning are identical to Fig. 3.

In terms of computational costs, one needs to consider two general computational overheads during runtime, compared to using a premixed grid. The first overhead comes from the handling of the individual  $k$ -tables, such as the interpolation, as compared to handling of just one premixed  $k$ -table. Secondly, obviously the computation of the mixing itself. We show the computation time needed to run the initial 100 days of the simulation in Table 2. The performance of the summation (ADD) method, in which individual  $k$ -tables are simply added up, is mainly constrained by the handling of the individual  $k$ -tables, since the cost of the summation can be neglected, whereas all the other methods are also subject to the computational cost of the mixing. Therefore, a significant fraction of the computational cost in the AEE, AEE\_we and DS simulations can be explained by the overhead of handling individual  $k$ -tables.

The slow runtime of the RORR mixing, renders the on-the-fly use of the RORR technique impossible. Even when using

an optimized sorting algorithm, which is now the standard in `petitRADTRANS`, and which significantly speeds up RORR, the RORR technique performs at least 6 times slower compared to the premixed case. One of the main reasons for the poor performance of RORR in our specific setup is the quadratic computational dependence on the number of  $g$  values. GCMs that use 8  $g$  values instead of 16 could therefore (at the cost of accuracy) have faster performance of the RORR method (for a discussion see Amundsen et al. 2017).

Due to the increased amount of scattering iterations, the adaptive equivalent extinction with weighting (AEE\_we) is generally the slowest of all the approaches. The poor convergence behavior of the AEE\_we method makes this method less reliable and less performant. This will be an even bigger issue, when scattering becomes non-negligible, which could be the case for lower temperatures or if clouds and hazes are included. However, when limiting the maximum amount of scattering iterations to

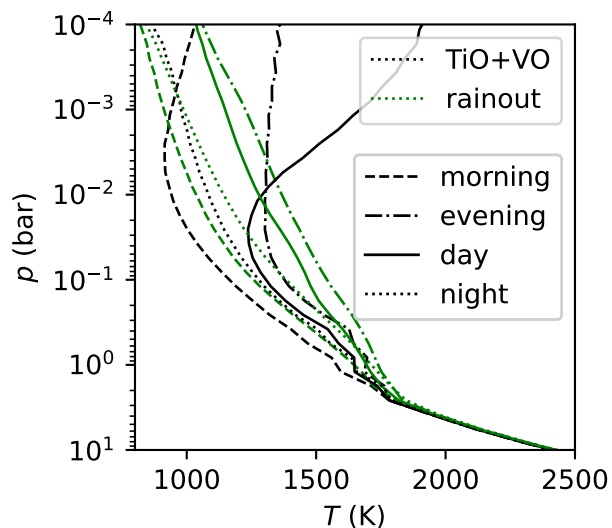


Fig. 5: The RORR model from Fig. 2 compared to a model with rainout of TiO and VO. The stratosphere disappears if rainout is included in the model.

two, as we have done here, it is fast enough to compete with the other methods and if the weighting were to be neglected (AEE), it would even be similarly fast as the ADD method and faster than the DS mixing. Performance-wise, we thus think that either of the ADD, AEE (with and without flux weighing), or DS mixing approach could be equally used in a setup that requires on-the-fly mixing. In our multi-stream setup, which requires us to iterate over the source function, we would however not recommend the use of the AEE<sub>we</sub> method, due to its numerical issues with the convergence of the source function.

## 5. Rainout

On-the-fly mixing is only relevant in GCMs if the chemical abundances are to be changed from chemical equilibrium. A sufficiently well resolved premixed table, will in most cases of chemical equilibrium be the fastest and most reliable method of choice. However, in the case of chemical disequilibrium, which is the case if photochemistry or chemical kinetics were to be considered, we can not simply premix k-tables, because the abundances of the relevant opacity species need to be changed during runtime. Another very simple scenario for such a situation could be the rainout of heavy refractory species such as TiO and VO. Due to their strong absorption in the UV, they absorb a significant fraction of the stellar flux at high altitude, thus heating the upper atmosphere significantly. Such a strong heating in the upper atmosphere can lead to a thermal inversion, where the atmosphere becomes hotter towards the top (e.g., Showman et al. 2009). However, due to vertical mixing and advection, condensed TiO and VO could gravitationally settle and therefore not be available in the gas phase at higher altitudes (e.g., Parmentier et al. 2013).

In Fig. 5 we show a model, in which all TiO and VO is removed by a simple rainout prescription (see Appendix C for details), which removes all TiO and VO from the gas phase, if it is condensed further down in the atmosphere. Using a premixed table, we would not be able to calculate the radiative effect of the rainout on the atmospheric structure, however, by using on-the-fly mixing we can trace the effect of the change in chemical abundances on the temperature. The lack of the strong

UV absorbers TiO and VO in the upper layers means that the upper layers get cooler, because less of the stellar flux is absorbed in those layers, whereas the intermediate pressure layers get warmer, where the bulk of the stellar flux is absorbed instead. As expected, we find that the thermal inversion caused by TiO and VO, as seen for example in Fig. 3, is self-consistently removed. By advection, this tendency swaps over to the night side.

## 6. Discussion and Conclusion

The correlated-k method is a useful approximation for rapid radiative transfer calculations with accuracies of a few percent (e.g., Amundsen et al. 2014; Leconte 2021; Schneider et al. 2022b), when used with resolutions typical for GCMs. We have demonstrated the performance and accuracy of several methods that could be used in GCMs to calculate the total opacity in the correlated-k assumption. We extended the work of Amundsen et al. (2017), who performed a similar analysis for the adaptive equivalent extinction (AEE) method. Furthermore, we have introduced two additional methods: The DeepSet (DS) mixing and a simple method in which k-tables are simply summed up (ADD). Whereas the work of Amundsen et al. (2017) only considered the accuracy of heating rates and fluxes, we incorporated the RORR, ADD, DS and AEE mixing methods into our GCM to demonstrate the performance in a real application. The DeepSet method turns out to be fairly accurate and flexible, leveraging machine learning, to calculate k-tables of gas mixtures. Overall, we find that

1. The random overlap with resorting and rebinning (RORR) method is too slow to be used in GCMs for on-the-fly mixing.
2. The AEE method requires a proper weighting to be accurate. Such a weighting, however, affects the numerical stability of the radiative transfer calculation, which will be especially important if scattering is non-negligible.
3. The ADD method and the AEE method are prone to systematic errors. This is especially problematic for the unweighted AEE and the ADD method, rendering a use of these methods questionable.
4. The DS method has minor statistical errors on fluxes and heating rates that do not seem to enhance the overall error, which seems to be an advantage of the DS method.

The DS mixing method provided by this work is accurate and open source<sup>4</sup>, and can be easily implemented in any radiative transfer package with no need to use complex libraries, as it only requires two matrix multiplications. Once trained, the network can perform on any composition and any set of opacity species. Although not strictly needed, we recommend training the network for a specific frequency resolution to maximize the accuracy. The amount of training data needed is small, and training can be performed within minutes on a standard personal computer. The provided open source package currently works with binned down `petitRADTRANS`-format k-tables, but can be easily extended to load any k-table format, and documentation is provided for how to achieve this.

The methods tested in this work, have been tested in terms of accuracy on the atmospheric structure and we think that these methods provided here will be key, when moving forward towards self-consistent transport of chemicals in the atmospheres of planets. We note, however, that these methods, do not translate to models that predict spectra. Future work is thus needed to test, if similar methods could also be used for atmospheric retrievals.

<sup>4</sup> <https://github.com/AaronDavidSchneider/opacmixer>



By implementing a simple chemical rainout procedure, we mimic the gravitational settling of TiO and VO, to demonstrate the DS method in a use-case of disequilibrium chemistry. By accounting for rainout in this way, we find that TiO and VO can be trapped in the deeper atmosphere, thus hindering the formation of a stratosphere. We note, that this approach is fairly simplified, and hope that our work enables future models to treat cold trapping self-consistently.

*Acknowledgements.* A.D.S., L.D., U.G.J. and C.H. acknowledge funding from the European Union H2020-MSCA-ITN-2019 under Grant no. 860470 (CHAMELEON). U.G.J. acknowledges funding from the Novo Nordisk Foundation Interdisciplinary Synergy Program grant no. NNF19OC0057374. The bibliography of this publication has been typeset using `bibmanager` (Cubillos 2020)<sup>5</sup>. The post-processing of GCM data has been performed with `gcm-toolkit` (Schneider et al. 2022)<sup>6</sup>

## References

- Aderoft, A., Campin, J.-M., Hill, C., & Marshall, J. 2004, *Monthly Weather Review*, 132, 2845
- Aderoft, A., Hill, C., & Marshall, J. 1997, *Monthly Weather Review*, 125, 2293
- Amundsen, D. S., Baraffe, I., Tremblin, P., et al. 2014, *A&A*, 564, A59
- Amundsen, D. S., Mayne, N. J., Baraffe, I., et al. 2016, *A&A*, 595, A36
- Amundsen, D. S., Tremblin, P., Manners, J., Baraffe, I., & Mayne, N. J. 2017, *A&A*, 598, A97
- Azzam, A. A. A., Tennyson, J., Yurchenko, S. N., & Naumenko, O. V. 2016, *MNRAS*, 460, 4063
- Barber, R. J., Strange, J. K., Hill, C., et al. 2014, *MNRAS*, 437, 1828
- Bergstra, J., Yamins, D., & Cox, D. D. 2012, arXiv e-prints, arXiv:1209.5111
- Carone, L., Baeyens, R., Mollière, P., et al. 2020, *MNRAS*, 496, 3582
- Chen, T. & Guestrin, C. 2016, arXiv e-prints, arXiv:1603.02754
- Chollet, F. et al. 2015, *Keras*, <https://keras.io>
- Coles, P. A., Yurchenko, S. N., & Tennyson, J. 2019, *MNRAS*, 490, 4638
- Cubillos, P. E. 2020, *bibmanager*: A BibTeX manager for LaTeX projects, Zenodo
- Deitrick, R., Heng, K., Schroffenegger, U., et al. 2022, *MNRAS*, 512, 3759
- Drummond, B., Hébrard, E., Mayne, N. J., et al. 2020, *A&A*, 636, A68
- Drummond, B., Mayne, N. J., Manners, J., et al. 2018, *ApJ*, 855, L31
- Edwards, J. M. & Slingo, A. 1996, *Quarterly Journal of the Royal Meteorological Society*, 122, 689
- Ehrenreich, D., Lovis, C., Allart, R., et al. 2020, *Nature*, 580, 597
- Feautrier, P. 1964, *Comptes Rendus Academie des Sciences (serie non specifiée)*, 258, 3189
- Goody, R., West, R., Chen, L., & Crisp, D. 1989, *J. Quant. Spectr. Rad. Transf.*, 42, 539
- Gustafsson, B., Bell, R. A., Eriksson, K., & Nordlund, A. 1975, *A&A*, 42, 407
- Gustafsson, B. & Jørgensen, U. G. 1994, *A&A Rev.*, 6, 19
- Helling, C. & Jørgensen, U. G. 1998, *A&A*, 337, 477
- Jørgensen, U. G. 1992, *Rev. Mexicana Astron. Astrofis.*, 23, 195
- Kesseli, A. Y. & Snellen, I. A. G. 2021, *ApJ*, 908, L17
- Kesseli, A. Y., Snellen, I. A. G., Casasayas-Barris, N., Mollière, P., & Sánchez-López, A. 2022, *AJ*, 163, 107
- Kingma, D. P. & Ba, J. 2014, arXiv e-prints, arXiv:1412.6980
- Lacis, A. A. & Oinas, V. 1991, *J. Geophys. Res.*, 96, 9027
- Leconte, J. 2021, *A&A*, 645, A20
- Lee, E. K. H., Parmentier, V., Hammond, M., et al. 2021, *MNRAS*, 506, 2695
- Lee, E. K. H., Tsai, S.-M., Hammond, M., & Tan, X. 2023, *A&A*, 672, A110
- Li, G., Gordon, I. E., Rothman, L. S., et al. 2015, *ApJS*, 216, 15
- Lodders, K. & Fegley, B. 2002, *Icarus*, 155, 393
- Marley, M. S., Saumon, D., Visscher, C., et al. 2021, *ApJ*, 920, 85
- McKemmish, L. K., Masseron, T., Hoeijmakers, H. J., et al. 2019, *MNRAS*, 488, 2836
- McKemmish, L. K., Yurchenko, S. N., & Tennyson, J. 2016, *MNRAS*, 463, 771
- Mollière, P., Stolker, T., Lacour, S., et al. 2020, *A&A*, 640, A131
- Mollière, P., van Boekel, R., Bouwman, J., et al. 2017, *A&A*, 600, A10
- Mollière, P., van Boekel, R., Dullemond, C., Henning, T., & Mordasini, C. 2015, *ApJ*, 813, 47
- Mollière, P., Wardenier, J. P., van Boekel, R., et al. 2019, *A&A*, 627, A67
- Mukherjee, S., Batalha, N. E., Fortney, J. J., & Marley, M. S. 2023, *ApJ*, 942, 71
- Ng, K. C. 1974, *J. Chem. Phys.*, 61, 2680
- Olson, G. L., Auer, L. H., & Buchler, J. R. 1986, *J. Quant. Spectr. Rad. Transf.*, 35, 431
- Parmentier, V., Showman, A. P., & Lian, Y. 2013, *A&A*, 558, A91
- Piskunov, N. E., Kupka, F., Ryabchikova, T. A., Weiss, W. W., & Jeffery, C. S. 1995, *A&AS*, 112, 525
- Polyansky, O. L., Kyuberis, A. A., Zobov, N. F., et al. 2018, *MNRAS*, 480, 2597
- Ronneberger, O., Fischer, P., & Brox, T. 2015, arXiv e-prints, arXiv:1505.04597
- Sainsbury-Martinez, F., Tremblin, P., Schneider, A. D., et al. 2023, *MNRAS*, 524, 1316
- Saxner, M. & Gustafsson, B. 1984, *A&A*, 140, 334
- Schneider, A. D., Baeyens, R., & Kiefer, S. 2022, `gcm_toolkit`, tools for GCM postprocessing
- Schneider, A. D., Carone, L., Decin, L., Jørgensen, U. G., & Helling, C. 2022a, *A&A*, 666, L11
- Schneider, A. D., Carone, L., Decin, L., et al. 2022b, *A&A*, 664, A56
- Showman, A. P., Fortney, J. J., Lian, Y., et al. 2009, *ApJ*, 699, 564
- Showman, A. P., Tan, X., & Parmentier, V. 2020, *Space Sci. Rev.*, 216, 139
- Sousa-Silva, C., Al-Refai, A. F., Tennyson, J., & Yurchenko, S. N. 2015, *MNRAS*, 446, 2337
- Tremblin, P., Amundsen, D. S., Mourier, P., et al. 2015, *ApJ*, 804, L17
- Wende, S., Reiners, A., Seifahrt, A., & Bernath, P. F. 2010, *A&A*, 523, A58
- Yurchenko, S. N., Amundsen, D. S., Tennyson, J., & Waldmann, I. P. 2017, *A&A*, 605, A95
- Yurchenko, S. N., Mellor, T. M., Freedman, R. S., & Tennyson, J. 2020, *MNRAS*, 496, 5282
- Zaheer, M., Kottur, S., Ravanbakhsh, S., et al. 2017, arXiv e-prints, arXiv:1703.06114
- Zamyatina, M., Hébrard, E., Drummond, B., et al. 2023, *MNRAS*, 519, 3129

<sup>5</sup> <https://bibmanager.readthedocs.io/en/latest/>

<sup>6</sup> <https://gcm-toolkit.readthedocs.io/en/latest/>

## Appendix A: Numerical considerations for the DeepSet implementation

One of the advantages of the DeepSet is its simplicity, which allows an easy naive implementation, since it only requires two matrix multiplication operations and one summation. There are several considerations regarding the performance of this approach. The first consideration is the use of the logarithm and, to reverse back to k-tables, the use of the exponential function. Both of these operations are unfortunately quite slow, albeit being needed as described in Sect. 2.5. Profiling the FORTRAN code used in the GCM shows that approximately half of the CPU time is spent on calculating the logarithm for the input scaling<sup>7</sup>. The other half of the CPU time is spent on the matrix multiplication, which can be optimized using hardware dependent compiler optimizations. For our architecture (using `ifort` on an intel CPU), we found that the build in `MATMUL` delivers the best performance compared to `MKL DGEMM`, `OpenBLAS DGEMM` and a manual implementation. The reason for the missing performance increase from the highly optimized `DGEMM` can most likely be found in the overhead of calling `DGEMM` compared to the small size of the matrix that is to be multiplied ( $16 \times 16$ ).

We note that faster and more efficient implementations could also be achieved by stacking the computations and deploying the computations to a GPU.

## Appendix B: Benchmarking RORR against PRE

In this work, we benchmarked several approximate k-table mixing methods to the RORR method (see Sect. 2.1). As discussed in Sect. 2.2, if the atmosphere is in chemical equilibrium, the abundances can be directly constrained as a function of pressure and temperature. It is then possible to use these abundances to create a grid of k-tables using the RORR method. This step can be done as a preprocessing step, so that the grid can be used as a lookup table to interpolate on during runtime. As mentioned in Amundsen et al. (2017), the accuracy of this approach will be greatly dependent on the resolution of the grid. We match the pressure coordinates with our premixed tables in `expeRT/MITgcm`, removing the need for interpolation in pressure and thereby resulting in a higher accuracy and faster runtime. We then use 1000 temperature grid points in this work and in Schneider et al. (2022b) to further maximize the accuracy.

In Fig. B.1 we show the temperature pressure profiles of the RORR simulation compared to the premixed (PRE) simulation. We find that both methods result in temperature profiles that overlap very well, with an error of less than 0.5%. Premixed tables will therefore stay the method of choice for setups, where equilibrium chemistry can be assumed, given their much faster performance (see Table 2).

## Appendix C: A simple algorithm to calculate rainout

In this work, we consider the rainout of TiO and VO in order to demonstrate a possible use case of the DS method (Sect. 5). We consider a species rained out, if, going upwards from the bottom of the computational domain, the species has, in local chemical equilibrium, first become available in the gas phase and has then disappeared further up in the atmosphere. If the species became more abundant in the gas in local chemical equilibrium, we would not allow this and instead keep the fixed abundance from

<sup>7</sup> The input scaling needs to be calculated  $N_s$  times more often than the reverse output scaling, see Fig. 1

---

**Algorithm 1** Procedure that takes in the pressure layers  $p$  (sorted from bottom to top) and mass mixing ratios  $\eta$  of an opacity species and returns the altered mass mixing ratios, under consideration of rainout.

---

```

procedure DO_RAINOUT( $p, \eta$ )
   $\tilde{\eta} \leftarrow \text{copy}(\eta)$ 
  species_appearing  $\leftarrow$  False
  species_cond  $\leftarrow$  False
  for  $i \leftarrow 2$  to  $\text{length}(p) - 1$  do  $\triangleright$  Loop from bottom to top
    grad  $\leftarrow -0.5 \times \left( \frac{\eta_{i-1} - \eta_i}{p_{i-1} - p_i} + \frac{\eta_i - \eta_{i+1}}{p_i - p_{i+1}} \right)$ 
    if (grad > 0) then  $\triangleright$  abundance increases
      species_appearing  $\leftarrow$  True
    end if
    if (species_appearing) then
      if (grad < 0) then  $\triangleright$  abundance decreases again
        species_cond  $\leftarrow$  True
      end if
    end if
    if (species_cond) then  $\triangleright$  Do the rainout
      if ( $\tilde{\eta}_{i-1} < \tilde{\eta}_i$ ) then
         $\tilde{\eta}_i \leftarrow \tilde{\eta}_{i-1}$ 
      end if
    end if
  end for
  if (species_cond) then  $\triangleright$  Treat the boundary
    if ( $\tilde{\eta}_{-2} < \tilde{\eta}_{-1}$ ) then
       $\tilde{\eta}_{-1} \leftarrow \tilde{\eta}_{-2}$ 
    end if
  end if
  return  $\tilde{\eta}$ 
end procedure

```

---

the layers below. These calculations are in practice performed by looking at the gradient of the mass mixing ratios, and we outline the algorithm used in this work in Algorithm 1. We would like to note that the algorithm presented here is likely oversimplified. Nevertheless, it provides mass mixing ratios that deviate from chemical equilibrium, and thus serves for demonstration purpose.

## Appendix D: Accuracy of fluxes and heating rates

The thermal forcing in the GCM is given by the thermodynamic heating rate  $H$ , which is calculated as (e.g., Amundsen et al. 2014; Showman et al. 2009)

$$H = -\frac{dF}{dz} = \frac{gp}{R_s T} \frac{dF}{dp}, \quad (\text{D.1})$$

where  $F$  is the net flux integrated over frequency, consisting of the bolometric stellar and planetary fluxes,  $g$ ,  $p$ ,  $z$ , and  $R_s$  are the surface gravity, pressure, geometric height, and specific gas constant respectively. For a deeper understanding of the accuracy of the different approaches, we compare the accuracy of the resulting fluxes and heating rates of the individual methods in comparison to the RORR method. Since all the simulations lead to different atmospheric states, we opted to compare the fluxes and heating rates in a separate set of models. These models start from the final output of the RORR models at 2000 d and run for 10 radiative time-steps (a total of 1000 seconds). Unlike in the 2000-day-long simulations mentioned above, we did not limit the amount of scattering iterations in these extra runs. Using one radiative time-step for the comparison is not good enough, since

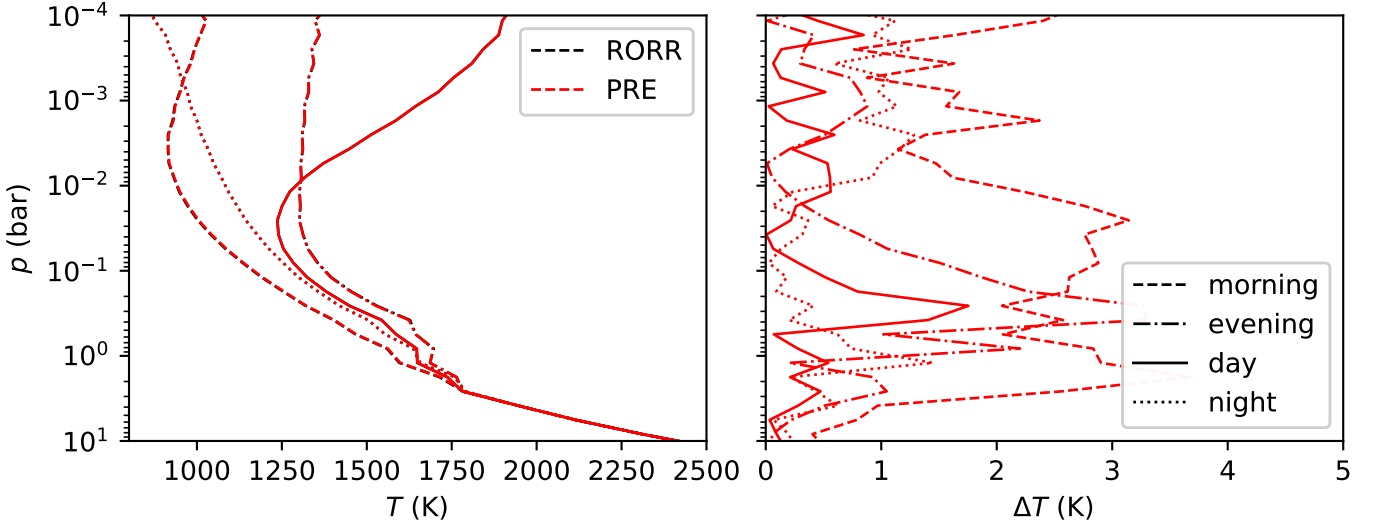


Fig. B.1: Temperature profiles of the RORR simulation compared to a simulation with premixing. The figure is in the same style as Figs. 3 and 4.

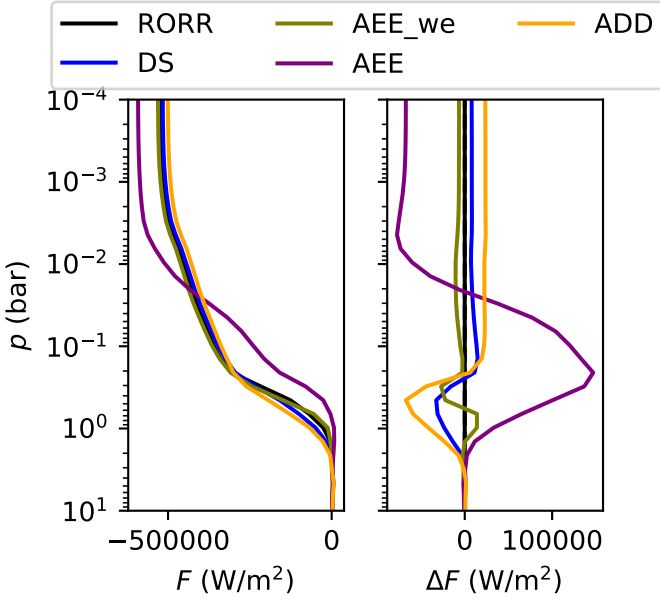


Fig. D.1: Total fluxes integrated over frequency at the substellar point (left) and difference to the RORR simulation (right).




Fig. D.2: Heating rates (see Eq. D.1) at the substellar point (left) and differences to the RORR simulation (right).

the AEE\_we method needs the flux from the previous time-step. The advantage of this approach is that the temperature profiles are identical in these models. The comparison of fluxes and heating rates is thus fairer. We show the fluxes and heating rates for the substellar point in Figs. D.1 and D.2 respectively, and those of the antistellar point in Figs. D.3 and D.4 respectively.

Looking at the day side (substellar point), we can see that AEE\_we and DS result in equally accurate heating rates and fluxes, followed by the ADD simulation, which also results in reasonable fluxes and heating rates. On the other hand, the non-weighted AEE simulations result in completely wrong fluxes and thus heating rates. The total flux at the substellar point is dominated by the stellar flux. We can see that the AEE method results in larger negative values of the stellar flux, again hinting to too little absorption in the upper parts of the atmosphere, as

discussed in Sect. 4. Conversely, due to the overestimation of  $\kappa_{\text{tot}}$  at small  $g$ , almost all radiation is absorbed within a fine pressure range. While these trends can also be seen in the weighted AEE method, it seems to cause much less of an error, if the major absorber is wisely picked.

Overall, we find that considering the substellar and antistellar fluxes and heating rates, the DS and AEE\_we method reproduce similarly accurate results. One of the reasons, why the DS method yields an overall better accuracy on the final atmospheric state might be found in the systematic error of the AEE method, which enhances  $\kappa_{\text{tot}}$  at small  $g$  and decreases  $\kappa_{\text{tot}}$  at large  $g$ , which is not found in the DS method, which exhibits a random noise on the different  $g$  values in each frequency bin. These diagnostics reinforce that adaptive equivalent extinction should only be used together with a weighted average opacity.




Fig. D.3: Total fluxes integrated over frequency at the antistellar point (left) and difference to the RORR simulation (right). The black dashed line in the right panel indicates the zero, which would mean no difference to RORR.




Fig. D.4: Heating rates (see Eq. D.1) at the antistellar point (left) and differences to the RORR simulation (right).