Spectral binning of precomputed correlated- k coefficients[★]

Jérémy Leconte

Laboratoire d'astrophysique de Bordeaux, Univ. Bordeaux, CNRS, B18N, allée Geoffroy Saint-Hilaire, 33615 Pessac, France
e-mail: jeremy.leconte@u-bordeaux.fr

Received 27 July 2020 / Accepted 26 October 2020

ABSTRACT

With the major increase in the volume of the spectroscopic line lists needed to perform accurate radiative transfer calculations, disseminating accurate radiative data has become almost as much a challenge as computing it. Considering that many planetary science applications are only looking for heating rates or mid-to-low resolution spectra, any approach enabling such computations in an accurate and flexible way at a fraction of the computing and storage costs is highly valuable. For many of these reasons, the correlated- k approach has become very popular. Its major weakness has been the lack of ways to adapt the spectral grid/resolution of precomputed k -coefficients, making it difficult to distribute a generic database suited for many different applications. Currently, most users still need to have access to a line-by-line transfer code with the relevant line lists or high-resolution cross sections to compute k -coefficient tables at the desired resolution. In this work, we demonstrate that precomputed k -coefficients can be binned to a lower spectral resolution without any additional assumptions, and show how this can be done in practice. We then show that this binning procedure does not introduce any significant loss in accuracy. Along the way, we quantify how such an approach compares very favorably with the sampled cross section approach. This opens up a new avenue to deliver accurate radiative transfer data by providing mid-resolution k -coefficient tables to users who can later tailor those tables to their needs on the fly. To help with this final step, we briefly present Exo_k, an open-access, open-source Python library designed to handle, tailor, and use many different formats of k -coefficient and cross-section tables in an easy and computationally efficient way.

Key words. planets and satellites: general – planets and satellites: atmospheres

1. Introduction

Despite the considerable increase in computing power, line-by-line radiative transfer is still considered a computationally intensive task in many cases of interest. One of the reasons for this state of affairs is that recent progress in spectroscopy has multiplied the size of molecular line lists by more than three orders of magnitude (Tennyson & Yurchenko 2012). We are at a point at which line list themselves, and a fortiori, high-resolution cross sections, represent a data volume that is not trivially handled and exchanged.

To face this while maintaining flexibility, some have opted for sampled or binned down cross-section tables¹ (Line et al. 2015; Waldmann et al. 2015). But Garland & Irwin (2019) showed that this approach is inaccurate if the resolution used is not sufficiently high, and too computationally expensive if the resolution is high enough (we briefly revisit this question in Sect. 2 whose results are summarized in Figs. 1 and 2). It seems that for typical resolutions below ~ 1000 , methods designed to group absorptions, such as the so-called correlated- k method (Liou 1980; Lacis & Oinas 1991), are more efficient while remaining accurate (Irwin et al. 2008; Amundsen et al. 2017; Garland & Irwin 2019; Zhang et al. 2020). It is especially true for

multidimensional models that need a very fast radiative transfer and use only 10–30 bins for the whole spectrum (Showman et al. 2009; Wordsworth et al. 2011; Leconte et al. 2013; Amundsen et al. 2017).

For these reasons, it is interesting to directly distribute reference k -coefficients for each molecule instead of high-resolution cross sections. This is currently done for the ExoMol project (Chubb et al. 2020). This approach decreases the data volume necessary to handle while keeping an optimal accuracy. A current drawback with precomputed k -coefficients, however, is that the person computing the tables needs to choose a spectral resolution (or wavenumber grid) without knowing exactly what the table will be used for, while k -coefficients are used at their highest potential when tailored to the needed resolution.

To circumvent this problem we present a way to accurately bin down precomputed k -coefficients to an arbitrary spectral grid without any significant additional loss in accuracy. This method is faster than computing the k -coefficients directly from line-by-line data. It therefore opens up a new way of disseminating accurate, reference molecular opacities: spectroscopists only need to compute k -coefficients once with a resolution that is higher than the highest resolution they envision their users need (e.g., the resolution of the instrument you want to compare your model to). Users can then easily bin down the data to their spectral grid of choice, possibly even on the fly.

To help with the latter, we developed Exo_k², an open-source Python 3 library that can directly load precomputed k -coefficient

[★] I wish to dedicated this paper to the memory of Adam P. Showman and Franck Hersant. You both were, each in your own way, an inspiration to me. Godspeed on your journey among the stars.

¹ Here, sampled cross sections specifically refer to cross sections that are computed monochromatically but at a resolution that does not necessarily resolve the lines or their shape. Binned-down cross sections are computed on a finely sampled spectral grid and then integrated over a coarser grid, thereby preserving area.

² See the online documentation for tutorials, example scripts and notebooks, and an extensive and up-to-date description of the (evolving) features of the library: http://perso.astronomy.u-bordeaux.fr/~jleconte/exo_k-doc/index.html

Fig. 1. Root mean square error on the transit spectrum of our fiducial hot Jupiter over the 1–2 μm window as a function of the final observed resolution. Dashed curves represent the error obtained with sampled cross-section tables at various sampling resolution (from top to bottom: $R_{\text{sp}} = \{1000, 3000, 10\,000, 30\,000\}$). The spectra are shown in Fig. 2. The black curve with dots represents the error using k -coefficients computed from high-resolution spectra directly at the final resolution. The gray curves show results when using binned down k -coefficients using natural (solid with stars) and non-integer (dotted) binning (see Sect. 3.4 for details).

and cross-section tables, change their spectral resolution as well as pressure-temperature grids, and save them back on disk. The library can handle many different formats used in various codes: for example, TauREx (Al-Refaie et al. 2019; both pickle and HDF5), petitRADTRANS (Mollière et al. 2019), Nemesis (Irwin et al. 2008), ARCIS (Min et al. 2020), Exo transmit (Kempton et al. 2017), LMD Generic global climate model (Wordsworth et al. 2011; Leconte et al. 2013). Several of these formats are provided by the latest ExoMol release (Chubb et al. 2020). Other formats are being implemented and new ones can be added on request. It is also possible to compute k -coefficient tables from high-resolution spectra. One can also combine the opacities of several molecules using the random overlap method (Lacis & Oinas 1991) to create a table for a given atmospheric mix, which is very useful for climate models, for example. Finally, the library includes methods to directly interpolate and combine the opacities, and even compute transmission and emission spectra for 1D planetary atmospheres. Leveraging on the Numba library, these operations are performed in a very computationally efficient way. So Exo_k can be imported directly into any radiative transfer code to handle molecular opacities easily and efficiently.

In this article, we first briefly compare the relative accuracy and efficiency of the sampled cross section and correlated- k approaches in Sect. 2. After introducing the necessary concepts and notations, Sect. 3 demonstrates rigorously why and how to bin down precomputed k -coefficients to an arbitrary resolution. Then, in Sect. 4, we validate the numerical algorithm presented in this work and implemented in Exo_k.

2. Sampled cross sections versus correlated- k tables

To accurately compare the efficiency of various radiative transfer methods when a given resolution and precision are needed, we briefly estimate the accuracy of atmospheric modeling when using the following types of input radiative data: sampled and binned cross section and k -coefficient tables.

We compare this for two different atmospheric configurations and observation geometries:

- Transmission spectrum of a hot Jupiter with a 1000 K isothermal $\text{H}_2/\text{H}_2\text{O}$ atmosphere in front of a Sun.

Table 1. Standard parameters used in our two fiducial cases.

Case		Hot Jupiter	Mars
Planet radius	R_{Jup}	1	–
Surface gravity	m s^{-2}	10	3.72
Surface temperature	K	1000	200
Stratospheric temperature	K	1000	100
Adiabatic index		–	0.22
Surface pressure	Pa	10^6	640
Volume mixing ratio of H_2O		10^{-3}	–
Stellar radius	$R_{\odot}$	1	–

Notes. A hyphen stands for parameters that are not needed for a given case.

- Emission spectrum of a pure CO_2 Mars-like atmosphere.

The parameters for these two cases are given in Table 1.

Each time, we start by computing a reference spectrum using high-resolution monochromatic cross sections (with a typical step of 0.01 cm^{-1}). For our CO_2 atmosphere, we use spectra from Turbet & Tran (2017). For our hot-Jupiter case, we use water spectra produced by the Exomol project with H_2 broadening (K. Chubb, private communication) based on the line list produced by Polyansky et al. (2018). Collision-induced-absorptions (CIA) are taken from the HITRAN database (Richard et al. 2011).

Then, the high-resolution cross sections are sampled at various resolutions ($R_{\text{sp}} = \{1000, 3000, 10\,000, 30\,000\}$) and a new set of spectra are computed at each of these resolutions.

The spectra obtained from the sampled cross sections are then binned (preserving area) at selected final resolutions ($R_{\text{fin}} = \{10, 20, 50, 100, 200, 500, 1000\}$). The transmission spectra computed in the hot-Jupiter case are shown in the left column of Fig. 2. Finally, these spectra are compared to the properly binned reference spectrum (right column of Fig. 2). The root mean squared (RMS) error is shown for each ($R_{\text{sp}}, R_{\text{fin}}$) in Fig. 1. The emission spectra computed in the Mars-like case are shown in Fig. 3.

For the correlated- k approach, k -coefficient tables are computed from the high-resolution spectra directly at the final desired resolution (computed with 20 Gauss-Legendre quadrature points in g -space) and compared to the reference spectrum binned at the same resolution. This gives the black dashed curves in Fig. 2. The RMS error is shown by the black curve in Fig. 1.

For all the computations described above we used a standard algorithm to describe the transmission and emission spectra of a 1D atmosphere that have been implemented in the Exo_k library².

We show that the correlated- k approach reaches an accuracy that is one to two orders of magnitude better than sampled cross sections, especially for resolutions bigger than 100. The relative drop in accuracy of the correlated- k method at $R_{\text{fin}} = 10$ is due to the low sampling of the CIA absorption, which is treated as constant in each spectral bin. This could be alleviated by applying the random overlap method to the CIA as well. To compare numerical efficiency, we focus on the $R_{\text{fin}} = 10$ case for which the $R_{\text{sp}} = 30\,000$ roughly manages to reach an accuracy equivalent to the correlated- k method. Then, considering that our k -coefficients were computed using 20 quadrature points, the computation is still 150 times faster with the correlated- k method compared to the sampled cross sections.

We conclude this comparison by mentioning binned-down cross sections – in which the finely sampled cross sections are integrated over a coarser grid, thereby preserving area. Although

Fig. 2. Effect of sampling resolution on the transit spectrum of a hot Jupiter. The various rows show the effect for different final observing resolutions (R_{fin} , specified in the top left corner of the panels). *Left column:* transit depth in the 1–2 μm window (with an offset of 11500 ppm). The various colors correspond to the four sampling resolutions (i.e., the resolution used in radiative transfer calculation before binning; R_{sp}) shown in the legend. The curve with dots shows the reference case computed from our very high-resolution cross sections ($\Delta\sigma \approx 0.01 \text{ cm}^{-1}$). The high-resolution reference spectrum before binning is shown as the shading in the *top left panel*. *Right column:* difference (in ppm) between the spectra of the left panel and the reference case. The difference between the calculation with k -coefficients and the reference case is shown with the black dashed curve; the two spectra would be indistinguishable in the left panels. The RMS standard deviation over the spectral region is shown in Fig. 1.

it seems the most natural way to bin cross sections, this method performs even worse than the sampled cross sections (by about one order of magnitude; Fig. 4). This rather counterintuitive result stems from the nonlinear nature of radiative transfer: the opacity dilution due to the binning procedure is not sufficient to

increase the transparency of the gas near optically thick lines, while decreasing it in the line wings. Overall, this systematically over estimates the opacity of the atmosphere and creates a systematic bias – on the order of 100 ppm in our hot-Jupiter test case.

Fig. 3. Effect of sampling resolution on the emission spectrum of our Mars-like case around the $15\mu\text{m}$ CO_2 band. The rows show the effect for different final observing resolutions (R_{fin}). *Left column:* flux. The various colors correspond to the four sampling resolutions (i.e., the resolution used in radiative transfer calculation before binning; R_{sp}) shown in the legend. The reference case is computed from our very high-resolution cross sections ($\Delta\sigma \approx 0.01\text{ cm}^{-1}$). The high-resolution reference spectrum before binning is shown as the shading in the *top left panel*. *Right column:* relative error between the spectra of the *left panel* and the reference case. The error between the calculation with k -coefficients and the reference case is shown with the black dashed curve; the two spectra would be indistinguishable in the *left panels*.

Fig. 4. Effect of sampling resolution on the transit spectrum of a hot Jupiter using cross sections that have been binned by preserving the total area. The panels can be compared to the results for sampled cross sections in Fig. 2 (see the caption for details). The binning systematically overestimates the opacity.

3. Method to bin down k -coefficients

3.1. Correlated- k formalism

Before demonstrating how to bin down k -coefficients, we briefly introduce the basic concepts we need. In the correlated- k method (see Liou 1980 for details), the wavenumber space, where ν is the wavenumber, is divided into N_b spectral bands of width $\Delta\nu^b \equiv \nu_+^b - \nu_-^b$. In each band, the transmission through a slab of matter can be written as

$$\mathcal{T}^b(u) = \frac{1}{\Delta\nu^b} \int_{\nu_-^b}^{\nu_+^b} e^{-k^\nu u} d\nu = \int_0^1 e^{-\hat{k}^b(g)u} dg \approx \sum_{i=1}^{N_g} w_i^b e^{-\hat{k}_i^b u}, \quad (1)$$

where N_g is the number of points (or abscissas) used to discretize the g -space and w_i^b are the associated set of weights for each band. The symbol $\hat{k}^b(g)$ denotes the inverse of the cumulative density function of the opacity within band b – the so-called k -distribution – and should not be confused with k^ν , the absorption coefficient, which is a function of the wavenumber. The so-called correlated- k coefficients (or simply k -coefficients) are the $\hat{k}_i^b$, which are the discretized version of the k -distribution ($\hat{k}^b(g)$).

An intuitive way to understand that is to say that now, in band b , the opacity can be described by N_g representative values ($\hat{k}_i^b$) for each of the N_g points in g -space (g_i), each of these values being affected a weight (w_i^b). Then any radiative transfer

calculation can be computed separately in the N_g bins to be later summed up with the proper weights.

An important, although often forgotten assumption of the correlated- k formalism is that both the spectral incoming flux impinging on our medium (F_{in}^ν) and the source function of the latter should be nearly constant within each spectral band (Lacis & Oinas 1991) so that it can be written $F_{\text{in}}^\nu = F_{\text{in}}^b / \Delta\nu^b$, where $F_{\text{in}}^b \equiv \int_{\nu_-^b}^{\nu_+^b} F_{\text{in}}^\nu d\nu$ is the total flux within the band. Simply put in the context of a purely absorbing medium, which can be generalized, this ensures that the outgoing flux after a path u is given by

$$\begin{aligned} F_{\text{out}}^b(u) &\equiv \int_{\nu_-^b}^{\nu_+^b} F_{\text{in}}^\nu e^{-k^\nu u} d\nu \approx \frac{F_{\text{in}}^b}{\Delta\nu^b} \int_{\nu_-^b}^{\nu_+^b} e^{-k^\nu u} d\nu \\ &= F_{\text{in}}^b \cdot \mathcal{T}^b(u). \end{aligned} \quad (2)$$

This assumption is important in our context because it is the only one that we need to make to be able to combine the k -coefficients of various bands. Finally, the bolometric flux can be obtained with

$$F = \sum_{b=1}^{N_b} F_{\text{out}}^b. \quad (3)$$

3.2. Combining k -coefficients of various bands

To reduce the resolution of precomputed k -coefficients, we essentially need to answer the following question: Can we determine the k -coefficients $\hat{k}_i^{\bar{b}}$ of a “super-band” $\bar{b}$ that is the union of two smaller, nonoverlapping “sub-bands”³, b_1 and b_2 , for which we know the k -coefficients, $\hat{k}_i^{b_1}$ and $\hat{k}_i^{b_2}$?

In this form, the answer is not trivial. However, when computing k -coefficients, the crucial mathematical object is not the k -distribution ($\hat{k}^b(g)$) but its inverse, the cumulative density function (CDF) of the opacity within the band (hereafter called the g -distribution, $\hat{g}^b(k)$) defined as

$$\hat{g}^b(k) \equiv \frac{1}{\Delta\nu^b} \int_{\nu_-^b}^{\nu_+^b} H(k - k^\nu(\nu)) d\nu, \quad (4)$$

where $H(x)$ is equal to 0 where x is positive and 1 elsewhere.

The question can thus be reframed to ask whether we can compute the g -distribution of the super-band, $\hat{g}^{\bar{b}}$, if we know the g -distribution within the two sub-bands ($\hat{g}^{b_1}$ and $\hat{g}^{b_2}$). This is rather straightforward, because as long as the two smaller bands do not overlap, we have $\Delta\nu^{\bar{b}} = \Delta\nu^{b_1} + \Delta\nu^{b_2}$ and

$$\begin{aligned} \hat{g}^{\bar{b}}(k) &\equiv \frac{1}{\Delta\nu^{\bar{b}}} \int_{\nu_-^{\bar{b}}}^{\nu_+^{\bar{b}}} H(k - k^\nu(\nu)) d\nu \\ &= \frac{1}{\Delta\nu^{\bar{b}}} \left(\int_{\nu_-^{b_1}}^{\nu_+^{b_1}} H(k - k^\nu(\nu)) d\nu + \int_{\nu_-^{b_2}}^{\nu_+^{b_2}} H(k - k^\nu(\nu)) d\nu \right) \\ &= \frac{\Delta\nu^{b_1}}{\Delta\nu^{\bar{b}}} \hat{g}^{b_1}(k) + \frac{\Delta\nu^{b_2}}{\Delta\nu^{\bar{b}}} \hat{g}^{b_2}(k). \end{aligned} \quad (5)$$

Therefore the g -distribution in the super-band is just the sum of the g -distributions in the sub-bands weighted by their spectral extent.

³ We restrict ourselves to the case of the combination of two bands. The generalization to an arbitrary number of bands is trivial.

This can be understood more intuitively when discussing in terms of probability. We want to know if $\hat{g}^{\bar{b}}(k)$, that is, the probability that a monochromatic ray of light falls in a spectral region where the opacity k^ν is lower than k . To compute this, we have to sum, for each band, the probability of the ray falling in the given sub-band; this probability is $\Delta\nu^b / \Delta\nu^{\bar{b}}$, as long as the incoming spectral flux is constant over the super-band, multiplied by the probability that $k^\nu < k$ within this sub-band, which is $\hat{g}^b(k)$. So we recover Eq. (5) as long as we keep the essential assumption of the correlated- k method that the incoming flux and sources functions do not significantly vary over the wavelength range that we wish to describe by a single distribution.

3.3. Numerical algorithm

As shown above, combining the g -distributions of various bands in a single g -distribution for a large band is straightforward. However, to use this in a radiative transfer code, we usually need the k -coefficients ($\hat{k}_i^b$) that are tabulated for fixed g -points (g_i) with the associated weights (w_i^b). The difficulty is that although g_i can be seen as the values of $\hat{g}^b$ in band b for the points located at $\hat{k}_i^b$ in opacity space, the distribution $\hat{g}^b$ is sampled at different locations in opacity space in each sub-band.

To circumvent this problem, we recommend the following approach, which is the one implemented in Exo_k. The whole process is illustrated in Fig. 5 with only two bands, but can be carried out for any number of bands. For the moment, however, we restrict ourselves to the case in which the super-band $\bar{b}$ is composed of an integer number of sub-bands (hereafter, natural binning). Once these sub-bands have been identified, we determine

$$k_{\min} = \min_{i, b \in \bar{b}} \{\hat{k}_i^b\} \quad \text{and} \quad k_{\max} = \max_{i, b \in \bar{b}} \{\hat{k}_i^b\} \quad (6)$$

and compute an evenly log-distributed grid of N_k points ($\tilde{k}_j$) between these two values. Then, for each sub-band, we interpolate $g_i(\hat{k}_i^b)$ on this grid, which yields N_k values per band, $\tilde{g}_j^b$, representing all the $\hat{g}^b$ on the same grid for all bands. It is thus now easy to compute the global g -distribution which is given by

$$\tilde{g}_j^{\bar{b}} = \sum_{b \in \bar{b}} \frac{\Delta\nu^b}{\Delta\nu^{\bar{b}}} \tilde{g}_j^b. \quad (7)$$

This generally yields an oversampled version of the distribution because we choose $N_k > N_g$ (the black curve in the middle panel of Fig. 5). Recomputing the k -coefficients is thus now simply a matter of resampling $\tilde{k}_j(\tilde{g}_j^{\bar{b}})$ onto our original g -grid (g_i) – going from the black to the gray curve in the right panel of Fig. 5 – as discussed in Lacis & Oinas (1991) and Amundsen et al. (2017). However, unlike those authors, we think that, to be consistent with the usual quadrature rules, the finely sampled k -distribution should not be weight-averaged over the final g -bins, but sampled at the precise abscissa of the g -point determined by the quadrature rule used, following Irwin et al. (2008).

As long as we resample on the original grid, we keep the original quadrature with N_g points, and the weights are left unchanged. However, it is still possible at this point to use different g -points, especially if we need to use less g -points for computational reasons (as is often the case with 3D climate models). For this, we can simply resample $\tilde{k}_j(\tilde{g}_j^{\bar{b}})$ on a different g -grid and use the relevant weights, which is an option of Exo_k. In general, Exo_k works for any choice of g -space sampling specified by the user.

Fig. 5. Example of the binning process for two bands. *Left:* k -coefficients on $N_g = 20$ Gauss-Legendre quadrature points for two bands ($\hat{k}_i^{b_1}$ and $\hat{k}_i^{b_2}$). *Middle:* g -distribution for the two bands resampled on $\tilde{k}_j$, a grid of $N_k = 100$ points between $k_{\min}$ and $k_{\max}$ ($\tilde{g}_j^{b_1}$ and $\tilde{g}_j^{b_2}$). The black line indicates the weighted sum of these two g -distributions, i.e., the g -distribution of the super-band ($\tilde{g}_j^b$). *Right:* colored lines are shown as in the left panel. The black line indicates the finely sampled k -distribution ($\tilde{k}_j(\tilde{g}_j^b)$). The gray line represents the k -distribution of the super-band re-sampled on the original g -grid with $N_g = 20$, i.e., the final k -coefficients ($\hat{k}_i^b$).

3.4. Non-integer binning

Up to now, we only considered the case in which the final super-band encompasses an integer number of sub-bands. This is required for our demonstration to be exact without any further assumptions.

In practice, it is still possible to bin k -coefficients on a spectral grid that does not exactly respect the limits of the sub-bands (hereafter, non-integer binning). To do that, we replace the spectral extent, $\Delta\nu^b$, of the two bands at the boundaries of the super-band in Eq. (7) by their spectral extent inside the super-band. This is equivalent to binning a high-resolution spectrum to a low-resolution grid using a weighted average. In theory, this would remain perfectly accurate if the statistical distribution of the absorption coefficient were constant within the sub-band. As this is not true in general, we test the accuracy of this approach in the next section.

4. Validation

To validate our approach, we compute spectra in the two cases detailed in Sect. 2. As we do not want to test the accuracy of the correlated- k approach itself, which has already been done before, but whether the binning introduces additional errors, the comparison proceeds as follows:

- Starting from very high-resolution spectra, we compute a first reference set of tables of k -coefficients with resolutions ranging from $R = 1000$ to $R = 5$ and a g -space sampled using 20 Gauss-Legendre quadrature points. The exact R values used do not matter. However, we make sure that the wavenumber points in the low-resolution grids can be found in the high-resolution grid to use only natural binning.
- Starting from the $R = 1000$ table of k -coefficients⁴ we just computed, we use the binning approach described above to compute a new set of ten tables of k -coefficients with the exact same resolutions and wavenumber grids as the reference set.
- The emitted or transmitted flux of our fiducial atmosphere is computed at each resolution with both the reference and

the binned k -coefficients and the results are compared two by two.

The results of the comparison are shown in Fig. 6. For the Mars emission case, we compare the RMS standard deviation of the relative difference between the two spectra in the 1–200 μm range. For the Hot-jupiter in transmission test case, the comparison metric used is the RMS standard deviation of the relative difference between the two spectra in the 1–2 μm window.

All the k -coefficients computed in this work use 20 Gauss-Legendre quadrature points in g -space. Because the number of intermediate points used to resample the g -distribution in opacity space (N_k ; see Sect. 3.3) is a free parameter, we tried different values of this parameter. We see that an insufficient intermediate sampling results in significant numerical errors, as could be anticipated. However, when N_k is a factor of a few larger than the initial number of g -points, the accuracy of the method asymptotes below 10^{-3} – 10^{-4} , which is itself much lower than the error introduced by the correlated- k method in the first place (see for example Amundsen et al. 2017). We note that although choosing a larger N_k may slightly increase the binning time, this does not affect the efficiency of the subsequent radiative transfer calculations in any way. We thus used results with $N_k = 5N_g$ as a baseline. We also found that using initial data with 40 Gauss-Legendre points did not significantly affect the performances of the binning procedure. Our algorithm is thus robust to the initial choice of g -space sampling.

The spectra obtained with this set of binned-down k -coefficients are also compared directly to the reference high-resolution spectrum as described in Sect. 2 to give the absolute difference shown in gray with stars in Fig. 1. Although not strictly equal because they are computed using a slightly different reference spectrum, the absolute difference shown in Fig. 1 and the relative difference from Fig. 6 can be shown to be consistent by multiplying the latter by the average transit depth as shown on the right axis of Fig. 6. In this particular case, we further wanted to test the impact of non-integer binning. For this a new set of k -coefficients were computed in the exact same way, but with a spectral grid that is slightly shifted in wavenumber to force non-integer binning. As shown by the dotted gray curve in Fig. 1, the accuracy is lower, although this method remains more accurate than sampled cross sections. In addition,

⁴ This choice is motivated by the resolution chosen by the ExoMol project for their k -coefficient tables (Chubb et al. 2020).

Fig. 6. Numerical relative error introduced by the binning process as a function of the binning factor (ratio of the initial resolution, R_{ini} , to the final one, R_{fin}) for four different values of the number of intermediate sampling points N_k (see Sect. 3.3; dotted: $N_k = N_g$; dashed: $N_k = 1.5N_g$; dot-dash: $N_k = 2N_g$; solid: $N_k = 5N_g$, where $N_g = 20$). *Top*: RMS relative top-of-atmosphere flux error in our Mars-like case. *Bottom*: RMS relative transit depth error in our hot-Jupiter case. The right axis gives the absolute error to be compared to the results of Figs. 1 and 2. With a sufficient sampling, the errors are under the 10^{-3} – 10^{-4} level depending on the geometry, which is below the error introduced by the correlated- k method in the first place.

the error decreases with the binning factor because the errors at the boundaries of our spectral bins are more diluted.

We therefore conclude that our binning method does not introduce any significant additional errors.

5. Conclusions

We have demonstrated mathematically how the k -distribution of the opacity in a large band can be obtained from the k -distributions inside each of the smaller bands that it encompasses. Then we have presented a numerical algorithm to use this property to bin down precomputed tables of k -coefficients to any arbitrary resolution. Finally, we showed, in two concrete cases, that the numerical error added by this procedure is relatively small compared to the errors due to the more general use of the correlated- k approach.

To facilitate the implementation of this new method, we developed Exo_k, an open-source Python library that has been designed to handle many different sources of radiative data (including, of course, k -coefficient tables, but also cross section and collision-induced-absorption tables). Because this library is still rapidly evolving, see our extensive online documentation for tutorials, example notebooks, and a complete description of the library's features².

We think the flexibility that this approach brings to the correlated- k method opens up a completely new way of disseminating reference radiative data for use in atmospheric models: precomputed k -coefficients can now be directly distributed at a fraction of the data volume while providing fast and accurate results for any user that does not need a higher resolution. For example, $R = 1000$ k -coefficients computed on 20 gauss points such as that distributed by the ExoMol project take about 50 times less space on disk than the high-resolution spectra used to produce them. This is a considerable compression factor, and that does not include the hundreds of CPU hours that would be needed to compute the spectra.

This of course makes us wonder what would be the optimum resolution and quadrature for a flexible and versatile correlated- k database. Our test show that $R = 1000$ k -coefficients computed on 20 gauss points, when combined with the flexibility brought by a tool like Exo_k, seem to be sufficient for most 1D-3D atmospheric model applications⁵ as well as modeling synthetic spectroscopic observations of planetary atmospheres up to that resolution. This method can thus be used to model HST, Spitzer, Ariel (Tinetti et al. 2018), and low-resolution JWST data. For high-resolution JWST observation modes, higher resolution radiative data will be needed.

Acknowledgements. I thank F. Selsis for being, as always, an outstanding scientific sparring partner and the TauREx team at UCL for pointing out many important Python tricks and libraries. I am also indebted to K. Chubb and M. Turbet for making some radiative data available before their publication. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement n° 679030/WHIPLASH).

References

- Al-Rfaie, A. F., Changeat, Q., Waldmann, I. P., & Tinetti, G. 2019, ArXiv e-prints [arXiv:1912.07759]
- Amundsen, D. S., Tremblin, P., Manners, J., Baraffe, I., & Mayne, N. J. 2017, *A&A*, **598**, A97
- Chubb, K. L., Rocchetto, M., Yurchenko, S. N., et al. 2020, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202038350>
- Garland, R., & Irwin, P. G. J. 2019, *MNRAS*, in press [arXiv:1903.03997]
- Irwin, P. G. J., Teanby, N. A., de Kok, R., et al. 2008, *J. Quant. Spectr. Rad. Transf.*, **109**, 1136
- Kempton, E. M. R., Lupu, R., Owusu-Asare, A., Slough, P., & Cale, B. 2017, *PASP*, **129**, 044402
- Lacis, A. A., & Oinas, V. 1991, *J. Geophys. Res.*, **96**, 9027
- Leconte, J., Forget, F., Charnay, B., et al. 2013, *A&A*, **554**, A69
- Line, M. R., Teske, J., Burningham, B., Fortney, J. J., & Marley, M. S. 2015, *ApJ*, **807**, 183
- Liou, K. N. 1980, *An introduction to atmospheric radiation* (Boston: Academic Press)
- Min, M., Ormel, C. W., Chubb, K., Helling, C., & Kawashima, Y. 2020, *A&A*, **642**, A28
- Mollière, P., Wardenier, J. P., van Boekel, R., et al. 2019, *A&A*, **627**, A67
- Polyansky, O. L., Kyuberis, A. A., Zobov, N. F., et al. 2018, *MNRAS*, **480**, 2597
- Richard, C., Gordon, I., Rothman, L., et al. 2011, *J. Quant. Spectr. Rad. Transf.*, **113**, 1276
- Showman, A. P., Fortney, J. J., Lian, Y., et al. 2009, *ApJ*, **699**, 564
- Tennyson, J., & Yurchenko, S. N. 2012, *MNRAS*, **425**, 21
- Tinetti, G., Drossart, P., Eccleston, P., et al. 2018, *Exp. Astron.*, **46**, 135
- Turbet, M., & Tran, H. 2017, *J. Geophys. Res. Planets*, **122**, 2362
- Waldmann, I. P., Tinetti, G., Rocchetto, M., et al. 2015, *ApJ*, **802**, 107
- Wordsworth, R. D., Forget, F., Selsis, F., et al. 2011, *ApJ*, **733**, L48
- Zhang, M., Chachan, Y., Kempton, E. M. R., Knutson, H. A., & Chang, W. H. 2020, *ApJ*, **899**, 27

⁵ Very cold and low pressure atmospheres, such as Pluto's, may of course require special care because of their very narrow lines, which usually call for tailored quadrature rules.