

The GRAPE aerosol retrieval algorithm

G. E. Thomas¹, C. A. Poulsen², A. M. Sayer¹, S. H. Marsh^{1,*}, S. M. Dean^{1,**}, E. Carboni¹, R. Siddans²,
R. G. Grainger¹, and B. N. Lawrence²

¹Atmospheric, Oceanic and Planetary Physics, University of Oxford, Oxford, UK

²Space Science and Technology Department, Rutherford Appleton Laboratory, Chilton, Didcot, UK

*present address: Department of Medical Physics and Bioengineering, Christchurch Hospital, New Zealand

**present address: National Institute of Water and Atmospheric Research, Wellington, New Zealand

Received: 12 March 2009 – Published in Atmos. Meas. Tech. Discuss.: 8 April 2009

Revised: 22 October 2009 – Accepted: 23 October 2009 – Published: 6 November 2009

Abstract. The aerosol component of the Oxford-Rutherford Aerosol and Cloud (ORAC) combined cloud and aerosol retrieval scheme is described and the theoretical performance of the algorithm is analysed. ORAC is an optimal estimation retrieval scheme for deriving cloud and aerosol properties from measurements made by imaging satellite radiometers and, when applied to cloud free radiances, provides estimates of aerosol optical depth at a wavelength of 550 nm, aerosol effective radius and surface reflectance at 550 nm. The aerosol retrieval component of ORAC has several incarnations – this paper addresses the version which operates in conjunction with the cloud retrieval component of ORAC (described by Watts et al., 1998), as applied in producing the Global Retrieval of ATSR Cloud Parameters and Evaluation (GRAPE) data-set.

The algorithm is described in detail and its performance examined. This includes a discussion of errors resulting from the formulation of the forward model, sensitivity of the retrieval to the measurements and *a priori* constraints, and errors resulting from assumptions made about the atmospheric/surface state.

1 Introduction

Despite the important role that atmospheric aerosols play in both climate forcing (both direct and through their interactions with clouds) (IPCC, 2007; Lohmann and Feichter, 2005) and air quality, there are relatively few long term data sets showing their spatial distribution and evolution through time. Imaging satellite instruments offer the ability to provide such measurements and many algorithms have been developed to exploit this ability for specific instruments

(Veefkind and de Leeuw, 1998; Mishchenko et al., 1999; Martonchik et al., 1998, 2002; von Hoyningen-Huene et al., 2003; Remer et al., 2005; Grey et al., 2006). In this paper an optimal estimation algorithm for the retrieval of aerosol loading from space-borne visible/near-infrared radiometers, is described and an indepth characterisation of its theoretical performance when applied to nadir-view measurements from the Along Track Scanning Radiometer 2 (ATSR-2) is given.

The retrieval of aerosol properties from satellite radiometers is a challenging problem for three main reasons:

1. Aerosol retrievals are highly sensitive to cloud contamination. The much greater particle size and optical thickness of clouds compared to background aerosol means that even a small amount of cloud contamination will greatly effect the retrieved aerosol properties.
2. It is difficult to disentangle the top-of-atmosphere (TOA) radiance contribution of aerosol from the surface contribution. This is particularly true over bright and heterogeneous land surfaces, where the TOA signal can dominated by the surface.
3. Even taking the above two factors as read, there are still many more factors effecting the TOA signal than current measurement systems can unambiguously distinguish. These factors include the gaseous composition of the atmosphere as well as properties of the aerosol itself (composition, variation with height, particle shape, mixing state, etc.)

In order to overcome these problems, a number of different approaches to aerosol retrieval from satellite radiometers have been developed. All of these algorithms rely on empirical thresholds of the measured radiance (be it magnitude, ratios at different wavelengths or spatial variability) to remove cloudy pixels (Ackerman et al., 1998; Gómez-Chova et al., 2007; Birks, 2007). The main parameter retrieved by most



Correspondence to: G. E. Thomas
(gthomas@atm.ox.ac.uk)

algorithms is the aerosol optical depth (AOD) at some measurement wavelength. Many also use the spectral variance of the AOD to give an indication of the size and composition of aerosols by using a range of representative aerosol types or mixtures of components (for which the size and composition is fixed) and picking the one that best reproduces the observed radiance.

Perhaps the biggest difference in the various algorithms is in how they separate the surface and atmospheric contributions to TOA signal. The approach taken is largely determined by the capabilities of the instrument being used. Where the instrument has relatively high spectral resolution, such as the MODerate resolution Imaging Spectroradiometer (MODIS) or the MEdium Resolution Imaging Spectrometer (MERIS), so called dark-target algorithms are used. These use an assumption of either the surface reflectance, or the extinction coefficient of aerosol, being effectively zero at a given wavelength. Using this assumption the AOD, or surface reflectance (depending on which assumption is used), can be unambiguously defined, allowing the optical depth and surface reflectance to be derived at different wavelengths on the basis of assumed aerosol properties and a model of the spectral reflectance of the surface. For example, the standard MODIS aerosol product makes use of both of these approaches. Over the ocean (Remer et al., 2005) the assumption is made that the surface reflectance at wavelengths greater than $0.66 \mu\text{m}$ is zero, thus allowing the clear sky TOA signal at such wavelengths to be entirely modelled as a product of aerosol and Rayleigh scattering, plus absorption. A predefined set of bimodal aerosol models can thus be fitted to the observed radiance in 6 channels, incorporating a model of the ocean reflectance for shorter wavelength channels, by varying the mixing ratio between the two modes. The retrieval thus gives the AOD, the ratio between the fine and coarse modes of the aerosol distribution and, by picking the best fitting of the 20 fine and coarse mode combinations, an indication of aerosol type. This basic approach is commonly used for aerosol retrievals over ocean by many other algorithms and instruments including MERIS (Antoine and Morel, 1999), ATSR (Veefkind and de Leeuw, 1998), AVHRR (Mishchenko et al., 1999).

Over land MODIS (Levy et al., 2007) uses the approximation that the AOD at $2.12 \mu\text{m}$ is close to zero, and thus the TOA signal is dominated by the surface contribution. This, along with assumptions of about the spectral variation of the reflectance of vegetated surfaces, the so called dark-dense vegetation (DDV) approximation, allows the AOD to be estimated from measurements at shorter wavelengths. MERIS also uses a DDV approach over land (Santer et al., 1999), although it lacks channels in the infrared, so has to rely solely on constraints on the spectral shape of the surface.

If measurements of the same air mass are available from different viewing geometries it is possible to add angular constraints to the retrieval of aerosol. The main instruments which provide such multi-angle measurements are the

Multi-angle Imaging Spectro-Radiometer (MISR), on board NASA's terra platform, and the ATSR series of instruments launched on ESA satellites. Both instruments not only provide measurements in the nadir direction but also viewing at an angle along the orbital track of the instrument. Due to the high orbital velocity of low Earth orbit satellites, such instruments thereby provide multiple measurements of the same surface point at different viewing zenith angles, separated by only a few tens of seconds. ATSR provides two views, one in the nadir and one centred at 55° along the orbital track, while MISR provides a total of nine views with maximum viewing zenith angles of approximately $\pm 70^\circ$. Over land, such measurements allow the separation of the atmospheric contribution to the TOA radiance from that from the surface because the ratio of the surface reflectance at different viewing geometries shows little angular dependence, since the size of the surface scattering elements are so much larger than the wavelengths being used (Flowerdew and Haigh, 1995). The simplest approach is to assume that this ratio is in fact constant with respect to wavelength, as is used by Veefkind and de Leeuw (1998) for ATSR retrievals. North (2002) shows how it is also possible to use the angular constraint to fit a simple empirical surface BRDF model to ATSR measurements, providing a retrieval of spectral surface reflectance. In both approaches, the atmospheric path radiance for each instrument channel and view is fitted using a range of aerosol types, thus providing an estimate of the AOD.

This angular constraint approach is also used in MISR retrievals (Dinner et al., 2005), but only as a first step in reducing the number of aerosol types made available to a second retrieval, based on a different assumption. In this second approach, described by Martonchik et al. (1998), the clear-sky atmosphere is assumed to be homogenous over a region of $17.6 \times 17.6 \text{ km}$ (containing 256 individual MISR pixels). Under the assumption that the atmospheric and surface signals are additive, a spatial scatter matrix can be constructed that doesn't depend on the atmospheric path radiance, by expressing the signal in each pixel as a bias relative to some reference pixel within the retrieval area. The TOA radiance for each pixel can then be reconstructed using a principal component decomposition of this matrix and an atmospheric path radiance calculated using a range of aerosol types.

The ORAC algorithm is somewhat different than the algorithms described above, in that it is not built around any particular method of separating the surface and atmospheric contributions to the TOA signal. Rather it makes use of an optimal estimation retrieval scheme to fit modelled radiances across a series of wavelength bands (channels) to radiances measured by a satellite instrument, as a function of aerosol optical depth, effective radius and surface reflectance, subject to constraint by *a priori* knowledge of these parameters. The algorithm was developed from the Enhanced Cloud Processor (ECP) (Watts et al., 1998) and the version described here is part of a unified cloud and aerosol algorithm. This algorithm uses a empirical cloud flagging scheme to

differentiate cloud and aerosol pixels. Pixels identified as cloud have the cloud retrieval algorithm applied, with the remaining pixels having the aerosol retrieval applied. The algorithm has been applied to ATSR-2, Advanced-ATSR and Spinning Enhanced Visible-InfraRed Imager (SEVIRI) data (Kokhanovsky et al., 2007; Thomas et al., 2007b), including the creation of the GRAPE global cloud and aerosol dataset derived from the full ATSR-2 data-set, spanning 1995–2001¹. For descriptions of the ATSR-2, AATSR and SEVIRI instruments the reader is referred to Mutlow et al. (1999), Llewellyn-Jones et al. (2001) and Aminou et al. (1997), respectively.

It is important to emphasize that this version of the algorithm, hence referred to as the GRAPE algorithm to distinguish it from the wider family of algorithms which carry the name ORAC, is primarily a cloud retrieval. The inclusion of a simple aerosol retrieval can be considered a way of adding value to the GRAPE cloud products and does not represent the state-of-the-art ORAC aerosol retrieval. There currently exist two further versions of the ORAC aerosol algorithm, which are more advanced stand alone aerosol retrievals. The GRAPE algorithm uses a relatively simple forward model, particularly in its treatment of the surface reflectance, which (in common with the cloud retrieval) is treated as Lambertian, and in its treatment of different aerosol types. This results in important limitations to the algorithm:

- The algorithm is only applicable over surfaces which can be reasonably approximated by a Lambertian reflectance, such as the ocean surface far from the sunglint region, and homogeneous land surfaces.
- The use of a Lambertian surface reflectance prohibits the inclusion of near simultaneous observations of the same scene at different viewing geometries, such as offered by the ATSR instruments' dual-view system. Even if a surface can be approximated with an effective Lambertian reflectance for a given viewing geometry, this value can be different for different viewing geometries.
- In the case of instruments such as ATSR-2 and SEVIRI, which have a small number of channels in the visible and near-infrared, the retrieval becomes highly dependant on good *a priori* knowledge of the surface reflectance. This is because the measurements do not contain enough information to decouple the surface and atmospheric components of the signal.

Despite these limitations, the GRAPE dataset has shown that the algorithm can provide aerosol properties of sufficient

¹GRAPE was a UK Natural Environment Research Council project. The full GRAPE data-set is freely available for download from the British Atmospheric Data Centre. See <http://badc.nerc.ac.uk/data/grape/> for further details.

accuracy (i.e. with a random error of $\lesssim 0.1$) to be scientifically interesting, particularly over the open ocean, and has been applied in the study of aerosol cloud interactions (Bulgin et al., 2008; Quaas et al., 2009). A validation of the GRAPE oceanic AODs against ocean and coastal AERONET sites, and comparisons with the Global Aerosol Climatology Project (GACP) AOD product derived from Advanced Very High Resolution Radiometer measurements, is presented by Thomas et al. (2009b). This study finds a correlation of 0.79 between the AERONET and GRAPE AODs over ocean pixels, with a best fit line of $\tau_G = (0.08 \pm 0.04) + (1.0 \pm 0.1)\tau_A$, where τ_G and τ_A are the 550 nm AODs from GRAPE and AERONET, respectively. Over land pixels, however, the GRAPE product AOD product has been found to be very noisy and has thus not been validated. The 0.08 offset in GRAPE oceanic AODs when compared to AERONET measurements could be partially due to the coastal location of most of the AERONET sites, as the ocean surface model used to set the *a priori* is most accurate when applied to the deep ocean. For an overview of all versions of the ORAC aerosol algorithm, including a multiview algorithm that makes use of a BRDF description of the surface reflectance, the reader is referred to Thomas et al. (2009a).

In Sect. 2 the ORAC forward model and numerical retrieval scheme are described. Sect. 3 is devoted to a discussion of the algorithm's performance and limitations, including the sensitivity of the algorithm and investigation of sources of error. The results are summarised and conclusions drawn from them in Sect. 4.

2 Forward model and retrieval algorithm

2.1 The forward model

The GRAPE algorithm aerosol forward model can be thought of as being composed of four components:

1. A model of aerosol radiative properties (single scatter albedo, extinction coefficient and phase function) at the required wavelengths as a function of aerosol size distribution and spectral refractive indices. Ideally the aerosol radiative properties should be computed across the channel bandpass functions of the instrument in question, but for most radiometers the width of each channel is small compared to the scale over which the aerosol radiative properties vary and thus the computation can be done for the weighted centre wavelength of each channel.
2. A model of gas absorption (or emission) over the bandpass of each channel.
3. A model or measurements of the surface reflectance, both of the land and ocean.

4. A radiative transfer model, which predicts the top of atmosphere radiance as a function of the radiative properties produced by the aerosol and gas models, the surface reflectance, and illumination and viewing angle.

For reasons of numerical efficiency, radiative transfer calculations are done off-line for a range of aerosol properties and viewing geometries to produce look-up tables, which are then used to predict the radiance during retrieval runs. The GRAPE algorithm uses Mie scattering to relate aerosol microphysical properties to their radiative properties. It should be noted that Mie scattering is only applicable to spherical particles and thus its use will introduce errors in the radiative properties of non-spherical aerosol particles (the most notable example being mineral dust). Although not included in this analysis, the development of non-spherical scattering code for use with the ORAC scheme is underway.

Microphysical properties are taken from published descriptions of typical atmospheric aerosol types, such as those found in the Optical Properties of Aerosols and Clouds (OPAC) database (Hess et al., 1998). Such models typically describe a range of atmospheric aerosol classes, each of which is made up of an external mixture of multiple components of a given refractive index, size distribution and mixing ratio. In external mixtures each component is assumed to exist separately from the others, with its own size distribution. For use with the GRAPE algorithm, each component is assumed to have a log-normal size distribution

$$n(r) = \frac{N_0}{\sqrt{2\pi}} \frac{1}{\ln S} \frac{1}{r} \exp \left[-\frac{(\ln r - \ln r_m)^2}{2 \ln^2 S} \right] \quad (1)$$

where the median radius r_m and spread S (where the standard deviation of $\ln r$ is $\ln S$) are taken from the aerosol database.

Mie scattering theory is then used to produce the single scattering albedo (ω_i), extinction coefficient (β_i^{ext}) and phase function (p_i) of each component, i , from which the values for the complete aerosol class can be calculated using:

$$\beta_{\text{ext}} = \frac{\sum_i \chi_i \beta_i^{\text{ext}}}{\sum_i \chi_i} \quad (2)$$

$$\omega = \frac{\sum_i \chi_i \beta_i^{\text{ext}} \omega_i}{\sum_i \chi_i \beta_i^{\text{ext}}} \quad (3)$$

$$p = \frac{\sum_i \chi_i \beta_i^{\text{ext}} \omega_i p_i}{\sum_i \chi_i \beta_i^{\text{ext}} \omega_i} \quad (4)$$

where χ_i is the mixing ratio of the i th component.

To enable the retrieval of aerosol size information, these quantities must be calculated for a range of different aerosol size distributions. The measure of the aerosol size used by the GRAPE algorithm is the effective radius, which is defined as the ratio of the third and second moments of the aerosol size distribution. For a log normal distribution this is

$$r_e = \frac{\sum_i \chi_i r_{m,i}^3 \exp(4.5 \ln^2 S_i)}{\sum_i \chi_i r_{m,i}^2 \exp(2 \ln^2 S_i)} \quad (5)$$

Each aerosol class will have a prescribed r_e (defined by the r_m and χ of each component). Optical properties are produced for a range of effective radii by changing the mixing ratios of the components. If an r_e is desired that is either smaller than that of the smallest component, or larger than that of the largest component, the aerosol class becomes a single component aerosol (since the r_e will have been minimised, or maximised, by setting the mixing ratio of the appropriate component to 100%). In such a case, the desired r_e is achieved by scaling the median radius of the one remaining component. It is obvious that by changing the effective radius in this fashion, the aerosol class will quickly lose its resemblance (not only in terms of size, but composition) to its source class as defined by the literature. Thus, the GRAPE algorithm works on the implicit assumption that our *a priori* knowledge of the aerosol we expect to observe is reasonably accurate, so that large changes to the effective radius are not needed to match the observed radiances. In practice, retrievals done using real data show that retrieved effective radii are very rarely so different from the *a priori* that the aerosol becomes a single component, at 0.01%–1% of retrieved states, depending on the aerosol class.

The modelling of gas absorption or emission over the wavelength bands of the instrument is performed using MODTRAN (Berk et al., 1998). As the retrieval uses visible/near-infrared wavelengths, the signal from atmospheric gases is dominated by Rayleigh scattering rather than absorption/emission. Thus temporal and spatial variation in the gas composition of the atmosphere can be neglected and standard-atmosphere gas concentrations can be used². The modelled gas absorption is then convolved with the instrument filter function at each channel to produce an optical depth due to gas absorption/emission.

Both the aerosol radiative properties and the gas optical depth are then passed to the DISORT (Stamnes et al., 1988) radiative transfer code, which models the radiance for each combination of aerosol effective radius and optical depth (determined by the number density of aerosol) over the plausible range of illumination and viewing angles for the given satellite. A 32 layer model of the atmosphere, extending to a height of 100 km is used for the DISORT calculations. In layers containing aerosol, the effects of molecular absorption, Rayleigh scattering and aerosol scattering and absorption are combined using the following expressions:

$$\tau_1 = \tau_a + \tau_R + \tau_g \quad (6)$$

$$\omega_1 = \frac{\tau_R + \tau_a \omega_a}{\tau_a + \tau_R + \tau_g} \quad (7)$$

$$p_1 = \frac{\tau_R p_R + \tau_a \omega_a p_a}{\tau_R + \tau_a \omega_a} \quad (8)$$

²This is particularly true for the ATSR series of instruments, which have channels designed specifically to avoid atmospheric gas spectral features. It has also been found to be an acceptable assumption for the SEVIRI instrument, despite its 0.87 μm channel's slight sensitivity to water vapour.

where τ_a , τ_R , τ_g are the optical depths due to aerosol, Rayleigh scattering and molecular absorption within the layer, which combine to give a total optical depth of τ_l for the layer. The single scattering albedo for each layer, ω_l is given as the average of the single scattering albedo of the aerosol, Rayleigh scattering (for which $\omega_R \equiv 1$) and molecular absorption (where $\omega_g \equiv 0$), weighted by their optical depths. Likewise, the phase function for each layer is also a weighted mean, with weights being the fraction of the optical depth for each process which is due to scattering (i.e. $\tau\omega$). All of these calculations are performed assuming a black surface and assume a constant illumination, producing a set of effective transmissions and reflectances for diffuse and direct-beam radiation for the atmosphere as a whole.

This allows radiative transfer through the Earth’s atmosphere to be treated as a series of partial reflections and transmissions. The first of these is a “reflection” of the solar beam off the atmosphere itself, described by a bidirectional reflectance $R_{bb}(\theta_0, \theta_v, \phi)$, as a function of solar and viewing zenith angles, θ_0 and θ_v and the relative azimuth angle ϕ . The subscript bb indicates that this reflectance function takes a direct-beam input and provides a direct-beam output. This is followed by a term composed of the transmission of the solar beam through the atmosphere ($T_{bt}^\downarrow(\theta_0)$, where bt indicates the transmission function accepts a direct beam as input and produces both beam and diffusely transmitted components), reflection off a Lambertian surface with reflectance ρ and transmission of the resulting diffuse radiation back through the atmosphere ($T_{tb}^\uparrow(\theta_v)$) to the satellite. Subsequent terms result from multiple reflectances between the surface and the atmosphere (described by a diffuse reflectance R_{dd}). The resulting geometric series:

$$\begin{aligned}
 R = & R_{bb}(\theta_0, \theta_v, \phi) + T_{bt}^\downarrow(\theta_0)\rho T_{tb}^\uparrow(\theta_v) \\
 & + T_{bt}^\downarrow(\theta_0)\rho^2 T_{tb}^\uparrow(\theta_v) R_{dd} \\
 & + T_{bt}^\downarrow(\theta_0)\rho^3 T_{tb}^\uparrow(\theta_v) R_{dd}^2 \\
 & + \dots
 \end{aligned}
 \tag{9}$$

can be simplified using the appropriate series limit to give

$$R = R_{bb}(\theta_0, \theta_v, \phi) + \frac{T_{bt}^\downarrow(\theta_0)\rho T_{tb}^\uparrow(\theta_v)}{1 - \rho R_{dd}}.
 \tag{10}$$

Lookup tables of the quantities $R_{bb}(\theta_0, \theta_v, \phi)$, $T_{bt}^\downarrow(\theta_0)$, $T_{tb}^\uparrow(\theta_v)$, and R_{dd} are created by DISORT as a function of aerosol optical depth, effective radius, solar zenith angle, satellite zenith angle and relative azimuth angle, as detailed in Table 1.

The final aspect of the forward model is a description of the reflectance of the Earth’s surface across each of the desired channels. The GRAPE algorithm retrieves the surface albedo at $0.55\mu\text{m}$ but the spectral shape (i.e. the ratio of the albedo at each channel to $0.55\mu\text{m}$) is fixed by the *a priori* spectral albedo. The determination of the *a priori* albedo is

Table 1. Details of the lookup tables of atmospheric transmission and reflection used in the GRAPE algorithm forward model. The dashes indicate the dimensions spanned by each LUT, for example, R_{dd} is a function of $\log_{10}(\tau_a)$ and $\log_{10}(r_e)$. The LUTs use an equally spaced grid in all dimensions.

LUT name	Function of				
	$\log_{10}(\tau_a)$ 9 values: –2–0.301	$\log_{10}(r_e)$ 20 values: –2–1	θ_0 10 values: 0–90	θ_v 10 values: 0–90	ϕ 11 values: 0–180
R_{bb}	–	–	–	–	–
T_{bt}^\downarrow	–	–	–	–	–
T_{tb}^\uparrow	–	–	–	–	–
R_{dd}	–	–	–	–	–

treated separately for land and ocean pixels. Over land the MOD43B Bi-directional Reflectance Distribution Function (BRDF) product from MODIS (Jin et al., 2003) is used to produce a Lambertian reflectance at each channel. For ocean pixels a sea surface reflectance model (Sayer, 2007) is used, based on Cox and Munk wave slope statistics as a perturbation to Fresnel reflection.

2.2 Retrieval algorithm

ORAC is an optimal estimation retrieval scheme, based on the formalism of Rodgers (2000). It uses a numerical minimisation to determine the maximum *a posteriori* solution, that is the solution that has the maximum probability of being the “truth” according to Bayesian statistics, to the problem of matching the observed and modelled radiances. This solution is given by minimising, via an iterative optimisation, the cost function:

$$\begin{aligned}
 J(\mathbf{x}) = & (\mathbf{Y} - \mathbf{F}(\mathbf{x}))^T \mathbf{S}_\epsilon^{-1} (\mathbf{Y} - \mathbf{F}(\mathbf{x})) \\
 & + (\mathbf{x} - \mathbf{x}_a)^T \mathbf{S}_a^{-1} (\mathbf{x} - \mathbf{x}_a)
 \end{aligned}
 \tag{11}$$

where \mathbf{Y} is the vector of measured radiances, $\mathbf{F}(\mathbf{x})$ is the corresponding forward-modelled radiances and \mathbf{S}_ϵ is the measurement covariance matrix, which defines the uncertainty of each measurement and the correlations between them. \mathbf{x} is the vector of state variables (i.e. the values we are retrieving), with the *a priori* state vector \mathbf{x}_a and corresponding covariance matrix \mathbf{S}_a , defining our knowledge of the state before the measurement and the uncertainty in this knowledge, respectively. The optimal estimation framework also allows for the explicit inclusion of known errors in the forward model, or parameters on which it depends, which both effect the value of $\mathbf{F}(\mathbf{x})$. This is achieved by modifying \mathbf{S}_ϵ by the addition of a further covariance matrix characterising the forward-modelling uncertainties.

The state vector for the GRAPE retrieval is made up of the aerosol optical depth at 550 nm, the aerosol effective radius and the surface reflectance at 550 nm:

$$\mathbf{x} = \begin{bmatrix} \tau_a \\ r_e \\ R_s(550) \end{bmatrix}. \quad (12)$$

The inclusion of prior knowledge of the state in a rigorous way is a particular strength of the optimal estimation approach. The retrieval can be viewed as a process of using the measurements to better define our prior knowledge of the state, whether that knowledge be based on a climatology or on independent measurements of the state.

The other advantage of the optimal estimation framework over other, more ad-hoc, retrieval methods is that it provides information on the quality of the retrieval. The first indicator of a successful retrieval is that it has converged. Successive iterations each produce their own estimate of the state and when the difference between these estimates becomes a small fraction of their estimated precision, the retrieval can be said to have converged. In practice, an upper limit is put on the number of iterations the retrieval can take (in the case of GRAPE this is set to 25): if a retrieval has not converged within this number of steps, it is deemed to have failed. In addition, the value of the cost function at the final solution and error estimates on the retrieved state provide indicators as to the quality of the retrieval. The former indicates how consistent the solution is with the measurement and *a priori* information, while the latter defines how well constrained the solution is.

To calculate uncertainties on the retrieved state, the assumption is made that for small changes in the state vector the forward model can be considered linear, so that it can be expressed as

$$\mathbf{F}(\mathbf{x}) - \mathbf{F}(\hat{\mathbf{x}}) = \hat{\mathbf{K}}(\mathbf{x} - \hat{\mathbf{x}}). \quad (13)$$

Here $\hat{\mathbf{x}}$ and \mathbf{K} are the state vector and weighting function matrix at the solution. The weighting function matrix contains the derivatives of the forward model with respect to the state vector:

$$K_{ij} = \frac{\partial F_i(\mathbf{x})}{\partial x_j}, \quad (14)$$

where the i and j subscripts denote elements of the measurements and state vectors, respectively. Under this assumption, we can relate the covariance of the measurements and *a priori* state to the uncertainty in the retrieved state $\hat{\mathbf{S}}$ via,

$$\hat{\mathbf{S}} = \left(\hat{\mathbf{K}}^T \mathbf{S}_\epsilon^{-1} \hat{\mathbf{K}} + \mathbf{S}_a^{-1} \right)^{-1}. \quad (15)$$

ORAC uses the Levenburg-Marquardt algorithm (Levenberg, 1944; Marquardt, 1963) to perform the minimisation of the cost function. This provides an extremely robust and numerically efficient method for minimisation which will converge successfully even for highly non-linear problems.

However it has no mechanism to avoid becoming trapped in local minima in the cost function. Furthermore, if the forward model is highly non-linear in the region surrounding the solution, the optimal estimation error propagation will no longer be valid and uncertainties reported on the solution will be inaccurate.

3 Retrieval performance

In this section we examine the theoretical performance of the ORAC aerosol retrieval algorithm in the configuration used in the GRAPE project, using the results of retrieval runs on simulated data that resembles measurements made by the ATSR-2 (or AATSR) instrument. Two questions are addressed:

1. How sensitive is the retrieval to changes in the state parameters? This gives us a measure of the best possible precision and accuracy we can expect the retrieval scheme to achieve, given the measurements we are providing.
2. How sensitive is the retrieval to errors in the forward model and assumptions made about aerosol and surface properties.

The configuration of the ORAC processor used in the GRAPE project uses channels 2–4 of the ATSR-2 instrument (centred on wavelengths of 0.67, 0.87 and 1.6 μm). The 0.55 μm channel is not used, as over the oceans it is either disabled, or in a *narrow-swath* (where the swath width of the instrument is reduced from 512 to 256 km), due to down-link bandwidth limitations of the ERS-2 satellite. The retrieval assumes an aerosol type based on the geographic location of each pixel and time of year, as shown in Fig. 1, with the possible types being maritime clean, continental average, desert dust, Arctic and Antarctic classes from the OPAC database.

The retrieval of parameters at 550 nm when the retrieval does not utilise a measurement at this wavelength may initially appear strange and it might be presumed that the τ_a and R_s are actually being retrieved at 670 nm (or some other wavelength) and then extrapolated to 550 nm. However, this is not the case. The optimal estimation retrieval fits all parameters with all available measurements, weighted by \mathbf{S}_ϵ , constrained by the assumed aerosol properties. Thus, although the inclusion of the 550 nm channel would undoubtedly better constrain the retrieval and lead to smaller uncertainties on the retrieved parameters, the retrieved state would remain consistent with that retrieved without it.

It should be noted that the assumed aerosol microphysical properties are a key source of uncertainty in all aerosol remote sensing applications and that describing the global distribution of aerosol with a small number of discrete classes is not a realistic reflection of its true diversity. Further more, our knowledge of typical aerosol properties




Fig. 1. The spatial distribution of aerosol types used in the GRAPE aerosol product. The left map shows the distribution used for Southern Hemisphere summer (October–March), the Northern Hemisphere summer (April–September) is given on the right.

in different regions is improving continuously, particularly thanks to Alumcantar retrievals from Aerosol Robotic Network (AERONET) measurements (Dubovik et al., 2002), and some of the classes within the OPAC database (particularly those which contain absorbing components) are now known to be inaccurate. However, the small number of channels used in the GRAPE retrieval means that there is not enough information available to the retrieval to enable different aerosol types to be distinguished. The magnitude of the errors introduced by this simplistic assumption of the aerosol optical properties is investigated in Sect. 3.3.

Absolute measurement uncertainties are set at fixed values of 0.005, 0.009, and 0.018 for channels 2, 3 and 4. For average cloud-free radiances these errors correspond to errors of approximately 2, 3 and 10% for each channel, respectively. These values were determined for cloud retrievals but are in line with the expected noise on ATSR-2 measurements (Smith et al., 2002), except for channel 4 where it is an overestimate. The larger uncertainty on channel 4 was originally set because it was not subjected to the same pre-launch calibration applied to the other channels. The use of fixed measurement errors simplifies the retrieval code somewhat, but subsequent applications of ORAC retrieval will implement a more rigorous scheme, based on the magnitude of the measured TOA reflectance and its spatial variability.

Additional forward model error to account for both the Lambertian surface reflectance approximation and the constraint of a fixed spectral shape to the surface reflectance is also included in S_ϵ at 20% of the *a priori* surface reflectance and correlations between the channels (i.e. off-diagonal elements in S_ϵ) of 0.4. Again, these values were optimized for the cloud retrieval.

The aerosol optical depth and effective radius are retrieved in \log_{10} space, while the surface reflectance is retrieved on a linear scale. The *a priori* and first guess surface albedo are set to the value defined by the sea-surface reflectance model over the ocean or the MODIS white-sky albedo over the land, with a 1σ error of 0.01. The *a priori* and first guess values for optical depth are set to $\log_{10}(\tau) = -1.0 \pm 1.0$, corresponding to $\tau=0.1$ with 1σ error bounds of $0.01 \leq \tau \leq 1.0$. The *a priori* and first guess effective radius is set to the value prescribed in literature for each aerosol type, with error bounds of ± 0.5 in $\log_{10}(r_e)$. Table 2 summaries the state variables and their associated *a priori*.

3.1 Retrieval sensitivity

The foremost limitation on the performance of a retrieval system is the information content of the measurements themselves and how sensitive each retrieved parameter is to perturbations in the state. This limitation can be investigated by performing the retrieval on simulated data, so that all sources of forward model and forward model parameter error can be removed from the problem. Figure 2 shows the aerosol state parameters (optical depth and effective radius) which have been used to produce simulated ATSR-2 radiances using the optical properties of the OPAC maritime clean aerosol class.

The retrieval has been run on these data, assuming the correct aerosol class and with the *a priori* surface reflectance set to the correct value of 0.02. The *a priori* effective radius for the maritime clean class is $0.832 \mu\text{m}$, giving 1σ error bounds of $0.26 < r_e < 2.63 \mu\text{m}$. The results of applying the retrieval to these data are given in Fig. 3, along with error estimates derived from the diagonal of the state covariance matrix, the value of the cost function at the solution and the number of iterations required for convergence. The first thing to notice is that some retrievals have failed to converge




Fig. 2. The optical depth and effective radius fields used to produce simulated radiances for the ORAC retrieval. The ranges have been chosen to span the full range of optical depths and effective radii covered by the ORAC look-up tables. Ideally, the retrieval should reproduce these plots when applied to the simulated radiances.

at the lowest effective radius (as indicated by white spaces in Fig. 3). This can be attributed to the optical depth being poorly constrained for the lowest size bin, as indicated by the large uncertainty estimates in its value (Fig. 3b). Poor performance of the retrieval in the lowest size bin is a recurring theme in the results presented in this paper for this reason, and thus will not be discussed for each individual set of retrievals.

Figure 3g shows that the cost function at the solution has a value which is always less than 6. The smooth dependence on the state exhibited by the cost function is due to the *a priori* portion of the cost function $((\mathbf{x} - \mathbf{x}_a)^T \mathbf{S}_a^{-1} (\mathbf{x} - \mathbf{x}_a))$: i.e. the retrieval is fitting the measurements extremely well for all states, with the cost function being essentially determined by the distance from the *a priori*. Statistically, one expects the cost function to follow a χ^2 distribution with a single degree of freedom³. This is not the case in this example, because the simulated measurements used in the retrieval did not have noise added to them, while the retrieval was run using the error covariance matrix described earlier in this section. Thus the forward model can consistently fit the measurements more accurately than predicted by the measurement covariance matrix, \mathbf{S}_ϵ and the retrieved uncertainties over estimate the error in the retrieval.

It is also very clear that the retrieval has failed to accurately estimate the true fields for states with low optical and large effective radius. To further explore why this is the case, we will examine the retrieval statistics further. Rodgers (2000) defines the averaging kernel for the maximum *a posteriori* solution as

$$\mathbf{A} = \mathbf{S}_a \hat{\mathbf{K}}^T (\hat{\mathbf{K}} \mathbf{S}_a \hat{\mathbf{K}}^T + \mathbf{S}_\epsilon)^{-1} \hat{\mathbf{K}}. \quad (16)$$

³The cost function reported by ORAC is normalised by the number of measurements. If it were not, the degrees of freedom would equal the number of measurements (Rodgers, 2000).

Table 2. Retrieval state parameters and *a priori* assumptions. *A priori* errors are given as 1σ .

State parameter	<i>a priori</i> value	<i>a priori</i> error
$\log_{10}(\tau_a)$	−1.0	1.0
$\log_{10}(r_e)$	literature value	0.5
R_s	MODIS over land, modelled over ocean	0.01

This matrix gives the sensitivity of the retrieved state to perturbations in the true state,

$$A_{ij} = \frac{\partial \hat{x}_j}{\partial x_i}, \quad (17)$$

where \hat{x}_j is the j th element of the retrieved state and x_i is the i th element of the true state. The diagonal elements of this matrix can be thought of as an indication of the fraction of the retrieved state which can be said to be determined by the true value of that quantity (with the rest being determined by the choice of *a priori* and the value of the other elements of the state). For a perfect retrieval system \mathbf{A} would be an identity matrix, while a value of zero on the diagonal indicates that the corresponding state element is entirely determined by the *a priori* and values of the other state elements. The trace of the averaging kernel also gives the degrees of freedom for signal for the given retrieval. This quantity gives the number of pieces of independent information retrieved. It is important to realise that d_s is not a direct estimate of the information content of the measurement, but rather indicates how much the measurement is able to improve our prior knowledge of the state. Figure 4 shows the diagonal elements of \mathbf{A} , as well the overall d_s corresponding to the retrieval shown in Fig. 3. The first thing to note is that $d_s \lesssim 2.5$, indicating that the retrieved state is always somewhat influenced by the *a priori*. Looking at Fig. 4a–c, we can see that




Fig. 3. The retrieval results when the GRAPE algorithm is applied to the simulated data field with the correct (OPAC maritime-clean class) aerosol optical properties. (a) and (b) show the retrieved optical depth and its uncertainty, respectively. Similarly, (c) and (d) show the retrieved effective radius and its uncertainty, while (e) and (f) show the surface reflectance and its uncertainty. (g) shows the value of the cost function, Eq. (11), at the solution and (h) shows the number of iterations required for convergence. White regions indicate where the retrieval failed to converge.




Fig. 4. The information content of the retrieval when applied to the simulated data field with the correct (OPAC maritime-clean class) aerosol optical properties. (a)–(c) show the diagonal elements of the averaging kernel for optical depth, effective radius and surface reflectance, respectively. (d) shows degrees of freedom for signal (i.e. the trace of the averaging kernel) for each retrieval.

in general optical depth is most sensitive to the measurement, as is effective radius at large optical depths, whereas the surface reflectance is mostly dominated by the *a priori* value.

It can also be seen that the regions of the domain where the retrieval has done least well (in particular, where optical depth is low and effective radius high) correspond to regions where d_s is low and the effective radius shows poor sensitivity to the true state. In such circumstances the effective radius is held at the *a priori* value and this results in an error in the retrieved optical depth, despite its good sensitivity to the true state.

Figure 5 shows the distribution of retrieved error estimates for each of the state elements, along with the distribution of the difference between the retrieved and true states. It can be seen that most retrievals provide optical depth to a precision between 0.05 and 0.15, while the majority of effective radii have a precision of less than 0.1. It is clear that for most states, the retrieved optical depth is more accurate than indicated by the retrieved error estimate. As with the low values of the cost function shown in Fig. 3g, this can be explained by the fact that no measurement noise was added to the simulated radiances used in the retrieval. The retrieved error estimates give the 1σ confidence interval on the retrieved values, given the measurement and *a priori* uncertainties: since the simulated measurements actually contain no error, the accuracy of the retrieval is substantially better than this estimate.

The effective radius errors show a similar pattern, although this is not apparent in Fig. 5b, due to the relatively coarse x -axis scale.

The surface reflectance error distribution shows somewhat different behaviour, since the retrieval used the correct value for the first guess and *a priori*. In addition, the surface reflectance is retrieved on a linear scale, while the \log_{10} of optical depth and effective radius are retrieved. Thus the maximum of the retrieved surface reflectance error distribution is defined by the *a priori* error (while the retrieved error for optical depth and effective radius depends on the value of the retrieved state and hence has a relative, rather than an absolute, maximum). It is clear from Fig. 5c that the retrieval has somewhat narrowed the confidence interval on the value of surface reflectance for many states, but many more have not been improved at all.

Taken together, Figs. 3, 4 and 5 tell us several things about the retrieval:

- Overall, the retrieval is working well as there is an improvement in our knowledge of almost all states (given by the narrowing of the uncertainties from their *a priori* values) and the retrieved states almost always agree with the true value within uncertainty estimates.
- The retrieval works best at high optical depths or effective radii between approximately 0.015 and 0.10.

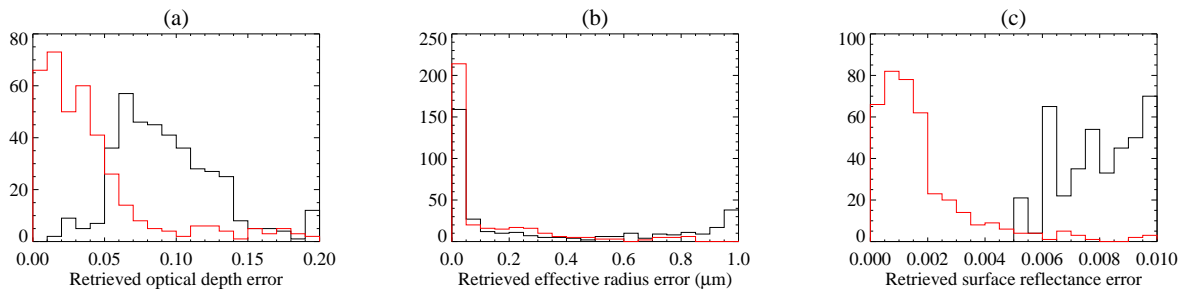


Fig. 5. The black line shows distribution of retrieved uncertainties for each of the state parameters (i.e. standard deviations taken from the state covariance matrix, \hat{S}). The red line shows the distribution of the actual errors (absolute difference between true and retrieved state). Note that 0.01 corresponds to the *a priori* error in the surface reflectance, which defines an upper limit on the retrieved uncertainty. The bin for the largest value includes all values greater than that bin, e.g. optical depths errors > 0.19 are included in the rightmost column of plot (a).

- Where the retrieval is working well, optical depth and effective radius are both retrieved with a precision of $\lesssim 0.1$.
- The measurement is adding between 1 and approximately 2.5 independent pieces of information to the system at this level of *a priori* constraint – i.e. the retrieval is under-constrained, since we are attempting to retrieve more quantities than we have pieces of information from the measurement.
- Surface reflectance is poorly retrieved, with the *a priori* accounting for 50% or more of the retrieved value. Optical depth and effective radius show good sensitivity to the true state throughout most of the range.
- The strange feature in the effective radius fields above a true effective radius of $\sim 1\mu\text{m}$ (particularly evident in Figs. 3d and 4b) indicates a region of state space where effective radius is poorly constrained by the measurements. Corresponding patterns in the surface reflectance fields indicate some degeneracy between these two variables. It is also notable that the optical depth retrieval remains quite stable throughout this region of state space.

The accuracy of the *a priori* surface reflectance also has a strong influence on the accuracy of the retrieved parameters. Although the GRAPE algorithm retrieves the magnitude of the surface reflectance, it is tightly constrained to the *a priori*. This constraint is required because of the limited amount of information available in the measurements – if the surface reflectance is not tightly constrained, the retrieval is prone to converging on highly unrealistic states: i.e. the cost function has multiple minima. The effects of a 0.01 error (i.e. equal to the *a priori* error) and a 0.03 error in the magnitude of the *a priori* and first guess surface reflectance have been investigated. A 0.01 error in surface reflectance produces results almost identical to those shown in Fig. 3, showing that the

retrieval is able to compensate for errors of this scale and fit the measurements as well and as quickly as when the correct *a priori* surface reflectance is used.

Fig. 6 shows the results of increasing the first guess and *a priori* surface reflectances to 0.05. In this case it is clear that the retrieval is unable to compensate for the larger discrepancy, particularly at either large, or the smallest effective radii, where the retrieved state is grossly different to the truth and the retrieved surface has not moved from the *a priori* value. Despite this the retrieved uncertainty estimates are similar to those in Fig. 3. Unsurprisingly, the retrieval deals with an error in the assumed surface reflectance better when the aerosol optical depth is high, as in this case the relative contribution of the surface to the overall signal is reduced. Figure 6g shows that for low effective radii, the cost function is slightly higher, suggesting a poorer fit to the measurements; however for much of state space the cost is either similar or smaller than when the correct *a priori* is used. This, along with the low retrieval-error estimates, demonstrates the degeneracy in the retrieval system: the effect of the incorrect surface reflectance can be compensated for by changes in optical depth and effective radius.

These results indicate that the retrieval is able to compensate for errors in the *a priori* surface reflectance of the order of the *a priori* error, but not much larger than this. Unfortunately the retrieved state is still consistent with the measurements and *a priori* (as indicated by very similar retrieval costs) in both of these cases, meaning that it would not be possible to detect instances of poor *a priori* surface characterisation from the retrieval itself. However, it should be noted that cases of grossly inaccurate surface characterisation will be evident from the retrieval results, since it will not be possible to correct for the surface reflectance by altering aerosol parameters. In such cases the retrieval will fail to converge, or will return an anomalously high cost.

3.2 Forward model errors

The term *forward model error* refers to inaccuracies that result from incomplete or incorrect modelling of the relevant physical processes by the forward model. In the case of ORAC, these can be divided into four main categories:

1. Errors in modelling the scattering from the aerosol itself; in particular, from the assumption of sphericity implicit in the use of Mie scattering.
2. Errors resulting in the discrete ordinates radiative transfer approach, including the assumption that the surface acts as a Lambertian reflector and the plane-parallel approximation.
3. Assumptions made in the formulation of the forward model expression (Eq. 10) used in the retrieval.
4. Interpolation errors due to the use of discrete look-up tables for the transmission and reflectance terms in Eq. (10).

As mentioned in Sect. 2, the first of these sources of error is not addressed in this paper, although work is ongoing to investigate its effects and incorporate non-spherical scattering into the ORAC system. The modelling of atmospheric gas absorption, emission and Rayleigh scattering is another potential source of error, but the authors are confident that this is a minor contribution, particularly in the case of the ATSR instruments and SEVIRI, where the radiance error due to the assumptions made in the forward model is significantly smaller than the random error on the measurements. Similarly, the error introduced by the plane-parallel approximation will be minor, as the retrieval is not run when the viewing geometry is such that this approximation is inaccurate. It should be noted that forward model error is unavoidable, especially if the retrieval algorithm is to be reasonably computationally efficient. As long as the sum of the forward model errors are kept well below the measurement noise level, however, their effects will be minimal.

The error due to the DISORT radiative transfer will be dominated by the Lambertian surface reflectance approximation, except at high zenith angles ($>75^\circ$) where the plane-parallel assumption of DISORT breaks down. This is a well known limitation of the DISORT method and is avoided by only running the retrieval on data which meets the plane parallel criterion. In order to investigate the effect of the Lambertian surface approximation, DISORT was used to model TOA reflectances (following a similar procedure to the calculation of the look-up tables) for both a bi-directional surface reflectance and an equivalent Lambertian reflectance. The MODIS BRDF product was used to provide a variety of surface reflectances which span the typical range for land surfaces (the ocean surface reflectance has been neglected in this analysis, but it is generally far more isotropic than land surfaces, except for areas effected by strong sun-glint).

Table 3. The range of viewing angles and particle properties used in comparing DISORT and the GRAPE algorithm forward model. The range of viewing angles has been chosen to be typical for the ATSR instruments.

Aerosol optical depth	Aerosol effective radius	Solar zenith angle	Instrument zenith angle	Relative azimuth angle
0.06–1.0	0.02–7.0 μm	36.0–72.0°	0.0–22.5°	0.0–162°

The calculation was repeated for a wide range of viewing geometries, aerosol loading and surfaces ranging from desert to dense forest. Results show that the Lambertian approximation generally over-estimates the directional surface reflectance (which will dominate the TOA signal from the surface for typical background aerosol loadings) with the mean difference being 1.9% (median 1.4%) of the reflectance, or, in absolute terms 0.002 (0.001). For a surface reflectance of ~ 0.25 (typical for a bright desert surface) the TOA reflectance using the Lambertian approximation is within 0.015 of the BRDF for all viewing geometries.

The latter two of the error terms listed at the beginning of this section can be quantified by comparing TOA reflectances computed directly from DISORT to those computed using the look-up tables and the forward model equation. Point 3 can be investigated by comparing the two qualities for viewing geometries and aerosol properties which correspond to points in the look-up tables (so that no interpolation is required). Point 4 can then be examined by doing the comparison for values which lie half-way between the look-up table points (where interpolation errors can be expected to be maximum). This procedure has been followed for 4500 points across a wide range of viewing geometries and particle states, as summarised in Table 3.

For values coincident with the look-up table points the forward model and DISORT agree to within $\pm 0.2\%$ for 95% of viewing geometry and aerosol loading combinations, and always agree to within $\pm 0.6\%$. There is, on average, a small positive bias apparent in the forward model expression, with the mean difference being 0.14% (median 0.14%). However, it is clear that for a Lambertian surface, Eq. (10) reproduces the DISORT modelled TOA radiances well.

Comparisons made for values where the effects of look-up table interpolation are maximised show that the interpolation error can far out-weigh that from the approximations made in Eq. 10. In some instances, it is possible for interpolation to introduce over 5% error into the modelled radiances, although for the vast majority of viewing angle/aerosol loading combinations the discrepancy is much lower than this. For this worst case ensemble of points the mean difference between the forward model and DISORT is -0.99% (median -0.72%).

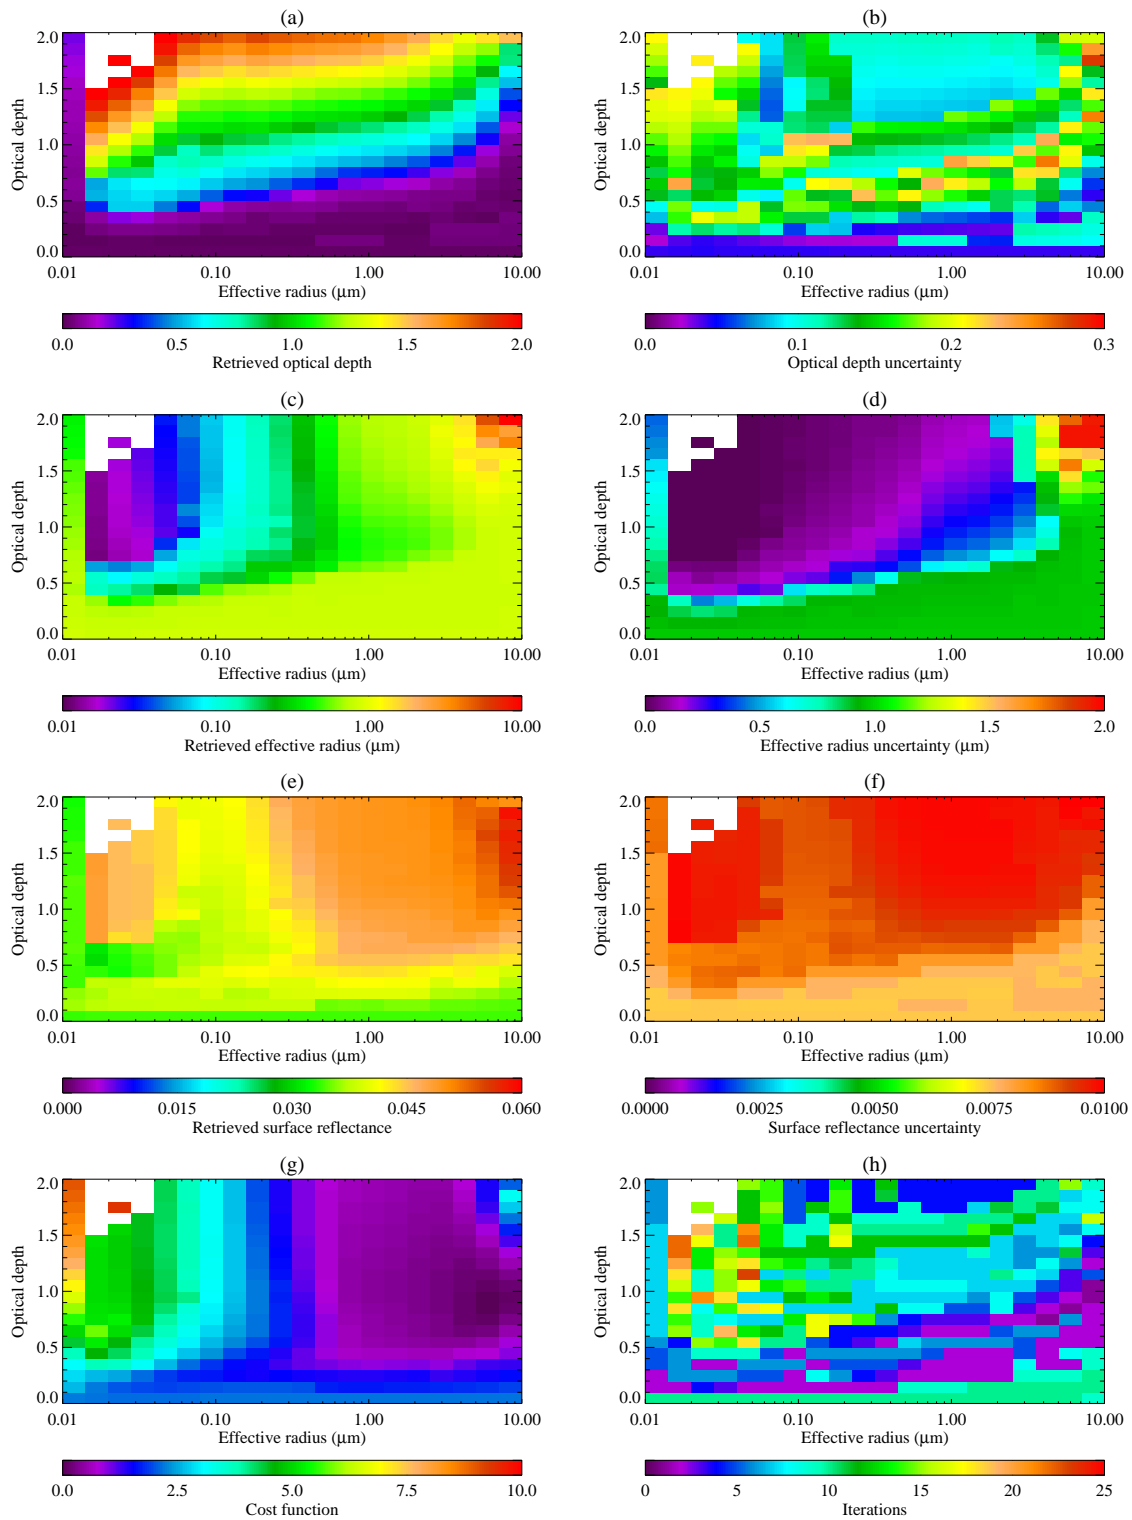


Fig. 6. Similar to Fig. 3. In this case the retrieval has been run using the correct aerosol optical properties, but the first guess and *a priori* surface reflectances were set to 0.05, while the simulated radiances were generated with a surface reflectance of 0.02.

Smith et al. (2002) quotes errors on ATSR-2 visible/near-infrared radiances as being between 1 and 5%, based on pre-launch calibration of the instrument and subsequent vicarious calibration of the visible channels using bright targets. Hence both the Lambertian surface approximation and look-up table interpolation can introduce errors which are larger than the measurement noise, with typical error values of the same order of magnitude.

3.3 Forward model parameter error

Forward model parameter error results from uncertainties in the parameters used in the computation of the forward model that are not included in the retrieval process. There are two main inputs into the modelled radiances which are most likely to significantly affect the output radiances:

1. The aerosol properties used in the calculation of the look-up tables (including the assumed size distribution, refractive indices and vertical distribution of the aerosol).
2. The spectral dependence of the surface reflectance.

In this section the effects on the retrieval of errors in a range of forward model parameters will be examined with a series of retrievals on simulated data, similar to that presented in Sect. 3.1. The effects to be examined are:

1. The effect of using an inappropriate aerosol class. The GRAPE algorithm does not include any ability to retrieve the composition of aerosol (aside from that implied by the change in composition that accompanies a change in effective radius) – it relies entirely on the accuracy of assumed optical properties, such as those provided by the OPAC database.
2. The effect of incorrect assumptions about the aerosol size distribution. Although the aerosol effective radius is retrieved, the form of the size distribution is fixed (to that defined in the OPAC database, for example).
3. The vertical distribution of aerosol used is an assumed, fixed profile. One might assume that the TOA radiance in the visible is largely insensitive to the height distribution of aerosol, but this should be quantified.
4. The effect of incorrect assumptions about the spectral dependence of the surface reflectance. Although the GRAPE algorithm is able to retrieve surface reflectance to a limited degree, it is only its magnitude which is permitted to vary. The spectral dependence is fixed at the *a priori* value.

Errors in any of these parameters will result in inaccuracies in the retrieved aerosol parameters. However, due to the non-linearity and complexity of the effects, and lack of knowledge about the accuracy of any one of them for a given retrieval, it is not practical to attempt to characterise them with

standard Gaussian error statistics. Indeed it is not even meaningful to attempt to define “typical” values for such errors, since their effects are likely to be so variable. Thus, the approach taken here is to test the retrieval in situations where the sources of error are completely known, in order to provide indications of the magnitudes of each effect.

3.3.1 Incorrect aerosol properties

Figure 7 shows retrieval results when the GRAPE algorithm is applied to simulated radiances produced using the maritime-clean aerosol class from the states shown in Fig. 2, but using the OPAC desert dust aerosol class in the retrieval. It is encouraging to see that the optical depth field still shows a reasonable agreement with the input field. Overall however, differences between the true and retrieved fields are much greater than when the correct aerosol properties are used in the retrieval, particularly in effective radius. In this particular case, the use of the incorrect aerosol parameters has resulted in differences in optical depth of between 10 and 30% for most states, while the effective radius shows typical differences of 100% or more. Despite this, the retrieved uncertainties suggest a higher degree of confidence in the retrieved values than was the case with the correct aerosol class, with the uncertainties on optical depth in particular being much smaller than in Fig. 3. Although this may seem a counter-intuitive result, it must be remembered that it simply indicates that the radiance shows a stronger dependence on aerosol optical depth for the desert aerosol class than for the maritime clean one. If the assumptions made in the retrieval are incorrect, the retrieved uncertainties cannot be assumed to be an accurate reflection of the actual error in the result.

Overall the effect of assuming desert aerosol in place of maritime can be summarised as resulting in dramatic errors in retrieved effective radius, while the optical depth shows smaller perturbations. It is clear from Fig. 7 that the retrieved state alone does not provide enough information to determine whether the aerosol class used in the retrieval is appropriate – the retrieval has converged to reasonable values, with low costs, for the majority of the states tested. One might expect that if the aerosol optical properties assumed within the retrieval are incorrect, the forward model would not be able to provide a good fit to the measurements, resulting in the cost function having a higher value at the solution. However, there is enough degeneracy in the system to allow incorrect assumptions about the optical properties of the aerosol to lead to a retrieval which is consistent with the measurements and *a priori*.

If the assumed aerosol properties are very far from the truth however, the retrieval breaks down completely. For instance, applying the highly absorbing OPAC urban aerosol class to the simulated radiances generated from the maritime-clean class, results in the retrieval failing to converge for most high optical depths and those states which are retrieved have very large uncertainty estimates (on the order of 100%

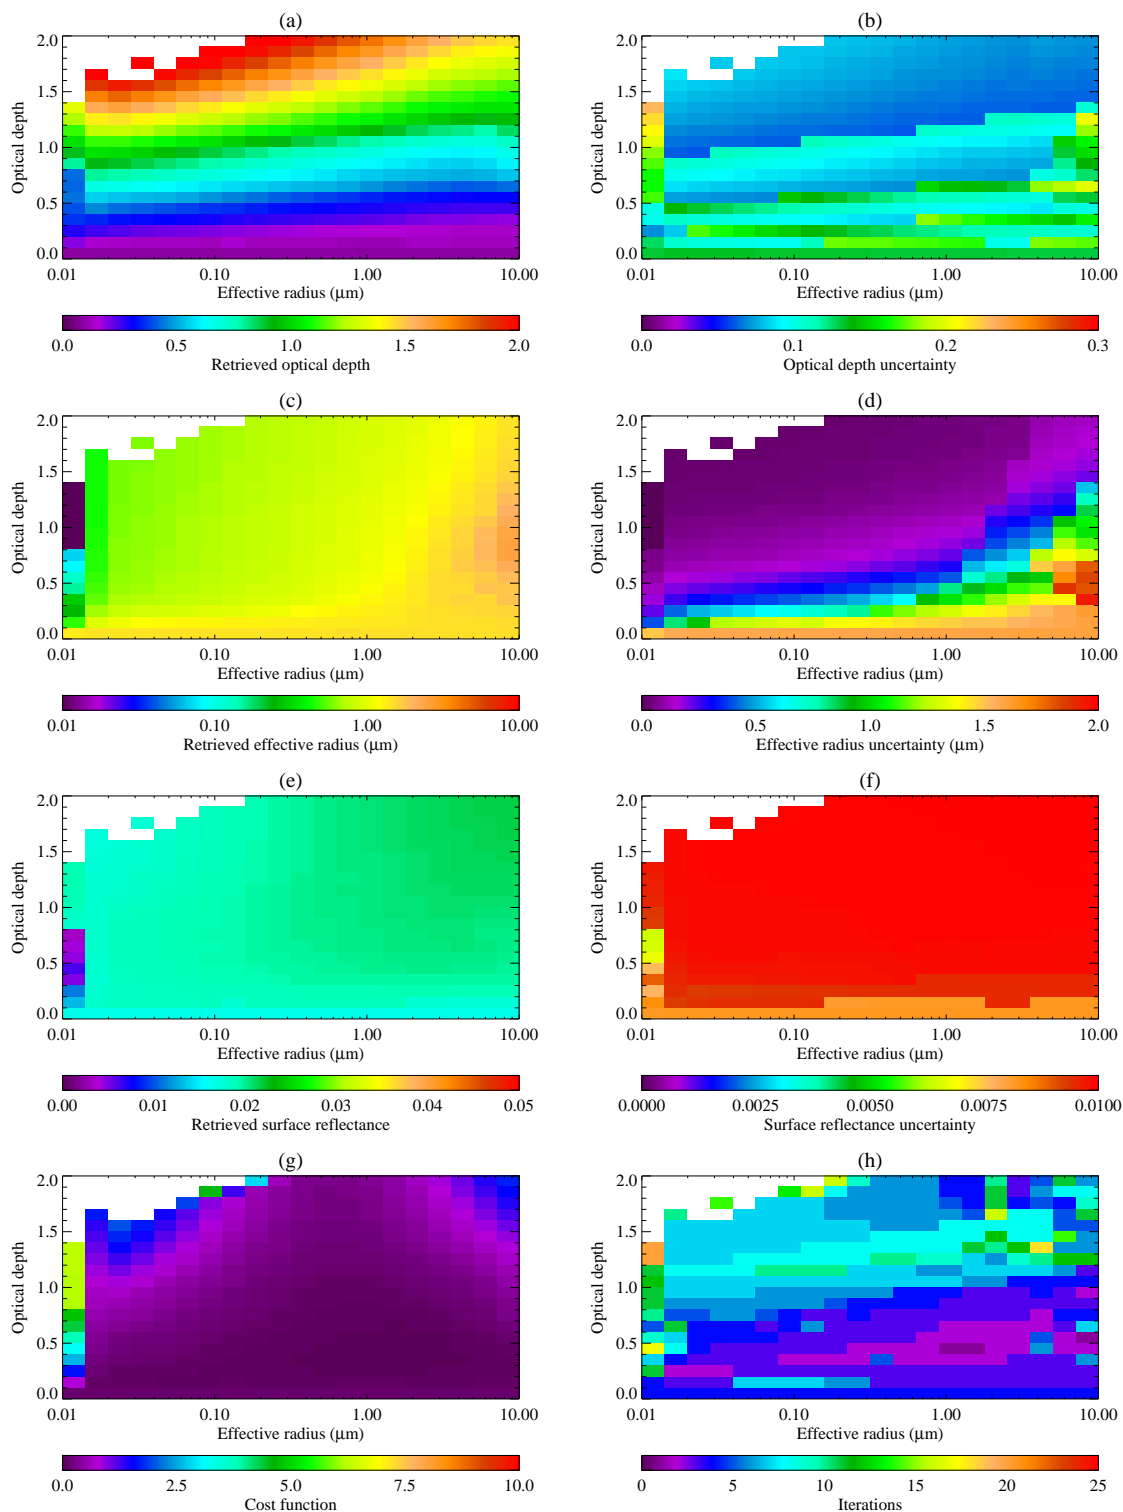


Fig. 7. Similar to Fig. 3, except in this case the retrieval has been done assuming the OPAC desert aerosol class rather than the maritime-clean one. Plots (a) and (b) show the retrieved optical depth and its uncertainty, (c) and (d) show the effective radius and its uncertainty, while (e) and (f) show the surface reflectance and uncertainty. (g) shows the cost function at the solution and (h) gives the number of iterations required for convergence.

or more in optical depth). However, even in such an extreme case, when the retrieval has converged it does so to a state with a low cost – again demonstrating the high degeneracy of the problem.

3.3.2 Aerosol size distribution

The size distribution of an aerosol population can have a great impact on its optical properties, particularly for populations dominated by larger particles. Incorrect assumptions about the size and number of modes in an aerosol distribution used to model the TOA radiance can thus be expected to have a significant impact on retrieved aerosol properties. To investigate the scale of such errors simulated ATSR-2 radiances were produced in the same way as in Sect. 3.1, but using an aerosol class consisting of a log-normal distribution of sea salt particles (as defined by the accumulation mode sea-salt component defined in OPAC). The retrieval was then run on these data using a bi-modal aerosol class consisting of two log-normal sea-salt distributions (with the second mode being defined from the OPAC coarse mode sea-salt component). The results of this retrieval are given in Fig. 8. It is clear that the effect on the retrieved state can be drastic, and the retrieval has failed to converge at all for states with a high optical depth and low effective radius. As might be expected it is the retrieval of effective radius itself which shows the greatest deviation from the truth: the retrieved field has become nearly flat. Again, for much of state space, the retrieved uncertainty estimates for the effective radius give little indication of the inaccuracy of the retrieval. Despite this however, the retrieval of optical depth shows a remarkable robustness: indeed for values less than approximately one, it is as accurate as the retrieval using correct assumptions.

The improvement in the retrieval performance at larger effective radii can be attributed to the coarse mode dominating the bi-modal distribution used in the retrieval, which effectively makes this distribution more mono-modal. However, the effective radius retrieval is still poor, and as both the sea-salt modes have the same prescribed width, this indicates that the retrieved effective radius is very sensitive to the form of the size distribution.

3.3.3 Aerosol height distribution

Figure 9 shows the difference between two retrievals using the same set of simulated ATSR radiances, created with the OPAC desert-dust aerosol class, but with differing assumptions about the height of the aerosol in the atmosphere. In calculating the radiances, the aerosol was assumed to lie between 0–2 km (corresponding to the lowest two levels of the DISORT forward model), with 60% of the aerosol optical depth lying between 0–1 km. Two retrievals have been performed on these radiances, the first assuming the correct height distribution, the second assuming the aerosol lay in a layer between 4–5 km altitude.

The two retrievals show strong similarities, with optical depth showing average differences of 4% and effective radius 7%. These differences indicate that the retrieval is slightly sensitive to the height distribution of the aerosol. The effect is most notable at large effective radii, but unlike all the other effects examined in this study, the perturbation to the optical depth is comparable that seen in the effective radius (with optical depth differing by as much as 20% at large effective radii). This is due to the interaction of the aerosol scattering with the background Rayleigh scattering by air molecules. At larger particle sizes the aerosol scattering phase function becomes more peaked in the forward direction, thus a greater proportion of the downwelling solar radiation is scattered towards the surface where, due to the low surface reflectance used in the simulations, it is likely to be absorbed. If the aerosol layer is elevated, there is a layer of Rayleigh scattering between the aerosol layer and the surface. Due to the isotropic nature of Rayleigh scatter, this will effectively increase the reflectance below the aerosol and hence produce higher TOA radiances. Thus if the forward model incorrectly makes the assumption of an elevated aerosol layer, it will over-predict the TOA radiance for a given optical depth, resulting in the retrieval under-predicting optical depth to compensate. This will only be noticeable at large aerosol effective radii, because at lower effective radii the aerosol phase function becomes more isotropic. This effect will also depend on viewing geometry, due to the angular dependence of the phase function of large particles. For example, for near-back scattering geometries, tests have shown the effect is not apparent, because the signal is dominated by the strong back scattering peak of the aerosol phase function.

Although the retrieval is not very sensitive to the aerosol height itself, Marsh et al. (2004) shows that the TOA radiance is dependent on the height distribution of aerosol properties. Variation of aerosol properties with height leads to effects similar to those described above, but to a much greater degree. An example analogous to the Rayleigh scattering effect discussed above would be that of a layer of absorbing aerosol overlying scattering aerosol compared to situation where the two aerosol are mixed in a single layer. In the former case the overlying absorbing layer will absorb a greater proportion of the incoming solar radiation and upwelling radiation scattered from the scattering aerosol below than would be the case if the aerosol were in a mixed layer, thus the TOA radiation will be reduced. The size of this effect will depend strongly on the amount and distribution of each aerosol and is thus difficult to quantify. However, Marsh et al. (2004) shows changes of up to 50% in TOA radiance over dark surfaces at high optical depths. It is clear that such effects could have drastic effects on the retrieval of aerosol properties. In extreme cases the retrieval would fail, as it would not be possible for it to reproduce the measurements, but for smaller perturbations, it is likely that the result would be a highly inaccurate retrieval of optical depth and effective radius.

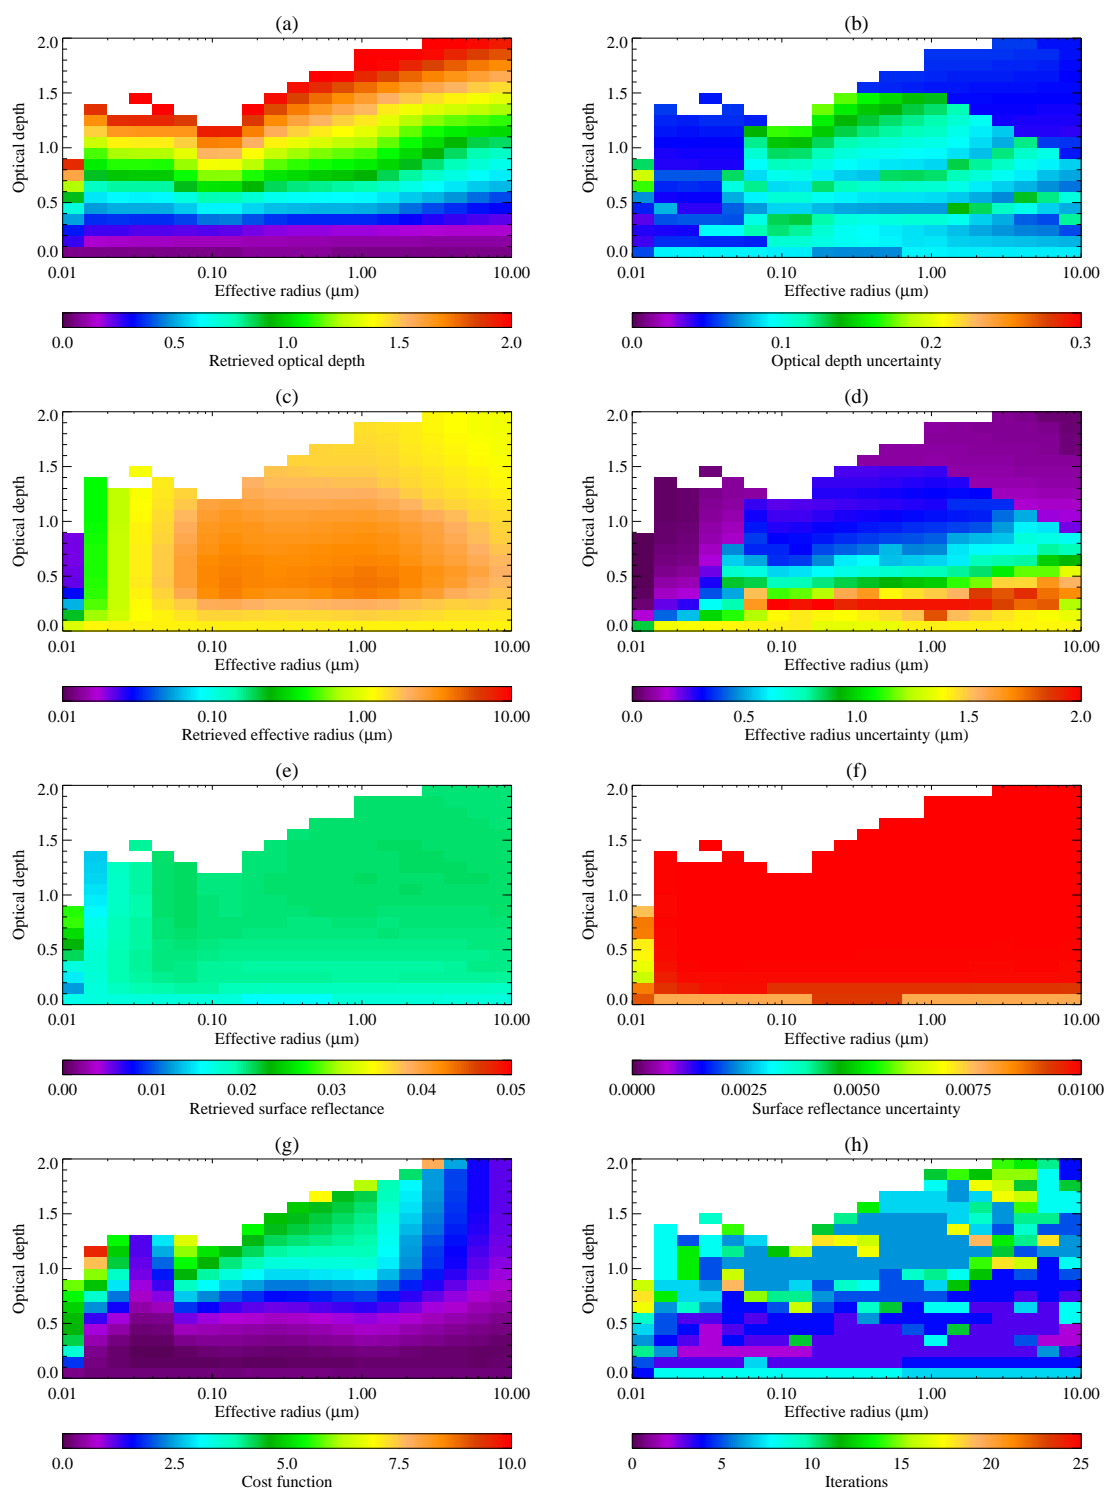


Fig. 8. Similar to Fig. 7, but here the simulated radiances have been produced using a pure sea-salt aerosol with a monomodal log-normal size distribution and retrieved assuming a bi-modal sea salt aerosol.

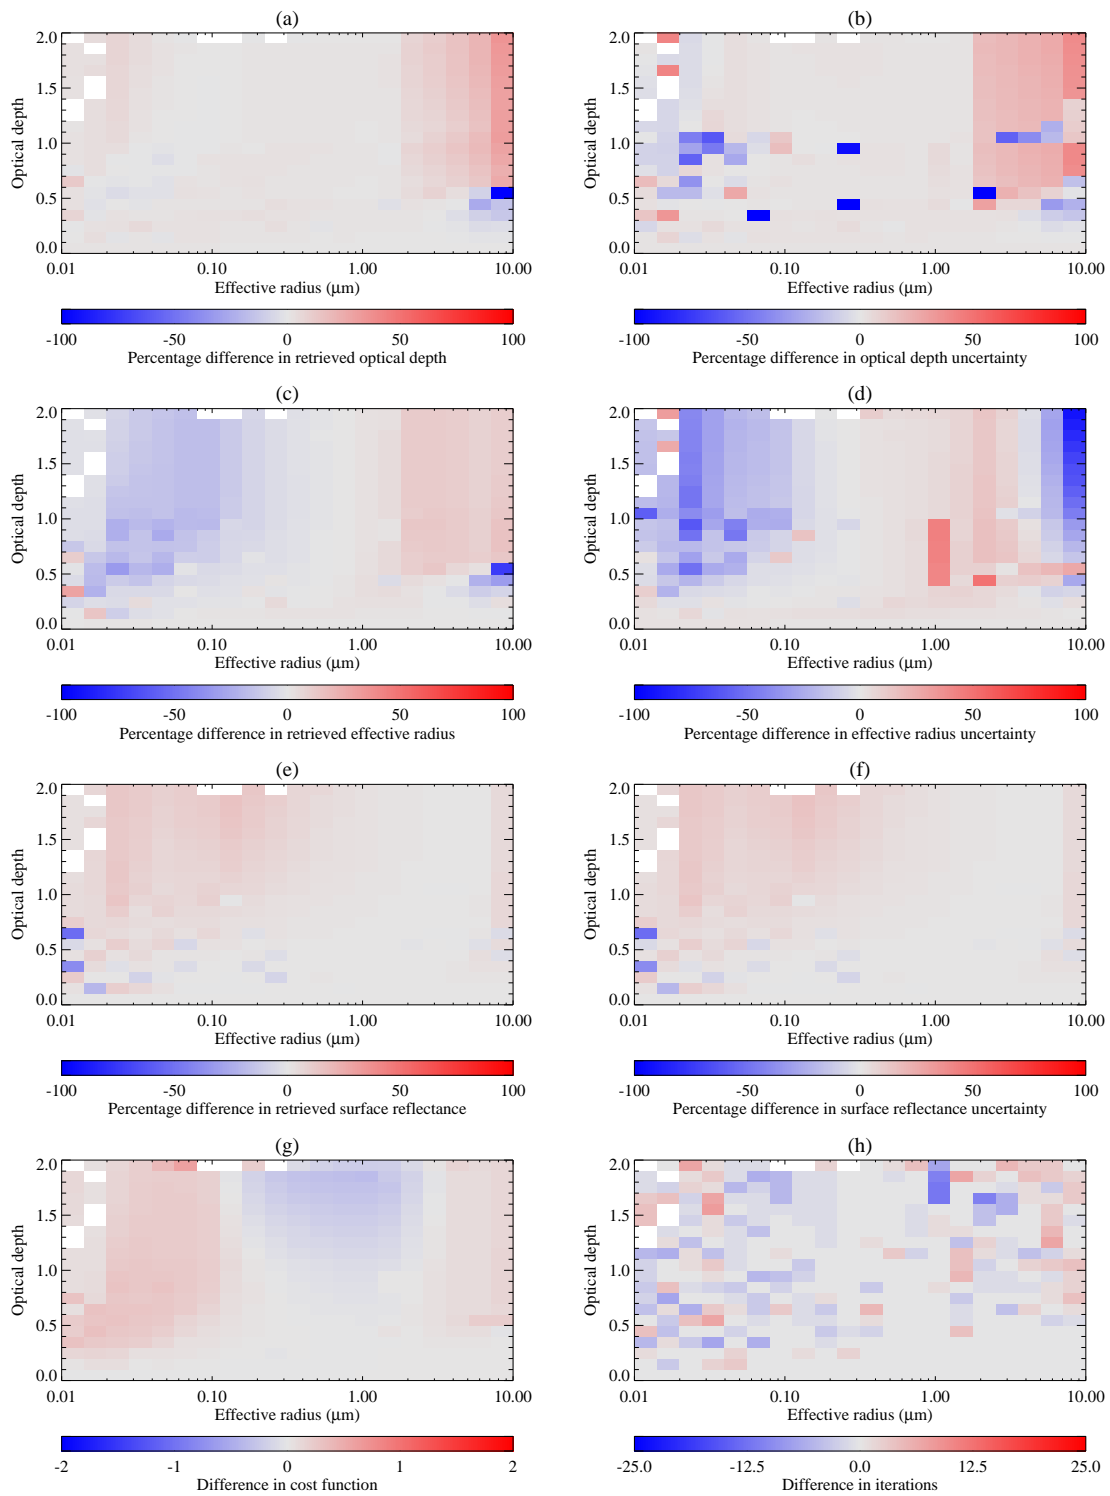


Fig. 9. Differences between a retrieval of desert dust where the aerosol is correctly assumed to lie near the surface and one in which it is erroneously assumed to be in an elevated layer.

3.3.4 Spectral surface reflectance

The final parameter to be examined is the spectral shape of the surface reflectance. This was tested by generating simulated AATSR radiances using the OPAC maritime clean aerosol class with surface reflectances of 0.02, 0.016, 0.012 and 0.008 at 550, 670, 870 and 1600 nm, respectively (i.e. the reflectance has been reduced by 20, 40 and 60% from the 550 nm value in the subsequent channels). The retrieval was then applied to the resulting data assuming a spectrally flat surface reflectance of 0.02 in all channels. Figure 10 shows that the retrieved optical depth and effective radius fields are almost as close to the input data as is the case when the retrieval is run with correct assumptions. The retrieved 550 nm surface reflectance shows that the retrieval has produced a lower reflectance at this wavelength in order to compensate for the over-estimate in reflectance at lower wavelengths produced by the assumed flat spectral shape. It should be noted that the sensitivity to the *a priori* 0.55 μm surface reflectance discussed in Sect. 3.1 will limit the ability of the retrieval to deal with very large discrepancies in the surface reflectance spectrum, as the algorithm is unlikely to retrieve a surface reflectance outside the 1σ *a priori* confidence interval.

4 Conclusions

An optimal estimation retrieval of aerosol properties from visible/near-IR satellite imagery, part of the Oxford-RAL Aerosol and Cloud retrieval scheme (ORAC), has been presented. The algorithm was developed to work within the ORAC cloud retrieval scheme for application to the ATSR series of instruments and SEVIRI, but should be applicable to the majority of near nadir viewing visible/near-infrared radiometers. The algorithm is based around a forward model using look-up tables calculated using the DISORT radiative transfer method with predefined aerosol properties. The retrieved parameters are aerosol optical depth at 550 nm, the aerosol effective radius and the surface albedo at 550 nm (the spectral shape of the surface is fixed).

Retrievals using simulated radiances have been used to assess the sensitivity of the retrieval, as well as its susceptibility to error in assumptions made in the forward model (forward model error) and non-retrieved forward model parameters (forward model parameter errors). The retrieval is under-constrained, requiring the surface reflectance to be tightly constrained by *a priori* information. If the assumed non-retrieved aerosol properties and *a priori* surface reflectances are correct, optical depth is retrieved to a precision of approximately 0.01 (with a range between approximately 0.05 and 0.14). The uncertainty on the retrieved effective radius depends strongly on its value, with typical errors being approximately 50%. Despite the tight *a priori* constraint, the precision of the surface reflectance is improved somewhat by the retrieval, with post retrieval error estimates being be-

tween 60 and 100% of the *a priori* values. It should be emphasised that these values do not represent the accuracy of aerosol properties derived from real measurements using this retrieval scheme, but represent how well constrained the retrieval of each of these parameters is, subject to the approximations and assumptions made in the retrieval. To gauge the accuracy of the GRAPE aerosol products comparisons of real data to ground-truth measurements is required. This analysis is presented by Thomas et al. (2009b).

The degrees of freedom for signal and the averaging kernel of the retrieval solution have been examined. At high aerosol optical depth and/or low aerosol effective radii, the scheme shows between 2–3 degrees of freedom, showing that the problem is under constrained. For the lowest optical depths, the degrees of freedom drops to approximately 1 and this poorly constrained region extends to higher optical depths for larger effective radii. Optical depth shows the greatest sensitivity to its true value for all states except those with the lowest optical depth, where it is the surface reflectance. Retrieved effective radius shows good sensitivity to the true value in the region of state space with 2–3 degrees of freedom, but this sensitivity drops to zero for the lowest optical depths. Surface reflectance shows relatively poor sensitivity, especially at higher optical depths. These patterns of measurement information content are reflected in the retrieved values of the state, and its dependence on the *a priori* constraints. Thus, errors in the *a priori* surface reflectance which are greater than the 1σ *a priori* error estimates cannot be corrected for by the retrieval and result in poor estimates of optical depth and effective radius.

Errors due to the approximations made in the forward model, such as the use of a Lambertian land surface reflectance in DISORT and the interpolation of the look-up tables used, have been found to result in errors in forward modelled radiance that are typically of the same order of magnitude as the measurement noise of the ATSR instruments ($\sim 1\text{--}5\%$), and can produce maximum errors of up to $\sim 5\%$. This is not ideal, as it could result in biases in the retrieved parameters. Although not implemented in time for processing of the version 3 GRAPE products, a new ORAC forward model has been created which uses a BRDF description of the surface reflectance (Thomas et al., 2009a), while new look-up tables with a finer spacing have been generated to address the interpolation errors. These improvements have reduced the associated errors to well below the measurement noise threshold of the ATSR instruments.

A more fundamental limitation of ORAC, which it shares with all satellite aerosol retrieval algorithms, is its dependence on assumptions regarding the aerosol state. A variety of perturbations to the assumed aerosol properties and the spectral shape of the surface reflectance have been tested with the GRAPE algorithm using simulated ATSR-2 data. Several conclusions can be drawn from the results of these tests. The aerosol optical depth is, in general, a far more robustly retrieved parameter than aerosol effective radius: in all

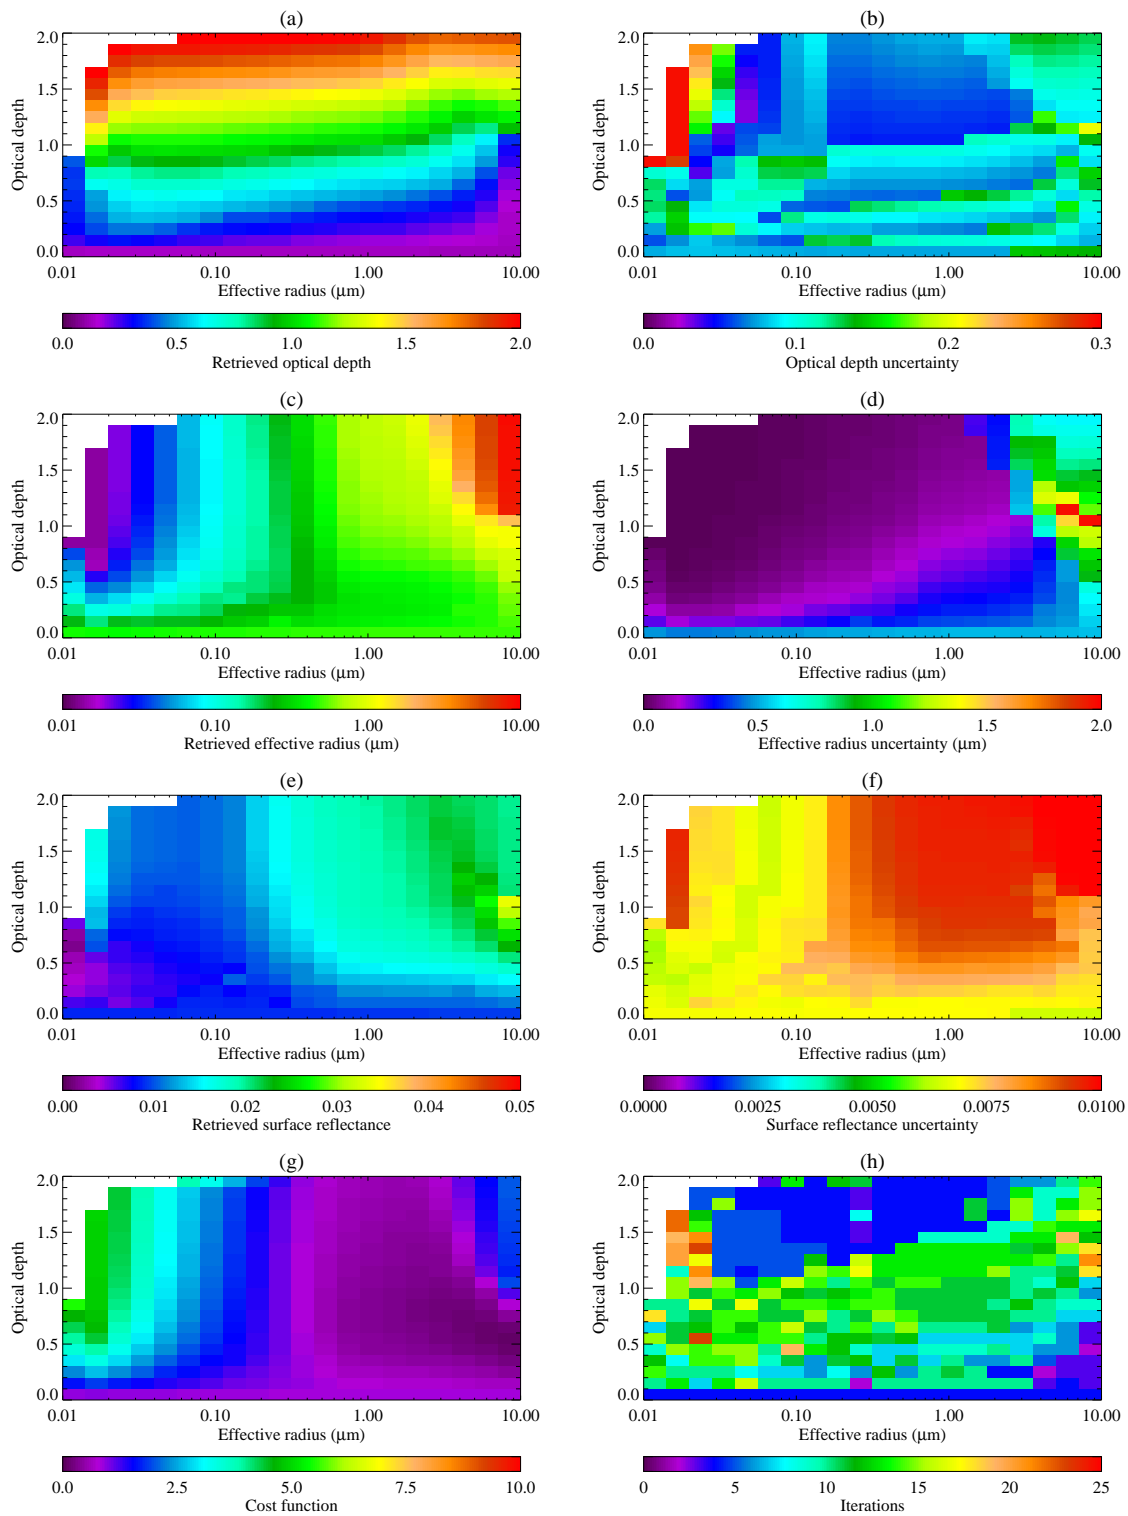


Fig. 10. Similar to Fig. 7, except in this case the simulated radiances have been produced using a surface reflectance that decreases with increasing wavelength, while the retrieval as assumed a spectrally flat surface reflectance.

tests except for perturbing aerosol layer height, the change in effective radius was far greater than in optical depth.

The retrieval was found to be somewhat sensitive to all changes tested, with the greatest effects being caused by changes to the aerosol size distribution and assumed aerosol class (a combination of differing size distribution and composition). The lowest sensitivities were found for changes to the spectral shape of the surface reflectances and the height of the aerosol layer. In the case of surface reflectance spectrum, it was found that the aerosol properties (optical depth and effective radius) were well retrieved, but that the retrieved surface reflectance was offset. Also, although the retrieval is not particularly sensitive to changes in the aerosol height directly, it will be sensitive to changes in aerosol composition with height, due to the large changes in TOA radiance that can result.

Overall, the GRAPE algorithm can retrieve aerosol optical depth and effective radius to a good degree of accuracy if the aerosol composition and size distribution are close to those assumed in the forward model, and if the equivalent Lambertian surface reflectance is known to within approximately 0.01. In practice, these provisos are the largest limitation to the accuracy of the retrieved products, as is the case with all aerosol products derived from satellite radiometers. The optical depth retrieved by the GRAPE algorithm shows a reasonable degree of robustness against moderate errors in *a priori* knowledge and forward model assumptions, but the effective radius is very sensitive to them. Thus, in the case of GRAPE aerosol products, the effective radius can only be considered to be accurate in conditions where the aerosol type and surface properties are likely to already be well defined, such as over the open ocean. The retrieved optical depth is likely to be a reasonable estimate for areas where aerosol type and surface are similar to those assumed in the retrieval, but should be treated with care when dealing with unusual or extreme aerosol events.

Due to the degeneracy of the retrieval scheme in the GRAPE version 3 configuration, the algorithm is often able to produce results that are consistent with the measurements even when the assumptions made in the retrieval are far from accurate. It is always necessary to keep the sensitivity of aerosol products derived from remotely sensed measurements to the assumptions made in the derivation in mind when dealing with such data-sets. It is possible to improve the sensitivity to the true aerosol state by adding more measurements to provide either better spectral characterisation of the atmosphere (as is the case with the MODIS instruments (Remer et al., 2005), for example) or angular sampling of the atmosphere (as is the case with the MISR instrument (Martonchik et al., 1998) and the ATSR series), however, the retrieval will still be constrained by the assumptions about aerosol composition, size distribution, etc. made in the forward model. Only in circumstances where aerosol properties are well defined beforehand can different remote sensing algorithms be expected to produce results which are consistent

with each other. At present, most satellite derived aerosol products show correlations with each other of ~ 0.5 or less (Kokhanovsky et al., 2007) – without a contrived agreement in the assumptions used in different algorithms, it is unlikely that this can be improved upon, and probably reflects the true accuracy of remotely sensed aerosol properties.

Acknowledgements. The work presented here was supported by the European Commission Framework 5 project PARTS the NERC UTGARD project and the ESA Globaerosol project. The NERC GRAPE project funded the creation of the GRAPE data-set.

Edited by: O. Torres

References

- Ackerman, S. A., Strabala, K. I., Menzel, W. P., Frey, R. A., Moeller, C. C., and Gumley, L. E.: Discriminating clear-sky from clouds with MODIS, *J. Geophys. Res.*, 103, 32141–32157, 1998.
- Aminou, D., Jacquet, B., and Pasternak, F.: Characteristics of the Meteosat Second Generation radiometer/imager SEVIRI, *Proc. SPIE, Europto Ser.*, 3221, 19–31, 1997.
- Antoine, D. and Morel, A.: A multiple scattering algorithm for atmospheric correction of remotely sensed ocean colour (MERIS instrument): principle and implementation for atmospheres carrying various aerosols including absorbing ones, *Int. J. Remote Sensing*, 20, 1875–1916, 1999.
- Berk, A., Bernstein, L. S., Anderson, G. P., Acharya, P. K., Robertson, D. C., Chetwynd, J. H. and Adler-Golden, S. M.: MODTRAN cloud and multiple scattering upgrades with application to AVIRIS, *Remote Sens. Environ.*, 65, 367–375, 1998.
- Birks A.: Improvements to the AATSR IPF relating to land surface temperature, European Space Agency Technical Note, 12 pp., available at: http://earth.esa.int/pes/envisat/aatsr/articles/LSTNote_IPFv6.0.pdf, 2007.
- Bulgin, C. E., Palmer, P. I., Thomas, G. E., Arnold, C. P. G., Campmany, E., Grainger, R. G., Poulsen, C. A., Siddans, R. and Lawrence, B. N.: Regional and seasonal variations of the Twomey indirect effect as observed by the ATSR-2 satellite instrument, *Geophys. Res. Lett.*, 35, L02811, doi:10.1029/2007GL031394, 2008.
- Diner, D. J., Martonchik, J. V., Kahn, R. A., Pinty, B., Gobron, N., Nelson, B. N., and Holben, B. N.: Using angular and spectral shape similarity constraints to improve MISR aerosol and surface retrievals over land, *Remote Sens. Environ.*, 94, 155–171, 2005.
- Dubovik, O., Holben, B. N., Eck, T. F., Smirnov, A., Kaufman, Y. J., King, M. D., Tanre, D., and Slutsker, I.: Variability of absorption and optical properties of key aerosol types observed in worldwide locations, *J. Atmos. Sci.*, 59, 590–608, 2002.
- Flowerdew, R. J. and Haigh, J. D.: An approximation to improve accuracy in the derivation of surface reflectances from multi-look satellite radiometers, *Geophys. Res. Lett.*, 22, 1693–1696, 1995.
- Gómez-Chova, L., Camps-Valls, G., Calpe-Maravilla, J., Guanter, L. and Moreno, J.: Cloud-screening algorithm for ENVISAT/MERIS multispectral images, *IEEE Trans. Geosci. Remote Sens.*, 45, 4105–4118, 2007.
- Grey, W. M. F., North, P. R. J., Los, S. O., and Mitchell, R. M.: Aerosol optical depth and land surface reflectance from

- multiangle AATSR measurements: global validation and inter-sensor comparisons, *IEEE Trans. Geosci. Remote Sens.*, 44, 2184–2197, 2006.
- Hess, M., Koepke, P., and Schult, I.: Optical Properties of Aerosols and Clouds: The software package OPAC, *Bull. A. Met. Soc.*, 79, 831–844, 1998.
- von Hoyningen-Huene, W., Freitag, M., and Burrows, J. P.: Retrieval of aerosol optical thickness over land surfaces from top-of-atmosphere radiance, *J. Geophys. Res.*, 108, D94269, doi:10.1029/2001JD002018, 2003.
- IPCC Climate Change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, edited by: Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K. B., Tignor, M., and Miller, H. L., Cambridge University Press, Cambridge and New York, 996 pp., 2007.
- Jin, Y., Schaaf, C. B., Woodcock, C. E., Gao, F., Li, X., Strahler, A. H., Lucht, W., and Liang, S.: Consistency of MODIS surface BRDF/Albedo retrievals: 1. Algorithm performance, *J. Geophys. Res.*, 108, D54158, doi:10.1029/2002JD002803, 2003.
- Kokhanovsky, A. A., Breon, F.-M., Cacciari, A., Carboni, E., Diner, D., Di Nicolantonio, W., Grainger, R. G., Grey, W. M. F., Höller, R., Lee, K.-H., North, P. R. J., Sayer, A., Thomas, G. and von Hoyningen-Huene, W.: Aerosol remote sensing over land: satellite retrievals using different algorithms and instruments, *Atmos. Res.*, 85, 372–294, 2007.
- Levenberg, K.: A method for the solution of certain nonlinear problems in least squares, *Quart. Appl. Math.*, 2, 164–168, 1944.
- Levy, R. C., Remer, L. A., Mattoo, S., Vermote, E. F. and Kaufman, Y. J.: Second-generation operational algorithm: Retrieval of aerosol properties over land from inversion of Moderate Resolution Imaging Spectroradiometer spectral reflectance, *J. Geophys. Res.*, 112, D13211, doi:10.1029/2006JD007811, 2007.
- Llewellyn-Jones, D., Edwards, M. C., Mutlow, C. T., Birks, A. R., Barton, L. J. and Tait, H.: AATSR: Global-change and surface-temperature measurements from Envisat, *ESA bulletin*, 105, 10–21, 2001.
- Lohmann, U. and Feichter, J.: Global indirect aerosol effects: a review, *Atmos. Chem. Phys.*, 5, 715–737, 2005, <http://www.atmos-chem-phys.net/5/715/2005/>.
- Marquardt, D. W.: An algorithm for least-squares estimation of nonlinear parameters, *SIAM J. Appl. Math.*, 11, 431–441, 1963.
- Marsh, S. H., Dean, S. M., Grainger, R. G., Quijano, A. L., Thomas, G. E. and Lawrence, B. N.: An optimal estimation aerosol retrieval scheme for ATSR-2, *Atmospheric, Oceanic and Planetary Physics Technical Memorandum*, University of Oxford, 53 pp., available at: <http://www.atm.ox.ac.uk/main/Science/technical/2004.2.pdf>, 2004.
- Martonchik, J. V., Diner, D. J., Kahn, R. A., Ackerman, T. P., Verstaete, M. M., Pinty, B., Gordon, H. R.: Techniques for the retrieval of aerosol properties over land and ocean using multiangle imaging, *IEEE Trans. Geosci. Remote Sensing*, 36, 1212–1227, 1998.
- Martonchik, J. V., Diner, D. J., Crean, K. A., and Bull, M. A.: Regional aerosol retrieval results from MISR, *IEEE Trans. Geosci. Remote Sensing*, 40, 1520–1531, 2002.
- Mishchenko, M. I., Geogdzhayev, I. V., Cairns, B., Rossow, W. B. and Lacis, A. A.: Aerosol retrievals over the ocean by use of channels 1 and 2 AVHRR data: sensitivity analysis and preliminary results, *Appl. Optics*, 38, 7325–7341, 1999.
- Mutlow, C. T., Murray, M. J., Bailey, P., Birks, A. R., and Smith, D. L.: ATSR-1/2 user guide issue 1, *ESA User Guide*, 29 pp., 1999.
- North, P. R. J.: Estimation of aerosol opacity and land surface bidirectional reflectance from ATSR-2 dual angle imagery: operational method and validation, *J. Geophys. Res.*, 107, D12, 4149, doi:10.1029/2000JD000207, 2002.
- Quaas, J., Ming, Y., Quaas, J., Ming, Y., Menon, S., Takemura, T., Wang, M., Penner, J. E., Gettelman, A., Lohmann, U., Bellouin, N., Boucher, O., Sayer, A. M., Thomas, G. E., McComiskey, A., Feingold, G., Hoose, C., Kristjánsson, J. E., Liu, X., Balkanski, Y., Donner, L. J., Ginoux, P. A., Stier, P., Feichter, J., Sednev, I., Bauer, S. E., Koch, D., Grainger, R. G., Kirkevåg, A., Iversen, T., Seland, Ø., Easter, R., Ghan, S. J., Rasch, P. J., Morrison, H., Lamarque, J.-F., Iacono, M. J., Kinne, S., and Schulz, M.: Aerosol indirect effects general circulation model intercomparison and evaluation with satellite data, *Atmos. Chem. Phys. Discuss.*, 9, 12731–12779, 2009, <http://www.atmos-chem-phys-discuss.net/9/12731/2009/>.
- Remer, L. A., Kaufman, Y. J., Tanré, D., Mattoo, S., Chu, D. A., Martins, J. V., Li, R.-R., Ichoku, C., Levy, R. C., Kleidman, R. G., Eck, T. F., Vermote, E., and Holben, B. N.: The MODIS aerosol algorithm, products and validation, *J. Atmos. Sci.*, 62, 947–972, 2005.
- Rodgers, C. D.: Inverse methods for atmospheric sounding: Theory and practice, *World Scientific*, Singapore, 238 pp., 2000.
- Santer, R., Carrere, V., Dubuisson, P. and Roger, J. C.: Atmospheric correction over land for MERIS, *Int. J. Remote Sensing*, 20, 1819–1840, 1999.
- Sayer, A.: A sea surface reflectance model suitable for use with AATSR aerosol retrieval, *Atmospheric, Oceanic and Planetary Physics Technical Memorandum*, University of Oxford, 31 pp., available at: <http://www.atm.ox.ac.uk/main/Science/technical/2007.2.pdf>, 2007.
- Smith, D. L., Mutlow, C. T. and Nagaraja Rao, C. R.: Calibration monitoring of the visible and near-infrared channels of the Along-Track Scanning Radiometer-2 by use of stable terrestrial sites, *Appl. Optics*, 41, 515–523, 2002.
- Stamnes, K., Tsay, S.-C., Wiscombe W., and Jayweera, K.: Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media, *Appl. Optics*, 24, 2502–2509, 1988.
- Thomas, G. E., Dean, S. M., Carboni, E., Grainger, R. G., Poulsen, C. A., Siddans, R., and Kerridge, B. J.: An optimal estimation aerosol retrieval scheme for (A)ATSR, *Atmospheric, Oceanic and Planetary Physics Technical Memorandum*, University of Oxford, 25 pp., available at: <http://www.atm.ox.ac.uk/main/Science/technical/2007.1.pdf>, 2007.
- Thomas, G. E., Poulsen, C. A., Curier, R. L., de Leeuw, G., Marsh, S. H., Carboni, E., Grainger, R. G. and Siddans, R.: Comparison of AATSR and SEVIRI aerosol retrievals over the Northern Adriatic, *Q. J. R. Meteorol. Soc.*, 133, 85–95, 2007.
- Thomas, G. E., Carboni, E., Sayer, A. M., Poulsen, C. A., Siddans, R., and Grainger, R. G.: Oxford-RAL Aerosol and Cloud (ORAC): Aerosol retrievals from satellite radiometers, in: *Aerosol Remote Sensing Over Land*, edited by: Kokhanovsky, A. A., and de Leeuw, G., Springer, Berlin, Germany, 193–225, 2009.

- Thomas, G. E., Poulsen, C. A., Siddans, R., Sayer, A. M., Carboni, E., Marsh, S. H., Dean, S. M., Grainger, R. G., and Lawrence, B. N.: Validation of the GRAPE single view aerosol retrieval for ATSR-2 and insights into the long term global AOD trend, *Atmos. Chem. Phys. Discuss.*, 9, 21581–21618, 2009, <http://www.atmos-chem-phys-discuss.net/9/21581/2009/>.
- Veefkind, J. P. and de Leeuw, G.: A new algorithm to determine the spectral aerosol optical depth from satellite radiometer measurements, *J. Aerosol Sci.*, 29, 1237–1248, 1998.
- Watts, P. D., Mutlow, C. T., Baran, A. J., and Zavody, A. M.: Study on cloud properties derived from Meteosat Second Generation observations, EUMETSAT ITT no. 97/181, 1998.