Atmospheric
Measurement
Techniques

Fast reconstruction of hyperspectral radiative transfer simulations by using small spectral subsets: application to the oxygen A band

A. Hollstein and R. Lindstrot

Freie Universität Berlin Department of Earth Sciences Institute for Space Sciences, Carl-Heinrich-Becker-Weg 6–10, 12165 Berlin, Germany

Correspondence to: A. Hollstein (andre.hollstein@fu-berlin.de)

Abstract. Hyperspectral radiative transfer simulations are a versatile tool in remote sensing but can pose a major computational burden. We describe a simple method to construct hyperspectral simulation results by using only a small spectral subsample of the simulated wavelength range, thus leading to major speedups in such simulations. This is achieved by computing principal components for a small number of representative hyperspectral spectra and then deriving a reconstruction matrix for a specific spectral subset of channels to compute the hyperspectral data. The method is applied and discussed in detail using the example of top-of-atmosphere radiances in the oxygen A band, leading to speedups in the range of one to two orders of magnitude when compared to radiative transfer simulations at full spectral resolution.

1 Introduction

Radiative transfer simulations are a key tool for the development of remote sensing algorithms in the field of earth observation. Depending on the spectral domain, a variety of techniques are used to solve the radiative transfer equation (RTE). Such simulated radiances can be compared to measurements and thus used for the inversion for the actual physical state of the atmosphere–surface or atmosphere–ocean (e.g., Thomas and Stamnes, 2008, and references therein).

Most radiative transfer models for the ultraviolet (UV) to short-wavelength infrared (SWIR) region that we are aware of solve a monochromatic version of the RTE and are used for channel-based radiative transfer simulations. Even if models are capable of treating problems with inelastic scattering, monochromatic RTE solvers with additional radiance

sinks and sources are generally used (e.g., Landgraf et al., 2004, and their treatment of atmospheric Raman scattering).

Therefore, for problems involving hyperspectrally resolved radiance measurements, the obvious approach is to perform large numbers of independent simulations. Various techniques for the increase of the computational efficiency of these radiative transfer simulations have been developed (e.g., Kokhanovsky, 2013, chapter 10 by Vijay Natraj). Two main approaches can be distinguished. Firstly, by relaxing the constraints on the accuracy of the solution of the radiative transfer equation, the computational time needed for each individual radiative transfer simulation is reduced. Examples are the use of two stream methods (e.g., Meador and Weaver, 1980), reduced order of scatterings (e.g., Natraj and Spurr, 2007), replacing line-by-line absorption calculations by exponential sum fitting techniques (ESFTs) (e.g., Wiscombe and Evans, 1977) and their advancements (e.g., Lacis and Oinas, 1991; Bennartz and Fischer, 2000; Doppler et al., 2013), or by relying on pre-computed databases (e.g., Wang et al., 2013). The error with respect to more rigorous solutions of the RTE in general depends on the scene and spectral band and can be estimated by carrying out rigorous simulations. Secondly, by making use of the inherent redundancy of line-by-line calculations by using data reduction techniques such as principal component analysis (e.g., Efremenko et al., 2013; Jolliffe, 2002), the number of individual radiative transfer simulations is reduced. Approaches have been published for the IR spectral region by Liu et al. (2006) and the VIS to SWIR region by Natraj et al. (2010, 2005) and Lindstrot and Preusker (2012).

In this paper, we propose a simple method in which only a small subset of the spectra is simulated and used to generate a

reconstruction matrix, based on principal component analysis, the expansion coefficients of the data, and the pseudo-inverse of a spectral subsample of the data. Then, for the bulk of the spectra to be simulated, only a relatively small spectral subset of the simulations is carried out, and hyperspectral results are constructed using multiplication with the reconstruction matrix. The method is explained and tested using the oxygen A band as test case. We provide numbers on the accuracy of the method vs. the achievable speedup. The method is independent of the underlying radiative transfer model and simple enough to be easily applied to existing radiative transfer schemes. The method is entirely based on post-processing techniques, so that the RT scheme itself does not have to be modified.

The oxygen A band was chosen since it allows a broad range of applications that could benefit from faster and more accurate radiative transfer simulations. Among others, possible applications are the retrieval of cloud-top pressure (e.g., Koelemeijer et al., 2001; Fischer and Grassl, 1991; Preusker and Lindstrot, 2009) and aerosol vertical distribution (e.g., Dubuisson et al., 2009; Sanghavi et al., 2012). Also, more and more hyperspectral radiance measurements will become available with current and future sensors such as TANSO-FTS on GOSAT (Suto et al., 2010), GOSAT-2 (Nakajima et al., 2012), TROPOMI on S-5 P (Veefkind et al., 2012), or OCO2 (e.g., Boesch et al., 2011), to name a few.

2 Notation

Throughout this paper a convenient matrix notation is employed, which allows denoting matrix elements and subsets of matrix elements. A real $a \times b$ matrix is defined as $\mathbf{M} \in \mathbb{R}^{a \times b}$, and $\mathbf{M}_{i,j} \in \mathbb{R}$, $0 < i, j \leq a, b$ identifies a single element of $\mathbf{M}$. Row and columnar subsets are accessed using a $-$ sign: $\mathbf{M}_{i,-} = (M_{i,1}, \ldots, M_{i,b}) \in \mathbb{R}^b$, as well as $\mathbf{M}_{-,j} = (M_{1,j}, \ldots, M_{a,j}) \in \mathbb{R}^a$. To access subsets of elements in rows and columns, a notation of index sets is introduced: $s^{n,l} = (s_1, \ldots, s_l)$ with $s_1 > 0, s_i < s_{i+1}, s_l \leq n, s_i \in \mathbb{N}$, and $n, l \in \mathbb{N}$. This notation denotes an ordered list of l unique elements between 1 and n. A specific index set is defined when its elements $(s_1, \ldots, s_l)$ are set. With the introduced matrix notation, such an index set can be used to access subsets of rows and columns: $\mathbf{M}_{i,s^{b,l}} = (M_{i,s_1}, \ldots, M_{i,s_l}) \in \mathbb{R}^l$.

3 Principal components as data reduction technique

The method was developed keeping in mind a radiative transfer forward operator based on a lookup table, although an application to other usage scenarios (e.g., forward radiative transfer simulations) is straightforward. Here we assume a forward radiative transfer operator (RT), which relates the state vector $x \in \mathbb{R}^{n_x}$ to the top-of-atmosphere radiance spectrum $y \in \mathbb{R}^{n_\lambda}$, using a spectral calibration vector

$\lambda = (\lambda_1, \ldots, \lambda_{n_\lambda}) \in \mathbb{R}^{n_\lambda}$ with $\forall i, j, i \neq j : \lambda_i \neq \lambda_j$:

$$y = \mathrm{RT}(x, \lambda). \tag{1}$$

For simplicity, the radiative transfer simulations are stored in a two-dimensional lookup table matrix $\mathbf{L} \in \mathbb{R}^{n_\lambda \times n_L}$, with $n_L = \prod_{j=1}^{n_x} \| \hat{x}^j \|$, where $\| \cdot \|$ denotes the number of elements of a vector, and the vectors $\hat{x}^j = (\hat{x}_1^j, \ldots, \hat{x}_{\|\hat{x}^j\|}^j)$ contain grid points for each state vector dimension. The corresponding parameter states are stored in the parameter matrix $\mathbf{X} \in (R)^{n_x \times n_L}$ such that for all rows i,

$$\mathbf{L}_{-,i} = \mathrm{RT}(\mathbf{X}_{-,i}, \lambda). \tag{2}$$

The matrices $\mathbf{X}$ and $\mathbf{L}$ could be used to construct a fast forward operator based on multivariate interpolation. A main problem of this approach is that even for small parameter spaces the matrix $\mathbf{L}$ can become large, especially if a very high spectral resolution is needed. This naturally leads to the employment of data reduction techniques like principal component analysis. The main idea is to replace the large matrix $\mathbf{L}$ with two much smaller matrices $\mathbf{P} \in \mathbb{R}^{n_\lambda \times n_P}$ and $\mathbf{C} \in \mathbb{R}^{n_P \times n_L}$ with $n_p \ll n_\lambda$, such that each spectrum $\mathbf{L}_{-,i}$ can be expressed as

$$\mathbf{L}_{-,i} = \mathbf{P} \cdot \mathbf{C}_{-,i} + O[n_P]. \tag{3}$$

The matrix $\mathbf{P}$ contains the principal components up to the order n_P, and the matrix $\mathbf{C}$ contains the expansion coefficients, which express the original spectra in the n_P dimensional space spanned by the principal components:

$$\mathbf{C}_{-,i} = \mathbf{P}^\mathsf{T} \cdot \mathbf{L}_{-,i}. \tag{4}$$

The principal components are orthogonal by construction, and thus $\mathbf{P}$ is a semi-orthogonal matrix with $\mathbf{P}^T = \mathbf{P}^{-1}$, which simplifies the computation of the expansion coefficients $\mathbf{C}_{-,i}$.

The reconstruction of a spectrum $\mathbf{L}_{-,i}$ is associated with an error $O[n_P]$, which in general decreases with increasing n_P. For a given problem, n_P has to be chosen such that the residual $O[n_P]$ is below the user requirement. The memory requirement for the matrix $\mathbf{C}$ is much smaller than for $\mathbf{L}$, and interpolation can be implemented much faster. There are two main downsides to this approach. Firstly, the singular vector decomposition of $\mathbf{L}$ becomes increasingly expensive with respect to computational time with increasing n_L. Secondly, to compute the matrix $\mathbf{C}$, the complete matrix $\mathbf{L}$ must be computed and stored in advance, which in general is time-consuming and costly in terms of storage and backup.

Both of these downsides are discussed in detail in this paper using the oxygen A band as an example. Strictly speaking, the presented analysis is thus only valid for this part of the spectrum, but we do not see any substantial obstacles to applying the technique to other parts of the spectrum.

4 Constructing the principal components from small parameter state space subsets of L

The computation of $\mathbf{P}$ becomes numerically more expensive with increasing size of the lookup table $\mathbf{L}$ and increasing spectral resolution. For actual computations we used the principal component analysis algorithm provided by the Python package Scikit-learn (Pedregosa et al., 2011). Throughout this paper a moderately small lookup table is discussed with the parameter state space sampling given in Table 1. The variability of the database is illustrated in Fig. 1. The lookup table was computed using the Matrix Operator Model (MOMO) radiative transfer model (see Fell and Fischer, 2001; Hollstein and Fischer, 2012), which is a matrix operator model widely used at Freie Universität Berlin. The gaseous absorption was computed using line parameters from the HITRAN spectral database (Rothman et al., 2009) and a modified scheme to compute the K distribution (Bennartz and Fischer, 2000). As shown in Table 1, the variation of the parameter space includes surface pressure, aerosol optical thickness, aerosol mean height, aerosol type, surface reflectance[1], and the viewing geometry. Aerosol optical models were implemented according to Levy et al. (2007). These models are also used by the MODIS aerosol retrieval and were specifically designed to fit observations for different locations on the globe. From the published optical properties, the urban, neutral, dust, continental, and absorbing types were implemented, and Mie calculations using the implementation provided by Wiscombe (1980) were used to compute phase functions, extinction, and single-scattering albedo. This state vector variability was set up to reproduce the variability of a clear-sky scene over land. The temperature profile was explicitly excluded from the state vector to keep it simple but realistic. As shown by Lindstrot and Preusker (2012), the spectral variability due to the temperature profile of the atmosphere can be accounted for by performing radiative transfer simulations only for a set of principal temperature components. The radiance spectrum can then be constructed as linear superposition of the simulated spectra for the principal temperature profile components, by using the expansion coefficients of the actual temperature profile in the space spanned by the principal temperature profile components.

The oxygen A band was simulated with a spectral sampling of 0.005 nm, which led to 4500 spectral channels. The parameter sampling as stated in Table 1 led to 121 500 different parameter states.

Numerical experiments show that, for this particular data set, the principal component matrix $\mathbf{P}$ can be computed without using the complete matrix $\mathbf{L}$, but by using a much smaller subset of the state space $\mathbf{L}_{-,r^{n_L,n_r}}$, where r^{n_L,n_r} is an index set with n_r randomly chosen members from the $n_L =$

---

[1]The surface reflectance spectrum is assumed to be linear and is modeled using a reflectance value at 755 nm and at 780 nm.

Fig. 1. The left panel illustrates the variability of the used spectral database by showing some randomly selected spectra. The right panel shares the reflectance axis with the left panel and shows logarithmic histograms for four selected wavelengths.

121 500 available states. The reconstruction signal-to-noise ratio (SNR) is used as simple scalar figure to test the reconstruction quality of each spectrum:

$$\mathrm{SNR}(r,t) = \mathrm{mean}(t)/\mathrm{stdev}(r-t). \qquad (5)$$

It is computed as the ratio of the mean value over all channels for a spectrum labeled as truth t and the standard deviation of the differences of t and a reconstructed spectrum r.

Figure 2 shows the mean and standard deviation of this value for all spectra in the lookup table $\mathbf{L}$ with respect to the size of the random sample that was used to compute the principal component matrix $\mathbf{P}$. Randomly selected parameter states are used to compute $n_P = 15$ principal components with n_P kept constant throughout this paper. The number of used principal components in general controls the reconstruction quality, where more components lead to a better reconstruction. Then, the reconstruction SNR is computed for all spectra in the lookup table, and these values are reduced to their mean value and standard deviation. This procedure was repeated 50 times to remove the effects of a particular realization of a random state sample and mean values are shown.

The data clearly show a SNR convergence to ≈ 2500 for randomly chosen sets larger then a few hundred. The standard deviation of the reconstruction SNR also converges to a finite number. This can be understood by analyzing the underlying histograms of the reconstruction SNR for various sample sizes, which is shown in Fig. 3. The analysis shows that the histograms for a sample size of 500 and 5000 are very similar. For these, the wide range of the reconstruction SNR is entirely controlled by the number of principal components and not the sample size that was used to construct the matrix $\mathbf{P}$. The case with a sample size of only 5 is clearly below that range and shows much smaller SNR values, which is caused by the small sample size.

Which of the state vectors in the lookup table are associated with the largest reconstruction errors is probably of

Table 1. State vector parameters of the discussed lookup table database. The state space contains physical parameters as well as the viewing geometry of a hypothetical instrument. The nadir viewing geometry is constrained to 0°–40° where the solar zenith angle and the relative viewing azimuth are less constrained. The aerosol optical models were implemented according to Levy et al. (2007). Here, a rather coarse resolution for all parameter states was chosen to simplify the presented analysis. The total number of $3 \cdot 3 \cdot 3 \cdot 5 \cdot 3 \cdot 3 \cdot 4 \cdot 5 \cdot 5 = 121\,500$ physical cases is considered. The spectral sampling is 0.005 nm, and 4500 channels within the O_2A band were simulated. Altogether this table represents $121\,500 \cdot 4500 \approx 5.5 \times 10^8$ radiative transfer simulations.

Surface pressure ρ	$n_\rho = 3$	$\rho = 800\,\text{hPa}, 950\,\text{hPa}, 1050\,\text{hPa}$
Aerosol optical thickness	$n_\tau = 3$	$\tau = 0.0, 0.1, 1.0$
Aerosol center height	$n_h = 3$	$h = 500\,\text{m}, 2500\,\text{m}, 4500\,\text{m}$
Aerosol type	$n_t = 5$	$t = 1, 2, 3, 4, 5$ (dust, urban, continental, neutral, absorbing)
Surface reflectance at 755 nm	$n_{\alpha_1} = 3$	$\alpha_1 = 0.1, 0.4, 0.7$
Surface reflectance at 780 nm	$n_{\alpha_2} = 3$	$\alpha_1 = 0.1, 0.4, 0.7$
Viewing zenith angle μ	$n_\mu = 4$	$\mu = 0.00, 13.63, 25.88, 38.10$ in deg
Solar zenith angle μ_S	$n_{\mu_S} = 5$	$\mu_S = 0.0, 13.63, 25.88, 38.10, 50.32, 62.53$ in deg
Relative azimuth angle ϕ	$n_\phi = 5$	$\phi = 0.00, 37.89, 75.79, 113.68, 151.58$ in deg
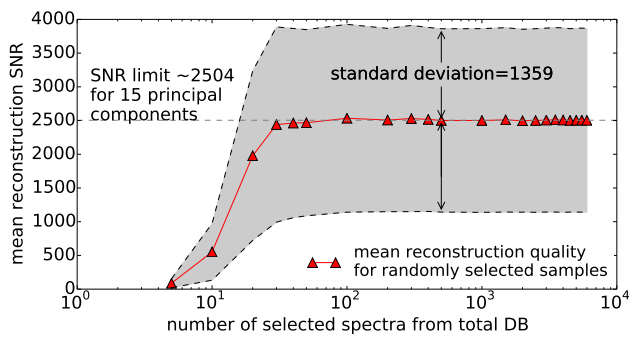
Fig. 2. Reconstruction SNR for a lookup table of top-of-atmosphere radiances in the oxygen A band (see Table 1 and Fig. 5 for reference). The reconstruction SNR was defined as the ratio of the standard deviation of the residual and the mean value of the truth where the residual was defined as the difference between truth and reconstruction. The considered index set was randomly chosen. For each step, computations were repeated 50 times, and the mean value is shown. Shown is the mean reconstruction SNR, which is the mean value over the reconstruction SNR for all cases in the lookup table. The grey shaded area shows the corresponding standard deviation. A histogram for the whole data set is shown in Fig. 3.

Fig. 3. Histograms of the reconstruction SNR for various sample sizes.

interest for users,. Figure 4 shows a histogram of state vector element values of those 10 % of the state vectors exhibiting the smallest reconstruction SNR, that is, the largest reconstruction errors. For this figure, 500 randomly selected states were used to compute the principal components. The states are clearly not equally distributed within this sample, but the over- and under-representation of some states over others is approximately within a factor of two. This fraction is dominated by cases with lower reflectivity where the reconstruction SNR is naturally smaller. This is clearly shown by the fact that surface reflectivity values of 0.7 are not represented within this sample and that the 0.1 case is largely over-represented. In a similar manner, this behavior is shown by the frequency of occurrence of the different viewing
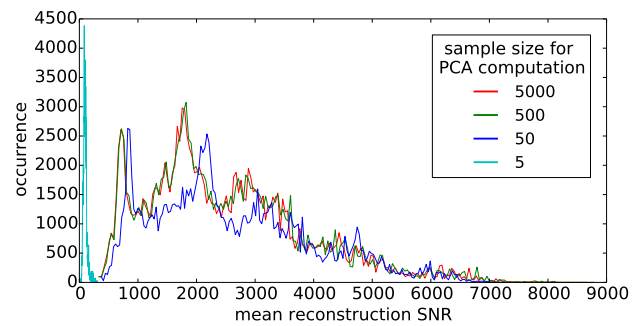
angles. The different aerosol types are almost equally distributed within this sample since their effect on the absolute reflectance is much smaller then the effect from the viewing geometry or surface reflectance.

Although not explicitly shown here, the states that have the smallest reconstruction error are likely among the set that was used to compute the principal components in the first place. This emphasizes the fact that these states should be representative for the total data set to achieve almost uniform distribution of the reconstruction error.

This analysis shows that a fairly small subset of states within $\mathbf{L}$ is sufficient to construct the principal components for the whole data set. This fact immediately raises the question of whether the complete lookup table $\mathbf{L}$ is needed to compute the coefficient matrix $\mathbf{C}$ as defined in Eq. (4), which is discussed in the next section.

5   Choosing small spectral subsets to construct the coefficient matrix C

Equation (4) states how the coefficient matrix $\mathbf{C}$ can be computed from the whole lookup table matrix $\mathbf{L}$ if the principal

Fig. 4. Normalized histograms for parameter state occurrence in the lowest 10 % fraction of a reconstruction SNR data set based on 500 samples. The occurring states are scaled to a range of [0, 1] such that they can be shown on the same axes. The corresponding values are shown in the legend. The occurrences are normalized by the sample size and multiplied by the number of unique parameter values (e.g., three for surface pressure and five for the aerosol type). A value of one indicates equal representation of this particular parameter value within the data set, while larger and smaller values represent over- and under-representation.

Fig. 5. The grey line shows an example spectrum from the lookup table. The red, green and blue lines are the result if only every 500th, 200th, or 100th channel of the original spectrum is considered.

Fig. 6. The grey line shows the first principal component for the complete spectral resolution ($\mathbf{P}_{0,-}$) while the colored lines show the corresponding rows for $\tilde{\mathbf{P}}$ for three sample sizes. The equal distribution sampling technique was used.

component matrix $\mathbf{P}$ is already known. In the previous section, it was shown that only a relatively small subset of the states within $\mathbf{L}$ is needed to compute the principal component matrix $\mathbf{P}$. In this section, it will be discussed how the expansion coefficients can be computed without requiring the knowledge of the total spectrum, such that far fewer radiative transfer simulations have to be carried out. Our approach is to assume that spectral subsamples $\lambda^n = \lambda^{n_\lambda,n}, n \ll n_\lambda$ exist, which are sufficient to compute the expansion coefficients up to an error $O[\lambda^n]$:

$$\mathbf{C}_{-,i} = \mathbf{P}^\mathsf{T} \cdot \mathbf{L}_{-,i} = \tilde{\mathbf{P}}^\mathsf{T}(\lambda^n) \cdot L_{\lambda^n,i} + O[\lambda^n], \tag{6}$$

where $\tilde{\mathbf{P}}(\lambda^n) \in \mathbb{R}^{n \times n_P}$ acts as an effective principal component matrix for a spectral subset, which, after multiplication with the simulations at the selected spectral subset points, produces the coefficients of the original expansion coefficients $\mathbf{C}_{-,i}$. In that respect, $\tilde{\mathbf{P}}$ is a function of a particular spectral subset, and the channel subset λ^n ideally identifies the channels carrying the information sufficient to reconstruct the complete spectrum. Such an idea of identifying the channels that carry the most information has been used in the past, for example, to specify the channel layout in the oxygen A band of the spaceborne remote sensing instrument MERIS on board the Envisat platform (e.g., Kollewe and Fischer, 1994). This led to the definition of three channels with moderate spectral width between 3.75 nm and 15 nm. Here, we focus on reconstruction of the complete hyperspectral data set using the channels that carry sufficient information.

By assuming that $O[\lambda^n]$ vanishes, $\tilde{\mathbf{P}}$ can be computed as the matrix product of the coefficient matrix and the pseudo-inverse of the lookup table matrix for the chosen subsets of states and spectral channels:

$$\tilde{\mathbf{P}}^\mathsf{T}(\lambda^n) = \mathbf{C}_{-,s^{n_s}} \cdot (L_{\lambda^n,n^{n_s}})^{-1}. \tag{7}$$

Here, $s^{n_s} = s^{n_L,n_s}$ is used as a shorthand notation for the state subset that was used to construct $\mathbf{P}$. Necessary for $\tilde{\mathbf{P}}$ to exist is the existence of the pseudo-inverse of the lookup table with respect to both the spectral and the state subset.

Equations (7) and (6) are the two main equations within this paper, since they define the effective principal component matrix $\tilde{\mathbf{P}}$ and show how it can be used to compute the hyperspectral expansion coefficients.

As a consequence, two points must be discussed. Firstly, how can spectral subsets be chosen optimally? And secondly, how many channels according to a selection scheme are required in order to reduce the reconstruction error $O[\lambda^n]$ to a level well below the user requirements?

Three different methods for the selection of spectral subsets are discussed and compared within the next subsections. Other, potentially more efficient methods might exist, as the following discussion is based entirely on heuristic approaches.

Fig. 7. Similar to Fig. 5, but the positions of the spectral channels were chosen with the optimization technique based on the minimization of **C** matrix residuals.

5.1 Equal sampling

Probably the most simple, while still reasonable, method to construct spectral subsamples is to define the size of the sample and spread the channels evenly throughout the spectral band. The reasoning behind this approach is to try to cover as much variability in the spectral band while also employing the high correlation of channels within a spectral interval.

Figure 5 shows this approach for three test cases that correspond to the selection of every 500th, 200th, and 100th channel from the original simulations. The resulting first principal component and the three first effective principle components are shown in Fig. 6.

5.2 Optimization based on random walks

While an equal sampling selection is simple and intuitive, better results might be achieved using optimization techniques. These can be distinguished by the minimized or maximized cost function and the technique employed for the optimization. While for this case it is straightforward to define various cost functions, choosing the optimization technique is more difficult. Commonly used techniques such as the Newton or Levenberg–Marquardt algorithm are based on computing first and second orders of partial derivatives of the optimized cost function. Here, the position of selected channels is to be optimized, and how to apply the standard techniques to this problem does not seem straightforward. The computation of derivatives is not necessary for random walk techniques, which we apply to this problem using two different cost functions.

The random walk starts from a selection of evenly distributed sampling points, and the random step is achieved by adding a randomly chosen step to each position of a selected channel. The maximum range of a random step is chosen to be half the difference of the initial distribution. If a random step leads to an improvement in the cost function, it is chosen as the basis for the next step; otherwise the previous state is used again for the next step. This procedure is repeated

Fig. 8. Similar to Fig. 6, but the optimization technique was used.
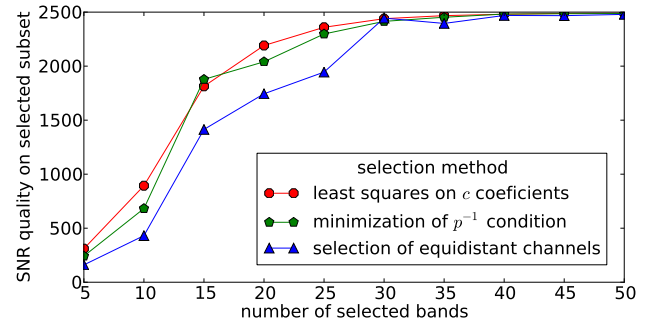
Fig. 9. Results of the mean reconstruction SNR for the three discussed spectral subset selection techniques.

several times, and a maximum number of possible step attempts is prescribed.

Two cost functions appear naturally in this framework:

1. maximization of the l_2 condition of $\mathbf{L}^{-1}_{\lambda^n, s^P}$ with respect to λ^n: $\max_{\lambda^n}(\kappa(\mathbf{L}^{-1}_{\lambda^n, s^P}))$, where the l_2 condition κ of a matrix $\mathbf{A}$ is defined as $\kappa(\mathbf{A}) = \| \mathbf{A}^{-1} \|_2 \cdot \| \mathbf{A} \|_2$, with $\| \mathbf{A} \|_2$ being the l_2 norm of $\mathbf{A}$, which is defined as the square root of the largest eigenvalue of the matrix product of the conjugate transpose of $\mathbf{A}$ and $\mathbf{A}$ itself;

2. minimization of the sum of squares of the difference between the original coefficients $\mathbf{C}_{-, S^P}$ and the reconstructed ones: $\min_{\lambda^n} \sum_{i,j}((\mathbf{C}_{-, S^P})_{i,j} - (\tilde{\mathbf{P}}(\lambda^n) \cdot L^{-1}_{\lambda^n, s^P})_{i,j})^2$.

Resulting spectral subsets using the minimization of **C** matrix residuals are shown in Fig. 7 and the resulting principal components in Fig. 8, while results for all three techniques are shown in Fig. 9.

5.3 Discussion of the reconstruction SNR

Figure 9 shows a comparison of the mean reconstruction SNR for the whole lookup table for the three discussed selection techniques. From the selection of 30 spectral bands

Fig. 10. Residuals between original and reconstructed spectra. The black line with the grey shaded area shows the residual between the original spectra and the one reconstructed using the full spectral information and 15 principal components. The red line shows the residual using only 9 spectral bands, which were selected using the equal sampling technique. The green line shows similar results, but the optimization of **C** matrix residuals was used to select the spectral sample.

Fig. 11. Similar to Fig. 10, but 23 spectral bands were used.

Fig. 12. The grey area shows the reconstruction SNR histograms for the original (i.e., using the full spectral information) reconstruction and two of the discussed reconstruction techniques using small spectral subsamples. The dashed lines show results for the equal sampling technique, while the lines show results for the selection based on the minimization of **C** matrix residuals. The colors indicate the different size of the spectral subsamples.

on, the techniques show only minor differences and are hence equally well suited to solve the problem. Furthermore, the achieved mean reconstruction SNR is equal to the results when using all available spectral bands (compare with Fig. 2). This shows that the highly correlated spectra, when represented with $n_P = 15$ principal components, can be equally well represented with using only 30 selected channels with already known hyperspectral principal components. This behavior is similar to that shown in Fig. 3. There, the number of randomly selected state vector samples approaches a number above which no increase of the reconstruction SNR could be gained by increasing the sample size. In that respect, the state vector sample and the spectral sample show quite similar behavior with respect to their sample size.
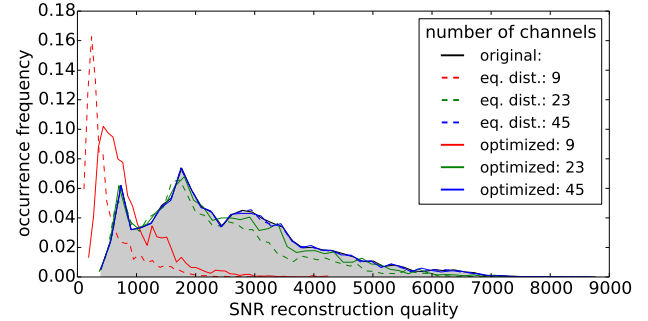
For sample sizes below 30, the techniques show differences, with the least-squares minimization of the **C** matrix residuals showing the best results. If one is willing to accept a larger error by choosing a smaller spectral subsample, the optimization methods are preferred. Using these techniques, similar ranges of the reconstruction SNR can be achieved by using only half of the spectral sample sizes.

One should note that the result for the equal sample selection technique could be optimized by using more suited borders of the spectral interval, but as a matter of fact the results of the optimization techniques are quite similar and might be of smaller importance for the solution of the overall problem.

The effect of choosing different selection techniques is demonstrated in Figs. 10 and 11, showing spectral residuals for the used reconstruction techniques. If 23 spectral bands are used, the results are similar, which corresponds to the results shown in Fig. 9. However, if fewer spectral channels are used, significant differences are apparent. The equal sample technique shows much larger residuals in the 761 nm and

slightly smaller errors elsewhere than the **C** matrix residual minimization technique. This can be nicely explained with the chosen spectral sampling (compare Figs. 5 and 7). The optimization technique has an additional channel moved to that spectral area to compensate for such errors.

Reconstruction SNR histograms for the discussed techniques are shown in Fig. 12. The optimum histogram is reached for the case using 45 channels for both techniques. A good representation of the optimal histogram is achieved by using only 23 channels. For both cases the selection techniques show similar results. Only for the case of using only 9 spectral bands, the optimized selection technique shows much better results.

The main goal of this technique is to reduce the number of radiative transfer calculations needed for a certain problem, which is clearly fulfilled. The achievable speedup however strongly depends on the problem. For a lookup-table-like problem as discussed here, the speedup depends on the number of states in the lookup table n_L, the number of states used for the computation of the principal components n_s, the number of original wavelengths n_λ, and the size of the spectral subsample n. A speedup (sp) can be computed:

$$\mathrm{sp}(n_L, n_s, n_\lambda, n) = \frac{n_L n_\lambda}{n_s \cdot n_\lambda + (n_L - n_s)n}. \qquad (8)$$
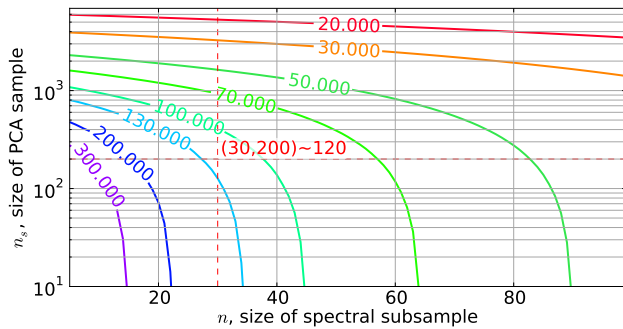
Fig. 13. Shown are lines of constant speedup using the presented technique for the lookup-table-like problem as defined in Table 1. The speedup is presented as a function of the size of the chosen spectral subsample and the size of the state space subsample used for the computation of the principal components. Equation (8) was used to create the figure, and the speedup for $n = 30$ and $n_s = 200$ is highlighted.

The speedup for the discussed lookup table is shown in Fig. 13. It becomes evident that speedups of two orders of magnitude are easily possible using this technique. A special point in this figure unfolds if the additional error introduced by this method is negligible. Figure 2 shows that, from a selection of several hundred spectra on, the mean reconstruction error reaches the limit controlled by the number of principal components. Furthermore, Fig. 9 shows that, from 30 channels on, the reconstruction error reaches its optimum. For the discussed application of a lookup table, a speedup of 120 comes at almost no cost of additional error. Further speedup can be easily achieved, but introduces errors, which have to be checked with the user requirements.

If the radiative transfer is used as a forward model in an inverse problem and the time spent for the computations of the hyperspectral principal is neglected, the speedup with almost no additional error is given by the ratio of the original spectral resolution and the number of channels needed to reconstruct it; here $n_\lambda / n = 4500/30 = 150$.

6 Conclusions

The presented analysis shows that the presented spectral subsampling technique could be employed to achieve major speedups for hyperspectral radiative transfer simulations. Its application to the oxygen A band showed that about 30 spectral channels are sufficient to reproduce the full hyperspectral data in a space spanned by 15 principal components. This number will in general depend on the spectral domain at which this method is applied, and for other domains the validity of the method has to be proven again. However, this is not a major drawback of the method. As it was shown, a number of hyperspectral simulations have to be performed to produce the principal components of the data set. From there, one can establish the validity of the method and benefit from

major speedups. If the validity cannot be established, one has to proceed using different techniques.

We want to highlight two of the main impacts of this technique for problems involving hyperspectral radiative transfer. Firstly, formerly used techniques to gain speedups can be revised. Such techniques could involve the reduction of the vertical resolution of the model atmosphere, neglecting polarization, or only including certain orders of scattering. This method could be used to avoid many of the simplifications made with respect to strict solutions of the radiative transfer equation, and thus to increase the radiative transfer calculation accuracy.

Secondly, this method enables a path to facilitate computationally costly radiative transfer simulations such as full 3-D simulations including polarization. These models have even more complexity in their state vector, including the description of a horizontally and vertically heterogeneous scene, and they are much more demanding in terms of computation time as compared to traditional 1-D methods. Hence, this method could be used to employ 3-D radiative transfer for applications that depend on fast hyperspectral radiative transfer.

As a last point, we want to highlight that this method is entirely based on post-processing techniques and requires no changes in the used radiative transfer code. These are often complex in their implementation and thus practically inaccessible for many users. The method is also simple enough to be implemented with ease using modern interpreted languages and high-level functions for the computation of pseudo-inverses and principal component analysis.

Edited by: S. Schmidt

References

Bennartz, R. and Fischer, J.: A modified k-distribution approach applied to narrow band water vapour and oxygen absorption estimates in the near infrared, J. Quant. Spectrosc. Ra., 66, 539–553, 2000.

Boesch, H., Baker, D., Connor, B., Crisp, D., and Miller, C.: Global Characterization of CO_2 Column Retrievals from Shortwave-Infrared Satellite Observations of the Orbiting Carbon Observatory-2 Mission, Remote Sensing, 3, 270–304, doi:10.3390/rs3020270, 2011.

Doppler, L., Preusker, R., Bennartz, R., and Fischer, J.: k-bin and k-IR: k-distribution methods without correlation approximation for non-fixed instrument response function and extension to the thermal infrared – Applications to satellite remote sensing, J. Quant. Spectrosc. Ra., online first, doi:10.1016/j.jqsrt.2013.09.001, 2013.

Dubuisson, P., Frouin, R., Dessailly, D., Duforêt, L., Léon, J.-F., Voss, K., and Antoine, D.: Estimating the altitude of aerosol plumes over the ocean from reflectance ratio measurements in the O_2 A-band, Remote Sens. Environ., 113, 1899–1911, doi:10.1016/j.rse.2009.04.018, 2009.

Efremenko, D., Doicu, A., Loyola, D., and Trautmann, T.: Optical property dimensionality reduction techniques for accelerated radiative transfer performance: Application to remote sens-

ing total ozone retrievals, J. Quant. Spectrosc. Ra., online first, doi:10.1016/j.jqsrt.2013.07.023, 2013.

Fell, F. and Fischer, J.: Numerical simulation of the light field in the atmosphere-ocean system using the matrix-operator method, J. Quant. Spectrosc. Ra., 69, 351–388, doi:10.1016/S0022-4073(00)00089-3, 2001.

Fischer, J. and Grassl, H.: Detection of cloud-top height from backscattered radiances within the oxygen A band. Part 1: Theoretical study, J. Appl. Meteorol., 30, 1245–1259, 1991.

Hollstein, A. and Fischer, J.: Radiative Transfer Solutions for Coupled Atmosphere Ocean Systems Using the Matrix Operator Technique, J. Quant. Spectrosc. Ra., 113, 536–548, 2012.

Jolliffe, I.: Principal Component, 2nd Edn., Springer, 2002.

Koelemeijer, R. B. A., Stammes, P., Hovenier, J. W., and de Haan, J. F.: A fast method for retrieval of cloud parameters using oxygen A band measurements from the Global Ozone Monitoring Experiment, J. Geophys. Res.-Atmos., 106, 3475–3490, doi:10.1029/2000JD900657, 2001.

Kokhanovsky, A. A. (Ed.): Light Scattering Reviews 8, Springer Praxis Books, 2013.

Kollewe, M. and Fischer, J.: Definition of the oxygen A-band channels of ENVISAT's Medium Resolution Imaging Spectrometer for cloud monitoring, Proc. SPIE, 2309, 205–213, doi:10.1117/12.196677, 1994.

Lacis, A. A. and Oinas, V.: A description of the correlated k distribution method for modeling nongray gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres, J. Geophys. Res.-Atmos., 96, 9027–9063, 1991.

Landgraf, J., Hasekamp, O., van Deelen, R., and Aben, I.: Rotational Raman scattering of polarized light in the Earth atmosphere: a vector radiative transfer model using the radiative transfer perturbation theory approach, J. Quantit. Spectrosc. Ra., 87, 399–433, doi:10.1016/j.jqsrt.2004.03.013, 2004.

Levy, R. C., Remer, L. A., and Dubovik, O.: Global aerosol optical properties and application to Moderate Resolution Imaging Spectroradiometer aerosol retrieval over land, J. Geophys. Res.-Atmos., 112, D13210, doi:10.1029/2006JD007815, 2007.

Lindstrot, R. and Preusker, R.: On the efficient treatment of temperature profiles for the estimation of atmospheric transmittance under scattering conditions, Atmos. Meas. Tech., 5, 2525–2535, doi:10.5194/amt-5-2525-2012, 2012.

Liu, X., Smith, W. L., Zhou, D. K., and Larar, A.: Principal component-based radiative transfer model for hyperspectral sensors: theoretical concept, Appl. Optics, 45, 201–209, doi:10.1364/AO.45.000201, 2006.

Meador, W. E. and Weaver, W. R.: Two-stream approximations to radiative transfer in planetary atmospheres: a unified description of existing methods and a new improvement, J. Atmos. Sci., 37, 630–643, 1980.

Nakajima, M., Kuze, A., and Suto, H.: The current status of GOSAT and the concept of GOSAT-2, Proc. SPIE, 8533, 853306-10, doi:10.1117/12.974954, 2012.

Natraj, V. and Spurr, R. J.: A fast linearized pseudo-spherical two orders of scattering model to account for polarization in vertically inhomogeneous scattering-absorbing media, J. Quant. Spectrosc. Ra., 107, 263–293, doi:10.1016/j.jqsrt.2007.02.011, 2007.

Natraj, V., Jiang, X., lie Shia, R., Huang, X., Margolis, J. S., and Yung, Y. L.: Application of principal component analysis to high spectral resolution radiative transfer: A case study of the band, J. Quant. Spectrosc. Ra., 95, 539–556, doi:10.1016/j.jqsrt.2004.12.024, 2005.

Natraj, V., Shia, R.-L., and Yung, Y. L.: On the use of principal component analysis to speed up radiative transfer calculations, J. Quant. Spectrosc. Ra., 111, 810–816, doi:10.1016/j.jqsrt.2009.11.004, 2010.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, J. Mach. Learn. Res., 12, 2825–2830, 2011.

Preusker, R. and Lindstrot, R.: Remote Sensing of Cloud-Top Pressure Using Moderately Resolved Measurements within the Oxygen A Band-A Sensitivity Study, J. Appl. Meteorol. Climatol., 48, 1562–1574, 2009.

Rothman, L., Gordon, I., Barbe, A., Benner, D., Bernath, P., Birk, M., Boudon, V., Brown, L., Campargue, A., Champion, J.-P., Chance, K., Coudert, L., Dana, V., Devi, V., Fally, S., Flaud, J.-M., Gamache, R., Goldman, A., Jacquemart, D., Kleiner, I., Lacome, N., Lafferty, W., Mandin, J.-Y., Massie, S., Mikhailenko, S., Miller, C., Moazzen-Ahmadi, N., Naumenko, O., Nikitin, A., Orphal, J., Perevalov, V., Perrin, A., Predoi-Cross, A., Rinsland, C., Rotger, M., Simecková, M., Smith, M., Sung, K., Tashkun, S., Tennyson, J., Toth, R., Vandaele, A., and Auwera, J. V.: The HITRAN 2008 molecular spectroscopic database, J. Quant. Spectrosc. Ra., 110, 533–572, doi:10.1016/j.jqsrt.2009.02.013, 2009.

Sanghavi, S., Martonchik, J. V., Landgraf, J., and Platt, U.: Retrieval of the optical depth and vertical distribution of particulate scatterers in the atmosphere using O_2 A- and B-band SCIAMACHY observations over Kanpur: a case study, Atmos. Meas. Tech., 5, 1099–1119, doi:10.5194/amt-5-1099-2012, 2012.

Suto, H., Kuze, A., Shiomi, K., and Nakajima, M.: On-orbit Status of TANSO-FTS on GOSAT, SPIE: Sensors, Systems, and Next-Generation Satellites XIV, 7826, 2010.

Thomas, G. E. and Stamnes, K.: Radiative Transfer in Atmosphere Ocean, Cambridge University Press, 2008.

Veefkind, J., Aben, I., McMullan, K., Förster, H., de Vries, J., Otter, G., Claas, J., Eskes, H., de Haan, J., Kleipool, Q., van Weele, M., Hasekamp, O., Hoogeveen, R., Landgraf, J., Snel, R., Tol, P., Ingmann, P., Voors, R., Kruizinga, B., Vink, R., Visser, H., and Levelt, P.: {TROPOMI} on the {ESA} Sentinel-5 Precursor: A {GMES} mission for global observations of the atmospheric composition for climate, air quality and ozone layer applications, Remote Sens. Environ., 120, 70–83, doi:10.1016/j.rse.2011.09.027, 2012.

Wang, C., Yang, P., Platnick, S., Heidinger, A. K., Baum, B. A., Greenwald, T., Zhang, Z., and Holz, R. E.: Retrieval of Ice Cloud Properties from AIRS and MODIS Observations Based on a Fast High-Spectral-Resolution Radiative Transfer Model, J. Appl. Meteorol. Clim., 52, 710–726, 2013.

Wiscombe, W. and Evans, J.: Exponential-sum fitting of radiative transmission functions, J. Comput. Phys., 24, 416–444, doi:10.1016/0021-9991(77)90031-6, 1977.

Wiscombe, W. J.: Improved Mie scattering algorithms, Appl. Optics, 19, 1505–1509, 1980.