Aerosol Remote Sensing Using AATSR

Andrew Mark Sayer

Submitted for the degree of Doctor of Philosophy in Physics
Michaelmas term 2008

Atmospheric, Oceanic and Planetary Physics
Department of Physics
University of Oxford

Acknowledgements

The sentiment of a set of acknowledgments is typically that espoused by one of John Donne's Meditations; namely, that 'No man is an island unto himself'. This is certainly true in scientific research, with each piece of work building upon that which has come before, and our ideas reaching maturity through discussion and debate. In this respect I have been fortunate to work under the guidance of Don Grainger, Chris Mutlow and Gareth Thomas and I am grateful to them for their advice and supervision over the past three years. I would also like to thank Anu Dudhia for his supervision at the start of my time in Oxford, during Don's sabbatical.

As well as my formal supervision, I have benefitted from being part of a close-knit research group in the form of the ORAC team, both at Oxford and RAL (Elisa Carboni, Elies Campmany, Caroline Poulsen and Richard Siddans, alongside my supervisors). It has been a pleasure to work with them and I hope that I am of more use to them now than I must have been when I began. Thanks also to those others at RAL working in Earth Observation with whom I have discussed my work (Brian Kerridge, Barry Latter, Tim Nightingale and Dave Smith). The excellent team of support staff within AOPP (namely Sarah Harrington, Monika Porada and Mark Munro) has also been instrumental in keeping everything running smoothly, so many thanks to them!

On a more personal note my parents, Christopher and Anne Sayer, have instilled in me the value of an education and I have always received support from all generations of my family. Returning again to Donne, without this support I would not be writing these words today. I could not ask for a better family.

I would like to thank Elizabeth Davidson, my wonderful girlfriend, for her constant love, for always being there for me, and for her patience. Not only has she made the past two years the happiest of my life, but her own tireless work ethic has been an inspiration to me. Here's to many more.

Through both AOPP and Trinity College I have had the good fortune to find myself within a great circle of friends. Thanks to Lewis Allan, Andrew Ng, Seth Wilson, Kate Rumbold, Matthew Johnston, Louise Hardwick, Seshadri Nadathur and others at Trinity for all the good times and not forgetting me when I'd bury myself in books for a week. Particular thanks to Kate, for proof-reading the thesis and catching all the typos my own tired eyes had missed. Clearly, I was fortunate in my choice of Trinity in this respect. I am grateful to Trinity College for awarding me a Graduate Scholarship during the 2007-2008 academic year, and for their provisions for postgraduate students. Annabel Ownsworth, Trinity's Academic Administrator, has always been pleasant, helpful and efficient throughout the many queries and the barrage of paperwork which I have subjected her to since 2005.

Thanks also to the community within AOPP—students, postdocs and staff alike—for welcoming me as one of their own. I look forward to continuing our interactions on both a personal and professional level in the years to come. Further back, thanks to those friends from my undergraduate days and before for pegging me as the bookish type and showing me that it's ok to enjoy your work.

Finally, I would like to thank the Natural Environment Research Council (NERC) for awarding the studentship that the work was funded under, without which I would not have been able to study for a DPhil.

Andrew Sayer, October 2008, Oxford

Aerosol Remote Sensing Using AATSR

Abstract: Andrew Mark Sayer, Trinity College

Submitted for the degree of Doctor of Philosophy in Physics, Michaelmas term 2008

This thesis describes a new algorithm for the remote sensing of atmospheric aerosol and surface properties from the Advanced Along-Track Scanning Radiometer (AATSR). Retrieved aerosol optical depth for September 2004 shows a correlation with AERONET measurements of 0.76 at 550 nm and 0.53 at 870 nm, and best-fit lines have slopes of 1.10 and 1.08 respectively. Weaker correlations of 0.46 over land and 0.50 over sea at 550 nm are found with monthly mean MISR data, although the datasets show similar spatial patterns. Statistical analysis of retrieved data show that the scheme performs well; 80 % of converged retrievals fit the measurements and *a priori* data to within the expected precision. As part of the new algorithm, the description of the reflectance of the Earth's surface is improved. The bidirectional reflectance of the sea varies between approximately 10^{-5} and 1, with a consistent shape along the swath due to sun-glint. Black-sky and white-sky albedo are less variable, generally from 0.01-0.10 (black-sky) and 0.05-0.07 (white-sky) at all wavelengths. Over land, a method is presented to correct surface reflectance data from MODIS to account for the different visible channel response functions of MODIS and AATSR. After correction the mean error in the *a priori* albedo arising from such differences is 0.0011 at 550 nm, 0.002 at 660 nm, 0.001 at 870 nm, and 0.0046 at 1.6 μm. Various applications of the algorithm are explored, including the creation of false-colour imagery, comparison of two months of data separated by 5 years and generation of a seasonal climatology of aerosol and surface properties in the Amazon. No statistically significant weekly cycle is found in aerosol optical depth over large cities in this region on local scales. Two estimates of deforestation in the region are made (125,000 km^2 and 156,000 km^2 from August 2002-August 2007) with good agreement to independent data (158,000 km^2).

Contents

List of Figures

List of Tables

Chapter 1

Atmospheric aerosol and satellite instrumentation

1.1 Atmospheric aerosol and its importance

1.1.1 What is atmospheric aerosol?

Formally defined, aerosols are particles composed of solid or liquid matter that are suspended in the atmosphere and have sizes in the region of 10^{-9} m to 10^{-4} m (Twomey, 1977). They encompass a broad array of different substances with varying shapes, sizes, sources and sinks, and physical and chemical properties. Their large temporal and spatial variability has made global modelling and prediction studies difficult, although over the past few decades the importance of aerosols in climate and weather systems has been realised, leading to many attempts—such as the one described in this thesis—to quantify and classify them.

The study of aerosols began with John Aitken's experiments with dust in the 1880s. He built apparatus to count particles in suspension by expanding and then cooling samples of air. Water condensed on the particles, allowing the resulting drops to be counted with the aid of a microscope. Aitken used this apparatus to measure a wide variety of samples, including sea salt produced from air blowing over the ocean. C. T. R. Wilson performed important work from 1897 onwards investigating the growth of water and dust particles, which resulted in the invention of the cloud chamber. From the 1920s to the 1950s work by Kohler, Howell and then Twomey established aerosols as cloud condensation nuclei (CCN) which led to the study of their implications for weather and climate.

Since this time a wide body of literature has grown around this field, and an overview of some relevant aspects is given in this first part of this thesis. An in-depth treatment of most aspects of atmospheric aerosol—albeit before the era of modern satellites—was given by Twomey (1977).

1.1.2 Impacts on climate: aerosol direct and indirect effects

It has long been acknowledged that by scattering light, aerosols exert a non-negligible influence on weather and climate (for example Charlson et al., 1992, Pueschel, 1996, Schwartz, 1996 or many more recent works such as Schulz et al., 2006). Much research on the subject has recently been collated and summarised by the Intergovernmental Panel on Climate Change (IPCC) in their 4$^{\text{th}}$ assessment report (Forster et al., 2007).

- They have the *direct effect* of altering the planetary albedo by scattering incoming shortwave solar radiation and absorbing outgoing longwave radiation. Over the globe as a whole, scattering predominates and this exerts a net negative forcing on the Earth, the anthropogenic component of which is suggested by Forster et al. (2007) to be between -0.9 Wm^{-2} and -0.1 Wm^{-2}, with a best estimate of -0.5 Wm^{-2}.

- They have several *indirect effects* of perturbing cloud cover and properties. These exert a mixture of negative and positive forcings, dependent on how the cloud properties (such as particle size, lifetime, height or albedo) are altered. These are ill-understood; Forster et al. (2007) states that the enhancement of cloud albedo due to anthropogenic aerosol leads to a radiative forcing of between -1.8 Wm^{-2} and -0.3 Wm^{-2} (with a best estimate of -0.7 Wm^{-2}) but understanding of other aerosol indirect effects is currently too low to state them with confidence.

The overall negative forcing of aerosols offsets some of the positive forcing of greenhouse gases, and is the least well-understood component of radiative forcing (Forster et al., 2007). While greenhouse gases persist for years, aerosol lifetimes are short and profiles show great spatial variability. We have no past record of global aerosol distribution or composition, and as a result understanding of historical aerosol forcing is poor. These unknowns mean it is of great importance to obtain an accurate picture of global aerosol distribution.

Forcing estimates have generally been derived from modelling studies (for example Schulz et al., 2006), although satellite data are now also being used for this purpose, such as in Bellouin et al.

(2008). Uncertainties in the forcing arise from diverse sources, from the amount of aerosol and its composition to the parametrisations of various interactions and feedback processes within the models used (Penner et al., 2006, Kinne et al., 2006, Stier et al., 2006, Stier et al., 2007 and others).

In most cases, the infrared impact of aerosols is minor because aerosols are mostly found low in the atmosphere where atmospheric temperature is close to surface temperature. The optical depth of aerosols also decreases at longer wavelengths. Volcanic eruptions can provide exceptions to this rule: the stratospheric ash produced absorbs infrared and visible radiation and heats the stratosphere (Grainger and Highwood, 2003). Over the long term, however, scattering of visible light by the more persistent sulphate aerosol predominates. There is a clear and sharp decrease in incoming radiation linked with the eruptions which can persist for several years; for example, the 1815 eruption of Tambora in Indonesia is thought (for example Pueschel, 1996) to be responsible for summer snowstorms and crop failures in mid-latitude regions.

Forcing effects are also dependent on the geographical location of the forcing substance. An increase in aerosol optical depth (and resultant increase in planetary albedo) has more impact over a dark surface (such as a cloud-free ocean) than over a bright one (cloud, ice or desert). Conversely, deposition of black carbon aerosol on snow leads to a decrease in surface albedo and hence has positive forcing effect (0.1 Wm^{-2} according to Forster et al., 2007).

1.1.3 Aerosol types and sources

It is usual practice to group aerosols by their origins: those produced as a result of naturally-occurring geological processes, those produced by biological processes ('biogenic'), and those produced as a result of human activity ('anthropogenic'). Distinction is also often made between tropospheric and stratospheric aerosol, as these often share different sources and have different radiative effects (Pueschel, 1996). Quantification of aerosol emission is less clear-cut than that of greenhouse gases or other pollutants. One source may produce particles of a wide variety of sizes or compositions, and so very different properties. Primary aerosols are those substances emitted directly as aerosols, while secondary aerosol forms in the atmosphere from various precursor substances. Combination of aerosol particles may also lead to new particles with different chemical and physical properties from those from which they were made, so it is important to consider not just aerosol sources but how these aerosols interact with each other and the atmosphere during their lifetime.

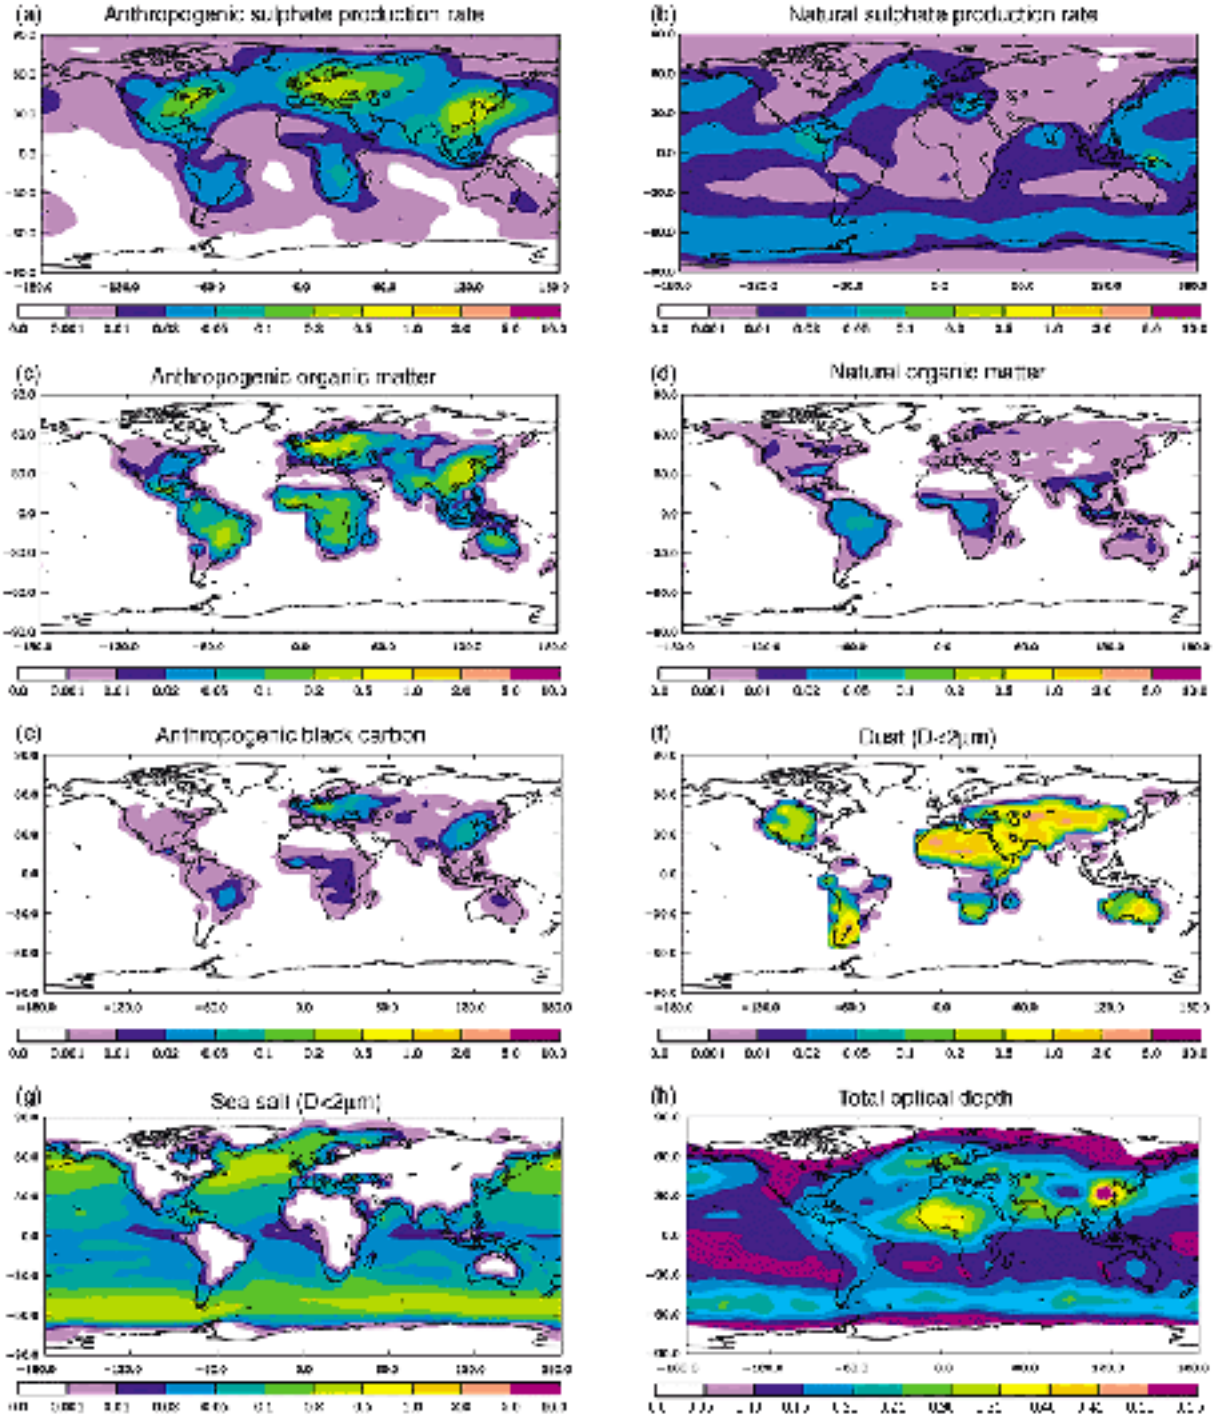

Figure 1.1: Annual average source strength in kg km^{-2} hr^{-1} for aerosol types considered in Penner et al. (2001) (a to g) with total aerosol optical depth (h). Shown are (a) the column average H_2SO_4 production rate from anthropogenic sources, (b) the column average H_2SO_4 production rate from natural sources (DMS and volcanic SO_2), (c) anthropogenic sources of organic matter, (d) natural sources of organic matter, (e) anthropogenic sources of black carbon, (f) dust sources for dust with diameters less than 2 μm, (g) sea salt sources for sea salt with diameters less than 2 μm, and (h) total optical depth. Taken from Penner et al. (2001).

Penner et al. (2001), the aerosol component of the third IPCC report, contains estimates for global annual emissions of aerosol and aerosol precursors, which is summarised in Figure 1.1. The discussion below follows the aerosol groupings from Penner et al. (2001).

Wind-driven sources: mineral dust and sea salt

Mineral dust and sea salt are primary aerosol. As shown in Figure 1.1 (f), mineral dust aerosol is particularly prevalent in tropical and sub-tropical regions where it is lofted by wind (or human activity) over deserts or vegetation-sparse land. Dust storm frequency has also been linked to climate cycles, such as the North Atlantic Oscillation (Ginoux et al., 2004). Wind blowing over the oceans results in the generation of sea salt aerosol. Global distribution is shown in Figure 1.1 (g). Dependence on wind speed means these aerosol types show strong temporal variability.

Industrial and volcanic ash

Since the Industrial Revolution human industry has been putting significant amounts of aerosol into the atmosphere. Industrial dust and ash is a primary anthropogenic aerosol. Volcanic ash, a primary geological aerosol, is ejected into the atmosphere by volcanic eruptions.

Organic carbon

Organic carbon aerosol is carbon emitted in volatile organic compounds (VOCs), distributed as shown in Figure 1.1 (c) and (d). Some is directly emitted into the atmosphere, either as products from incomplete combustion from industry or biomass burning or as compounds (e.g. terpenes) emitted by plants, while some secondary organic aerosol (SOA) is produced by modification (typically oxidation) of these VOCs. As a result, organic carbon can be both a primary and a secondary, as well as an anthropogenic or biogenic, aerosol. Many of these compounds contain polar groups (such as hydroxyl or carboxylic acids) meaning they are water soluble, and so readily take part in cloud formation, as discussed in Section 1.1.4.

Although organic carbon is produced both naturally and anthropogenically, its oxidation depends on compounds such as NO_x, O_3 and the hydroxy radical $OH^{\cdot}$. Levels of these, and hence the rate of SOA production, are strongly influenced by human activities creating these pollutants.

Black carbon

Black carbon is so called because it is composed of strongly light-absorbing carbonaceous compounds. Quantification of carbon aerosol is important because of the very different effects of black carbon (primarily absorbing) and organic carbon (primarily scattering). A primary aerosol, black

carbon tends to be insoluble and is produced as the result of incomplete combustion. Its distribution is shown in Figure 1.1 (e).

Other primary biogenic aerosol

Other sources of primary biogenic aerosol include airborne viruses and bacteria, as well as pollen and spores and other debris. Little is known about precise spatial or temporal distribution of this class of aerosol, but it would be expected to vary with land cover.

Sulphate

Sulphate aerosol is produced from precursor molecules, either anthropogenic (from industry or transportation, including SO_2, OCS and CS_2) or natural (such as H_2S or SO_2 from volcanic eruptions, see e.g. Grainger and Highwood, 2003, Bingen et al., 2004). Over the oceans DMS (dimethyl sulphide, CH_3SCH_3) emitted by plankton is a major source. In the atmosphere, these precursors are oxidised and combine with water vapour to produce aerosols of mixed H_2SO_4 and H_2O. Sulphate production is shown in Figure 1.1 (a) and (b), although there are large uncertainties in emissions (Penner et al., 2001). DMS emittted by marine plankton is an important source of sulphur although quantifying release is difficult as it is problematic to measure directly, and there have been problems correlating DMS to proxies such as the chlorophyll content of a body of water (Andreae and Crutzen, 1997), although more recent work by Kloster et al. (2005) has made progress to improve the representation of DMS release in models. Furthermore, the complexity of reaction pathways involved in the generation of sulphate aerosol means that the precise proportion of precursors which end up as H_2SO_4 is uncertain.

Nitrate

Atmospheric nitrate aerosol is linked to NH_3 and H_2SO_4 abundance as these two participate together in neutralisation reactions. The global mean ratio is thought to be close to 1, but in regions where there is surplus ammonia its oxidisation to nitrate aerosol is possible. Anthropogenic NO_x is also important in polluted areas, although Penner et al. (2001) suggests the total forcing due to anthropogenic nitrate is small.

Like sulphate, nitrate chemistry involves many complex redox cycles. Nitrate aerosols may be

effectively removed by formation of the less reactive HNO_3, either by water reacting with N_2O_5 or via the reaction of HCl with $ClONO_2$. Both these reactions are prevalent in the stratosphere, particularly on ice-crystal polar stratospheric clouds (PSCs) in the Antarctic winter, where the redox processes are linked to ozone destruction (Pueschel, 1996).

Particle growth and size distributions

Direct emissions of aerosols and their precursors tend to be small molecular clusters. These quickly coagulate to give larger particles on length scales of 10^{-9} m which are said to be in *nucleation* mode (Junge, 1955). Coagulation readily occurs as particles are often coated in a film of moisture at typical humidities; alternatively, growth may occur through gas condensation onto existing particles. The mechanics and statistics of both of these are discussed in detail by Twomey (1977).

Not all aerosol emission is in the nucleation mode. Dependent on soil type and wind speed, dust particles of microns in size can be lofted, although the largest will typically settle soon afterwards (see for example Marticorena and Bergametti, 1995). Sea spray from bursting water droplets on the crests of waves can be up to 50 μm on emission (Zhang et al., 2005).

Aerosols of around 10^{-8} m to 10^{-7} m grow less rapidly (because they are larger, more coagulation or condensation events are required to increase their size, despite their larger surface area), and are known as *accumulation* mode aerosol. Particles of this size are still fairly buoyant due to Brownian motion, and are comparatively long-lived because of this buoyancy and slow growth (Twomey, 1977). Larger aerosol particles of size scales of 1 μm and greater are called *coarse* mode particles. Their lifetimes tend to be comparatively short, for two main reasons:

1. Particles of this size have sufficient mass to be affected by gravitational settling. A spherical particle of diameter 1 μm will fall at a speed of 1 mm every 5 seconds or so, with falling speed increasing quadratically with radius.

2. The largest particles act as cloud condensation nuclei (CCN) and so will eventually fall as rain or snow (typically as droplets of size 10^{-3} m, i.e. 1 mm).

As particles will be emitted at different sizes and grow at different rates, when considering samples of aerosol it is of use to define a distribution which accurately represents them. Junge (1955) showed that the sizes of many natural aerosols follow inverse power distributions, and as a result this

is also sometimes called the Junge distribution. With radius r and for some constants α and c, these distributions take the form:

$$r^{\alpha}\left(\frac{\mathrm{d}n(r)}{\mathrm{d}\log r}\right) = c \tag{1.1}$$

This is solved to give a cumulative distribution for the number density $n(r)$ of particles of size greater than some radius r, which is proportional to $r^{-(\alpha+1)}$. Junge (1955) placed the value of α at approximately 3. When integrating this distribution to obtain total particle number the integral diverges to infinity for all values of α. In practice some minimum radius is chosen as a cutoff. As a result other measurements, such as effective radius (see Section 2.1.1) or total volume, are more useful when describing the population than total particle number or number density.

More recent studies such as Davies (1974) and Hess et al. (1998) have suggested the lognormal distribution to be most appropriate for individual aerosol components. The aerosol classes used in the retrieval scheme presented here consider aerosols to be mixtures of lognormally-distributed components. For an aerosol with modal radius r_{i} and total number density N the number density $n(r)$ of particles of size r in a distribution with spread σ is given as follows:

$$n(r) = \frac{N}{\sqrt{2\pi}\, r \log \sigma \ln 10} \exp\left[-\frac{1}{2}\left(\frac{\log r - \log r_{\mathrm{i}}}{\log \sigma}\right)^2\right] \tag{1.2}$$

Size is not the only property of interest as a particle's composition may alter as the particle develops. This can occur through collision with different particles, condensation of water vapour or other gases, or chemical reactions within the aerosol particle, with the effect on its shape and microphysical properties depending on both the composition and the mixing state of its constituents (see, for example, Hãnel, 1976 or Wang and Martin, 2007). These changes are known as ageing processes.

1.1.4 Aerosol transport and sinks

Vertical transport

Most aerosol is produced at fairly low altitudes and remains confined to the troposphere, due to the lack of large-scale vertical mass exchange between the troposphere and stratosphere. Large particles

in particular rarely make their way from troposphere to stratosphere because they are affected by gravitational settling.

The most violent volcanic eruptions are able to inject aerosol and precursor gases directly into the stratosphere. Although the ash quickly settles, other aerosol (typically sulphate) is able to persist for longer timescales of a year or more (Pueschel, 1996) than tropospheric aerosol because of the high stability of the stratosphere. Aircraft exhaust also provides a source of stratospheric aerosol.

Even in periods of volcanic inactivity there is a persistent background stratospheric aerosol layer (Bingen et al., 2004). It is not thought to show strong temporal variability or be dense enough to affect climate: the main climatic perturbations of stratospheric aerosol come from periodic volcanic eruptions (Pueschel, 1996, Grainger and Highwood, 2003).

Horizontal transport

As small particles, aerosols are easily transported by winds, and thus the aerosol present in a region may well come from a distant source. This is important in pollutant analysis as the impact of the pollutant will not be confined to the emitting source. Saharan dust is also thought to contribute to fertilisation of the Amazon basin, as transport across the Atlantic Ocean has been observed (Formenti et al., 2001 and others). Measurement of aerosol transport is difficult as source regions may be large and inhomogeneous, and the distances involved are often great (on the order of 10^3 km). Nevertheless, various methods have been developed to attempt to qualify and quantify the long-range transport of aerosols.

One commonly-used method (for example Rahn and Lowenthal, 1984, Formenti et al., 2001, Borbély-Kiss et al., 2004) is to calculate the abundances of several trace elements in the aerosol sample, and then attempt to map these data to a particular source region. Possible source locations may be narrowed down using wind back-trajectory analysis. Elemental analysis may be performed with proton-induced X-ray emission methods, although this method is not without problems. One criticism from Thurston et al. (1985) is that as it measures total mass it cannot account for mass observed in one region coming from different sources, which may alter the apparant relative abundances of the elements analysed.

Transport may also be inferred by monitoring the temporal evolution of large aerosol masses. Light scattering apparatus such as LIDAR and nephelometers are often used in conjunction with

satellite data and wind fields to track the movement of aerosol (especially dust clouds). Such methods also often provide some detail about size distribution and vertical profiles and are frequently described in the literature (Dulac et al., 2001, Hamonou and Chazette, 1998), sometimes combined with elemental analysis (Formenti et al., 2001).

Aerosol sinks

In common with the wide variety of processes leading to the generation of atmospheric aerosol, there exists a wide variety of processes removing aerosol from the atmosphere. The following discussion on removal processes is based on the treatment by Twomey (1977), and may be summarised in three forms:

- Wet removal processes occur when aerosol is removed in precipitation (water or ice).
- Dry removal processes remove aerosol without the involvement of precipitation.
- Cloud formation also removes aerosol from the atmosphere, and eventually provides the mechanisms for wet removal processes to occur.

Wet removal processes

Aerosol particles may be rained out by collision with falling raindrops (inertial removal). The efficiency of this process is highest for heavy particles. Removal is also possible through diffusion of aerosols into falling drops, which is most efficient for small raindrops and particles less than 10^{-8} m in size (Twomey, 1977). The strong size dependence of these removal processes means that particles in the region of 10^{-8} m to 10^{-7} m are not efficiently removed, although as they are typically the most numerous particles a large total number is lost. For particles of all sizes the total aerosol amount removed through wet processes does, of course, depend on how much precipitation there is.

Dry removal processes

The largest difference between wet and dry removal is that while wet removal can occur throughout the atmosphere wherever there is precipitating water, dry removal relies on the particles being transported near to the surface by turbulence or gravity. The movement of small particles will be dominated by Brownian motion, while heavy particles will feel the effects of gravitational settling.

Diffusive removal is again only important for very small particles, and inertial only for larger particles, with low rates of dry removal of particles in the range from 10^{-8} m to 1 μm (Twomey, 1977).

Cloud formation and aerosol removal processes

As previously mentioned, large aerosol particles are able to act as cloud condensation nuclei (CCN). Cloud removal processes tend to be efficient over all aerosol size ranges due to the high surface area of cloud droplets, as well as the rapid condensation of small particles into larger ones which takes place. Clouds are able to form and rain within an hour, which can lead to rapid removal of aerosol.

Nucleation can only occur on aerosol particles larger than a critical radius of approximately $10^{-8}S_c^{-2/3}$ m (where S_c is the supersaturation in percent; Twomey, 1977). Following nucleation a period of rapid growth leads to the droplet expanding to the order of microns within a few seconds, provided supersaturation is maintained. Further growth is slower. Sufficiently large droplets fall out as rain, removing both directly-incorporated aerosol as well as any caught by wet removal processes. As a result, clouds are considered to be the most important means of removing aerosol from the atmosphere, with small particles being caught up diffusively and larger ones serving as CCN. Water-insoluble aerosol is affected less by wet removal and cloud processes than soluble, although insoluble particles may still act as cloud nucleation sites.

Work by Junge (1955) showed that, due to the high efficiency of wet removal processes, in principle a few millimetres of rain could effectively clear the atmosphere of aerosol. This suggests that the residence time of aerosol in the atmosphere is determined more by the frequency of rain, and the height at which clouds form, than the total amount of rain that falls.

1.2 The historical context of AATSR, and its capabilities

The Advanced Along Track Scanning Radiometer (AATSR) is an instrument aboard the ESA satellite Envisat, launched in March 2002. It is the successor of the earlier similar instruments ATSR-1, launched on ERS-1 in July 1991 and in operation until March 2000, and ATSR-2, launched on ERS-2 in April 1995 and in operation until February 2008 (with regular attempts to restart since then). Together the series provides a 17-year data set with many applications in climate studies, aside from their main purpose of monitoring sea surface temperatures.

While ATSR-1 measured radiance at one wavelength in the near-infrared and three in the thermal infrared part of the spectrum, ATSR-2 and AATSR have an additional three channels in the visible region. It is these visible channels which are key to the instruments' ability to provide data suitable for aerosol retrievals. AATSR is an improvement on ATSR-2 in this regard as ATSR-2 suffers from restrictions on the amount of data which can be downlinked while AATSR is able to provide continuous data from all channels.

1.2.1 Satellite orbital parameters

Envisat is in a sun-synchronous polar orbit with a mean local solar time (MLST) of 10:00 am for the descending node. It is at an altitude of approximately 800 km and 98.55° inclination to the equator. The reference orbit repeats once every 35 days, and global coverage is achieved once every 3 days as a result of orbital precession. A single Envisat orbit takes about 101 minutes. The orbit is controlled to keep it as precise as possible: within a maximum deviation of ± 1 km on the ground and ± 5 minutes on the MLST at the equator (ESA, 2007).

1.2.2 AATSR design and data acquisition

Channel	Central wavelength	Bandwidth
0.55 μm	0.555 μm	20 nm
0.66 μm	0.659 μm	20 nm
0.87 μm	0.865 μm	20 nm
1.6 μm	1.61 μm	0.3 μm
3.75 μm	3.70 μm	0.3 μm
11 μm	10.85 μm	1.0 μm
12 μm	12.00 μm	1.0 μm

Table 1.1: AATSR channel locations and bandwidths.

The basic structure of the instrument can be seen in Figure 1.2. AATSR and its predecessors are unique in that they use two views (near-simultaneous in time) with differing path lengths to discriminate between radiance from the surface and radiance from the atmosphere. The radiometer measures radiance at nadir and along track (the 'forward view') at seven channels, with positions and bandwidths summarised in Table 1.1.

Light is reflected off a scan mirror onto a paraboloid mirror, after which it is focussed and reflected

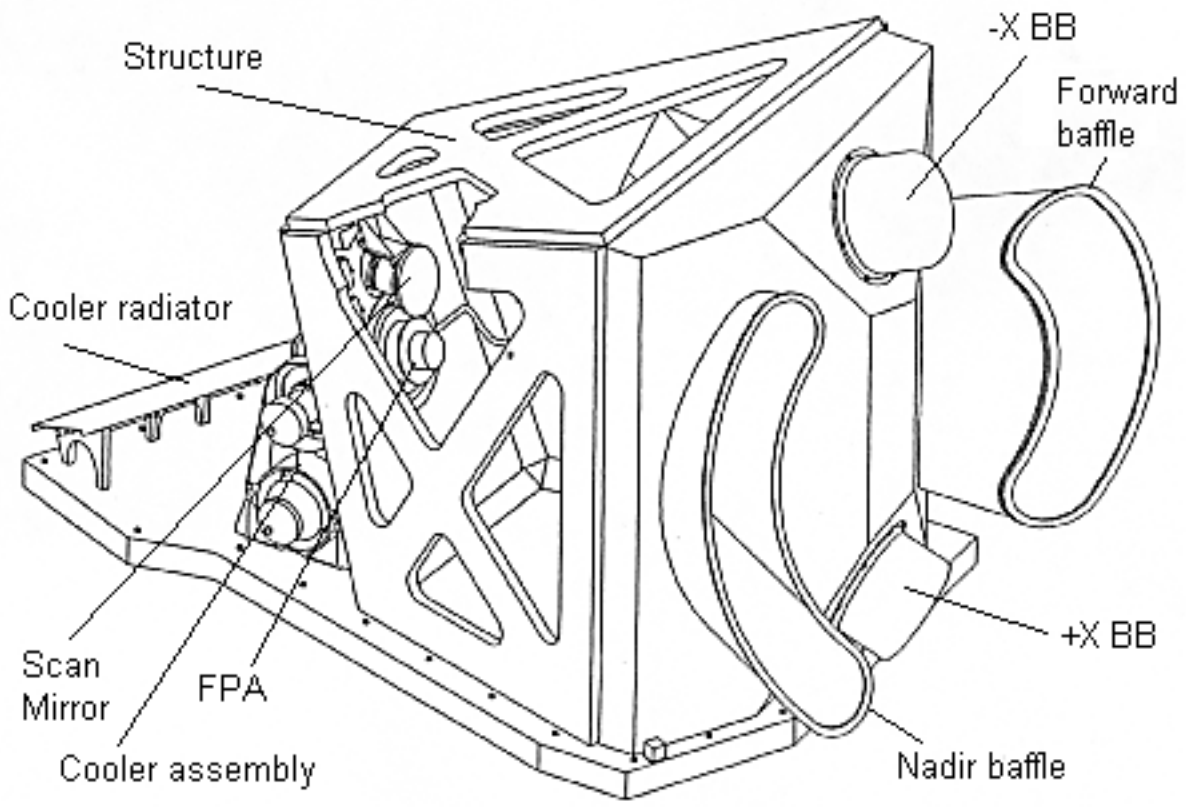

Figure 1.2: AATSR instrument structure. The FPA is the focal plane assembly, and the ±X BB the hot and cold black bodies respectively, used in calibration of the infrared channels. Taken from ESA (2005).

into the visible and infrared focal plane assemblies (FPAs). Here the radiance is converted into electrical signals, which are digitised and then sent to data-downlinking systems to be transmitted back to Earth. The baffles shield the nadir and forward windows from direct sunlight, ensuring radiance picked up is only from the region being scanned. The FPA for the infrared channels is cooled to about 80 K whilst that for the visible channels is maintained at ambient temperature.

The viewing geometry of the radiometer is shown in Figure 1.3. The region sampled consists of two curved swathes around 500 km wide: a nadir view, looking down at zenith angles from 0°-22°, and a forward view inclined between 53°-55° to the normal to the surface. There are 555 pixels across the nadir swath (with a size of about 1 km^2 at the centre) and 371 across the forward swath (with a size of about 1.5km^2 at the centre). During each scan cycle the satellite moves 1 km onward with respect to the Earth's surface. The scan cycle repeats 6.6 times per second (ESA, 2007).

With an altitude of around 800 km, the forward view samples approximately 1,000 km in front of the nadir at any given time. Thus after around 150 seconds the satellite has moved such that nadir view samples the same region, giving two views of the scene with differing path lengths. While this

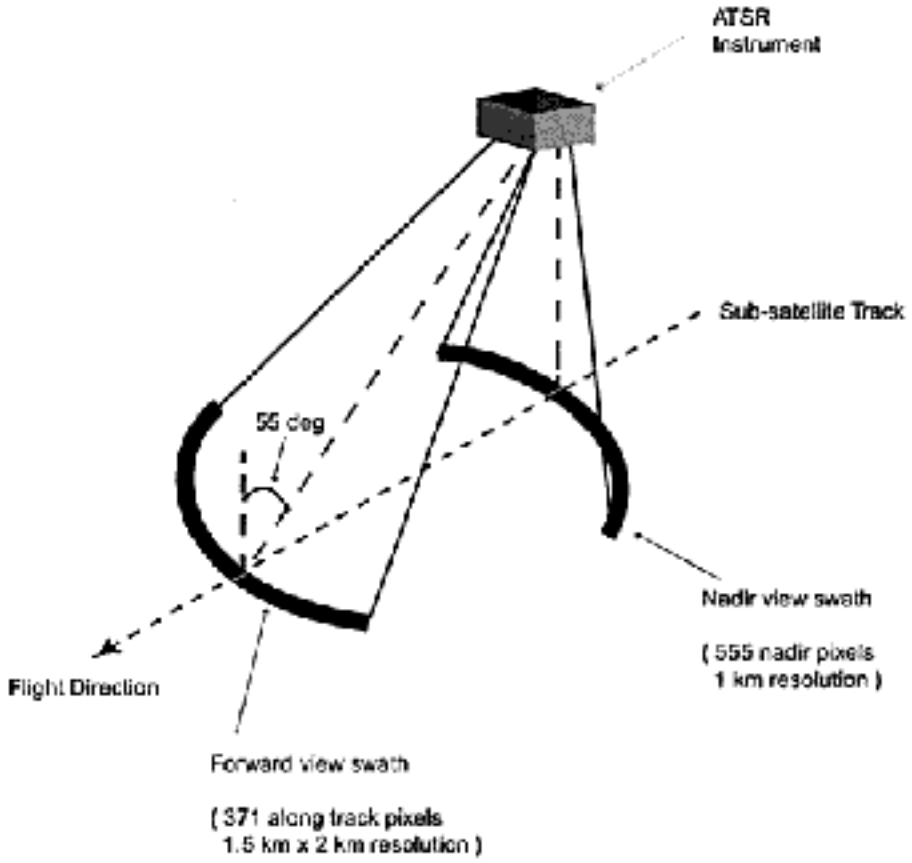

Figure 1.3: The (A)ATSR viewing geometry, taken from ESA (2005).

time delay means the views are not *strictly* simultaneous, it is still valid to treat them as such because this timescale is short compared with that of atmospheric processes. Were it not, then the dual-view approach of the (A)ATSR instruments would not be able to be exploited in this way.

The FPA is occasionally warmed up to outgas. This is because contaminants condense onto the cold surfaces of the FPA and detectors, degrading instrument operation due to signal attenuation, as well as increasing the burden on the FPA coolers. These contaminants do not affect signal calibration because they affect both views of the Earth and of calibration targets. During outgassing, which occurs for around 2 days at a time each 3 months, no data are collected from the IR or 1.6 μm channels. As the aerosol retrieval uses the 1.6 μm channel, aerosol retrievals cannot be performed during outgassing.

1.2.3 Level 1 data

AATSR level 1 data (TOA reflectances) is provided at full nadir-view resolution. Nadir and forward-view pixels are colocated with an accuracy of 1 pixel; the forward view's coarser resolution means

that some of the colocated forward-view data consist of 'cosmetic fill' pixels created via interpolation of 'true' forward-view measurements. Quality control parameters and ancillary data such as pixel ground coordinates and a cloud flag (for both nadir and forward-view pixels) are also provided.

1.2.4 Visible channel precision

VISCAL calibration of measurements and VISCAL precision

During operation, the AATSR instrument measures TOA radiance at the visible channel wavelengths which are integrated and digitised to give a 12-bit count taking values from 0 to 4095 (ESA, 2007). The on-board visible calibration system (VISCAL) consists of an opal diffuser which views the Sun once per orbit; a calibration detector count C_{VISCAL} is measured from this. In contrast to the view of the Sun, a view of space once per orbit gives a reference count C_{dark}. As the instrument flies the scene count, C_{scene}, is recorded for each pixel. The VISCAL reflectance factor, r_{VISCAL} (measured pre-launch), can be used in combination with C_{scene} and the solar zenith angle θ_0 to obtain the sun-normalised reflectance of the scene, r_{scene}, by the following method (Smith et al., 2001, Smith, 2005):

$$r_{\mathrm{scene}} = \frac{r_{\mathrm{VISCAL}}}{\cos(\theta_0)} \cdot \frac{C_{\mathrm{scene}} - C_{\mathrm{dark}}}{C_{\mathrm{VISCAL}} - C_{\mathrm{dark}}} \tag{1.3}$$

The retrieval uses the measurements of Sun-normalised reflectance. This is a dimensionless quantity which takes values from 0 to (typically) around 1, with 0 corresponding to no signal received at the detectors and 1 to a radiance equal to a view of the Sun.

Using this method, the calibration is not dependent on the absolute instrumental response for a pixel. This on-board automatic calibration is a unique feature of the ATSR instruments and ensures high data quality. Noise on the visible channels is small and random (as opposed to systematic). Smith (2007) stated the uncertainty on each of C_{VISCAL}, C_{dark} and C_{scene} to be one digitised count; however, periodic buildups of ice and outgassings to clear them affect r_{VISCAL}, giving rise to a temporal variation in the calibration of the results. The following factors, then, can lead to error (systematic or random) in an AATSR measurement:

- The instrument receiving radiance from a location outside of the ground pixel. The instrumental baffles are designed to minimise this contribution and so it is expected to be very small.

- Detector nonlinearities and degradation. These are expected to either be small or well-modelled.

Wavelength	Percentage error from Smith (2005)
550 nm	2.4 %
660 nm	3.2 %
870 nm	2.0 %
1.6 μm	3.3 %

Table 1.2: Instrumental noise on measured TOA reflectance for AATSR visible channels. These are percentage errors on r_{VISCAL} from pre-launch calibration discussed by Smith (2005).

- Uncertainty in the conversion of the measured radiance into a 12-bit count. Assuming the counts to be accurate to within 1, use of Equation 1.3 permits the estimation of the uncertainty arising on the reported scene reflectance. True values of the instrumental counts are not available, but it is likely that this error will be smaller than 1 %. This source of error diminishes with higher C_{scene} and greater contrasts between C_{VISCAL} and C_{dark}.

- Buildup of ice leading to changes in r_{VISCAL} and its precision, and hence an associated error in the reported r_{scene}. The ice acts as an interference film and causes oscillations in the signal. These change depending on the instrumental outgassing cycle. Smith (2004) states that r_{VISCAL} itself is minimally affected after additional ground-based calibration results (Smith et al., 2002) are applied; the signal to noise ratio, however, will change during this cycle.

The above points mean that any systematic error in the visible channels is minimal and the precision of r_{VISCAL} (i.e. random noise) is the dominant error source.

The accurate characterisation of instrumental noise is important as it affects both the progress of the retrieval and the error estimates of retrieved quantities at the solution. The accuracy of the measurements dictates how accurate the forward model needs to be, and sets a limit on how much the retrieval can be improved by refinements to the forward model. In a general model of noise N, we will have terms corresponding to the signal strength S and a random error:

$$N = e_1 S + e_2 \tag{1.4}$$

The relative magnitude of these coefficients e_1 and e_2 is important. If e_2 dominates then we may consider the noise to be approximately constant across the range of instrumental signal strengths; if not then the retrieval needs to be able to cope with different uncertainty estimates for different pixels.

Pre-launch noise tests on the visible channels were carried out by Smith (2005), giving a percent-

age error on r_{VISCAL}, as shown in Table 1.2. Much of the information within Smith (2005) was also published prior to Envisat's launch in Smith et al. (2001). Although in-flight performance may differ, these values are used as error estimates on the measured reflectances in the retrieval as information on the temporal variation of AATSR's precision is not available. It is possible to estimate whether these values are of an appropriate magnitude by considering the standard deviation of groups of nearby pixels, as undertaken below.

Methodology of noise estimation via standard deviation

As will be discussed in Chapter 2, the averaging of nearby pixels in the retrieval (known as *superpixelling*) is done with the assumption that on small scales an image is homogeneous so averaging data pixels will reduce noise. Looking at this another way, the variability of reflectance within superpixels is an estimate of the noise on the measurements. Two caveats must be borne in mind:

1. For very dull scenes the measurement may be affected by the digitisation of the signal, which will in turn affect the standard deviation.

2. The assumption of homogeneity will not hold true for all pixels. Coastal pixels, most land pixels and pixels in or near cloud may show particularly high variability. These will also tend to be brighter pixels, which limits how good an estimate can be made for high signal strengths.

To test this, an orbit of nadir-view data from September 5^{th} 2004 was superpixelled at 2×2 pixel resolution and resulting mean signal and standard deviation recorded for channels used in the retrieval. Forward-view data were not used as cosmetic fill pixels would artificially decrease the standard deviation of the superpixel. In groups of 100,000 superpixels (ordered by brightness, from dullest to brightest) the mean reflectance, mean standard deviation and modal standard deviation range (in increments of 10^{-4}) were calculated. This was done for the following pixel classifications:

- Dull superpixels (sun-normalised reflectance less than 0.05).

- Slightly brighter superpixels (sun-normalised reflectance up to 0.1).

- All pixel brightnesses.

A linear fit between mean reflectance and the most common standard deviation range gives gradient and intercept values (coefficients e_1 and e_2) as shown in Table 1.3. An example of this process for the first category of points is shown in Figure 1.4.

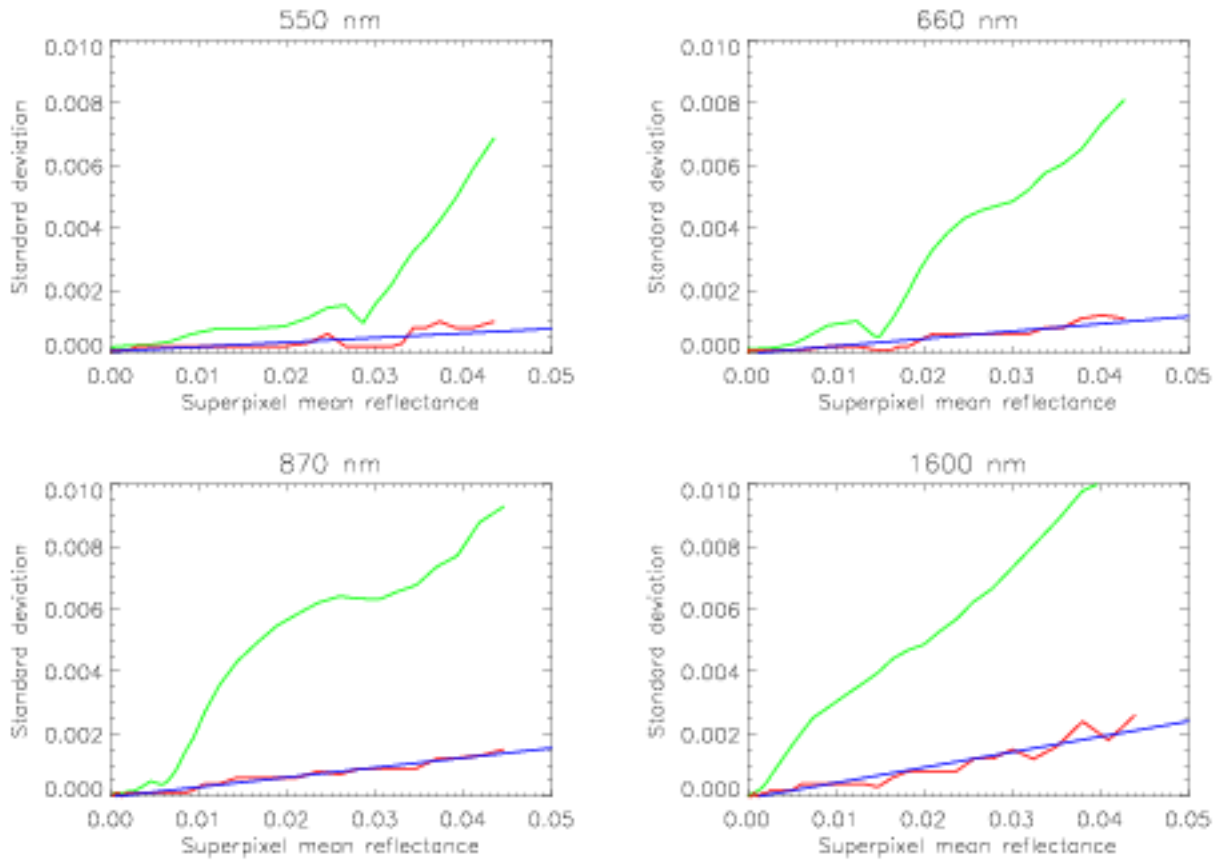

Figure 1.4: Superpixel mean sun-normalised reflectance against standard deviation for 2×2 pixel superpixels. Red lines show the modal standard deviation range (in increments of 0.0001) and green the mean standard deviation for groups of 100,000 points. Blue lines are a linear fit of the signal to the modal standard deviation.

The modal standard deviation range is more appropriate to perform a linear fit to than the mean because of the assumption of homogeneity. The standard deviation of a truly homogeneous superpixel will be the noise on the measurement. Any inhomogeneity will increase the mean standard deviation; the mode will be less affected provided most pixels are homogeneous. Evidence supporting this assumption is given by the values in Table 1.3; for brighter scenes (which are, as discussed previously, likely more inhomogeneous) there is an increase in the gradient of the noise.

Replication of the calculations using a different day of data or different bin size leads to similar results. Using larger superpixels leads to steeper gradients, which is as expected because the larger superpixel size means the scene observed is likely to be less homogeneous.

Comparison of estimates with pre-launch values

Previously, the ORAC algorithm assumed flat values of instrumental noise for each of the visible channels equal to 0.0048 at 550 nm, 0.0058 at 660 nm, 0.009 at 870 nm and 0.018 at 1.6 μm (in

Channel	Duller than 0.05		Duller than 0.1		All points	
wavelength	Gradient	Intercept	Gradient	Intercept	Gradient	Intercept
550 nm	0.014	0.00004	0.046	-0.00044	0.058	-0.00071
660 nm	0.024	0.00003	0.047	-0.00029	0.067	-0.00068
870 nm	0.031	-0.00001	0.058	-0.00038	0.072	-0.00058
1.6 μm	0.050	-0.00006	0.057	-0.00020	0.052	0.0

Table 1.3: Coefficients for noise estimates obtained by considering the standard deviations of reflectance in small superpixels. Given are intercepts and gradients (in units of sun-normalised reflectance) of linear fits of mean superpixel sun-normalised reflectance to standard deviation for each of the four vis/nIR wavelengths. Fits were obtained using three thresholds of brightness of mean sun-normalised reflectance: all superpixels duller than 0.05, all duller than 0.1 and finally all superpixels.

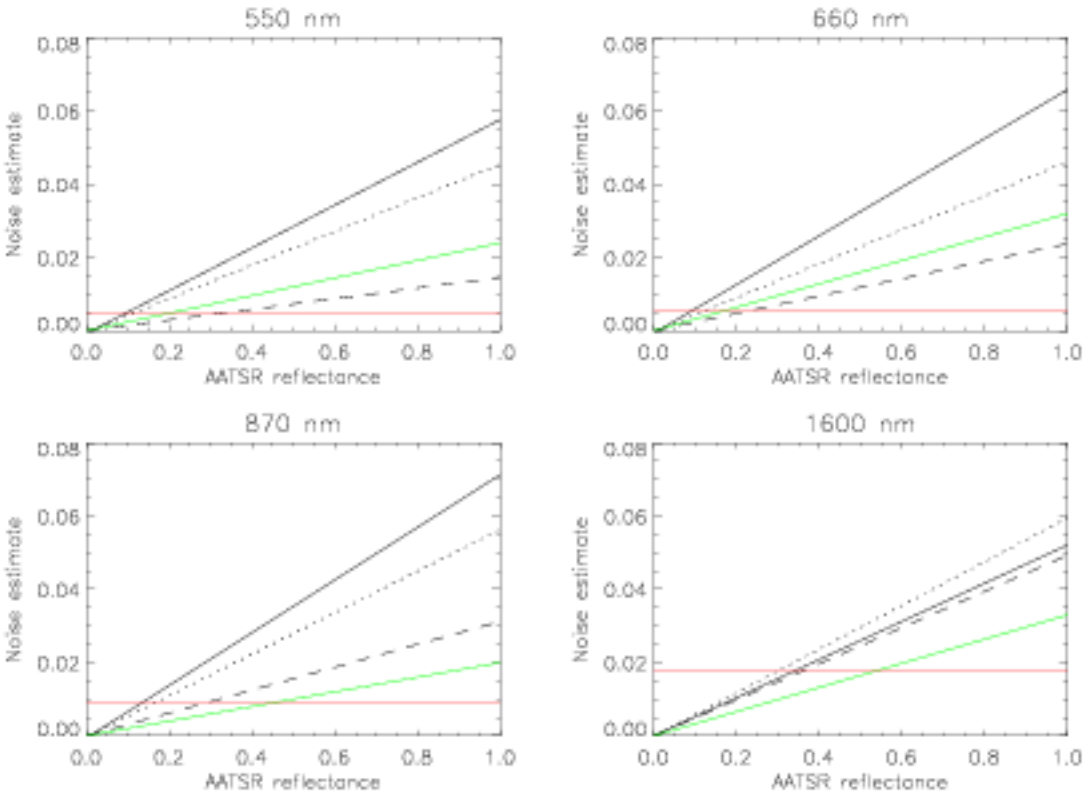

Figure 1.5: Noise estimates as a function of signal strength based on standard deviation measurements (see Table 1.3). Solid lines show data from all points, dotted using reflectances duller than 0.1 and dashed duller than 0.05. Red lines indicate the flat value previously taken as instrumental noise by the ORAC retrieval. Green lines show noise assuming uncertainties dependent on signal strength according to pre-launch calibration results (Smith, 2005), as currently used in the retrieval scheme.

units of sun-normalised reflectance). Figure 1.5 shows calculated noise profiles for AATSR visible channels based upon the coefficients found for Table 1.3, with the pre-launch calibration estimates, and the flat values previously used, for comparison. In all cases there is a crossover point at which the previous flat values change from being an overestimate to an underestimate as compared to the standard deviation-based measurements and pre-launch results currently used. This crossover point is

above the observed reflectance for typical cloud-free scenes, suggesting that the previous flat values were too high.

The standard deviation-based estimates are generally higher than the pre-launch calibration percentage uncertainties. This positive bias is expected as a result of the assumption of homogeneity of superpixels. The discrepancy may also be due to changes in instrumental performance since launch. It is encouraging that the pre-launch results are closest to the estimates based on dull superpixel standard deviation, as that is where this method is expected to work best. The negative intercepts found in Table 1.3 cannot translate to a negative magnitude of uncertainty, but the fact that they are all small suggests that the random term of the noise (Equation 1.4) is almost zero.

This analysis suggests that the pre-launch calibration results are appropriate for use as error estimates on reflectances measured in-flight. Orbit-based measurement noise values would be the optimal choice from a mathematical point of view, but in their absence it is encouraging that results from the pre-launch tests seem to hold.

Minimum threshold of noise

The percentage values of noise must have a lower limit corresponding to errors in the digital counts (Equation 1.3). Assuming C_{VISCAL} and C_{scene} to be large means that this will only be a factor for very dark pixels. For these points it would be necessary to abandon a percentage estimate of noise and use a flat value. Inspection of the data from the previous analysis reveals that such a lower limit of standard deviation is observed for the dullest superpixels (with r_{scene} less than around 0.002). This lower limit is about 0.0005 at 550 nm and 0.0003 at longer wavelengths. This suggests a threshold of pixel sun-normalised reflectance below which digitisation error dominates of around 0.025, which is recorded only rarely. Therefore in most cases the count error is not important, assuming these results are typical of the mission as a whole.

These lower-limit values are adopted in the retrieval whenever the percentage error estimate would result in a value for noise smaller than these threshold values. This prevents the retrieval algorithm from being overconstrained.

1.2.5 Infrared calibration and long-term performance

The infrared channels of AATSR are calibrated during each instrument scan cycle. Between each nadir and forward view scan, one of a pair of black-body targets (visible in Figure 1.2) are viewed to calibrate the IR channels. One is maintained at 305 K while the other is at the ambient instrument chamber temperature (256 K). This range encompasses all the expected brightness temperatures that the instrument measures. The black bodies are cylindrical, insulating and non-reflective with a temperature monitoring system.

The performance of AATSR to date was discussed most recently by Smith (2008) and Smith et al. (2008). The instrument was found to be fulfilling its original specifications, despite having passed its nominal lifetime of 5 years. For the work described in this thesis, the most recent calibration data are applied throughout.

1.3 Non-aerosol uses for AATSR data

1.3.1 Retrieval of sea surface temperature

As mentioned, the primary objective of the ATSR series of instruments is to provide a continuous, long term global record of sea surface temperature (SST). This was originally determined from infrared data, using the 1.6 μm channel to aid in cloud-flagging (Závody et al., 1995). Measuring SST accurately is important in climate research; high precision is required so any long-term trend is not hidden by noise. As a result, the AATSR mission requirements state that the instrument should be able to retrieve SST with an accuracy of better than 0.3 K (1σ limit), and up to 0.1 K under favourable conditions. In addition, there are strict requirements on instrumental calibration, to assess any long-term drift and provide continuity with the ATSR-1 and ATSR-2 data sets.

Further improvements to the instruments, such as the inclusion of the additional visible channels, have been accompanied by improvements in the algorithm (Merchant et al., 1999 and Merchant and Le Borgne, 2004). However as a correction is made for the effects of aerosol, rather than a full retrieval, there is scope for providing an improved SST dataset through extension of the algorithm presented here to include the infrared measurements.

Satellites offer several advantages over other SST data sources (such as ships and buoys), notably a much greater coverage. Furthermore as it provides a single dataset AATSR is self-consistent

whereas ships and buoys, being many datasets collected in different ways, are not. It should, however, be noted that ship and buoy measurements are still useful (for example for validation of results from satellite data).

1.3.2 Land remote sensing

A secondary objective of AATSR was to continue ATSR-2's records of data for land remote sensing of properties such as land surface temperature (LST) and Normalized Difference Vegetation Index (NDVI), which provides information about surface cover. A difficulty associated with remote sensing over land is that the land surface reflectance often shows a strong dependence on viewing geometry and can be highly heterogeneous over the scale of the instrument's resolution. Despite these difficulties, progress in land remote sensing including retrieval of LST (Sória and Sobrino, 2005) and surface reflectance (simultaneously with aerosol properties) has been made with (A)ATSR data (North et al., 1999 and North, 2002a). The algorithm presented here makes use of prescribed surface bidirectional reflectance (BRDF) models but retrieves the surface albedo, which is an important parameter in climate models and can be an indicator of changing land-use. The NDVI can be derived from this.

Burning vegetation can also be monitored using the same method as retrievals of LST, although the fact that fire temperatures exceed that of the hot calibration body (305 K) means that the temperatures retrieved are not as precise as those for SST or normal LST. Still, the ATSR series has become an important tool in the monitoring of biomass burning (Huang and Siegert, 2004).

1.3.3 Retrieval of cloud properties

The GRAPE project (Grainger, 2008) provided the foundation for the work described in this thesis. This ongoing joint project between the University of Oxford and the Rutherford Appleton Laboratory (RAL) involves the development of the optimal estimation retrieval scheme known as ORAC (Oxford-RAL Aerosol and Cloud) to derive cloud and aerosol parameters from ATSR-2. The algorithm presented here is a greatly updated version of the GRAPE aerosol algorithm. For both aerosol and cloud, GRAPE currently makes use of only the nadir-view data.

The principles behind aerosol and cloud retrievals are similar in that aerosol may be considered to be a very optically thin (and often low-lying) layer of cloud. The cloud algorithm makes use of information in the visible and infrared channels to retrieve cloud optical depth, effective radius,

cloud-top temperature, the fraction of cloud within a pixel or superpixel, and the phase of the cloud particles (Poulsen et al., 2008, Sayer et al., 2008). Cloud-top temperature may be converted to height and pressure.

1.4 Other instruments in use for aerosol characterisation

A wide variety of instrumentation is in current or recent use to characterise atmospheric aerosols. These can be broadly divided into two categories: satellite instruments similar to AATSR for remote sensing, and terrestrial instruments (either on the ground or aboard ships and aircraft).

1.4.1 Satellite-borne

MERIS

Like AATSR, the MEdium Resolution Imaging Spectrometer (MERIS) is aboard Envisat. Both are radiometers, with MERIS measuring at 15 wavelengths in the visible region. MERIS also has a higher resolution (approximately 300 m) and wider swath (1,150 km). MERIS only has one viewing geometry, which gives AATSR an advantage in terms of making corrections for the aerosol contribution to signal. MERIS is primarily used for ocean colour measurements and vegetation remote sensing. A lack of channels in the IR can cause problems with cloud contamination.

Being jointly aboard ENVISAT offers opportunities for synergy between MERIS and AATSR data. There is currently no operational synergystic product, although a synergistic retrieval of aerosol and surface properties is being developed by North et al. (2008). MERIS data are used along with (A)ATSR and SEVIRI in the GlobAerosol project (Thomas et al., 2008a), although the instruments are processed individually, meaning this is not a true synergystic retrieval.

MISR

The Multi-angle Imaging Spectro-Radiometer (MISR) flies on NASA's TERRA satellite, launched in 2000. Designed and used for aerosol retrieval (for example Abdou et al., 2005 or Keller et al., 2007), like AATSR it measures in the visible (over 4 channels), at 9 widely-spaced angles. The multi-angle measurements provide some information about the sphericity and size of the aerosol. MISR is capable of a very high spatial resolution of 275 m (250 m at nadir) although in practice for

a more reasonable data transmission rate it tends to work with a resolution of between 1 × 1 km and 4 × 4 km, adjustable for each detector. However, cloud clearing of scenes can be difficult as it lacks infrared channels. Its narrow swath of 360 km limits coverage.

MODIS

An instrument carried on two satellites (NASA's TERRA and AQUA, launched in 2000 and 2002 respectively), the MODerate resolution Imaging Spectroradiometer (MODIS) measures at 36 bands in the visible and infrared. Resolution varies between 250 × 250 m and 1 × 1 km dependent on band, and it has a wide swath (55.5°, giving a 2,330 km overpass). Remer et al. (2008) have recently published a global aerosol climatology from MODIS, which is an important step in quantifying aerosol abundance. As well as retrieving aerosol, a primary use of MODIS is land surface classification and a MODIS product is used to set the *a priori* surface albedo used in the AATSR retrieval (described in Section 2.3.2 and Chapter 4). Over the ocean MODIS is used for ocean colour and SST measurements.

The easy availability of MODIS data and its excellent spatial coverage mean it is often used in validation and intercomparison studies. However, the single viewing geometry means that it lacks sufficient information to retrieve aerosol and surface parameters simultaneously from a single overpass without assumptions about the spectral shape of the surface or characteristics of the aerosol. Although a MODIS surface reflectance product is used in the retrieval scheme described in this work, it uses data from 16 days of orbits to obtain a sufficient atmospheric correction.

Due to the simplifying assumptions required by having only a single view of measurements, the assorted products are produced with different algorithms, and in some cases multiple algorithms are used for the same product. For example, the regular land MODIS aerosol algorithm (Abdou et al., 2005, Remer et al., 2008) often fails over bright desert surfaces, where the *Deep Blue* algorithm (Hsu et al., 2004, Hsu et al., 2006) is used instead. An advantage of the scheme presented here is that it is consistently and globally applicable.

SEVIRI

Unlike the other satellite instruments mentioned here, the Spinning Enhanced Visible and Infra-Red Imager (SEVIRI) is in a geostationary orbit. Located above equatorial Africa aboard Meteosat 8 and

9, each images the visible portion of the globe every 15 minutes using 3 channels in the visible/near-infrared region and 8 in the infrared at 3 km by 3 km resolution. Launched on 29 January 2004 and 21 December 2005 respectively, they are expected to have a lifetime of 7 years. As well as aerosol detection, SEVIRI is used in remote sensing of cloud and surface properties, and to provide clear-sky radiance measurements for numerical weather prediction models.

SEVIRI deserves special mention in that it is used in the GlobAEROSOL project alongside (A)ATSR and MERIS (Thomas et al., 2008a). The same principles apply to both (A)ATSR and SEVIRI aerosol retrievals in the project, although as previously mentioned the SEVIRI viewing geometry is different, as are the precise channel wavelengths and bandwidths.

AVHRR

The Advanced Very High Resolution Radiometers (AVHRR) are a series of instruments which have flown on the American satellites TIROS-N and NOAA-6 onwards (with improvements made since NOAA-15). They are primarily used for surface temperature and cloud measurements, but have also been used to obtain aerosol properties (Wetzel and Stowe, 1999).

AVHRR measures at 5 bands in the visible and infrared regions at a resolution adjustable to either 1.1× 1.1 km or 4×4 km. Aerosol properties over the land are less reliable as simplifying assumptions about the surface reflectance must be made, but over the ocean it provides a long time-series of data (October 1978 to present).

This long time-series has been put to use by Mischenko and Geogdzhayev (2007) and Zhao et al. (2008) to study long-term trends in aerosol optical depth over global oceans. Confidence in these trends, however, is limited by the quality of the calibration of the AVHRR data. The ATSR series of instruments have an advantage here in that their increased number of measurements and dual-viewing geometries mean retrieval over land is reliable, and their on-board calibration should mean that results from ATSR-2 and AATSR are more directly comparable.

OMI

A joint venture between the American NASA, the Netherlands' Agency for Aerospace Programs (NIVR) and Finnish Meteorological Institute (FMI), the Ozone Monitoring Instrument (OMI) aboard AURA was designed as a successor to the Total Ozone Mapping Spectrometer (TOMS). OMI is a

nadir viewing radiometer which measures backscattered solar radiation from 0.270 to 0.500 μm, with a spectral resolution of about 0.5 nm. It has a 2600 km swath, resulting in daily global coverage, and a resolution of 13$\times$ 24 km in normal operational mode. In addition to O_3, OMI has been used to retrieve aerosol properties (Braak, 2006).

CALIPSO

Another joint venture (between NASA and the French CNES), the Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observation (CALIPSO) satellite was designed specifically to monitor cloud and aerosol (Kaufman et al., 2003). Launched on April 28 2006, CALIPSO carries 3 co-aligned instruments measuring at nadir. The Cloud-Aerosol Lidar with Orthogonal Polarization (CALIOP) measures backscattered light intensity at 532 nm and 1064 nm. This can provide information on the vertical distribution of cloud and aerosol, although coverage is limited to the sub-satellite track.

1.4.2 Terrestrial

AERONET

The Aerosol Robotic Network (AERONET, Holben et al., 1998) is a global system of sun photometers providing data on aerosol properties and precipitable water. A diverse range of environments are covered, as shown in Figure 1.6. Data are available from 1993 onwards (although not all sites have been active since then).

Because of the good spatial and temporal coverage AERONET sites are often chosen to validate satellite retrievals, although large areas of land (and open ocean) lack stations so true global coverage is not possible. Ground-based photometers have an advantage over satellites in that measurement can be taken continually at a specific location if continuity is required, although as they are point measurements they are less useful than satellites for estimating global coverage. Being physically located on the ground also makes maintenance of such instrumentation simpler than that of a satellite, and their calibration is aided by the presence of the Sun as a well-defined light source to sample. As such they are generally considered accurate data sources and good to validate against (Holben et al., 1998, Ichoku et al., 2002).

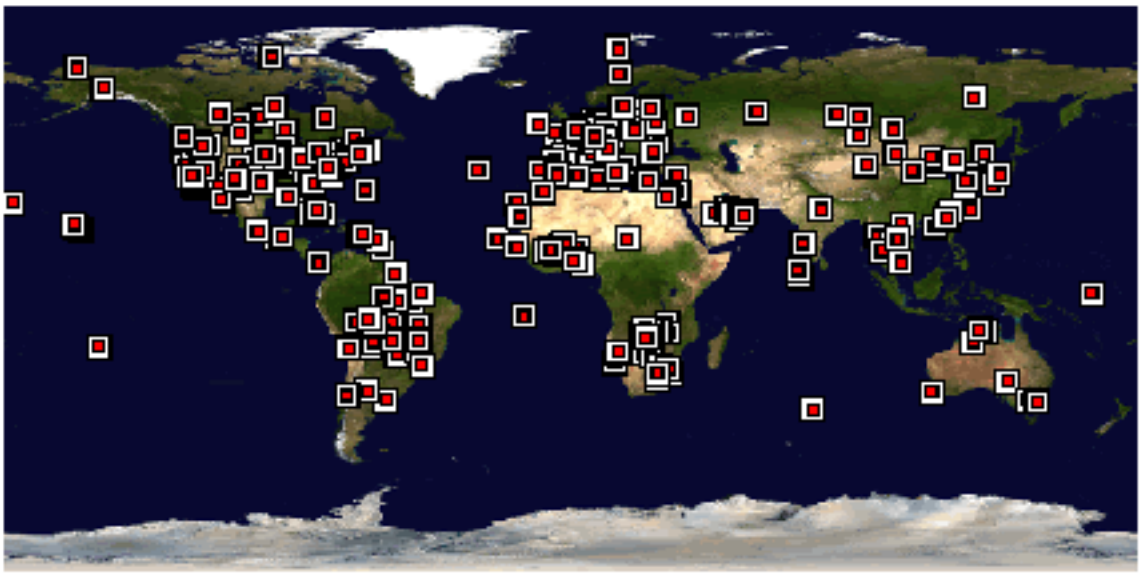

Figure 1.6: AERONET global coverage. Red squares indicate the presence of an AERONET station. Image taken from NASA (2008).

Other land, sea and air campaigns

The literature shows a wealth of one-off or ongoing campaigns of aerosol characterisation based upon data collected on land (Lu et al., 2006), at sea (Ramachandran, 2005) and increasingly by aircraft (such as Inomata et al., 2006, Targino and Noone, 2006 and Yoon et al., 2006), although most of these are local as opposed to worldwide in scope. A variety of techniques are typically employed, ranging from photometry and LIDAR to mass spectrometry.

Chapter 2

Aerosol retrieval algorithm

2.1 The inverse problem

The field of retrieval theory is dedicated to the solution of the inverse problem, defined and investigated thoroughly by Rodgers (2000). There are systems in which we are unable to directly measure the quantity we are interested in, but are able to measure some other quantity related to or resulting from it. From these measurements $\mathbf{y}$ we wish to find out something about the state $\mathbf{x}$ the system is in. Given $\mathbf{y} = \mathbf{F}(\mathbf{x})$ for some function $\mathbf{F}$ known as the 'forward function', approximated by the 'forward model' $\mathbf{f}(\mathbf{x})$, solving the inverse problem means being able to find the inverse function $\mathbf{x} = \mathbf{f}^{-1}(\mathbf{y})$ and so obtain our best estimate $\hat{\mathbf{x}}$ of the state $\mathbf{x}$ from measurement $\mathbf{y}$.

The premise behind performing an aerosol retrieval is that aerosol in the atmosphere will affect the light received at the satellite through a combination of absorption and scattering. The top-of-atmosphere reflectances AATSR measures form the measurement vector $\mathbf{y}$ and the state information vector $\mathbf{x}$ to be retrieved consists of various atmospheric aerosol and surface properties (namely aerosol optical depth, effective radius, and surface albedo at the four AATSR visible/near-IR wavelengths). Aerosol radiative properties depend on factors such as aerosol particle size, shape, chemical composition and vertical profile so a forward model must take such factors into account. The use of satellite data in this way is an example of remote sensing (or remote sounding), as the measurements are taken at a distance.

The algorithm developed and used here is based on the nadir-view GlobAEROSOL algorithm for retrieval of aerosol properties, described by Thomas et al. (2006) and Thomas et al. (2008b),

and developed at Oxford and the Rutherford Appleton Laboratory (RAL). The nadir-view algorithm has itself been developed from an earlier algorithm designed for ATSR-2 (Marsh et al., 2004). The scheme is known as the Oxford-RAL retrieval of Aerosol and Clouds (ORAC) and is applicable to both ATSR-2 and AATSR date. Continuing the chain further back, ORAC was built from the Enhanced Cloud Processor (ECP) described by Watts et al. (1998), which was employed to retrieve cloud properties from ATSR-2.

The main development from the previous ORAC scheme to the one described here is that the former only uses the instrument's nadir view channels, while this scheme makes use of both nadir and forward views (and hence is known as a dual-view retrieval). This doubles the number of measurements in use, significantly improving the characterisation of aerosol, and allows the retrieval of surface albedo at all four wavelengths of interest (as opposed to just 550 nm). Additionally, the error budget of the algorithm has been thoroughly investigated and revamped, and the treatment of the surface has been overhauled by use of bidirectional reflectance distribution functions, as described in succeeding chapters.

The mathematical notation used in this work is collated in Appendix A for ease of reference.

2.1.1 Retrieved aerosol parameters

A sample of atmospheric aerosol can be described by its number density, size and shape distribution, vertical profile and chemical make-up (which determines refractive index). Assuming the particles are spherical allows computation of radiative properties from the theory of Mie (1908).

Aerosol optical depth, $\tau(\lambda)$ (for wavelength λ), is the first quantity retrieved in the ORAC scheme, and describes the extinction of incoming irradiance by the atmosphere due to absorption and scattering by the aerosols. Optical depth at 550 nm is retrieved directly during the iteration, and later related to optical depth at 870 nm using lookup tables (based upon knowledge of the aerosol model used and the retrieved effective radius, discussed later). Optical depth is defined as:

$$\tau(\lambda) = \int_0^\infty \beta_e(z,\lambda)\,dz = \int_0^\infty (\beta_a + \beta_s)(z,\lambda)\,dz \tag{2.1}$$

The total extinction coefficient β_e, is the sum of the extinction due to absorption (β_a) and scattering (β_s). Radiative characteristics are described by the vertical profiles (z) of β_a and β_s along with the

scattering phase function $P(\theta)$, which determines the angular distribution of the scattered radiation, where θ is the scattering angle between source and observer. Other useful optical properties include the single scattering albedo ω_0, which is the ratio β_s/β_e, and the asymmetry parameter g which is the cosine-weighted integral of $P(\theta)$ from 0 to 180°. For a given aerosol shape, size, and refractive index, β_e is proportional to the aerosol concentration while $P(\theta)$ is not. Mie theory shows that the extinction coefficient β_e is given by:

$$\beta_e(z,\lambda) = \int_0^\infty Q_e(z,\tilde{n},x)\,\pi r^2\,n(z,r)\,dr \tag{2.2}$$

Q_e is the Mie extinction efficiency factor, dependent on the Mie size parameter $x = 2\pi r/\lambda$ and the refractive index of the particles ($\tilde{n}= n_r - i\kappa$). $n(r)$ is the aerosol number size distribution, dependent on radius r.

The most appropriate distribution to use to characterise atmospheric aerosol is the log-normal one (Davies, 1974). This is described by its modal radius r_i, standard deviation (in log space) σ and total number density N as described previously in Equation 1.2. The second parameter retrieved provides information about aerosol particle size. This is the effective radius, r_e, defined as the ratio of the 3rd to 2nd moments of the size distribution:

$$r_e = \frac{\int_0^\infty r^3 n(r)dr}{\int_0^\infty r^2 n(r)dr} \tag{2.3}$$

Vertical distribution also in principle affects radiative properties, although this is not retrieved as the visible channels do not provide sufficient information on it for AATSR visible measurements (shown in more detail in Section 2.4.3). As a result the aerosol height distribution at height h and for scale height Z (dependent on aerosol type) is fixed and described as follows:

$$N(h) = N(0)e^{-h/Z} \tag{2.4}$$

No information on aerosol shape distribution is retrieved and aerosol is assumed to be spherical. This is largely due to difficulties in the modelling of radiative characteristics of non-spherical particles, although some progress has recently been made in this field by using spheroid models (such as Dubovik et al., 2006). The spherical approximation holds for many classes of aerosol, although solid

aerosol particles like desert dust may be nonspherical.

2.1.2 Retrieved surface parameters

A large proportion of the radiance arriving at the satellite is frequently due to direct reflection of sunlight off the surface, particularly over land. The next four retrieved parameters are the white-sky surface albedo at AATSR visible channel wavelengths (550 nm, 660 nm, 870 nm and 1.6 μm). This is defined as the ratio of outgoing to incoming irradiance integrated over all solar and viewing angles, and as such is dimensionless and takes values from 0 to 1. The treatment of the surface reflectance in the retrieval is described in details in Section 2.2.3 and the following two chapters of this document.

If, for whatever reason, a single-view retrieval is performed instead of a dual-view retrieval then only the albedo at 550 nm is retrieved and the spectral shape from the *a priori* model is used to adjust the albedo at the other wavelengths. This is because, with only one view, there is not enough information to retrieve the albedo at all wavelengths simultaneously: the problem is underdetermined.

2.2 The forward model

To save on computational time, the aerosol forward model used to compute radiance for a given atmospheric state uses a series of lookup tables (LUTs) rather than performing radiative transfer calculations on-the-fly. The LUTs are generated using assumed values of aerosol refractive index, size distribution, phase function and vertical distribution, as well as assumptions of the absorption of atmospheric gases. They are generated for a 32 layer plane parallel atmosphere of 100 km in height, with a 1 km spacing from 0 to 25 km, stretching to 30 km for the top layer.

The Optical Properties of Aerosols and Clouds (OPAC) database of Hess et al. (1998) provides information on physical (e.g. size and vertical distribution) and optical (e.g. wavelength-dependent complex refractive index) aerosol properties, which are used to build up profiles for the characteristic aerosol types. Information on various aerosol components is used to generate four aerosol classes: continental clean, desert, maritime clean and urban. Additionally, a model for biomass burning aerosol taken from Dubovik et al. (2002) is used as a fifth aerosol class. Information on the components of the first four classes can be found in Table 2.1 and details of the assumed composition of the classes in Table 2.2. Increased relative humidity provides an increase in the amount of atmospheric

Component	r_i (μm)	σ	Wet Mode Radius
Insoluble	0.47	2.51	indep.
Water-soluble	0.0212	2.24	*
Soot	0.0118	2.00	*
Sea Salt (Accumulation)	0.209	2.03	*
Sea Salt (Coarse)	1.75	2.03	
Mineral (Nucleation)	0.07	1.95	
Mineral (Accumulation)	0.39	2.00	indep.
Mineral (Coarse)	1.90	2.15	

Table 2.1: Aerosol components, modal radius, and standard deviation for a log-normal size distribution and an assumed relative humidity of 0 %. Entries in the fourth column refer to the publications where the values of the wet mode radii can be obtained: * refers to d'Almeida et al. (1991) or Hess et al. (1998). The label regarding the wet mode radius is set to 'indep.' if the aerosol component is not hygroscopic.

water vapour which can condense onto the particles, altering their size and refractive index. This effect is taken into consideration utilising the method of Hänel (1976).

- The water-insoluble component of aerosol particles consists mostly of soil particles containing organic material. The water-soluble substances are a mixture of sulphate, nitrate, and organic compounds, the mixing ratio of which varies with the type of sources (such as biogenic, fossil fuel combustion, volcanic, biomass burning).

- The term 'soot' represents absorbing black carbon that is either a product of direct particle emission into the atmosphere by combustion processes that form highly absorbing graphitic along with weakly absorbing organic particles, or a transformation within the atmosphere of combustion related gases to particles.

- Sea salt particles comprise the various kinds of salt contained in seawater. Two modes are given to allow for a different wind-speed-dependent increase of particle number for particles of different sizes.

- Desert dust is produced in arid regions and consists of a mixture of quartz and clay minerals. It is modelled with three modes to allow for increasing relative amounts of large particles with increasing turbidity. The size of mineral aerosol particles is assumed to be insensitive to relative humidity effects.

Class	Components
Continental clean	Water-soluble, water-insoluble
Desert dust	Water-soluble, mineral (nucleation, accumulation & coarse modes)
Maritime clean	Water-soluble, sea salt (accumulation & coarse modes)
Urban	Soot, water-soluble, water-insoluble

Table 2.2: Composition of aerosol classes used in the retrieval for classes derived from Hess et al. (1998). A 50 % relative humidity is assumed for the hygroscopic components. Mixing ratios depend upon the distribution effective radius.

Mie theory is used to calculate the scattering phase function for the central wavelength of each channel over a range of 20 effective radii r_e (from 0.02 μm to 20 μm). The exponentially-decreasing vertical profile described previously is then used to give a vertical profile of scattering. Phase function generation requires the following pair of assumptions:

- The radiative properties of the aerosol are linear across the spectral width of each instrument channel. This is a reasonable assumption because the features of aerosol extinction spectra are very broad in comparison with gas absorption or emission features.

- The components' size distributions are, as previously mentioned, assumed log-normal. Changes in effective radius are achieved by adjusting the proportions of the differently-sized aerosol modes. If effective radius for an aerosol type equals that of the smallest or largest component in it, then the type becomes a single-component aerosol. If the size lies outside this range then the modal radius is shifted (keeping the distribution width constant)—in such situations the assumed size distribution becomes less reliable.

For each layer at which the aerosol distribution is defined, the extinction coefficient β_e, single scattering albedo ω_0 and the coefficients of a Legendre expansion of the scattering phase function are calculated for each channel and effective radius.

2.2.1 Modelling atmospheric gas absorption

Following on from aerosol scattering properties, gas absorption over the instrument band passes is calculated using the MODTRAN database (version 3.5-v1.1) outlined by Brown et al. (2004) and convolved with the instrument filter transmission functions. This database provides tropical, mid-latitude summer and winter, subarctic summer and winter, and US Standard Atmosphere climatolog-

ical atmospheres for H_2O, CO_2, O_3, N_2O, CO, CH_4, plus single profiles for HNO_3, NO, NO_2, SO_2, O_2, N_2, NH_3 and CFCs.

ORAC LUTs are generated using the mid-latitude summer atmosphere only. This simplification can be made as gas absorption is weak compared to aerosol extinction in the visible. Additionally, the instrumental channels used are located in so-called 'atmospheric window' regions of the spectrum and so are free from strong absorption features of gases which show large spatial and temporal variability (most notably H_2O).

2.2.2 Modelling atmospheric transmission and reflectance

To conclude LUT generation, atmospheric transmission and bidirectional reflectance values are computed from aerosol phase function and gas absorption profiles previously calculated. The DIscrete Ordinates Radiative Transfer (DISORT) software package (Stamnes et al., 1988, Olesen, 2005) is used for this purpose. DISORT is a well-documented and often used general purpose algorithm for time-independent multiple-scatter transfer calculations. It solves the equation for the transfer of monochromatic light at wavelength λ as described by:

$$\mu \frac{dL_\lambda(\tau_\lambda, \mu, \phi)}{d\tau} = L_\lambda(\tau_\lambda, \mu, \phi) - L^S_\lambda(\tau_\lambda, \mu, \phi) \tag{2.5}$$

In Equation 2.5, $L_\lambda(\tau_\lambda, \mu, \phi)$ is the intensity along direction (μ, ϕ), where μ is the cosine of the zenith angle and ϕ is the azimuth angle, at optical depth τ_λ, measured perpendicular to the surface of the medium. $L^S_\lambda(\tau_\lambda, \mu, \phi)$ is the source function.

The surface albedo is set to zero for LUT creation. The equation is solved for L to give tabulated transmission or reflectance (depending on the LUT) for each of the 20 effective radii, 20 logarithmically-spaced 550 nm optical depths (between approximately 0.01 and 6) and Sun/satellite geometry (specified by 10 equally-spaced zenith angles and 11 equally-spaced azimuth angles). During the retrieval, linear interpolation is used for values falling in between the points. Five LUTs for each aerosol type/instrument channel combination are produced:

- Bidirectional reflectance of the aerosol layer, R_{BD}
- Reflectance of the aerosol layer to diffuse radiance, R_{FD}

- Diffuse transmission of the incident beam, T_{BD}
- Direct transmission of the beam, T_{DB}
- Transmission of diffuse incident radiance, T_{D}

Molecular absorption and Rayleigh scattering in the aerosol layer

The effects of molecular absorption and Rayleigh scattering are included in DISORT by an adjustment of the aerosol layer's optical depth, the particle's single scattering albedo and phase function in the following manner:

$$\tau_{\mathrm{l}} = \tau_{\mathrm{a}} + \tau_{\mathrm{R}} + \tau_{\mathrm{g}} \tag{2.6}$$

$$\omega = \frac{\tau_{\mathrm{R}} + \omega_0 \tau_{\mathrm{a}}}{\tau_{\mathrm{g}} + \tau_{\mathrm{R}} + \tau_{\mathrm{a}}} \tag{2.7}$$

$$P(\theta) = \frac{\tau_{\mathrm{a}}\omega_{\mathrm{a}} P_{\mathrm{a}}(\theta) + \tau_{\mathrm{R}} P_{\mathrm{R}}(\theta)}{\tau_{\mathrm{a}}\omega_{a} + \tau_{R}} \tag{2.8}$$

Here τ_{a}, τ_{R} and τ_{g} are the contributions to the total layer optical depth τ_{l} from aerosol extinction, Rayleigh scattering and gaseous absorption within each layer respectively. The aerosol single scattering albedo is denoted ω_0. For each layer bounded by lower and upper pressure levels p_{l} and p_{u} respectively and ground level pressure p_0, τ_{R} is calculated as:

$$\tau_{\mathrm{R}} = \frac{\tau_{\mathrm{RT}}[p_{\mathrm{u}} - p_{\mathrm{l}}]}{p_0} \tag{2.9}$$

τ_{RT}, the wavelength-dependent Rayleigh scattering optical depth for a column of atmosphere extending from the ground surface to the top of the atmosphere (TOA), is obtained from Justus and Paris (1985):

$$\tau_{\mathrm{RT}}(\lambda) = \frac{p_0}{p_{\mathrm{s}}} \times \frac{1}{117.03\,\lambda^4 - 1.316\,\lambda^2} \tag{2.10}$$

This takes the value of p_{s} (standard pressure) as 1013.25 mbar, and requires the ground pressure p_0 in mbar and wavelength λ in μm.

2.2.3 Computing top of atmosphere radiances

The ORAC retrieval allows for two different types of surface treatment: in the first case, a simple Lambertian surface forward model, and in the second a surface described by a bidirectional re-

flectance distribution function (BRDF), which allows the surface reflectance to vary as a function of geometry. The Lambertian treatment was the original forward model of ORAC; the BRDF forward model of Thomas (2007) allows for a more thorough treatment of real-world surfaces. In practice the BRDF forward model is used unless *a priori* surface albedo information (from sources described in succeeding chapters) is missing, in which case default values and a Lambertian surface are assumed.

For completeness and to aid an appreciation of the differences between the two approaches, both forward models are presented here, after Thomas (2007).

The Lambertian forward model

The Lambertian forward model uses the LUTs to predict a top of atmosphere (TOA) reflectance using a radiative transfer model, shown diagrammatically in Figure 2.1. The solar beam is incident on the aerosol layer (described by many layers in DISORT).

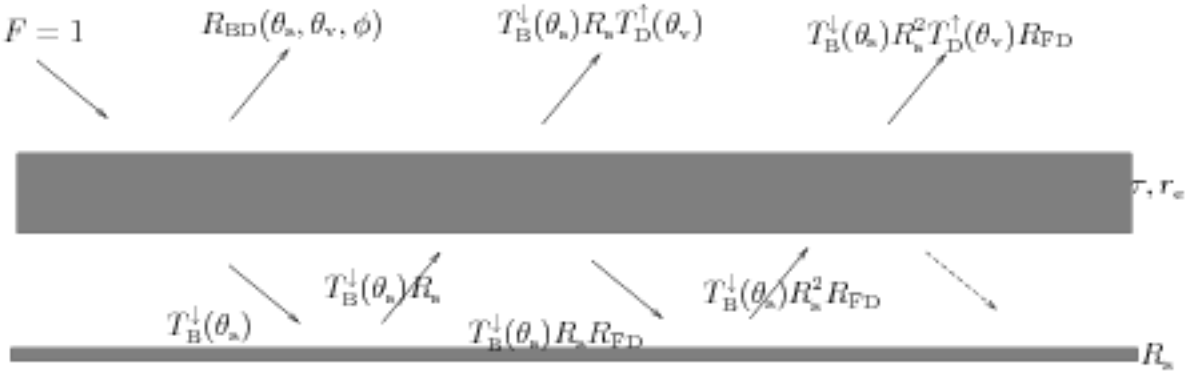

Figure 2.1: Aerosol layer - surface interactions in the Lambertian forward model.

The first contribution to the TOA reflectance received at the satellite is the direct bidirectional reflectance (R_{BD}) of the aerosol layer. Transmission through the layer is partly by direct transmission of the beam ($T_{\mathrm{DB}}^{\downarrow}$) and partly by diffuse transmission of scattered radiance ($T_{\mathrm{BD}}^{\downarrow}$). These can be combined to give the total downward transmission through the aerosol layer $T_{\mathrm{B}}^{\downarrow} = T_{\mathrm{BD}}^{\downarrow} + T_{\mathrm{DB}}^{\downarrow}$, as a Lambertian surface is assumed (meaning any directionality in the radiance is lost on reflection).

The reflected radiance is assumed to be diffuse and is partially transmitted ($T_{\mathrm{D}}^{\uparrow}$) by the aerosol layer into the view direction, thereby giving the second term. The term $T_{\mathrm{D}}^{\uparrow}$ is the total transmission of the atmosphere (the intensity at the top of the atmosphere where the incoming flux at the bottom is equal to one); hence, the diffuse only component of the transmitted radiation is $T_{\mathrm{D}}^{\uparrow} - T_{\mathrm{DB}}^{\uparrow}$, the total transmission minus the directly-transmitted component.

The aerosol layer also reflects downwards, giving rise to a series of multiple reflections and transmissions giving with rapidly decreasing contributions to the TOA reflectance. This multiple-scattering process is represented as follows:

$$\begin{aligned} R =& R_{\mathrm{BD}}(\theta_{\mathrm{s}}, \theta_{\mathrm{v}}, \phi) \\ &+ T^{\downarrow}_{\mathrm{DB}}(\theta_{\mathrm{s}}) R_{\mathrm{s}} T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) + T^{\downarrow}_{\mathrm{BD}}(\theta_{\mathrm{s}}) R_{\mathrm{s}} T D^{\uparrow}(\theta_{\mathrm{v}}) \\ &+ T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) R_{\mathrm{s}}^2 T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) R_{\mathrm{FD}} \\ &+ T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) R_{\mathrm{s}}^3 T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) R_{\mathrm{FD}}^2 \\ &+ \cdots \end{aligned} \tag{2.11}$$

Combining T_{DB} and T_{BD} gives the following:

$$\begin{aligned} R =& R_{\mathrm{BD}}(\theta_{\mathrm{s}}, \theta_{\mathrm{v}}, \phi) \\ &+ T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) R_{\mathrm{s}} T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) \\ &+ T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) R_{\mathrm{s}}^2 T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) R_{\mathrm{FD}} \\ &+ T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) R_{\mathrm{s}}^3 T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) R_{\mathrm{FD}}^2 \\ &+ \cdots \end{aligned} \tag{2.12}$$

Here θ_{s} is the solar zenith angle, θ_{v} is the satellite (instrumental view) zenith angle and ϕ is the relative (solar to satellite) azimuth angle. This expression can be simplified to give:

$$R = R_{\mathrm{BD}}(\theta_{\mathrm{s}}, \theta_{\mathrm{v}}, \phi) + T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) R_{\mathrm{s}} (1 + R_{\mathrm{s}} R_{\mathrm{FD}} + R_{\mathrm{s}}^2 R_{\mathrm{FD}}^2 + \cdots) \tag{2.13}$$

Further simplification as a geometric series limit gives the equation used to calculate TOA reflectance seen by the satellite:

$$R = R_{\mathrm{BD}}(\theta_{\mathrm{s}}, \theta_{\mathrm{v}}, \phi) + \frac{T^{\downarrow}_{\mathrm{B}}(\theta_{\mathrm{s}}) T^{\uparrow}_{\mathrm{D}}(\theta_{\mathrm{v}}) R_{\mathrm{s}}}{(1 - R_{\mathrm{s}} R_{\mathrm{FD}})} \tag{2.14}$$

The surface BRDF forward model

The incorporation of a surface BRDF forward model into the retrieval necessitates the definition of several additional terms not present in Equation 2.14. Firstly, the instead of just being described by the Lambertian albedo it is necessary to define three types of surface reflectance:

1. The surface bidirectional reflectance, R_{SBD}. This describes the reflection of the direct solar beam into the viewing angle, and is a function of both solar and viewing angles.

2. The black-sky albedo, R_{SLB}. This describes the diffuse reflection of the direct beam over the whole hemisphere, and is a function of the solar angle.

3. The white-sky albedo, R_{SLW}. This describes the reflection of diffuse downwelling radiation, assumed isotropic. Hence it is independent of the geometry, and is equivalent to the surface albedo used in the Lambertian forward model described above.

As as further consequence of the definition of these three types of surface reflectance it is no longer possible to combine $T^{\downarrow}_{\mathrm{DB}}$ and $T^{\downarrow}_{\mathrm{BD}}$ into $T^{\downarrow}_{\mathrm{B}}$, as they are subject to different surface reflectances. Hence Equation 2.12 takes on a more complicated form:

$$\begin{aligned} R = R_{\mathrm{BD}} &+ T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SBD}} T^{\uparrow}_{\mathrm{DB}} + T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}} (T^{\uparrow}_{\mathrm{D}} - T^{\uparrow}_{\mathrm{DB}}) \\ &+ T^{\downarrow}_{\mathrm{BD}} R_{\mathrm{SLW}} T^{\uparrow}_{\mathrm{D}} \\ &+ T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}} R_{\mathrm{FD}} R_{\mathrm{SLW}} T^{\uparrow}_{\mathrm{D}} + T^{\downarrow}_{\mathrm{BD}} R_{\mathrm{SLW}} R_{\mathrm{FD}} R_{\mathrm{SLW}} T^{\uparrow}_{\mathrm{D}} \\ &+ T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}} R_{\mathrm{FD}} R_{\mathrm{SLW}} R_{\mathrm{FD}} R_{\mathrm{SLW}} T^{\uparrow}_{\mathrm{D}} \\ &+ T^{\downarrow}_{\mathrm{BD}} R_{\mathrm{SLW}} R_{\mathrm{FD}} R_{\mathrm{SLW}} R_{\mathrm{FD}} R_{\mathrm{SLW}} T^{\uparrow}_{\mathrm{D}} \\ &+ \cdots \end{aligned} \tag{2.15}$$

The $T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{S}} T^{\uparrow}_{\mathrm{D}}$ term present in the Lambertian forward model is replaced by two new terms:

- $T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SBD}} T^{\uparrow}_{\mathrm{DB}}$ is the direct reflection of the solar beam into the viewing angle at the surface. Because of this, it uses the transmission of the direct beam for both the downward and upward paths through the atmosphere.

- $T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}} (T^{\uparrow}_{\mathrm{D}} - T^{\uparrow}_{\mathrm{DB}})$ is the diffuse reflection of the direct solar beam.

Analagously to the Lambertian forward model, the term $T^{\downarrow}_{\mathrm{BD}} R_{\mathrm{SLW}} T^{\uparrow}_{\mathrm{D}}$ represents the reflected diffusely-transmitted solar radiation. Succeeding terms represent multiple surface-atmosphere scattering.

The following two approximations are made for this BRDF forward model to hold:

1. It is assumed that all directionality in the diffusely-scattered radiation from the direct beam (represented by $T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}}$) is lost as it transmitted through the atmosphere.

2. When taken as a pair the surface and atmosphere act as Lambertian reflectors, so that any directionality left in the reflected beam ($T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SBD}}$) is lost in that proportion which is reflected back towards the ground. Hence after the first pair of surface-atmosphere reflections, the R_{SLW} and $T^{\uparrow}_{\mathrm{D}}$ may be used for subsequent reflections.

Because of these approximations, the BRDF forward model should be thought of as a 'simplified' BRDF rather than full BRDF forward model. However these simplifications greatly reduce the complexity of the resulting equations, and are more realistic than the Lambertian surface approximation. The equivalent of Figure 2.1 for a visual representation of scattering in the BRDF forward model is presented in Figure 2.2.

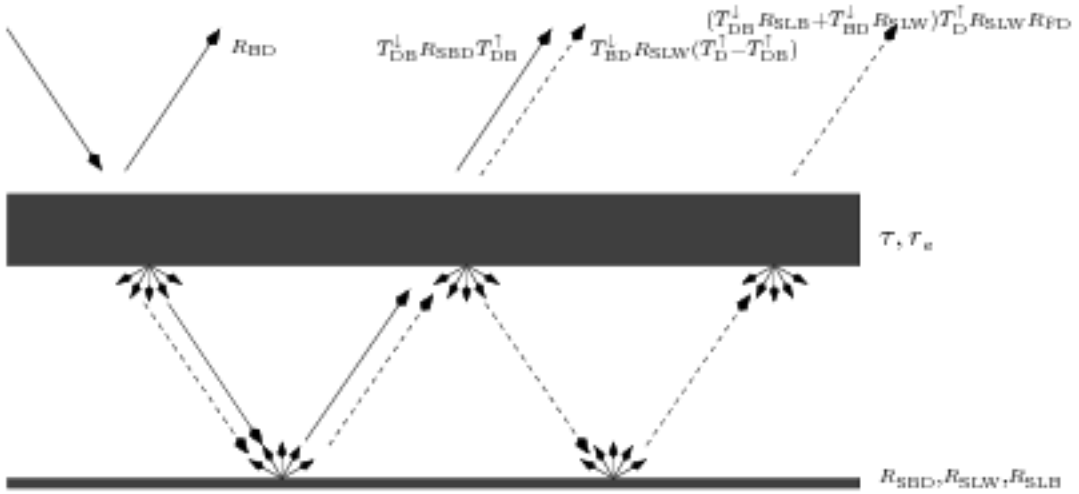

Figure 2.2: Aerosol layer - surface interactions using the simplified BRDF forward model, after Thomas (2007).

Analagously to the Lambertian forward model, the scattering can be simplified to the following:

$$\begin{aligned} R =& R_{\mathrm{BD}} + T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SBD}} T^{\uparrow}_{\mathrm{DB}} - T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}} T^{\uparrow}_{\mathrm{DB}} \\ &+ \left(T^{\downarrow}_{\mathrm{DB}} R_{\mathrm{SLB}} + T^{\downarrow}_{\mathrm{BD}} R_{\mathrm{SLW}} \right) T^{\uparrow}_{\mathrm{D}} \left(1 + R_{\mathrm{SLW}} R_{\mathrm{FD}} + R^2_{\mathrm{SLW}} R^2_{\mathrm{FD}} + ... \right) \end{aligned}$$

Again, by applying the same series limit as before the TOA reflectance is given by:

$$R = R_{\mathrm{BD}} + T^{\downarrow}_{\mathrm{DB}}(R_{\mathrm{SBD}} - R_{\mathrm{SLB}})T^{\uparrow}_{\mathrm{DB}} + \frac{\left(T^{\downarrow}_{\mathrm{DB}}R_{\mathrm{SLB}} + T^{\downarrow}_{\mathrm{BD}}R_{\mathrm{SLW}}\right)T^{\uparrow}_{\mathrm{D}}}{1 - R_{\mathrm{SLW}}R_{\mathrm{FD}}}, \tag{2.16}$$

Equation 2.16 is used to calculate the TOA reflectance with the surface BRDF forward model. In the case of the three surface reflectances R_{SBD}, R_{SLB} and R_{SLW} being equal (i.e. the surface being Lambertian), Equation 2.16 simplifies to Equation 2.14.

An example of predicted TOA reflectance

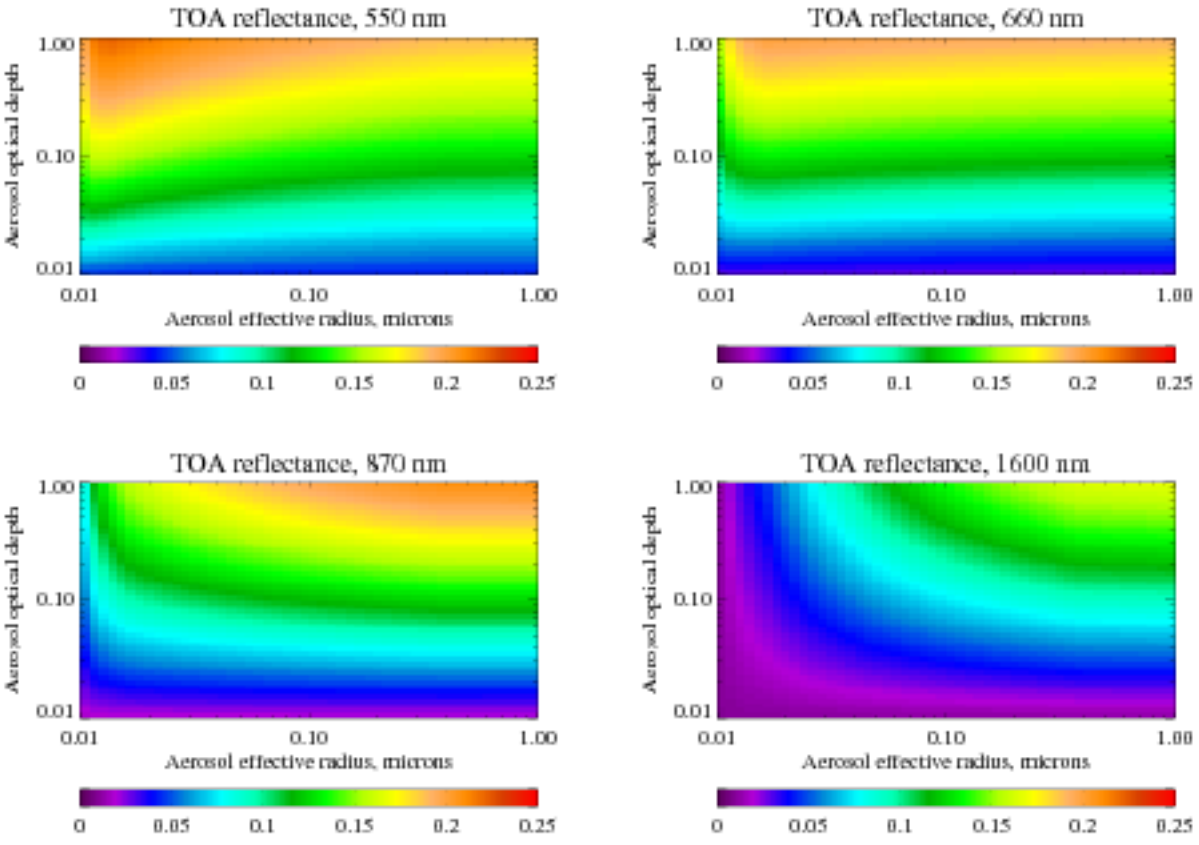

Figure 2.3: TOA reflectance arriving at satellite channels as predicted by the forward model for θ_{s} =60°, θ_{v} =55°, a Lambertian surface albedo of 0.02 at all wavelengths and the maritime aerosol class.

An example of TOA reflectance predicted by the (in this case Lambertian) forward model as a function of τ and r_{e} is shown in Figure 2.3. The range of optical depth and effective radii scanned was

chosen to include commonly-encountered values for this aerosol type (maritime). The measurements provide the most information when contours on different channels are orthogonal. Similar patterns are encountered for other geometries and aerosol models. The change in TOA reflectance with state variables is larger than the uncertainties in the forward model (discussed later), showing that the retrieval is sensitive to the state variables and so the measurements contain useful information.

Forward model radiance gradient

The gradient of the forward model ($\partial \mathbf{y}/\partial \mathbf{x}$) is required for the inversion of the non-linear reflectance model by the Levenberg-Marquardt algorithm, as discussed later.

The BRDF forward model gradient with respect to optical depth (or effective radius, with ∂r_e replacing $\partial\tau$) is as follows:

$$\begin{aligned}\frac{\partial R}{\partial \tau} =& \frac{\partial R_{\mathrm{BD}}}{\partial \tau} + (R_{\mathrm{SBD}} - R_{\mathrm{SLB}}) \left(T_{\mathrm{DB}}^{\downarrow} \frac{\partial T_{\mathrm{DB}}^{\uparrow}}{\partial \tau} + T_{\mathrm{DB}}^{\uparrow} \frac{\partial T_{\mathrm{DB}}^{\downarrow}}{\partial \tau} \right) \\ &+ \frac{\left(T_{\mathrm{DB}}^{\downarrow} R_{\mathrm{SLB}} + T_{\mathrm{BD}}^{\downarrow} R_{\mathrm{SLW}} \right) T_{\mathrm{D}}^{\uparrow} R_{\mathrm{SLW}} \frac{\partial R_{\mathrm{FD}}}{\partial \tau}}{\left(1 - R_{\mathrm{SLW}} R_{\mathrm{FD}}\right)^2} \\ &+ \frac{\left(T_{\mathrm{DB}}^{\downarrow} R_{\mathrm{SLB}} + T_{\mathrm{BD}}^{\downarrow} R_{\mathrm{SLW}} \right) \frac{\partial T_{\mathrm{D}}^{\uparrow}}{\partial \tau} + \left(R_{\mathrm{SLW}} \frac{\partial T_{\mathrm{BD}}^{\downarrow}}{\partial \tau} + R_{\mathrm{SLB}} \frac{\partial T_{\mathrm{DB}}^{\downarrow}}{\partial \tau} \right) T_{\mathrm{D}}^{\uparrow}}{\left(1 - R_{\mathrm{SLW}} R_{\mathrm{FD}}\right)}\end{aligned} \tag{2.17}$$

The derivative with respect to surface reflectance requires expression of the derivatives of R_{SBD} and R_{SLB} in terms of a derivative of R_{SLW}. Since R_{SBD} and R_{SLB} both depend linearly on R_{SLW} for a given viewing geometry, the following expressions may be formed:

$$\frac{\partial R}{\partial R_{\mathrm{SBD}}} = \frac{\partial R}{\partial R_{\mathrm{SLW}}} \frac{\partial R_{\mathrm{SLW}}}{\partial R_{\mathrm{SBD}}} = \frac{1}{\alpha} \frac{\partial R}{\partial R_{\mathrm{SLW}}} \tag{2.18}$$

$$\frac{\partial R}{\partial R_{\mathrm{SLB}}} = \frac{\partial R}{\partial R_{\mathrm{SLW}}} \frac{\partial R_{\mathrm{SLW}}}{\partial R_{\mathrm{SLB}}} = \frac{1}{\beta} \frac{\partial R}{\partial R_{\mathrm{SLW}}} \tag{2.19}$$

In the above the parameters α and β are taken as the ratio of R_{SBD} and R_{SLB} to R_{SLW} respectively. The derivative may then be expressed as the following:

$$\begin{aligned}\frac{\partial R}{\partial R_{\mathrm{SLW}}} =& T_{\mathrm{DB}}^{\downarrow} (\alpha - \beta) T_{\mathrm{DB}}^{\uparrow} \\ &+ \frac{T_{\mathrm{DB}}^{\downarrow} \beta T_{\mathrm{D}}^{\uparrow} + T_{\mathrm{BD}}^{\downarrow} T_{\mathrm{D}}^{\uparrow}}{1 - R_{\mathrm{SLW}} R_{\mathrm{FD}}} + \frac{\left(T_{\mathrm{DB}}^{\downarrow} R_{\mathrm{SLB}} + T_{\mathrm{BD}}^{\downarrow} R_{\mathrm{SLW}} \right) T_{\mathrm{D}}^{\uparrow}}{\left(1 - R_{\mathrm{SLW}} R_{\mathrm{FD}}\right)^2}\end{aligned} \tag{2.20}$$

Equivalent expressions are easily obtained from these equations for the special case of a Lambertian surface by setting R_{SBD}, R_{SLB} and R_{SLW} to be equal.

In the retrieval, these gradients are used to form the weighting function matrix $\mathbf{K}_x$. This has dimensions n_y by n_x, where n_y and n_x are the number of measurements and state variables respectively.

2.3 *A priori* information

The retrieval algorithm makes use of *a priori* information on the state variables $\mathbf{x}$. This acts as a constraint on the retrieval and so make the process less likely to retrieve unphysical values for $\mathbf{x}$. The *a priori* values are also used as a first guess at the solution. Although it is not necessary for the *a priori* and first guess to be the same, it is convenient for it to be. The precise treatment of the *a priori* by the retrieval algorithm is discussed in Section 2.6.

2.3.1 Aerosol properties

Each different aerosol class in the retrieval has its own associated *a priori* optical depth and effective radius, with associated errors, dependent on its composition. These are listed in Table 2.3. The *a priori* uncertainties on τ and r_e are assumed to be independent.

2.3.2 Surface albedo

Creating an *a priori* model of the surface for a pixel depends on whether it is a land or sea pixel. In both cases R_{SBD}, R_{SLB} and R_{SLW} are defined. The bidirectional reflectance is integrated over all satellite viewing zenith θ_v and relative azimuth ϕ_r angles to obtain the black-sky albedo for a given

Aerosol Class	ORAC label	Log optical depth $\log(\tau)$	Error on $\log(\tau)$	Log effective radius $\log(r_e)$, μm	Error on $\log(r_e)$
Continental clean	A00	-1.2	1.0	-0.64	0.5
Desert	A03	-0.54	1.5	0.14	0.5
Maritime clean	A05	-1.0	1.0	-0.08	0.5
Urban	A09	-0.19	2.0	-0.82	0.5
Biomass burning	A11	-0.19	2.0	-0.87	0.3

Table 2.3: *A priori* information for the aerosol classes used in the retrieval scheme. Quantities are logarithmic as this is the form in which the retrieval code deals with them.

solar zenith angle θ_s:

$$R_{\mathrm{SLB}}(\theta_s) = \frac{\int_0^{2\pi}\int_0^{\pi/2} R_{\mathrm{SBD}}(\theta_s,\theta_v,\phi_r)\cos\theta_v \sin\theta_v d\theta_v d\phi_r}{\int_0^{2\pi}\int_0^{\pi/2}\cos\theta_v \sin\theta_v d\theta_v d\phi_r} = \frac{1}{\pi}\int_0^{2\pi}\int_0^{\pi/2} R_{\mathrm{SBD}}(\theta_s,\theta_v,\phi_r)\cos\theta_v \sin\theta_v d\theta_v d\phi_r \quad (2.21)$$

This is then integrated over all solar angles to obtain the white-sky albedo:

$$R_{\mathrm{SLW}} = \frac{\int_0^{\pi/2} R_{\mathrm{SLB}}(\theta_s)\cos\theta_s \sin\theta_s d\theta_s}{\int_0^{\pi/2}\cos\theta_s \sin\theta_s d\theta_s} = 2\int_0^{\pi/2} R_{\mathrm{SLB}}(\theta_s)\cos\theta_s \sin\theta_s d\theta_s \quad (2.22)$$

R_{SBD} is calculated as described in Chapter 3 for sea pixels and as described in Chapter 4 for land pixels. Over the sea, Gaussian quadrature integration with four quadrature points has been found to yield results correct to three significant figures in calculation of R_{SLB} and R_{SLW}; this provides an *a priori* value for R_{SLW} for each wavelength. For land pixels an analytic expression described in Chapter 4 is used instead. This process also provides the ratios between R_{SBD}, R_{SLB} and R_{SLW}, which are fixed in the retrieval.

Sun-glint

The specular reflection of sunlight from the ocean surface (Sun-glint) poses a problem for nadir measurements over the ocean. Sun-glint is less of an issue for the forward viewing geometry, for two reasons:

- The forward geometry has a longer atmospheric path length, so the relative aerosol contribution to the signal is magnified.
- Because the instrument orbits at an approximately constant local solar time, the Sun's position in relation to the instrument is always less favourable for the nadir view than the forward view. In other words, the forward view is less likely to be affected by sun-glint because it is not pointing at the Sun-glint region.

Use of a BRDF forward model means that the differences in surface reflectance between forward and nadir viewing geometries may be accurately accounted for, enabling dual-view retrievals in the

nadir sun-glint region. This is one advantage of multiple-view sensors such as AATSR.

2.4 Discussion of forward model errors

It is not feasible to model accurately the full range of variation in aerosol composition, size, shape and vertical distributions, atmospheric composition and surface properties which affect the TOA reflectance. Even if all factors could be modelled the measurements do not provide sufficient information to retrieve them all: mathematically, the problem is ill-posed. The resulting approximations and assumptions in the forward model will lead to errors in the retrieved state parameters and it is important that these errors be quantified.

2.4.1 Surface reflectance and albedo

Over the land, except in circumstances of extremely high aerosol loading (or cloud), the instrumental signal will be dominated by surface reflection, with the aerosol signature essentially adding a perturbation. This is true for both views, although more so for the nadir view (which has a shorter path length). Although the surface albedo is a retrieved quantity, the ratios between R_{SBD}, R_{SLB} and R_{SLW} are fixed and it is this which leads to a forward model error. This is discussed in more detail in Section 2.6.2. As a result, accurate knowledge of the angular shape of the surface reflectance is of key importance in determining the overall accuracy of the retrieval. Several points can be made about surface modelling:

- Desert, snow and ice surfaces have a very high albedo, making the surface contribution to radiance even greater. This increases the impact of any problems in the assumed ratios on the forward model.

- Surface albedos at the four wavelengths are not independent, so correlations are applied between the relevant elements of the *a priori* covariance matrix as described by Chapters 3 and 4. It is important that these correlations are realistic and not too strict, however, as an overly-constrained surface would lead to errors in calculation of the spectral dependence of the aerosol scattering.

2.4.2 Aerosol size distribution

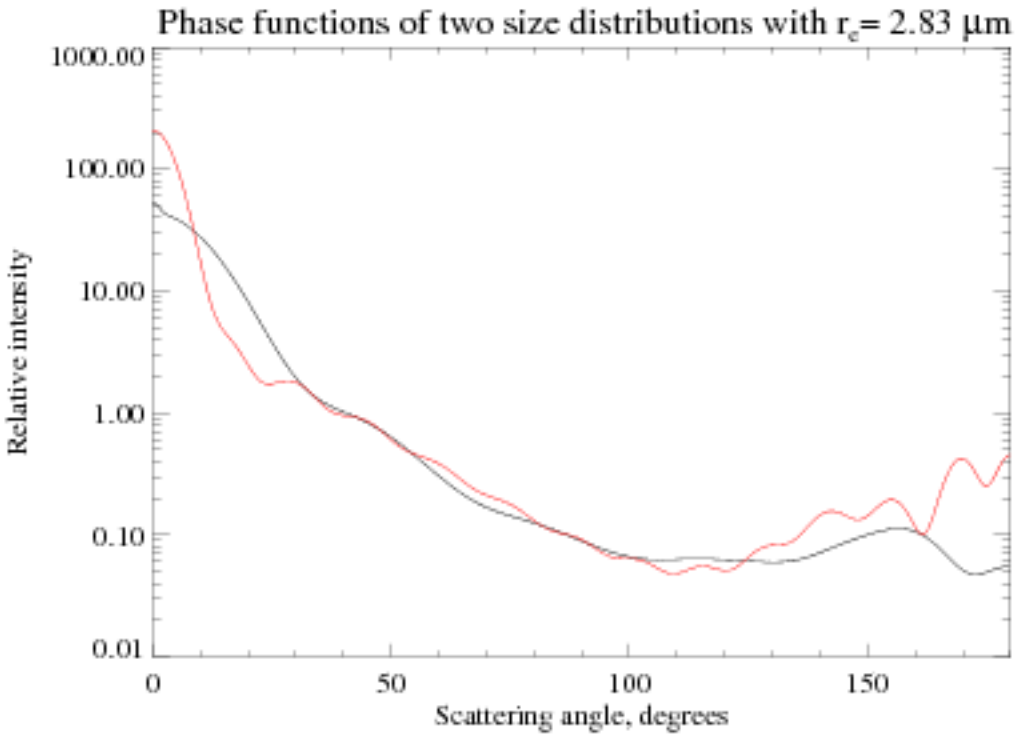

Figure 2.4: Comparison of the phase functions at 550 nm of two size distributions (shown in Figure 2.5 with the same colour key), both with an effective radius of 2.83 μm.

Figure 2.4 gives an example of the difference between the scattering phase functions that correspond to two grossly different size distributions (shown in Figure 2.5), which have the same effective radius. Although the difference can approach an order of magnitude for angles smaller than 30° or larger than 160°, AATSR observes at scattering angles between approximately 50° and 150°, where the phase functions are similar. As a result this disparity will have little impact for the bulk of scenes which are optically thin (where single scattering, as opposed to multiple scattering, dominates). It is likely, therefore, that this method of setting the effective radius does not introduce significant errors in the modelled radiances. Furthermore as the distributions used in the retrieval scheme are drawn from the literature (Hess et al., 1998, Dubovik et al., 2002) it is expected that they would be representative of real aerosol, hence extreme deviations of the assumed distribution from the true distribution are unlikely.

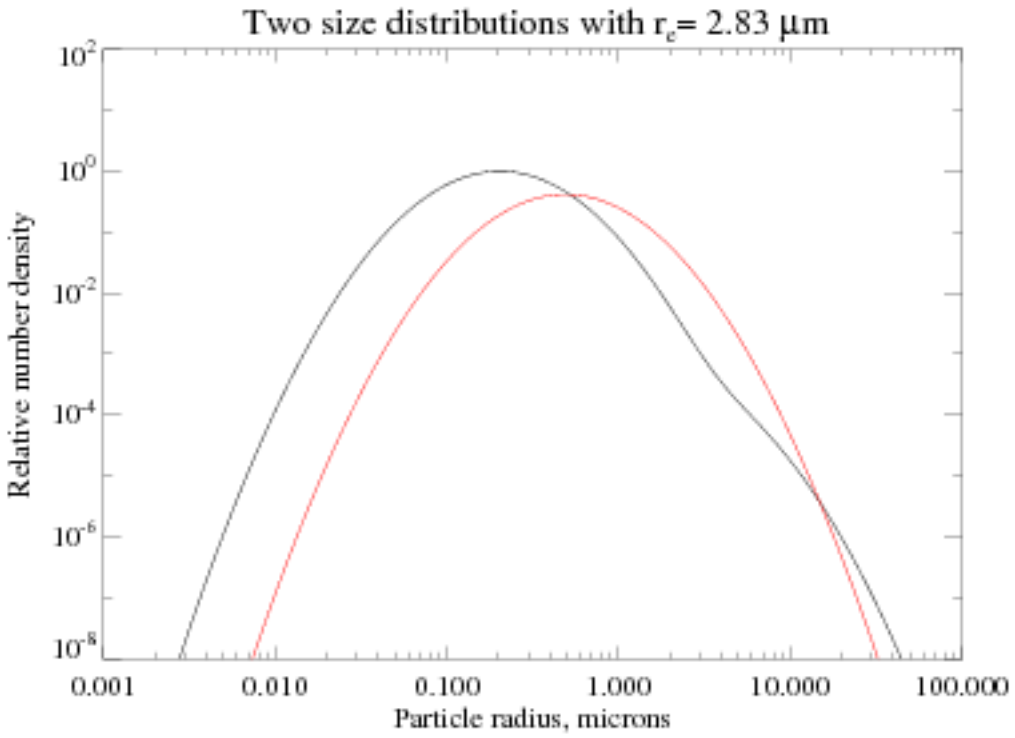

Figure 2.5: Aerosol size distributions used for comparison of phase functions shown in Figure 2.4. The first distribution (black) consists of the accumulation mode and coarse mode sea salt components from the OPAC database (Hess et al., 1998) at 50 % relative humidity, with a mixing ratio of 300:1, while the second distribution (red) is the accumulation mode alone, with modal radius increased by a factor of 2.4 to give it the same effective radius as the combined distribution (2.83 μm).

2.4.3 Sensitivity to vertical profiles of aerosol parameters

A simple 5-layer forward model from the ground to 2 km has been used to test whether the TOA reflectance varies as a function of the aerosol number distribution, ω_0 and g. This analysis was performed using a previous version of the lookup table generation that used the asymmetry parameter g to describe the angular distribution of the aerosol scattering, rather than a Legendre polynomial expansion of the phase function (Thomas et al., 2006). The TOA reflectance was investigated for a model with ω_0 =0.8 and g =0.61 for all layers. Surface albedo was set to zero. The following four vertical profiles were tested for each of aerosol number distribution, ω_0 and g:

- Exponentially increasing with height
- Exponentially decreasing with height
- Linearly increasing with height
- Linearly decreasing with height

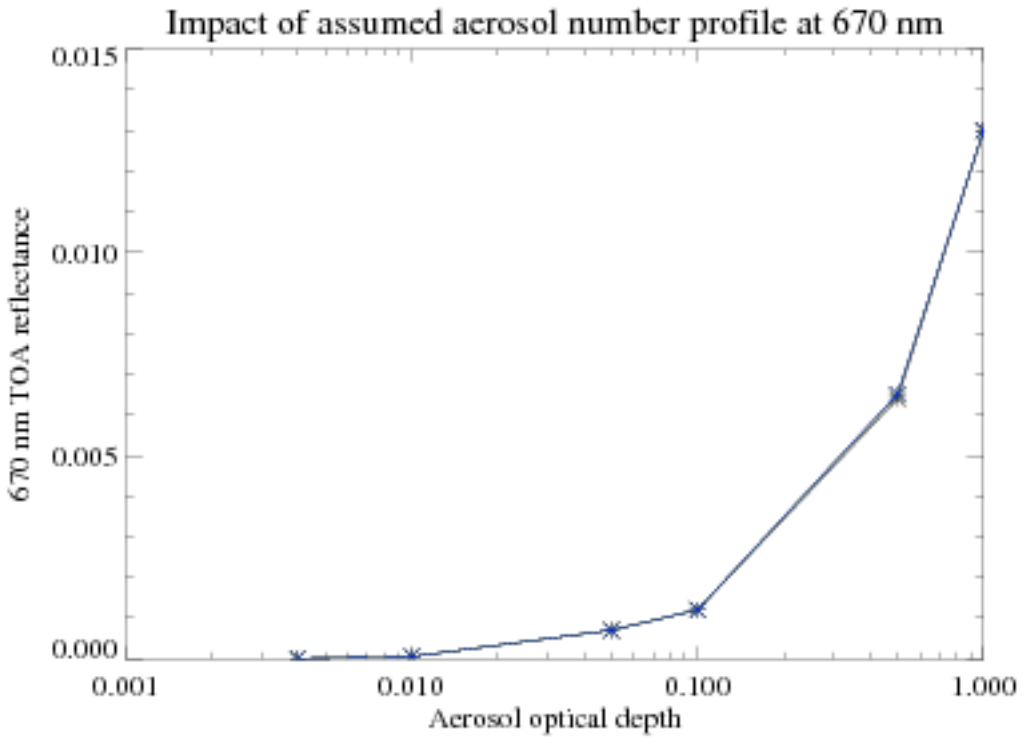

Figure 2.6: TOA reflectance values for each of the four vertical aerosol N distributions at 670 nm. Black indicates an exponential increase with height, red an exponential decrease, green a linear increase and blue a linear decrease. Adapted from Thomas et al. (2006).

The TOA reflectances for the vertical profiles are shown in Figure 2.6. This figure shows that the radiance values for all profiles are identical—if ω_0 and g were kept constant with height, the forward model would not be sensitive to the aerosol profile. Varying ω_0 instead produces the set of graphs in Figure 2.7. It can be seen that there is very little difference between the four plots when $\tau < 0.1$; however, for $\tau > 0.1$ the TOA reflectance is dependent on whether ω_0 increases or decreases with height. This can be explained by more of the incoming radiation being absorbed high in the model atmosphere when ω_0 decreases with height. For the set of graphs in Figure 2.8 the only parameter which is permitted to vary with height is the aerosol asymmetry parameter g. Again it can be seen that the TOA reflectance varies with its vertical distribution.

The conclusion to be drawn from these tests is that the forward model is sensitive to the vertical distribution of aerosol—although it is not directly sensitive to the vertical number distribution N it *is* sensitive to the vertical profiles of ω_0 and g. Equations 2.6, 2.7 and 2.8 show that values for τ, ω_0 and g within the layers of the forward model are dependent on τ_a and hence on N.

Due to an absence of information about vertical aerosol distribution at a given place and time this dependence cannot be accounted for by some parameterization, and the measurements do not

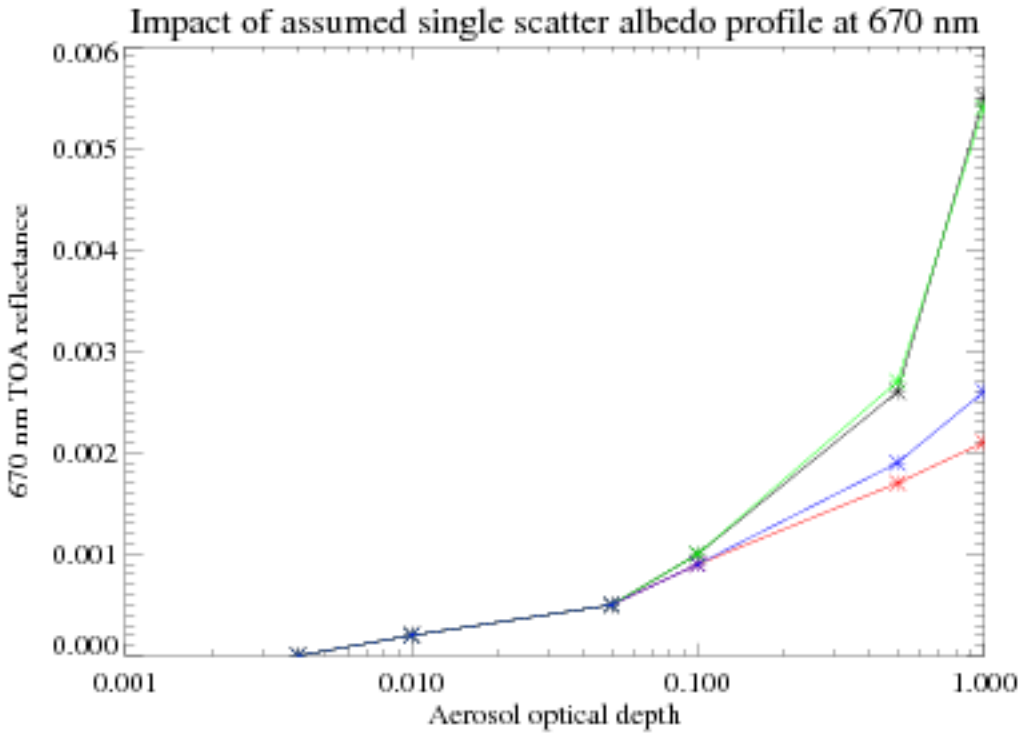

Figure 2.7: TOA reflectance values for each of the four vertical aerosol ω_0 distributions at 670 nm. Black indicates an exponential increase with height, red an exponential decrease, green a linear increase and blue a linear decrease. Adapted from Thomas et al. (2006).

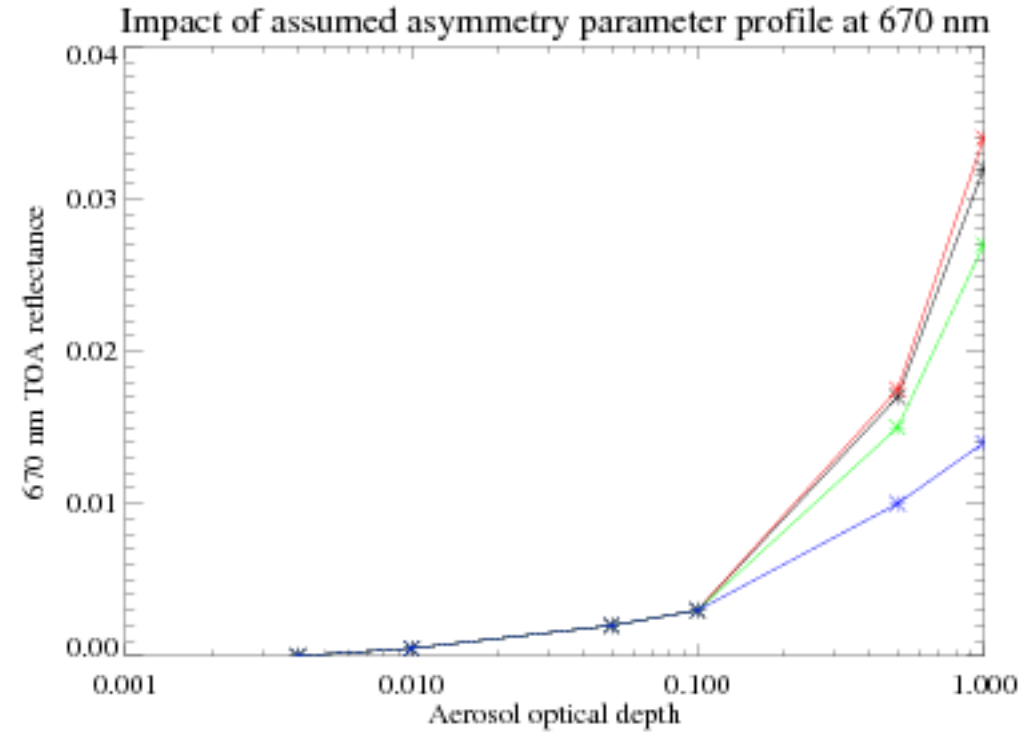

Figure 2.8: TOA reflectance values for each of the four vertical aerosol g distributions at 670 nm. Black indicates an exponential increase with height, red an exponential decrease, green a linear increase and blue a linear decrease. Adapted from Thomas et al. (2006).

provide enough information to allow the retrieval of the vertical distribution, as the differences in TOA reflectance between likely profiles are small except in the case of very high aerosol loading. As a result this dependence contributes to forward model error in the retrieval. It should be noted that in reality, other aerosol distibutions (such as a peaked distribution or multiple discrete aerosol layers) are possible. The examples chosen serve to illustrate a general sensitivity of the measurements to the aerosol vertical distribution for high aerosol loadings. If the assumed distributions were inaccurate, this would manifest as a high value of the retrieval cost (see Section 2.6.1).

2.5 Data preprocessing

The ORAC processor expects all input data to be presented in unformatted binary files. The primary role of the preprocessing is to read in the required data in its own format (generally more portable standards), manipulate it and output it in unformatted form. The use of preprocessing code in this way enables the FORTRAN retrieval code to be treated almost as a 'black box', while the preprocessing code can be written in a higher level language (here Interactive Data Language, IDL, by ITT, 2008) which is much better suited to the quick development of new routines for reading different data. This provides a stable, highly optimised retrieval scheme, while maintaining the flexibility needed to allow the easy application of the scheme to new input and ancillary data.

IDL was selected as the language for the majority of the preprocessing code as there are readily available I/O routines for most scientific data formats and because of its inbuilt routines for data manipulation. Analysis of output data is also typically performed using IDL routines. The preprocessing additionally performs some important functions required by the retrieval, as described below.

Land flagging

The ORAC scheme uses the operational ATSR land flag database, described in Zavody et al. (1994). It consists of vector descriptions of land, with outlines of rivers and political boundaries at approximately 1 km resolution.

Cloud flagging

Retrievals are only performed on cloud-free pixels. Over ocean the ESA cloud flag product (ESA, 2007) produced in the ESA SST retrieval is used. Over land, the algorithm of Birks (2004), which was developed to address problems with the original operational cloud flag, is used instead.

Cloud flagging is performed at the full instrumental resolution regardless of the spatial resolution required by the retrieval, which minimises errors in cloud flagging due to sub-pixel clouds. Cloud flags in forward and nadir views may differ (due to both cloud movement in the time between views and differences in apparent position due to parallax) and a pixel is only flagged as cloud-free if both forward and nadir views register as cloud-free.

Superpixelling

As well as running at full instrumental resolution of 1 km $\times$ 1 km the retrieval allows spatial averaging of data ('superpixelling'). If run at reduced resolution, radiances not flagged as cloud are averaged to give a mean and variance for use during the retrieval. The fraction of pixels within each superpixel containing cloud is also calculated, to give a cloud fraction product. An advantage of running at reduced resolution is that it cuts computational time by a factor of about n^2 for an n km $\times$ n km superpixelling as compared to running at full resolution. It also decreases the noise on the measurement by a factor of the square root of the number of cloud-free pixels within the superpixel. However, very large superpixelling may impact the accuracy of the retrieval, given factors such as varying geometry and scene inhomogeneity.

Superpixelling is normally performed on a sinusoidal grid with 4008 cells along the equator, resulting in approximately 10 km $\times$ 10 km superpixels. The advantage of a sinusoidal grid over an instrumental one is that the locations of the instrumental grid cells are not fixed with time, while the geographical sinusoidal grid cells are. This makes it simpler to, for example, compare to ground-based targets (as the point on the ground will always lie in the same sinusoidal grid cell).

Driver file generation

The behaviour of the ORAC retrieval code is controlled by a driver, generated in IDL preprocessing. It is a text file which passes various pieces of information to the FORTRAN code such as file paths, *a priori* and first guess settings, convergence criteria, and active instrumental channels.

2.6 The retrieval algorithm

The ORAC retrieval is an optimal estimation (OE) scheme which uses the Levenburg-Marquardt iteration algorithm, detailed by Levenberg (1944) and Marquardt (1963). The basic principle of optimal estimation is to maximise the probability of the retrieved state based on the value of the measurements and *a priori* information. Formally, it is required to maximise the conditional probability $P = P(\mathbf{x}|\mathbf{y}, \mathbf{x}_a, \mathbf{b})$ with respect to the values of the state vector $\mathbf{x}$, where $\mathbf{x}_a$ is the *a priori* value of the state vector and $\mathbf{b}$ are all other parameters not modelled by the forward model. The full state vector $\mathbf{x}$ is defined as follows:

$$\mathbf{x} = [\tau_{550}, r_e, R_{\mathrm{SLW}(550)}, R_{\mathrm{SLW}(660)}, R_{\mathrm{SLW}(870)}, R_{\mathrm{SLW}(1600)}]$$

The terms, as previously mentioned, are the optical depth at 550 nm, effective radius, and surface albedo at the four wavelengths. The previous nadir-view algorithm was only able to retrieve surface albedo at one wavelength, and assumed a fixed surface spectral shape, so $R_{\mathrm{SLW}(550)}$ was the only surface term in $\mathbf{x}$.

2.6.1 Definition of retrieval cost *J* and method of minimisation

The assumption is made that errors in the measurements, *a priori* and model parameters have Gaussian distributions with zero mean and variance given by $\mathbf{S}_y$, $\mathbf{S}_x$ and $\mathbf{S}_b$ respectively. Then, the conditional probability takes on the quadratic form:

$$\begin{aligned} P(x) \propto & \exp\left[-(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m)\mathbf{S}_y^{-1}(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m)^T\right] \\ & \times \exp\left[-(\mathbf{x} - \mathbf{x}_a)\mathbf{S}_x^{-1}(\mathbf{x} - \mathbf{x}_a)^T\right] \\ & \times \exp\left[-(\mathbf{b}_t - \mathbf{b})\mathbf{S}_b^{-1}(\mathbf{b}_t - \mathbf{b})^T\right] \end{aligned} \tag{2.23}$$

The three terms present in the equation represent weighted deviations from measurements, the *a priori* state and the model parameters respectively. Here $\mathbf{y}(\mathbf{x})$ refers to the values of $\mathbf{y}$ predicted by the model forward model from the current value of the state vector; for clarity, the measurement vector is denoted by $\mathbf{y}_m$. $\mathbf{b}_t$ has been written for the (unknown) true values of the model parameters. Maximising probability is equivalent to minimising the negative logarithm, so that the quantity to be

minimised J, known as the *retrieval cost*, is equivalently minimised for the minimum of:

$$\begin{aligned} J(x) = &(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m)\mathbf{S}_y^{-1}(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m)^T \\ &+ (\mathbf{x} - \mathbf{x}_a)\mathbf{S}_x^{-1}(\mathbf{x} - \mathbf{x}_a)^T \\ &+ (\mathbf{b}_t - \mathbf{b})\mathbf{S}_b^{-1}(\mathbf{b}_t - \mathbf{b})^T \end{aligned} \tag{2.24}$$

The minimisation is done with respect to the state variable $\mathbf{x}$, so that the derivative of J is independent of the third term and b therefore cannot be part of the solution. There are many techniques used in the minimisation of J. Essentially any method of finding the minimum is acceptable, with the caveat that in an operational context it must be robust and fast. The particular characteristics of this problem are that:

1. First and second derivatives of J (with respect to $\mathbf{x}$) exist and are continuous. This implies descent algorithms that make use of the local gradient are possible, and these are generally faster than methods that do not.

2. Multiple minima have been found to be unlikely (cost functions have only one distinct region of low values). This implies that excessive domain searches to avoid them are probably not required.

3. J is likely to be approximately quadratic in the region of the solution and far from quadratic elsewhere. This means that quick convergence from a poor starting position is unlikely.

The cost function to be minimised is a reduced form of Equation 2.24 since there is no explicit dependence on b:

$$\begin{aligned} J(x) = &(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m)\mathbf{S}_y^{-1}(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m)^T \\ &+ (\mathbf{x} - \mathbf{x}_a)\mathbf{S}_x^{-1}(\mathbf{x} - \mathbf{x}_a)^T \end{aligned} \tag{2.25}$$

The first and second derivatives of J with respect to $\mathbf{x}$ are given by:

$$J' = \frac{\partial J}{\partial \mathbf{x}} = \mathbf{K}_x^T \mathbf{S}_y^{-1}(\mathbf{y}(\mathbf{x}) - \mathbf{y}_m) + \mathbf{S}_x^{-1}(\mathbf{x} - \mathbf{x}_a)$$
$$J'' = \frac{\partial^2 J}{\partial \mathbf{x}^2} = \mathbf{K}_x^T \mathbf{S}_y^{-1} \mathbf{K}_x + \mathbf{S}_x^{-1} \tag{2.26}$$

The matrix $\mathbf{K}_x$ is the weighting function (forward model gradient), as previously mentioned. The expression for J'' is a commonly-used approximation in that $\mathbf{K}_x$ is assumed to be independent of $\mathbf{x}$, i.e. the radiative transfer is linear in $\mathbf{x}$. This is only strictly true near the solution (in the region where J is quadratic) but since J'' is only employed near the solution (as discussed in the next section) the approximation is acceptable.

When the retrieval solution is found, the retrieved state is denoted by $\hat{\mathbf{x}}$ while $\hat{\mathbf{y}}$ indicates the final estimate of the measurement vector $\mathbf{y}$ as computed from $\hat{\mathbf{x}}$.

The Levenburg-Marquardt algorithm

The minimum is found by starting at a 'first guess' state $\mathbf{x}_0$ (here set equal to the *a priori* value) and then proceeding to make steps $\delta\mathbf{x}_n$, based on some algorithm. Assuming the value of J decreases at each step, the updated $\mathbf{x}$ vector is taking the process towards the cost function minimum. To be consistent with the three points made above, the Levenburg-Marquardt algorithm is used to define the value of $\delta\mathbf{x}_n$. This algorithm uses a weighted combination of steepest descent and Newtonian descent according to the characteristic of the cost function.

When the cost function is near quadratic (generally near the solution) the efficiency of the Newtonian scheme is employed, while when the cost function is far from quadratic (generally when far from the solution) the robustness of the steepest descent algorithm is favoured.

The steepest descent algorithm is intuitively the simplest. The vector $-J'$ defines the 'downward' direction of the local steepest gradient. A move $\delta\mathbf{x} = -J'$ is almost certainly at least approximately in the direction of the minimum (although it may be too far or hardly far enough). The step is therefore usually scaled, $\delta\mathbf{x} = -\alpha J'$ for some variable α. If J is found to be decreasing (so the step is not far enough) α can be increased to move faster; if J increases (meaning the step was too far in this direction) then α is reduced until J decreases. J must eventually decrease with this method otherwise there is an error in the calculation of J'.

Steepest descent can be very slow to converge, especially near the solution where the gradient becomes small. In contrast, Newtonian descent is very fast near the solution because it will find it in one iteration if J is quadratic. Newton's method finds the root of an equation and is therefore applied here in the form to find the root of $J' = 0$. The Newton step is therefore defined as $\delta\mathbf{x} = -J'/J''$. The problem with Newtonian descent is that away from the solution J can be very non-quadratic; the J'' can easily be the 'wrong' sign leading the step away from the solution. This is not a problem scaling can solve.

A combination of both methods as used in this retrieval achieves the best results. Before each step is taken, the resulting cost is checked to see whether it decreases. If so it is taken and the next step is adjusted to be more Newtonian. If instead J increased, the step is not made and an adjustment is made towards more 'steepest descent' character. The resulting Levenburg-Marquadt procedure becomes steepest descent far from the solution, and Newtonian iteration close to it. It can be described formally by Equation 2.27. Here I is the unit matrix (with size equal to n_x, the number of elements in $\mathbf{x}$) and α the control variable described. When α is large (compared to J'') the step tends to steepest descent; when small the step approximates Newtonian.

$$\delta\mathbf{x} = -(J'' + \alpha I)^{-1} J' \tag{2.27}$$

For the initial step, α is set proportional to the average of the diagonals of J'' to obtain a reasonable value:

$$\alpha_0 = MQ_{\text{start}} \times \overline{Trace(J'')} \tag{2.28}$$

After a successful step with decreasing J, α is decreased as follows:

$$\alpha_{n+1} = \frac{\alpha_n}{MQ_{\text{step}}} \tag{2.29}$$

Following an unsuccessful step attempt α is increased:

$$\alpha_{n+1} = \alpha_n \times MQ_{\text{step}} \tag{2.30}$$

In the ORAC retrieval the control parameters MQ_{start} and MQ_{step} have default values of 0.001 and 10 respectively.

Boundaries

The state space in which the solution is to be found is typically bounded by physical constraints and the retrieval needs to ensure that values are not retrieved outside of these boundaries. These bounds are optical depths τ_{550} of between 0.01 and 5, effective radii r_e between 0.01 μm and 10 μm, and surface albedoes of between 0 and 1. If an iterative step would overstep these boundaries, the state vector element is instead set to the corresponding limit.

Convergence criteria

The iteration process proceeds until the decrease in J between consecutive iterations (δJ_n) is so small as to be considered negligible. The default threshold is 0.05. This prevents a parameter which is not well constrained from oscillating or becoming unstable (because the cost function is 'flat' in that direction). Such an instability makes it appear that the solution is not yet found whereas, in fact, the cost is minimal and cannot decrease further. If convergence has not occurred after 25 iterative steps, the retrieval is stopped and failure to converge noted.

2.6.2 Value of the matrix $\mathbf{S}_y$

As mentioned in Section 2.6.1, the covariance matrix $\mathbf{S}_y$ deals with errors on the measured TOA reflectances, and as such is often referred to as the *measurement covariance matrix*. As the retrieval forward model is imperfect, however, additional terms are added so that this matrix more properly represents the precision to which the retrieval algorithm should be able to fit the measured TOA reflectances. Thus $\mathbf{S}_y$ includes the following errors:

- The measurement error. As outlined in Table 1.2 and Section 1.2.4, there is an uncertainty on the measured TOA reflectance. Superpixelling decreases this error by a factor of the square root of the number of cloud-free pixels in a superpixel, down to a minimum digitisation limit. The decrease in error for the forward-view channels is only 2/3 of the value for nadir-view channels, due to approximately 1/3 of forward view pixels being cosmetically-filled data.

- The numerical precision of the forward model.

- Errors in the TOA reflectance arising from an incorrect aerosol model.

- Errors in the TOA reflectance arising from errors in the fixed forward model parameters (surface reflectance ratios $R_{\mathrm{SLW}} : R_{\mathrm{SBD}}$ and $R_{\mathrm{SLW}} : R_{\mathrm{SLB}}$).

The first source of error has been detailed in Section 1.2.4. The others are dealt with below.

Numerical error in the forward model

AATSR wavelength	Mean absolute % error	Standard deviation of absolute % error	Mean % error	Standard deviation of % error
550 nm	0.81	0.65	-0.22	1.02
660 nm	0.67	0.72	0.14	0.98
870 nm	0.66	0.96	0.62	0.98
1.6 μm	0.68	0.99	0.06	1.20

Table 2.4: Error arising from LUT interpolation in the ORAC forward model. Defined as a percentage of the DISORT-predicted sun-normalised reflectance. Results for the continental aerosol model.

The LUT-based forward model equation described in Section 2.2.3 is an approximation of the more accurate forward model which would be obtained by running a full DISORT simulation. The 'fast' forward model using LUTs is necessary as running the full radiative transfer model is much slower and would be impractical for processing the AATSR data, where many such calls to the forward model are necessary. As a result, is is important to quantify how closely the forward model matches the full DISORT output; any persistent trends or biases in the ORAC implementation may lead to biases in retrieved data. Assuming DISORT to be a perfect model of the atmosphere, a forward model error can be defined as the difference between the sun-normalised reflectance as predicted by ORAC and DISORT. There are two sources of disagreement between the ORAC fast forward model and the full DISORT code:

1. Simplifying assumptions in the ORAC forward model. The impact of these has been investigated by comparing the signal predicted by ORAC and DISORT for geometric and aerosol conditions lying exactly on LUT points. For these cases, the mean absolute percentage error is of the order of 0.1 % of the predicted sun-normalised reflectance for all four visible channels, with no strong biases or correlations with atmospheric or geometric parameters. The maximum disagreement is 0.6 %. This shows that the ORAC implementation is a reliable approximation of the DISORT forward model, for conditions corresponsding to LUT points. Figures of these results are omitted for brevity.

2. LUT interpolation error. The ORAC forward model linearly interpolates between LUT points; if the LUT value does not vary linearly between points then there will be an error arising in the modelled sun-normalised reflectance from the error in the interpolation.

The impact of LUT interpolation error has been quantified by generation of an ensemble of 30,000 sets of random atmospheric, surface and geometric conditions and using ORAC and DISORT to predict the sun-normalised reflectance seen by the satellite. For the continental aerosol model, this leads to the results shown in Table 2.4. Similar results are obtained for the other aerosol classes. Table 2.4 shows that the mean absolute percentage errors in the forward model are smaller than the instrumental noise. This indicates that for the majority of conditions the forward model is sufficiently accurate. Additionally, the mean percentage errors are (aside from at 870 nm) significantly smaller than the mean absolute percentage errors, hence any biases due to interpolation are small. Figures of these results are omitted for brevity.

In this analysis the correlation between ORAC forward model percentage error and state or geometric variables was small (always lower in magnitude than 0.25 and generally under 0.15), except in the case of surface albedo which showed a correlation of -0.79 at 550 nm, -0.72 at 660 nm, -0.56 at 870 nm and -0.54 at 1.6 μm. This can be explained in terms of Equation 2.14; over dull surfaces, the R_{BD} term is most important, and since it depends on all geometric and atmospheric state variables it is the most likely to suffer from interpolation error. As the surface albedo increases the TOA reflectance increases, and although this increases the interpolation error this error increases more slowly than the TOA reflectance, resulting in a decrease of the relative forward model error.

These small correlations show that the ORAC forward model is a good approximation of DISORT for all sets of atmospheric and geometric conditions, and so is unlikely to introduce any biases in retrieval results. These negative correlations with surface albedo are not generally significant; as the albedo is quite strongly constrained in state space by the *a priori* it might be expected that the retrieval cost (discussed in Section 2.6) would be lower over land surfaces than ocean as the fast forward model would be better able to fit the observed reflectance over brighter land pixels.

LUT interpolation error could be further mitigated by either an increase in the number of LUT points or use of higher-order interpolation. Both of these methods have their downsides, however; increased LUT size means more storage space is required, and processing time increases as the LUTs take longer to read in. Moving to quadratic interpolation similarly increases processing time. The

benefit gained from it is unclear as, although tests show it can decrease the mean forward model interpolation error, the improvements are seen for the more severe cases of interpolation error (where ORAC-DISORT disagreement decreases typically by approximately a third as compared to linear interpolation) while, for the bulk of cases with small interpolation errors, there is no overall benefit.

While further studies could determine a more optimal set of LUT points to minimise ORAC forward model disagreement with DISORT, as the forward model error is generally small when compared to the measurement error the forward model is considered to be of sufficient accuracy.

Aerosol model parameter error

No extra error term is added to $\mathbf{S}_y$ to deal with errors in the assumed aerosol model (such as particle size, refractive index, shape and height). Such errors will manifest themselves in a high retrieval cost where the aerosol model is inappropriate; attempting the retrieval for all aerosol models should then allow an attempt at aerosol speciation by picking the aerosol model which retrieves with the lowest cost. Conversely, it is possible that in some situations multiple aerosol models will be able to fit the data within retrieval precision. Because of this degeneracy, and as aerosol parameter error is not the only source of error in the retrieval, lowest-cost speciation is not expected to always be reliable, particularly in low aerosol loadings where the aerosol component of the signal is small.

Surface reflectance ratio parameter error

As discussed in the succeeding two chapters, BRDF models of surface reflectance are integrated to provide black-sky and white-sky albedos, and the ratios $R_{\mathrm{SLW}} : R_{\mathrm{SBD}}$ and $R_{\mathrm{SLW}} : R_{\mathrm{SLB}}$ are fixed in the retrieval. If these ratios are incorrect, then, this will lead to an inability to fit the measurements.

The magnitude of this contribution to the total error has been estimated by numerical simulation of data with a large (2,000 member) random ensemble of different representative geometric and atmospheric conditions. This requires a knowledge of the precision of the BRDF and black-sky albedo. For sea pixels, this was obtained via the same method as used to estimate an appropriate uncertainty on the *a priori* white-sky albedo, as detailed in Section 3.5.4. For land pixels, this was estimated using the quoted uncertainty of 0.02 on the MODIS white-sky albedo (see Chapter 4) and the differing noise-amplification factors for the BRDF, black-sky albedo and white-sky albedo detailed by Lucht and Lewis (2000).

View	Channel	Percentage error, land pixels	Percentage error, sea pixels
Nadir	550 nm	2.5 %	1.4 %
	660 nm	3.2 %	1.3 %
	870 nm	2.3 %	3.6 %
	1600 nm	1.9 %	7.8 %
Forward	550 nm	1.9 %	0.16 %
	660 nm	2.5 %	0.37 %
	870 nm	2.1 %	0.90 %
	1600 nm	1.8 %	1.6 %

Table 2.5: Forward model errors in TOA reflectance arising from incorrect model parameters (the fixed ratios $R_{\mathrm{SLW}} : R_{\mathrm{SBD}}$ and $R_{\mathrm{SLW}} : R_{\mathrm{SLB}}$). Values are given as percentages of the measured TOA reflectance, for each view and for both land and sea pixels.

Values for the resulting error in sun-normalised reflectance, expressed as percentages of the measured signal, are given in Table 2.5. The forward model shows different sensitivities to the reflectance ratios for nadir and forward-view pixels because of the different atmospheric path length and, over the ocean, the different position of the sun-glint region along the instrumental swath. Generally, aside from at long wavelengths over the sea (where the aerosol contribution is weak and the surface glint stronger), the sensitivity is lower for sea pixels than land.

The errors tabulated in Table 2.5 were found to be generally independent of geometry and surface conditions (aside from whether the satellite zenith angle was from the nadir or forward viewing geometry), with correlation coefficients generally smaller than 0.2 between error and conditions. The exception to this is that a stronger anticorrelation (of between -0.4 and -0.8 dependent on wavelength) was found between the error due to incorrect surface reflectance ratio parametrisation and the aerosol optical depth. This can be explained by higher optical depths decreasing the relative contribution of the atmosphere to the TOA reflectance. However, for simplicity and as there is no *a priori* way to tell what the aerosol optical depth for a given scene is, this dependence of the error is ignored.

Combining the errors

The three contributing factors to $\mathbf{S}_y$ (measurement error, forward model precision and model parameter error) are assumed to be independent of each other and independent between wavelengths. Thus each element of $\mathbf{S}_y$ is created by simply adding the variances due to the three terms (the squares of the errors described above). As a result of the assumption of indepedence of errors at different wavelengths, $\mathbf{S}_y$ is a diagonal matrix. Although in reality errors may be correlated between wave-

lengths, independence is assumed for computational simplicity and an exhaustive determination of the appropriate correlations has not yet been undertaken due to time constraints.

Typical values of the total error arising from these terms are of the order of 4 %-5 % of the measured TOA reflectance. Future work could look into a further improvement in the forward model precision and an improved characterisation of the sensitivity of the forward model to errors in the fixed model parameters, but such work is unlikely to have a marked impact on the retrieval algorithm as, with a small improvement to either term, the total error going into $\mathbf{S}_y$ is likely to remain of the order of 4 %-5 % of the measured TOA reflectance and other parts of the retrieval error budget are likely to be more important.

Comparison to treatment in previous versions of ORAC

The previous nadir-view version of ORAC had a simpler treatment of $\mathbf{S}_y$:

- Measurement uncertainty in the visible channels was treated as a flat value, as discussed in Section 1.2.4. No adjustment was made to account for the decrease in error resulting from superpixelling. This resulted in an overestimate of the measurement error in most cases.

- Forward model precision was not taken into account.

- Sensitivity to errors in the surface albedo was assumed to be 20 % of the surface albedo's value, with a correlation of 0.4 between wavelengths. It should be noted that at this point the retrieval was a nadir-only Lambertian retrieval, retrieving only the albedo at 550 nm and assuming the spectral shape of the *a priori* was correct, so this term also encompassed errors due to a non-Lambertian surface and errors in the *a priori* spectral shape. This high proportional value for the error led to low costs over land, and the positive correlation led to biases in the residuals (as a lower-cost solution would be found when the residuals were all of the same sign).

The prior treatment of errors resulted in part from the fact that the retrieval developed as a cloud retrieval algorithm, where the surface albedo was assumed fixed, so estimated accuracies in the surface albedo were propagated through into $\mathbf{S}_y$. As a result this work is the first time that a thorough error analysis tailored to the visible-channel aerosol retrieval has been performed as part of the ORAC retrieval scheme.

2.6.3 Quality control information

This is a very important aspect of any retrieval as there needs to be an effective way of checking the validity of the result. Fortunately, OE provides diagnostics which allow a reasonably strict quality control to be applied, principally:

- The cost, J, a check of the 'goodness of fit' of the solution to its measurement and *a priori* constraints. This should identify situations where the scene is not well-modelled.
- The retrieved error field. Scenes that are modelled successfully but have larger than expected errors can indicate problems such as a lack of information in the measurements on a particular state variable.
- The number of iterations. A failure to converge indicates no solution has been found, and instant may convergence suggests that the signal does not contain information on the state (or that the first guess at the solution was the correct state).

Model fit

The value of J at the solution indicates whether the result is good to within the statistical accuracies assumed for measurements and *a priori* data. If, at the solution, none of the measurements deviate from the calculated values (i.e. $\hat{\mathbf{y}} - \mathbf{y}$) by significantly more than their expected noise ($\mathbf{S}_y$) and no state variables deviate from their *a priori* values (i.e. $\hat{\mathbf{x}} - \mathbf{x}_a$) by significantly more than the *a priori* error ($\mathbf{S}_x$), then J will be of order $n_y + n_x$ (where n_y is, as before, the number of measurements used and n_x is the number of state variables retrieved). J_y and J_x are used to denote the measurement and *a priori* components respectively. J is divided by n_y on output; as J_y will then have an expected value of 1, this provides a quick check of the goodness-of-fit to measurements.

Because the state variables are bounded there are actually fewer degrees of freedom than $n_y + n_x$. A state variably loosely bound by $\mathbf{S}_x$ over its range, such as τ_{550} here, cannot contribute fully to J. Therefore if there are n_b state variables that are bounded but without significant *a priori* then an acceptable solution will have J of order $n_y + n_x - n_b$.

The retrieval keeps track of J_y and J_x separately (where the total cost is, from Equation 2.25 $J = J_x + J_y$). In the case of a high cost, the elements of $\hat{\mathbf{y}} - \mathbf{y}$ and $\hat{\mathbf{y}} - \mathbf{x}_a$ can be examined to find the offending measurement or state variable.

Because J depends on the estimation of values for $\mathbf{S}_y$ and $\mathbf{S}_x$ it is likely that J will not initially be of order $n_y + n_x - n_b$ as expected. Values too low imply an overestimation of (probably) measurement error; values too high imply either underestimation of noise levels or convergence criteria that are too loose.

2.6.4 Linear error analysis

OE allows rigorous error analysis of the retrieved state vector $\hat{\mathbf{x}}$. It is based on the assumption that in the vicinity of the solution the radiative transfer model is linear. In this case, it can be shown that the solution for the minimum of J (Equation 2.24) is found for:

$$\hat{\mathbf{x}} = (\mathbf{S}_x^{-1} + \mathbf{K}_x^T \mathbf{S}_y^{-1} \mathbf{K}_x)^{-1} (\mathbf{S}_x^{-1} \mathbf{x}_a + \mathbf{K}_x^T \mathbf{S}_y^{-1} \mathbf{y}) \tag{2.31}$$

We can define an 'inversion operator' $\mathbf{D}_y = \partial / \partial \mathbf{y}$ at the solution, which will be given by:

$$\mathbf{D}_y = (\mathbf{S}_x^{-1} + \mathbf{K}_x^T \mathbf{S}_y^{-1} \mathbf{K}_x)^{-1} \mathbf{K}_x^T \mathbf{S}_y^{-1} \tag{2.32}$$

$\mathbf{D}_y$ is sometimes also known as the *gain matrix* and termed $\mathbf{G}$. From Rodgers (2000), three error sources can be identified. These are known as 'null space' (or *a priori smoothing*) error, measurement mapping error and model parameter mapping error and are expressed as covariances $\mathbf{S}_N$, $\mathbf{S}_M$, and $\mathbf{S}_S$ respectively.

Null space error

This expresses error arising from a lack of information in the measurement system (hence the term 'null space'). It is calculable as follows:

$$\mathbf{S}_N = (\mathbf{D}_y \mathbf{K}_x - \mathbf{I}) \mathbf{S}_x (\mathbf{D}_y \mathbf{K}_x - \mathbf{I})^T \tag{2.33}$$

Measurement mapping error

This expresses the mapping of the errors in the measurements onto the solution and is defined by:

$$\mathbf{S}_M = \mathbf{D}_y \mathbf{S}_y \mathbf{D}_y^T. \tag{2.34}$$

Model parameter mapping error

This expresses the mapping of the errors in the forward model parameters onto the solution. The term is frequently neglected from error analyses on the grounds that $\mathbf{S}_b$ is either small (all parameters are well-known) or that all uncertain parameters are included in the state vector $\mathbf{x}$, and errors are therefore manifest through $\mathbf{S}_N$ and $\mathbf{S}_M$. It is also common to include the forward model error into $\mathbf{S}_y$, as is done in this scheme.

$$\mathbf{S}_S = (\mathbf{D}_y\mathbf{K}_b)\mathbf{S}_b(\mathbf{D}_y\mathbf{K}_b)^T \tag{2.35}$$

Error on the retrieved state $\hat{\mathbf{x}}$

The total of the null space and measurement errors gives the total error on $\hat{\mathbf{x}}$:

$$\hat{\mathbf{S}}_T = (\mathbf{S}_x^{-1} + \mathbf{K}_x^T\mathbf{S}_y^{-1}\mathbf{K}_x)^{-1} \tag{2.36}$$

Note that this expression is the same as that for J'' in Equation 2.26. All error terms are covariances and the correlations implied by the off-diagonal terms are important for interpreting the information content of the measurements. However, it is unrealistic to output the full covariances for all three error sources in an operational code environment and these are only available as special breakpoint variables for diagnosis. The ORAC nominal output contains only the square root of the diagonals for the total error, $\mathbf{S}_T$ (or optionally $\mathbf{S}_T + \mathbf{S}_S$). The following points are worth noting about the error estimate:

1. It applies to each state parameter separately—a high expected error in one parameter does not necessarily imply a high error in another.

2. It is conditional upon the correct solution being found.

3. It is conditional upon the fixed model parameters (such as the aerosol model and BRDF angular shape) being correct.

2.7 Aerosol speciation

The fact that the retrieval contains error means that it is sometimes possible for multiple aerosol classes to give a successfully-converging retrieval with low cost. This means that attempting a re-

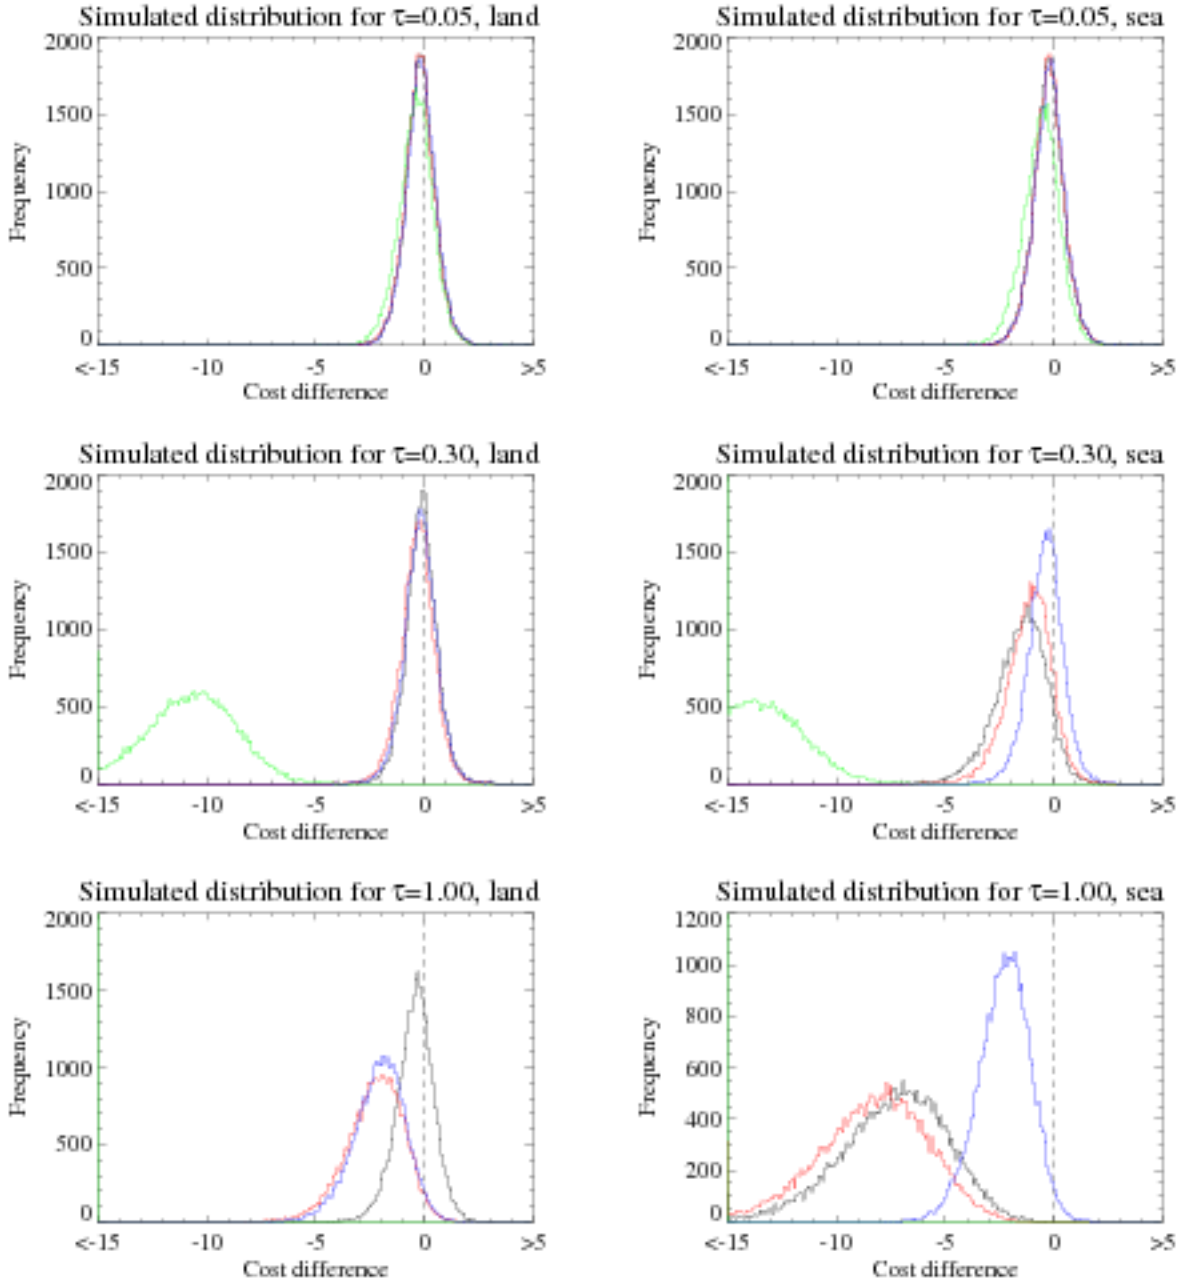

Figure 2.9: Cost difference distributions for simulated continental aerosol data using all 5 aerosol classes. The cost difference is calculated such that negative values mean the continental model retrieves with a lower cost. The method of calculation is described in the text. The left-hand plots are for a land surface and the right-hand side a sea surface. The top row are for $\tau_{550} = 0.05$, with the middle row $\tau_{550} = 0.3$ and the bottom $\tau_{550} = 1$. In all figures, the black line is the distribution for the desert dust aerosol class, the red line maritime, the green line urban, and the blue line biomass burning.

trieval with all 5 aerosol classes (Tables 2.2 and 2.3, along with Dubovik et al., 2002) and using the lowest retrieval cost as a tool for aerosol speciation will not always be reliable, particularly at low aerosol optical depths or over land where the aerosol contribution to the signal is smaller.

Figure 2.9 illustrates this point. In this case the forward model was used to generate TOA re-

flectance for the continental aerosol class and typical geometries (satellite zenith angles of 10° for nadir and 55° for forward view, a solar zenith angle and relative azimuth angle both of 45°) and state parameters (an effective radius close to the *a priori* value of 0.23 μm and surface albedos of 0.10, 0.20, 0.30 and 0.30 over land and 0.06, 0.06, 0.05 and 0.05 over sea for 550 nm, 660 nm, 870 nm and 1.6 μm respectively). For 30,000 random samples, a typical forward model error of 4 % of the TOA reflectance was used to generate an ensemble of simulated measurements for this state, and retrieval costs calculated for this state assuming each of the different aerosol classes in turn. This exercise was performed for τ_{550} of 0.05, 0.3 and 1.0 for land and sea pixels.

The cost difference in Figure 2.9 is calculated as the difference between the cost for the continental model, and that for each other model, such that negative values mean the continental model retrieves with a lower cost. For reliable discrimination between different aerosol models these distributions would be distinct and have most of their area in negative cost differences (indicating that the lowest cost was found with the correct aerosol model–in this case continental). The proportion of the distributions where the cost difference is positive is significant, particularly for lower aerosol loadings and over land surfaces. This is an indication that the measurements are insufficient to distinguish between the different types of aerosol, and there is a strong likelihood of aerosol being misidentified. As the continental, desert and maritime classes are all predominantly scattering aerosols, it is particularly difficult to distinguish between them. Similar distributions occur for other geometries, aerosol classes and particle sizes.

The distributions above assumes that the correct retrieved state is obtained with the wrong aerosol class. Another possible issue would be the retrieval converging on the *wrong* state with the wrong aerosol class; this is ameliorated by the *a priori* constraints, which should result in such situations having a high retrieval cost, but it is likely that for some pixels the incorrect state (albeit plausible within the precision of the forward model) will be retrieved. These would be best spotted via manual inspection of the data, such as urban aerosol in the clean open ocean.

Improving constraints on measured and *a priori* data may aid speciation by making these cost distributions more distinct. Addition of thermal infrared channels to the retrieval may also provide some additional contrast between the different aerosol classes. Due to the limitations of the instrument, however, it is not possible to use J alone as an unambiguous tool for aerosol speciation. Future generations of instruments will be required for this purpose; Hasekamp and Landgraf (2007) and

others have discussed that multi-angle multi-spectral instruments measuring both total light intensity and polarization, ideally with both more viewing geometries and measurement wavelengths than AATSR or other present instruments provide, are required to accurately retrieve aerosol properties of interest without recourse to significant *a priori* information.

Aerosol speciation is further discussed in Chapter 5.

2.8 Calculation of the Ångstrøm coefficient

As described the first two pieces of state information provided by the ORAC retrieval scheme are the optical depth at 550 nm and effective radius for the best-fit aerosol model. As the aerosol properties used in generation of forward model LUTs are known, additional LUTs are used after the retrieval to provide an estimate of the optical depth at 870 nm, τ_{870} . This is then used to estimate the Ångstrøm coefficient, which is defined by:

$$\text{Å} = \frac{\mathrm{d}\,\log\left[\tau(\lambda)\right]}{\mathrm{d}\,\log\left[\lambda\right]} \tag{2.37}$$

By assuming that τ varies linearly with λ in the visible, it can be directly calculated from the two optical depths:

$$\text{Å}_{550,870} = \frac{\log\left[\tau_{870}/\tau_{550}\right]}{\log\left[870/550\right]} \tag{2.38}$$

Information on the error on τ_{550} and r_e is propagated through to calculations of τ_{870} and $\mathring{A}$.

2.9 Differences between this and other AATSR algorithms

Two other algorithms, designated in this work as AATSR-Swansea and AATSR-TNO (referencing where they have been developed, to contrast with AATSR-ORAC), have been developed and applied for aerosol retrieval from AATSR. Similarities and differences between these algorithms are briefly discussed here. To date, no intercomparison of results from the assorted algorithms has been carried out. This is because none of the algorithms yet produce output data routinely, and a thorough intercomparison is planned for a future date. Two studies, Thomas et al. (2007) and Kokhanovsky et al. (2007), have involved the ORAC algorithm and one each of the others. In the former the nadir-view algorithm was in use, and the latter involved an earlier version of the current dual-view algorithm, so neither provided a comparison of the most recent versions of all the algorithms.

2.9.1 AATSR-Swansea

This algorithm is described by primarily by North et al. (1999) and North (2002a), with validation in Grey et al. (2006b). A more recent adoption of LUTs to increase the speed of the algorithm was detailed in Grey et al. (2006a). It has been applied to both ATSR-2 and AATSR data. In contrast to AATSR-ORAC, AATSR-Swansea is primarily an algorithm to retrieve the surface BRDF, with the aerosol optical depth retrieved as a byproduct of the atmospheric correction.

North et al. (1999) developed a simple BRDF model using a total of 7 parameters, which represent the angular and spectral reflectance properties of the surface. The 8 measurements are used to retrieve 6 of these parameters (the final one having been found to be approximately constant over natural surfaces), with the aerosol optical depth as an additional parameter.

The retrieval scheme, like AATSR-ORAC, finds a solution by least-squares fitting of modelled and observed data. However, instrumental precision is not taken into account and the optimal estimation framework is not used, which means that robust quality control and error measurements in the manner of AATSR-ORAC are not provided. Like ORAC-AATSR, a range of aerosol methods are used and the one resulting in the best fit to the observed data is adopted as the dominant aerosol type for a pixel. Grey et al. (2006b) also note that, like AATSR-ORAC, the measurements are not always sufficient to identify the correct aerosol model.

An advantage of this scheme over AATSR-ORAC is that it does not require *a priori* information on the surface reflectance. However, AATSR-ORAC offers several advantages:

- The assumptions made for BRDF model developed in North et al. (1999) do not hold for the sea surface, so the published literature on AATSR-Swansea only involves retrievals over land. Results have recently been obtained over the sea but are at present unpublished (North and Bevan, 2008).

- Aerosol models used in AATSR-Swansea assume fixed size distributions; by retrieving the aerosol effective radius, AATSR-ORAC allows for more flexibility in the retrieved aerosol.

- Lack of the optimal estimation framework means AATSR-ORAC provides more information about the quality of the retrieval.

- AATSR-Swansea retrieves the surface BRDF at AATSR wavelengths and view directions; it does not, however, provide information about the BRDF at other viewing directions, or provide

the surface white-sky albedo (which is a useful data product) like AATSR-ORAC does.

Comparison of the AATSR-Swansea retrieved optical depths with AERONET is discussed by Grey et al. (2006b) and results are similar in quality to those for AATSR-ORAC (discussed in Section 5.1). The retrieved BRDF data also compare well to MODIS and MISR-derived BRDF data. Optical depth derived from AATSR-Swansea matched well with that from MISR and less well with that from MODIS and TOMS. The correlation with MISR data was larger than that observed between AATSR-ORAC and MISR (see Section 5.4) but this is likely due to the fact that instantaneous observations were compared, rather than a monthly mean dataset, as presented later in this work.

2.9.2 AATSR-TNO

The second aerosol algorithm considered was developed for ATSR-2 imagery, but has been applied to AATSR data too. It is described and validated by Veefkind and de Leeuw (1997), Veefkind and de Leeuw (1998), Veefkind et al. (1998), Veefkind et al. (1999) and Veefkind et al. (2000); this information is collected together in Veefkind (1999). Although the literature contains more recent applications of the algorithm, such as Robles González et al. (2003) and Thomas et al. (2007), there have been no further published algorithm refinements. Development of the algorithm has, however, recently resumed (de Leeuw, 2008). Over both land and sea the algorithm shows a reasonable match with other data sources, although no extensive global validation is presented.

An advantage of AATSR-ORAC over AATSR-TNO is that the former applies a consistent algorithm over land and ocean. AATSR-TNO, however, applies different algorithms for land and ocean pixels. Additionally, AATSR-TNO lacks the statistical framework provided by optimal estimation.

In both the AATSR-TNO land and sea algorithms the effect of Rayleigh scattering and absorption by atmospheric ozone are accounted for in a simpler way than that of Section 2.2. They are instead treated as independent contributions to TOA reflectance from the aerosol, and these contributions are subtracted from the TOA reflectance before the aerosol retrieval is performed. In this respect AATSR-ORAC provides a more physically-accurate representation of atmospheric scattering and absorption.

Retrieval over ocean

The aerosol retrieval algorithm over water is first described by Veefkind and de Leeuw (1997) and Veefkind and de Leeuw (1998). It uses only one viewing geometry, although it can be applied to either, allowing for a consistency check by comparing the two. The view to be used is dependent on the geographic location, as the nadir and forward views will be in the sun-glint region at different points in their swath. Although the dull sea surface means a single view is often sufficient to characterise aerosol (as shown by, for example, AVHRR, MODIS and the previous version of AATSR-ORAC), use of the dual-view as in AATSR-ORAC allows for increased precision and the retrieval of more state parameters.

Surface reflectance over the ocean is fixed in the aerosol retrieval, based upon a model drawing upon similar sources to that described in Chapter 3. TOA reflectance is assumed to be a linear combination of that due to the surface and that due to aerosol; subtraction of the surface component yields the radiance arising from aerosol.

A least-squares fit is performed of the resulting radiance to that which would arise from a combination of externally-mixed aerosol models (corresponding to anthropogenic and sea salt aerosol); the retrieved parameters are the relative amounts of each of these classes present for a scene, assuming that the total radiance for an externally-mixed combination of the two aerosol types is simply the weighted total TOA reflectance for both of them. A linear relationship between the radiance due to aerosol and the aerosol optical depth is assumed, and used together with the weighting parameters derived to obtain the spectral aerosol optical depth.

This simple approach has the advantage of being rapid, but makes many approximations which will result in decreased precision as compared to AATSR-ORAC or AATSR-Swansea, and additionally will be sensitive to errors in assumed wind or ocean pigment fields as the surface reflectance is fixed. It also allows for mixed aerosol types in a pixel, which AATSR-ORAC does not, instead using models containing assumed mixtures of components as described previously.

Retrieval over land

Over land, a dual-view approach developed in Veefkind et al. (1998) is adopted. This approximates the TOA reflectance in a similar way to the Lambertian forward model (Equation 2.14).

The assumption is made that the ratio of the surface reflectance between the forward and nadir

viewing geometries is the same for all wavelengths. North et al. (1999) examine this approximation but dismiss it in favour of the BRDF model developed there (although they note that, for certain surface types, it holds true). Additionally it is assumed that aerosol contribution to the signal at 1.6 μm is negligible, so the ratio between views of TOA reflectance obtained at this wavelength is used as the assumed ratio for all wavelengths.

The strong influence of vegetation at 870 nm means that the ratio assumption does not hold here, so this channel is not used in the retrieval.

These approximations are used to eliminate the need for knowledge of the surface albedo, and allow the fitting of a weighted combination of two aerosol models to the TOA reflectance analagously to the procedure performed over the ocean. The spectral aerosol optical depth at 550 nm, 660 nm and 1.6 μm is then computed in the same way, and may be used to derive the surface reflectance.

This approach has the advantage, like AATSR-Swansea, that it does not require *a priori* information on the surface reflectance. However the aerosol models must be carefully picked for each location, limiting global applicability, and meaning the predictive power of the algorithm to spot an aerosol type where it does not typically occur is limited. Additionally, strong aerosol events or certain aerosol types may have a non-negligible impact on TOA reflectance at 1.6 μm, meaning the assumptions used to separate surface and atmospheric components will break down.

Chapter 3

A new sea surface reflectance model

As previously detailed in Section 2.3.2 the retrieval forward model requires that, for each wavelength used, three surface parameters are defined: the bidirectional reflectance R_{SBD}, black-sky albedo R_{SLB} and the white-sky albedo R_{SLW}. R_{SBD} is integrated over all satellite geometries to obtain R_{SLB}, which is then integrated over all solar angles for R_{SLW}. This white-sky albedo is one of the retrieval products, although as the ratio between these different types of reflectance is fixed during the retrieval it is important that it is well-modelled.

This chapter describes a new algorithm for the calculation of sea surface bidirectional reflectance, which can then be used to obtain the black-sky and white-sky albedos. Methods used for many ocean-colour sensors (such as the 6S model of Vermote et al., 1997) are typically concerned with the region from 400 nm to 700 nm. AATSR's visible channels are from 550 nm to 1.6 μm, hence new values of parameters must be researched, and previous assumptions re-evaluated, to create a model suitable for AATSR.

3.1 Definition and components of reflectance

3.1.1 Definition of terms

The reflectance of a body is defined as the ratio of outgoing to incoming radiance. Hence a perfectly white body has a reflectance of 1, a perfectly black body a reflectance of 0, and it is a dimensionless quantity. The distinction between reflectance and albedo is that the former is defined for a particular viewing direction, while the latter is not (either being defined for a solar angle in the case of R_{SLB}

or independent of geometry as for R_{SLW}). The sea surface bidirectional reflectance R_{SBD} may be parametrised in terms of the following factors:

- Solar zenith angle θ_{s} (measured from vertical) and azimuth angle ϕ_{s} (measured from north).
- Viewing zenith angle θ_{v} (measured from vertical) and azimuth angle ϕ_{v} (measured from north).
- Relative azimuth angle, $\phi_{\mathrm{r}} = \phi_{\mathrm{s}} - \phi_{\mathrm{v}}$.
- Wavelength λ.
- Complex refractive index $\bar{n} = n - i\kappa$ with real component n and extinction coefficient κ.
- Wind speed w, measured in ms^{-1}.
- Wind direction. Rotating clockwise from north by the wind azimuth ϕ_{w}, direction $\chi = \phi_{\mathrm{s}} - \phi_{\mathrm{w}}$.
- The total concentration C of phytoplankton pigments, principally chlorophyll-a, in $\mathrm{mg\ m}^{-3}$.
- The absorption coefficient a_{CDOM} at 550 nm due to detritus and coloured dissolved organic matter (CDOM), in m^{-1}.

3.1.2 The three components of R_{SBD}

The total reflectance R_{SBD} is defined by Koepke (1984) as being composed of three terms, representing different sources of upwelling irradiance. Firstly, light can be reflected off whitecaps in the rough ocean surface; secondly, it can be reflected off the foam-free portion of the surface. The contributions from these two factors will depend on the roughness of the sea surface, which is determined by the wind speed. Thirdly, light penetrating the surface can be scattered back up into the atmosphere by molecules within the body of water. The combination of these terms leads to the following relationship:

$$R_{\mathrm{SBD}} = W\rho_{\mathrm{wc}} + (1 - W)\rho_{\mathrm{gl}} + (1 - W\rho_{\mathrm{wc}})\rho_{\mathrm{ul}} \tag{3.1}$$

In the above $W\rho_{\mathrm{wc}}$ is the contribution from whitecaps (where W is the fractional cover of whitecaps and ρ_{wc} their reflectance); ρ_{gl} the sun glint; and ρ_{ul} the 'underlight' term from radiance reflected just below the surface of the water. These three components are dealt with individually here. ρ_{gl} is weighted against by a factor of (1-W) because specular reflection and whitecap reflection are dealt

with separately. Similarly, $\rho_{\rm ul}$ is weighted by a factor of (1-$W\rho_{\rm wc}$) because light reflected by whitecaps cannot be re-reflected as underlight. $\rho_{\rm ul}$ includes a correction to account for light lost due to glint reflection at the surface (see Section 3.4).

3.2 Whitecaps

Whitecaps are where the ocean appears white due to the action of wind creating a foam. The most simple of the three components of $R_{\rm SBD}$, their only dependence is on wind speed and wavelength.

3.2.1 Calculation

From Equation 3.1, the contribution of whitecaps to reflectance is the product of the proportion W of the surface covered by whitecaps and their average reflectance $\rho_{\rm wc}$. Older work such as Koepke (1984) treated $\rho_{\rm wc}$ in the visible region as constant with wavelength, although noting that in the near-infrared it might be expected to decrease due to absorption by water molecules. More recent coastal (Frouin et al., 1996) and open ocean (Nicolas et al., 2001) work suggests a decrease of about 40 % at 850 nm and 85 % at 1.65 μm, with a reflectance of about 0.4 at shorter wavelengths. These ratios have been adopted here for use at the nearby AATSR visible channel wavelengths, with reflectance at 550 nm and 660 nm assumed equal to 0.4.

The whitecap fraction W is here parameterised in terms of wind speed by a simple power law according to the method of Monahan and Muircheartaigh (1980). With the caveat that W cannot be greater than 1, the fraction of whitecaps is given by:

$$W = 2.951 \times 10^{-6} w^{3.52} \tag{3.2}$$

It should be noted that determination of W is complicated and various formulations based on wind speed and other environmental factors have been developed. An overview of some of these methods is given by Anguelova and Webster (2006). The method of Monahan and Muircheartaigh (1980) is used as it has been widely adopted (for example by Koepke, 1984 and more recently Vermote et al., 1997) and requires only wind speed data. In the ORAC retrieval scheme, 6-hourly winds at 1 degree resolution from the European Centre for Medium-range Weather Forecasts (ECMWF, 2008), linearly interpolated in space and time, are used throughout.

3.2.2 Magnitude of contribution

Wind speed w, ms^{-1}	Whitecap fraction W	Contribution to reflectance at 550 nm $W \times \rho_{\mathrm{wc}}$
0	2.951×10^{-6}	1.18×10^{-6}
5	8.23×10^{-4}	3.29×10^{-4}
10	9.77×10^{-3}	3.91×10^{-3}
15	1.96×10^{-2}	7.84×10^{-3}
20	1.12×10^{-1}	4.48×10^{-2}

Table 3.1: Values of W and $W \times \rho_{\mathrm{wc}}$, the whitecap fraction and contribution to sea surface reflectance at 550 nm, calculated for a range of representative wind speeds w from Equation 3.2.

The contribution of whitecap reflectance to R_{SBD} for several wind speeds is listed in Table 3.1. For wind speeds of approximately 10 ms^{-1} and higher the whitecap fraction and contribution to the reflectance are considerable.

3.2.3 Uncertainties

There are several sources of uncertainty with this section of the algorithm:

- There is a large uncertainty of up to 50 % in ρ_{wc} (Frouin et al., 1996, Nicolas et al., 2001), which will map into a corresponding uncertainty into its contribution to R_{SBD}.

- Anguelova and Webster (2006) reveal that different parameterisations of W can lead to estimates differing by up to an order of magnitude, as a simple dependence on wind speed is inadequate to explain observed variability. This may be a significant source of error in high-wind, low-chlorophyll environments away from the sun-glint region but over most of the open ocean the low values of W mean that the whitecap contribution is masked by the glint and underlight terms.

3.3 Glint reflectance

The contribution ρ_{gl} results from rays of light striking the sea surface and being specularly reflected in the observer's direction. It is calculated using the Fresnel equations, modified to account for the roughness of the wind-ruffled sea surface according to statistics developed by Cox and Munk

(1954a), (1954b). More complicated than ρ_{wc}, ρ_{gl} depends strongly on geometry, wind speed and wind direction, and weakly on wavelength.

3.3.1 Calculation

Slope distribution

The algorithm defines a coordinate system (P, X, Y, Z) such that P is the observed point on the surface and Z the altitude with PY in the direction of the Sun and PX in the direction perpendicular to the Sun's plane. The surface slope is defined by the following two components:

$$Z_x = \frac{\partial Z}{\partial X} = \sin(\alpha)\tan(\beta) \tag{3.3}$$

$$Z_y = \frac{\partial Z}{\partial Y} = \cos(\alpha)\tan(\beta) \tag{3.4}$$

In the above α is the azimuth of the ascent (clockwise from the Sun) and β the tilt. Z_x and Z_y are related to the incident and reflected directions as follows, where $\theta_s < \pi/2$ and $\theta_v > 0$:

$$Z_x = \frac{-\sin(\theta_v)\sin(\phi)}{\cos(\theta_s) + \cos(\theta_v)} \tag{3.5}$$

$$Z_y = \frac{\sin(\theta_s) + \sin(\theta_v)\cos(\phi)}{\cos(\theta_s) + \cos(\theta_v)} \tag{3.6}$$

In reality, the slope distribution will be anisotropic and dependent on wind direction χ. We can rotate the axes clockwise from the north by χ to define a new coordinate system (P, X', Y', Z) where PY' is parallel to the wind direction. Then the slope components may be re-expressed:

$$Z'_x = \cos(\chi)Z_x + \sin(\chi)Z_y \tag{3.7}$$

$$Z'_y = -\sin(\chi)Z_x + \cos(\chi)Z_y \tag{3.8}$$

The slope distribution is described by a Gram-Charlier series:

$$\begin{aligned} P(Z'_x, Z'_y) = &\frac{1}{2\pi\sigma'_x\sigma'_y} e^{(-\frac{\zeta^2+\eta^2}{2})}\Big[1 - \frac{1}{2}C_{21}(\zeta^2 - 1) \\ &- \frac{1}{6}C_{03}(\eta^2 - \eta) + \frac{1}{24}C_{40}(\zeta^4 - 6\zeta^2 + 3) \\ &+ \frac{1}{4}C_{22}(\zeta^2 - 1)(\eta^2 - 1) + \frac{1}{24}C_{04}(\eta^4 - 6\eta^2 + 3)\Big] \end{aligned} \tag{3.9}$$

Terms in this series come from Cox and Munk (1954b) and, for a clean surface, are defined as follows:

- $\zeta = Z'_x/\sigma'_x$ and $\eta = Z'_y/\sigma'_y$ where σ'_x and σ'_y are the root mean square values of Z'_x and Z'_y respectively. The value of σ'^2_x is given as 0.003+0.00192w$\pm$0.002 and σ'^2_y =0.00316w$\pm$0.004.
- C_{21} and C_{03} are the skewedness coefficients. These have values of 0.01-0.0086w$\pm$0.03 and 0.04-0.033w$\pm$0.12 respectively.
- C_{40}, C_{22} and C_{04} are the peakedness coefficients. These have values of 0.40$\pm$0.23, 0.12$\pm$0.06 and 0.23$\pm$0.41 respectively.

The Fresnel reflection coefficient

Wavelength λ	Real component n	Imaginary component κ
550 nm	1.341	1.96×10^{-9}
660 nm	1.338	2.23×10^{-8}
870 nm	1.334	3.91×10^{-7}
1.6 μm	1.323	8.55×10^{-5}

Table 3.2: Refractive indices $\tilde{n}$ for water at relevant wavelengths. Both real n and imaginary κ components are given, although only the real part is used in determining the Fresnel coefficient. Adapted from Hale and Querry (1973) and Quan and Fry (1995).

The Fresnel reflection coefficient R describes the proportion of light hitting the surface that is reflected back. Cox and Munk (1954a) provide a method to compute R that neglects the complex part of the index of refraction (a valid approximation because this is a small term, as shown in Table 3.2). For unpolarised light, R is calculated as follows:

$$R = \frac{1}{2}\left(\left[\frac{\sin(\Theta - \Theta')}{\sin(\Theta + \Theta')}\right]^2 + \left[\frac{\tan(\Theta - \Theta')}{\tan(\Theta + \Theta')}\right]^2\right) \tag{3.10}$$

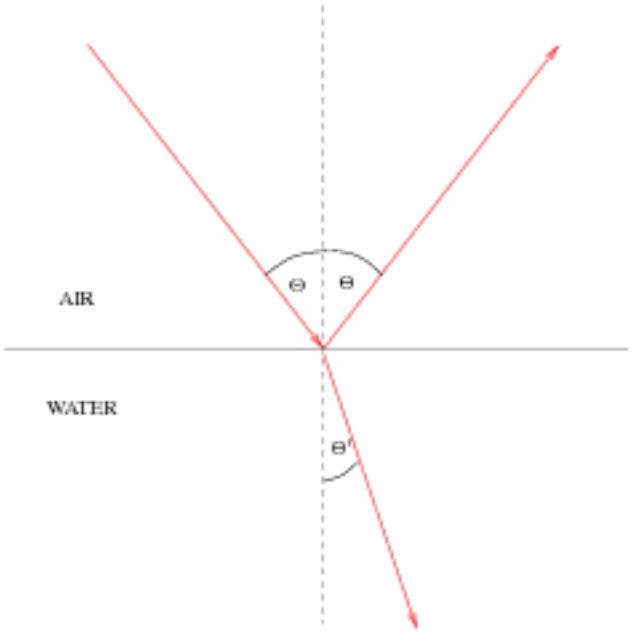

Figure 3.1: Geometry of Fresnel reflectance: the angle Θ is the angle of incidence to the surface and angle of reflection from it, both measured relative to the normal. Θ' is the angle of the refracted beam in the water. As water has a higher refractive index than air, by Snell's Law $\Theta > \Theta'$. Red lines represent paths of light rays.

Here Θ is the scattering angle while Θ' is the angle of the refracted ray passing into the water, as shown in Figure 3.1. Snell's Law describes the relationship between these angles:

$$n_{\mathrm{air}} \sin\Theta = n_{\mathrm{water}} \sin\Theta' \tag{3.11}$$

The real component of the refractive index of air, n_{air}, is taken as 1.00029 for all wavelengths. For water, real components of refractive indices were calculated at 550 nm and 670 nm using the method of Quan and Fry (1995) assuming a typical temperature of 15°C and salinity of 35 parts per thousand, but are correct to four significant figures over the range of typical temperatures and salinities.

This model extends only to 700 nm, so for the longer wavelengths values for pure water from Hale and Querry (1973) were used. At shorter wavelengths there was an offset of around 0.0065 between the refractive index as predicted for pure water and that of salinity typical for the sea, so this adjustment was also applied to the pure water data used at 870 nm and 1.6 μm. Values for the imaginary part of the refractive index were likewise taken from Hale and Querry (1973). The final data used are shown in Table 3.2. More recent optical properties are not readily available.

The angle Θ can be calculated from Equations 3.3, 3.4, 3.5 and 3.6, and is described in the α and β coordinate system by Cox and Munk (1954a). The following result can be obtained in terms of the known values θ_s and θ_v, which indirectly provides Θ:

$$\cos(2\Theta) = \cos(\theta_v)\cos(\theta_s) + \sin(\theta_v)\sin(\theta_s)\cos(\phi_r) \tag{3.12}$$

From this and the previous equations R may be calculated at the wavelengths of interest.

Combination of terms

The total contribution ρ_{gl} is given by the following equation (Cox and Munk, 1954a):

$$\rho_{gl} = \frac{\pi P(Z'_x, Z'_y) R}{4\cos(\theta_s)\cos(\theta_v)\cos^4(\beta)} \tag{3.13}$$

The tilt β may be calculated from:

$$\cos(\beta) = \frac{\cos(\theta_s) + \cos(\theta_v)}{\sqrt{2 + 2\cos(2\Theta)}} \tag{3.14}$$

3.3.2 Magnitude of contribution

Over the range of Θ sampled (from 0 to around 60°) the reflection coefficient R shows slow variation up to 40° and a consistent spectral shape. This is shown in Figure 3.2. P varies sharply with wind speed: for very low winds it is of the order 10^{-5} or smaller, leading to a very low ρ_{gl}, although for winds of about 5 ms^{-1} or more it can reach values approaching 10. P shows strong angular dependence; the shape of ρ_{gl} for a given angle owes as much to P as it does to R, as shown in Figure 3.3. The presence of multiple points at the same Θ but different ρ_{gl} is due to the geometric dependence of Equation 3.13.

In summary, the magnitude of ρ_{gl} can vary widely, with common values in the region 10^{-4} to 10^{-2} and approaching 1 toward the centre of the glint region. There is a weak dependence on wavelength but a strong one on geometry and wind speed.

As discussed in Section 2.3.2, the sea surface BRDF is integrated using Gaussian quadrature with 4 points to obtain R_{SLB} and R_{SLW}. The glint contribution requires a large number of points to calculate accurately; as a result a precalculated LUT (using 360 quadrature points) of the integrated

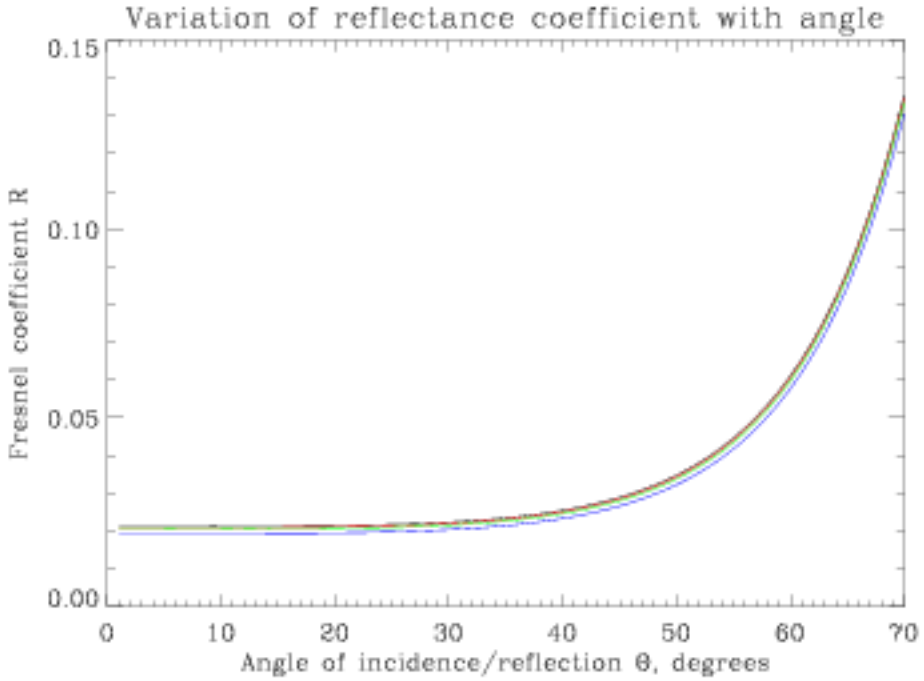

Figure 3.2: Variation of the Fresnel reflection coefficient R with Θ, the angle of incidence and reflection to the normal of the surface. Black is 550 nm, red 660 nm, green 870 nm and blue 1.6 μm.

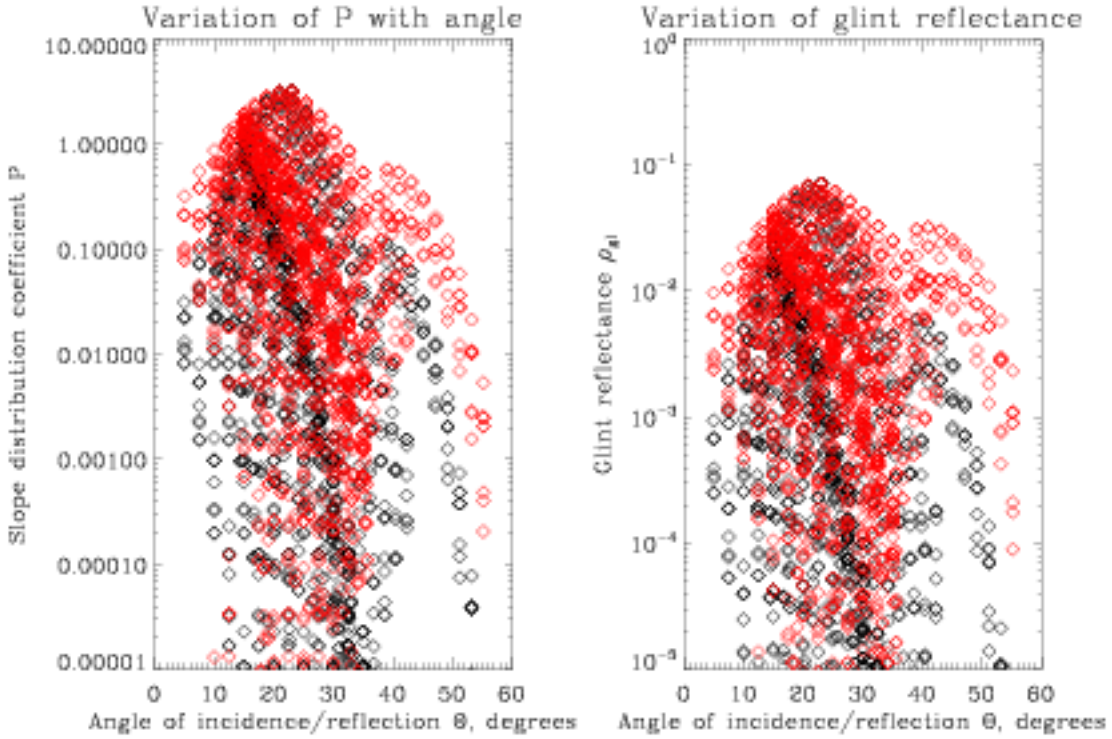

Figure 3.3: Left: Variation of slope coefficient P with scattering angle Θ. Right: Variation of ρ_{gl} at 550 nm with Θ. In both cases, black diamonds represent values for a wind speed of 5 ms^{-1} while red are for 15 ms^{-1}.

ρ_{gl}, in terms of wind speed (for R_{SLW}) and wind speed and solar zenith angle (for R_{SLb}) is used in the retrieval for speed.

3.3.3 Uncertainties

Over a range of typical conditions, the uncertainties in the coefficients used in the Gram-Charlier series lead to a variability of around 10 % in P, causing a corresponding uncertainty in ρ_{gl} of the same amount. This variability decreases at higher wind speeds, and is unlikely to be something that can be improved on.

3.3.4 Prior treatment by ORAC

Before this project began, the ORAC scheme used a similar methodology to the ρ_{gl} calculations presented here to generate the sea surface albedo. Although the equations calculate a directional reflectance, previously this was used as the Lambertian albedo. Hence this sea surface reflectance model, which correctly accounts for directionality as well as the effects of whitecaps and dissolved pigments, represents a large improvement over the original version of the algorithm.

3.4 Underlight

Underlight is upwelling irradiance from just below the surface of the ocean. As such, ρ_{ul} is influenced strongly by pigment concentration and wavelength, and weakly by geometry. The model described here is designed for Case I waters, following the nomenclature of Morel and Prieur (1977). In Case I waters (typically open ocean) the chlorophyll concentration is high compared to the scattering coefficient; in Case II waters (typically coastal and so additionally shallow) scattering by inorganic particles dominates. The semi-empirical relationships between ocean constituents and surface reflectance developed by Morel and Prieur (1977) and later work are different. This work focusses on Case I waters because they cover the majority of the Earth's surface, and scattering in Case II waters is less well-understood.

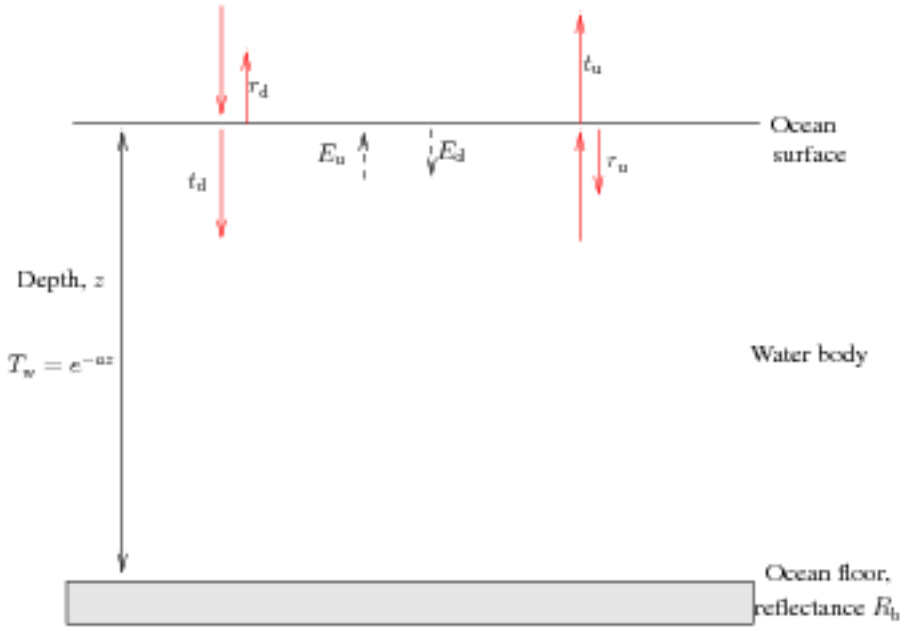

Figure 3.4: Terms important in calculation of underlight. Red lines represent light rays.

3.4.1 Calculation

The underlight contribution may be calculated using the method presented by Austin (1974), with the work begun in Morel and Prieur (1977) and developed further in Morel and Gentili (1991) for calculation of the water body reflectance R_w. The upwelling irradiance ρ_{ul} is described by the following equation:

$$\rho_{ul} = \frac{n_{air}^2}{n_{water}^2} \frac{t_d t_u (R_w + \frac{T_w^2 R_b}{1 - R_w R_b})}{1 - \bar{r}_u (R_w + \frac{T_w^2 R_b}{1 - R_w R_b})} \tag{3.15}$$

The factor of n_{air}^2/n_{water}^2 arises as the flux from a solid angle underneath the surface is refracted over a larger solid angle above it. The various parameters t_d, the transmittance coefficient for downwelling radiation at the surface, r_d, the reflectance coefficient for downwelling radiation at the surface, t_u, the transmittance coefficient for upwelling radiation at the surface, $\bar{r}_u$, the mean value of r_u, the reflectance coefficient over all observation angles for upwelling radiation at the surface, and R_w, the water body reflectance, are described in succeeding sections of this document. The parameters T_w and R_b are the transmittance through the body of water and the reflectance of the bottom of the sea respectively. These terms are visually represented in Figure 3.4.

Several approximations may be made involving these last two terms. Firstly, as R_w is likely to

be small (discussed later) and R_{b} is necessarily less than 1, $R_{\mathrm{w}}R_{\mathrm{b}}$ may be taken as very small, and so the denominators of Equation 3.15 where the term appears will be very close to 1. Secondly, the transmittance T_{w} may be calculated as follows:

$$T_{\mathrm{w}} = e^{-a_{\mathrm{w}}z} \tag{3.16}$$

In the above a_{w} is the absorption coefficient of the water, and z the depth. For pure water, a_{w} can be calculated from the complex part of the refractive index:

$$a_{\mathrm{w}} = \frac{4\pi}{\lambda}\kappa \tag{3.17}$$

Values of κ were tabulated in Table 3.2 and, assuming pure water, may be used to calculate a_{w} and hence T_{w} for a variety of depths. Over all wavelengths of interest, and even in shallow water (with $z = 100$ m), the transmittance T_{w} is very small (with a maximum of 10^{-2} for z =100 m at 550 nm, and orders of magnitude smaller for deeper water or longer wavelengths). These calculations are for pure water, and substances in seawater would further decrease T_{w}. As a result, $T_{\mathrm{w}}^2 R_{\mathrm{b}}$, the proportion of light transmitted through the water, reflected off the bottom and then transmitted up through the water body, may be neglected as almost zero. Then, with the first approximation of $R_{\mathrm{w}}R_{\mathrm{b}}$ being small, Equation 3.15 may be simplified as follows:

$$\rho_{ul} = \frac{n_{\mathrm{air}}^2}{n_{\mathrm{water}}^2}\frac{t_{\mathrm{d}}t_{\mathrm{u}}R_{\mathrm{w}}}{1-\bar{r}_{\mathrm{u}}R_{\mathrm{w}}} \tag{3.18}$$

This is the expression used to calculate ρ_{ul} in this scheme.

Downwelling transmittance coefficient, t_{d}

The term t_{d} in Equation 3.18 represents the transmittance of downwelling radiation, and assuming the underlight contribution is Lambertian and a flat sea surface may be calculated as follows after Vermote et al. (1997):

$$t_{\mathrm{d}} = 1 - \int_0^{2\pi}\int_0^{\pi/2} R_{\mathrm{sw}}(\theta_{\mathrm{s}},\theta_{\mathrm{d}},\phi)\cos(\theta_{\mathrm{d}})\sin(\theta_{\mathrm{d}})\,\mathrm{d}\theta_{\mathrm{d}}\,\mathrm{d}\phi \tag{3.19}$$

In the above, θ_{d} represents the zenith angle of a solar beam reflected from the rough water surface.

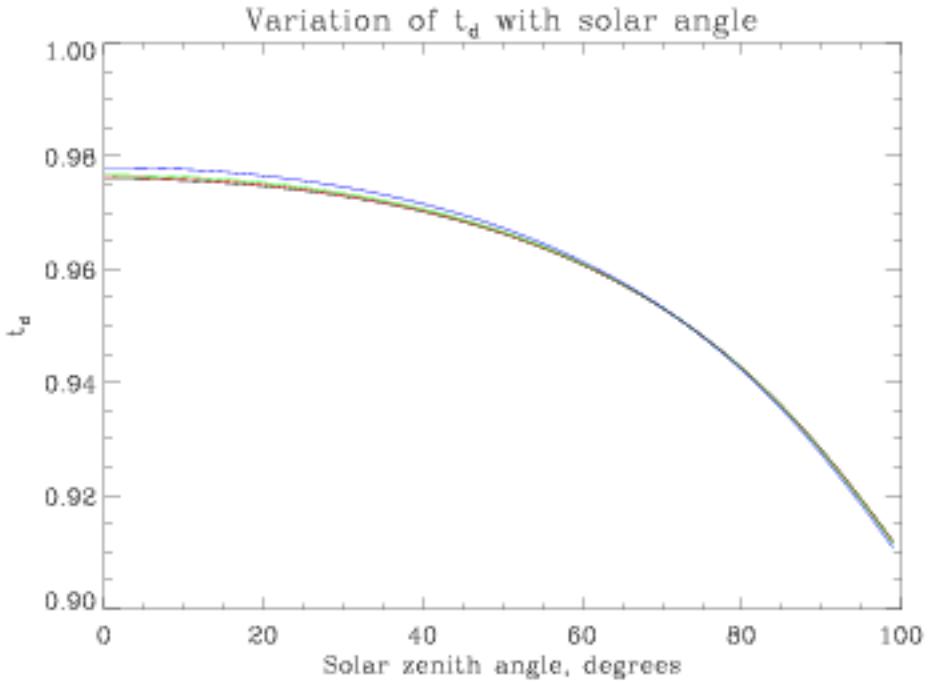

Figure 3.5: Downwelling transmittance coefficient t_{d} as a function of solar zenith angle θ_{s}. Black is 550 nm, red 660 nm, green 870 nm and blue 1.6 μm.

As integration is carried out over all θ_{d} and azimuth angles ϕ, t_{d} is dimensionless. The coefficient R_{sw} is the Fresnel reflection coefficient for a ray incoming in at zenith angle θ_{s} and reflected at θ_{d} with relative azimuth angle ϕ between them. R_{sw} is calculated as described in Section 3.3.1, using θ_{d} for θ_{v} and ϕ for ϕ_{r} in Equation 3.12.

The variation of t_{d} with θ_{s} is shown in Figure 3.5. The range of variability is not high, and there is a comparatively weak dependence on wavelength. A lookup table is generated at 1° intervals and linearly interpolated during the retrieval preprocessing.

Upwelling transmittance coefficient, t_{u}

Analagous to t_{d}, the parameter t_{u} represents the transmittance of rays from below the water surface up to the observer at zenith angle θ_{v}. It is calculated as follows:

$$t_{\mathrm{u}} = 1 - \int_0^{2\pi} \int_0^{\pi/2} R_{\mathrm{ws}}(\theta_{\mathrm{v}}, \theta_{\mathrm{u}}, \phi) \cos(\theta_{\mathrm{u}}) \sin(\theta_{\mathrm{d}}) \, \mathrm{d}\theta_{\mathrm{u}} \, \mathrm{d}\phi \tag{3.20}$$

The relevant geometry is shown in Figure 3.6. The angle θ_{u} represents the zenith angle *in the water* of an upwelling light beam refracted to the observer in the air at θ_{v}. The coefficient R_{ws} is

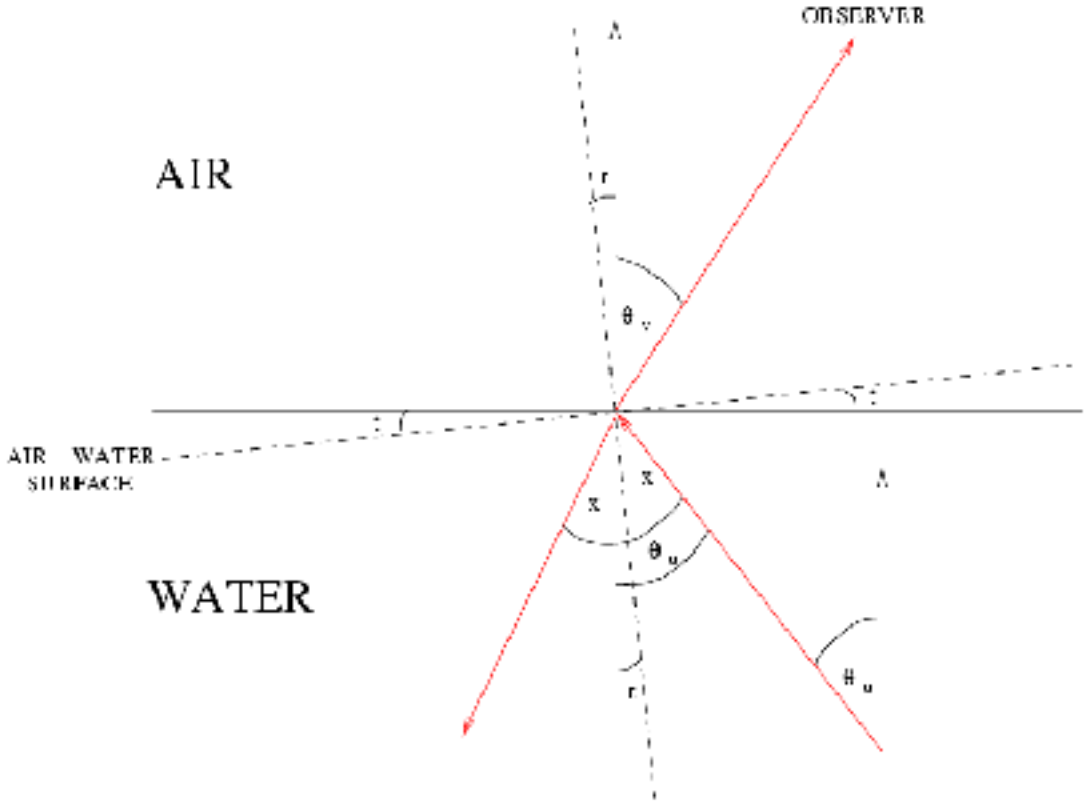

Figure 3.6: Geometry system defined for calculation of t_u. Solid black lines are the verticals zenith is measured against and the normal to it, while dotted are the inclincation and normal to the true sea surface. Red lines represent paths of light rays.

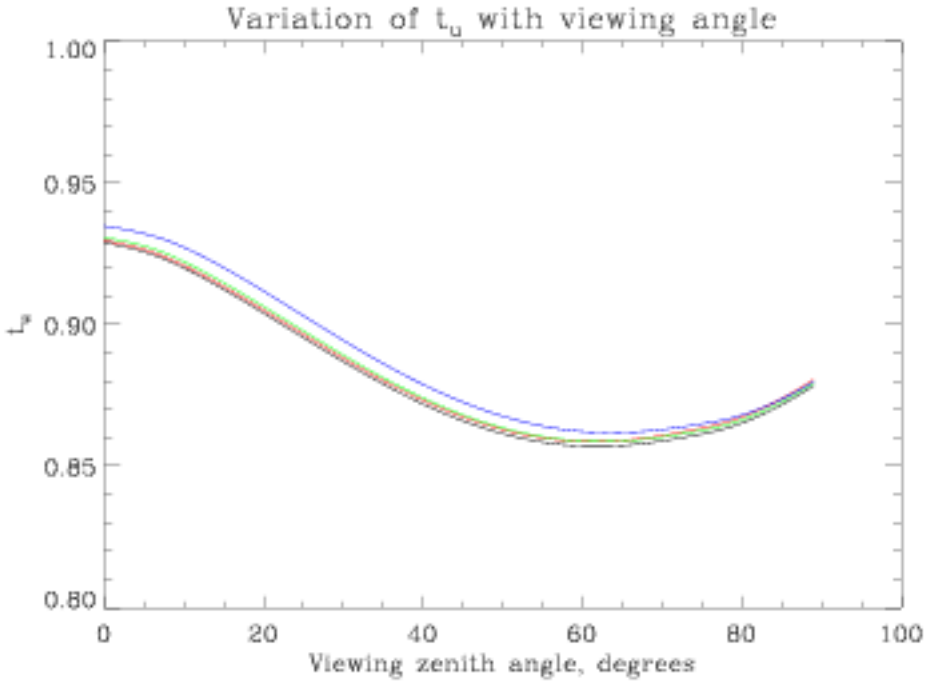

Figure 3.7: Upwelling transmittance coefficient t_u as a function of instrument zenith angle θ_v. Black is 550 nm, red 660 nm, green 870 nm and blue 1.6 μm.

the Fresnel coefficient for a beam of zenith angle θ_u coming up from the water and refracted to θ_v: hence, θ_u cannot simply be used directly in Equation 3.12 as we instead need to know the angle x this incoming beam makes with the water-air interface, which can be obtained as follows:

$$\tan(x) = \frac{1}{\frac{n_{water}}{n_{air}\sin(\theta_v+\theta_u)} - \cot(\theta_v + \theta_u)} \tag{3.21}$$

Equation 3.12 may then be used to calculate R_{wa} using x in place of θ_u, $(\theta_v + \theta_u - x)$ in place of θ_v and ϕ for ϕ_r. Care must also be taken when using Equation 3.11 to calculate Θ' as light is travelling from water to air, rather than air to water: Θ and Θ' are effectively interchanged. Figure 3.7 shows the variation of t_u with θ_v. Again, variability is small with both angle and wavelength, so an angular lookup table has been generated and is linearly interpolated during use.

Mean upwelling reflectance coefficient, $\bar{r}_u$

The final geometric term $\bar{r}_u$ is the mean reflectance coefficient for upwelling radiance at the water-air boundary, and is defined as follows:

$$\bar{r}_u = 1 - \int_0^{\pi/2} t_u(\theta_v)\cos(\theta_v)\sin(\theta_v)\,\mathrm{d}\theta_v \tag{3.22}$$

Rather than calculate this explicitly, Vermote et al. (1997) adopt a constant value of 0.485 based on values obtained by Austin (1974). This is done as calculation of r_u is computationally expensive, the value does not vary greatly with conditions and, as will be discussed in the next section, the generally low values R_w takes means that ρ_{al} is not affected by small changes in r_u. Therefore the value of 0.485 is also adopted here.

Water body reflectance, R_w

Perhaps the most important of terms in Equation 3.15, R_w depends on the optical properties of the water. It is known as the water body reflectance, and is defined as the ratio of upwelling radiance from just below the surface $E_u(\lambda)$ to downwelling radiance just above it $E_d(\lambda)$:

$$R_w = \frac{E_u(\lambda)}{E_d(\lambda)} \tag{3.23}$$

Wavelength λ	Absorption coefficient of water a_w, m^{-1}
550 nm	0.064
660 nm	0.410
870 nm	5.65
1.6 μm	672

Table 3.3: Absorption coefficient a_w (m^{-1}) for water at AATSR visible channel wavelengths. See Morel and Prieur (1977), Table 3.2 and Equation 3.17.

The method of calculation is based on the method presented first in Morel and Prieur (1977), and further developed and discussed on many occassions (for example in Morel, 1988 or Morel and Gentili, 1991). R_w is calculated from the optical properties of the water as follows:

$$R_w = f\frac{b_b(\lambda)}{a(\lambda)} \tag{3.24}$$

This describes the colour of the water as the ratio of the total backscattering coefficient $b_b(\lambda)$ to the absorption coefficient $a(\lambda)$, multiplied by some factor f.

Absorption coefficient

A more thorough treatment can be given to the absorption coefficient of water than the approximation made previously. The total absorption coefficient a of seawater can be thought of as the sum of the absorption due to pure water, a_w (as in Equation 3.17), that due to phytoplankton pigments a_{ph}, and a_{CDOM}, the absorption due to detritus and coloured dissolved organic matter (CDOM), also known as gelbstoff (literally, 'yellow stuff'):

$$a(\lambda) = a_w(\lambda) + a_{ph}(\lambda) + a_{CDOM}(\lambda) \tag{3.25}$$

The absorption coefficients used for water are tabulated in Table 3.3. Values for 550 nm and 660 nm are taken for seawater from Morel and Prieur (1977); for longer wavelengths data are unavailable so a_w at 870 nm and 1.6 μm is estimated using imaginary components of the refractive index from Table 3.2 with Equation 3.17. Use of this approximation is justified as the underlight contribution to R_{SBD} is small at these wavelengths, as discussed later, so any error is minimal in effect.

For a_{ph}, the two-component model outlined by Sathyendranath et al. (2001) is used, with coeffi-

cients from analysis of multiple measurement campaigns in different regions as described by Devred et al. (2006). This relates the absorption due to phytoplankton to the concentration C of chlorophyll-a in mg m^{-3}, assuming a mixed population of two phytoplankton types, and reproduces observed absorption with a RMSE of just under 30%, by the following equation:

$$a_{\mathrm{ph}}(\lambda) = U(1 - e^{-S_{\mathrm{chl}}C}) + a_2^*(\lambda)C \tag{3.26}$$

The parameter U is defined as follows:

$$U(\lambda) = C_1^m(a_1^*(\lambda) - a_2^*(\lambda)) \tag{3.27}$$

In the above C_1^m is the maximum chlorophyll-a concentration associated with phytoplankton population 1 in mg m^{-3}, and a_1^* and a_2^* are the specific absorption coefficients in m^{-1} (mg chl-a)$^{-1}$ of the two populations at the wavelength of interest. S_{chl} describes the nonlinearity of absorption and has units of m^3 (mg chl-a)$^{-1}$.

For a global dataset, Devred et al. (2006) found $C_1^m = 0.62$ mg m^{-3}, $a_1^* = 0.0109$ m^{-1} (mg chl-a)$^{-1}$ at 550 nm and 0.0173 m^{-1} (mg chl-a)$^{-1}$ at 660 nm, $a_2^* = 0.0064$ m^{-1} (mg chl-a)$^{-1}$ at 550 nm and 0.0085 m^{-1} (mg chl-a)$^{-1}$ at 660 nm, and $S_{\mathrm{chl}} = 1.61$ m^3 (mg chl-a)$^{-1}$. Absorption by pigments is neglected at 870 nm and 1.6 μm; the very strong absorption of the water at these wavelengths (Table 3.3) means this has approximation has negligible impact. This model results in the profile of a_{ph} shown in Figure 3.8.

It is not appropriate to consider a_{ph} as simply the product of C and a reference pigment absorption coefficient, as a_{ph} is observed to vary non-linearly with C, due to dependencies on plankton cell size and structure and the presence of accessory pigments (see, for example, Fujiki and Taguchi, 2002).

Operationally, data for both chlorophyll concentration and CDOM/detritus absorption are obtained from the GlobColour project (Barrot et al., 2006). This provides global values of various ocean colour parameters from merged satellite (MERIS, SeaWIFS and MODIS) datasets. Monthly mean values on an approximately 25 km $\times$ 25 km grid are used, with gaps filled using an annual mean (shown for 2004 in Figure 3.9).

Figure 3.9 shows mean pigment (chlorophyll-a and phaeophytin-a, normally abbreviated as just 'chlorophyll') concentrations over 2004. The large spatial variability, as well as range of concen-

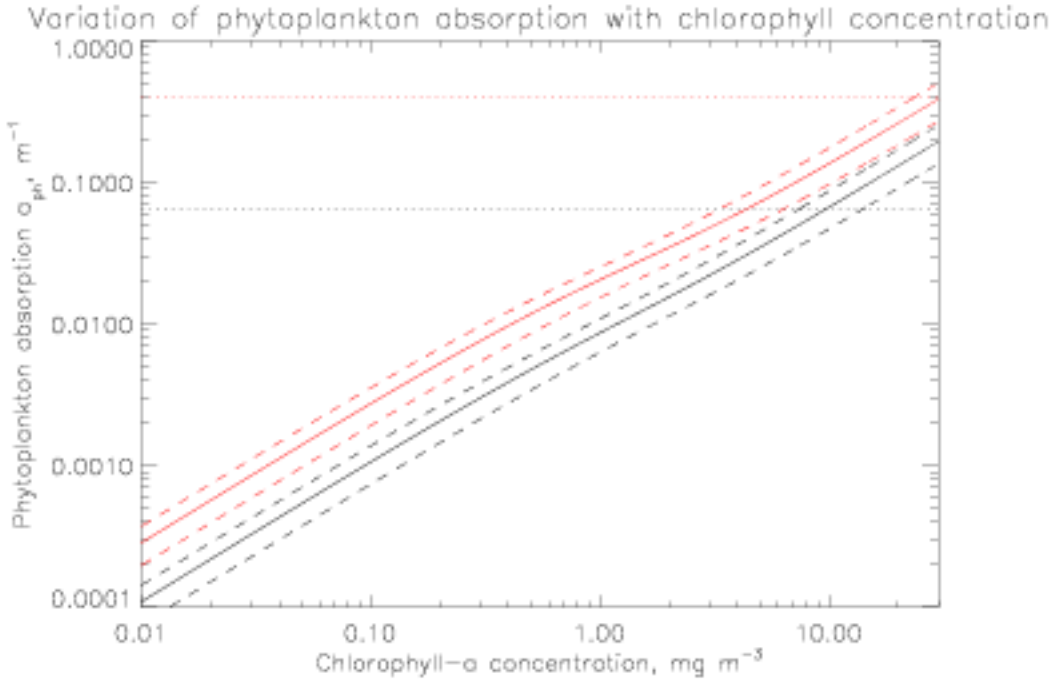

Figure 3.8: Absorption due to phytoplankton pigments a_{ph} (m^{-1}) at 550 nm (black) and 660 nm (red) as a function of chlorophyll-*a* concentration C (mg m^{-3}). Solid lines indicate a_{ph} and dashed lines the error bounds corresponding to the uncertainty in C from Barrot et al. (2006). Dotted horizontal lines are the absorption a_{w} of water at these wavelengths.

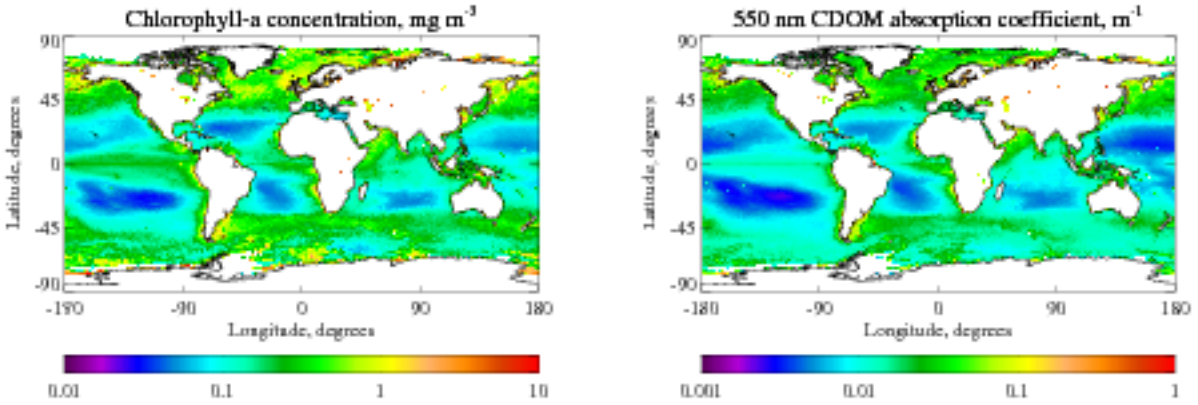

Figure 3.9: Left: GlobColour-derived chlorophyll-*a* concentration (mg m^{-3}). Right: CDOM and detritus absorption coefficient at 550 nm (m^{-1}), derived from GlobColour 443 nm absorption coefficient data. Both maps show annual mean data for 2004.

trations spanning several orders of magnitude, is evident. In the open ocean, typical values are in the range 0.05 to 1 mg m^{-3} but near the coast the concentration can reach 10 mg m^{-3} or higher. Chlorophyll content is also important for calculating the backscattering coefficient, as discussed in the next section.

According to Roesler et al. (1989), absorption from detritus and CDOM can be treated as one

parameter due to their similar spatial distributions and absorption properties; hence, the quantity retrieved by ocean colour algorithms is the total absorption coefficient for both substances. GlobColour provides the absorption coefficient at 443 nm which is related to absorption at longer wavelengths by the following equation:

$$a_{\mathrm{CDOM}}(\lambda) = a_{\mathrm{CDOM}}(443)e^{-S(\lambda-443)} \tag{3.28}$$

In Equation 3.28 the parameter S describes the spectral slope of the absorption. Roesler et al. (1989) found different values (generally in the range 0.011 to 0.018) worked well for different regions of the world, with 0.014 a good value for global studies. This value of 0.014 is used here, as well as assorted other studies (such as Chen et al., 2003).

The CDOM absorption coefficient, as shown in Figure 3.9, tends to covary with chlorophyll concentration. Typical ocean values are in the range 0.001 to 0.01 m^{-1} but again higher values, generally up to 0.1 m^{-1}, can be observed in productive or coastal waters.

This algorithm only takes a_{CDOM} into account at 550 nm. At 660 nm the value of S means that CDOM absorption is only around a fifth as strong as at 550 nm. Combined with the fact that absorption by water and chlorophyll increases by roughly an order of magnitude, the CDOM contribution to the total absorption coefficient is negligible. At longer wavelengths this effect is even more pronounced. Taking into account the decreasing importance of ρ_{ul} with increasing wavelength, this approximation has minimal effect on results.

Backscattering coefficient

The term $b_{\mathrm{b}}(\lambda)$ represents the total backscattering coefficient, defined after Morel and Prieur (1977):

$$b_{\mathrm{b}}(\lambda) = b_{\mathrm{bw}}(\lambda) + b_{\mathrm{bp}}(\lambda) \tag{3.29}$$

This equation simply states that the total backscattering coefficient b_{b} is the sum of backscattering due to molecules, b_{bw}, and particles, b_{bp}. These terms may be further parameterised; firstly, b_{bw} may be represented as $b_{\mathrm{w}}/2$ (molecular scattering is forward-back symmetric so the backscatter coefficient b_{bw} is half of the molecular scattering coefficient b_{w}). Secondly, in Case I waters scattering due to particles may be regarded as the product of the particle backscattering probability $\tilde{b}_{\mathrm{b}}$ and particle

backscattering coefficient b, leading to Equation 3.30:

$$b_{\mathrm{b}}(\lambda) = \frac{1}{2} b_{\mathrm{w}}(\lambda) + \tilde{b}_{\mathrm{b}}(\lambda) b \tag{3.30}$$

Values for b_{w} for pure water from 380 nm to 700 nm were given by Morel and Prieur (1977). Morel (1974) tabulated values from 350 nm to 600 nm for both pure water and typical seawater. The data were shown to fit a power law with a dependence on $\lambda^{-4.32}$, with seawater scattering around 1.30 times as much as pure water. This relationship has been used to extrapolate this data to the 660 nm, 870 nm and 1.6 μm channels. The values obtained for b_{w} are 1.93×10^{-3} m^{-1} at 550 nm, 8.77×10^{-4} m^{-1} at 660 nm, 2.66×10^{-4} m^{-1} at 870 nm and 1.91×10^{-5} m^{-1} at 1.6 μm. The value predicted for 660 nm is in good agreement with that given for pure water at 660 nm in Morel and Prieur (1977) multiplied by the aforementioned factor of 1.30. Another point to note is that b_w is generally small (both in absolute terms and when compared to $\tilde{b}_{\mathrm{b}}$) at longer wavelengths, meaning any error in this extrapolation is minor in terms of influence on R_{w}.

The second parameter in Equation 3.29, $\tilde{b}_b(\lambda)$, is the backscattering probability: the ratio of the backscattering to scattering coefficients of the pigments. It is related to the total concentration C of chlorophyll a and pheophytin a, measured in mg m^{-3}, and wavelength λ, measured in nm, by the following expression:

$$\tilde{b}_{\mathrm{b}}(\lambda) = 0.002 + 0.02(0.5 - 0.25\log_{10}C)\frac{550}{\lambda} \tag{3.31}$$

The final term in the backscatter component of Equation 3.29, b is calculated as:

$$b = 0.3C^{0.62} \tag{3.32}$$

The relationship between b and C was derived by Morel (1988) for data at 550 nm; the wavelength-dependence of particle backscattering is taken into account by the λ^{-1} factor in Equation 3.31. It should be noted that although parametrised in terms of C, the models were developed to account for scattering from suspended organic matter as well as pigment.

Ratio multiplier f and combination for water body reflectance R_{w}

Morel and Prieur (1977) initially gave f, the multiplier of the ratio of total backscattering to total

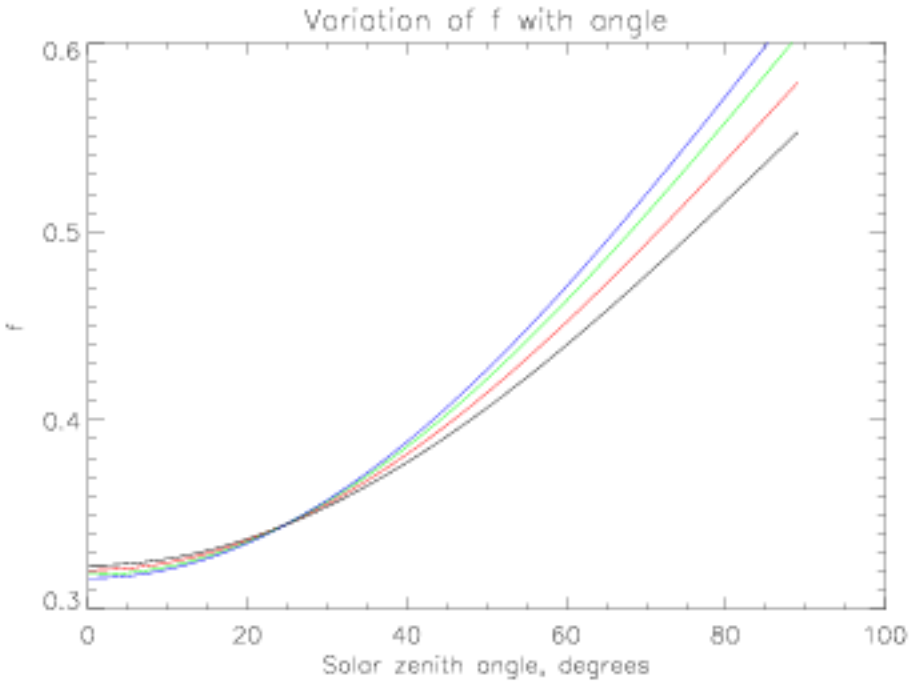

Figure 3.10: Variation of total backscattering:absorption ratio multiplier f with solar zenith angle and backscattering ratio η_b, based on a pigment concentration of 0.3 mg m^{-3}. The black line shows values for 550 nm, red for 660 nm, green for 870 nm and blue for 1.6 μm.

absorption used to calculate the water body reflectance R_w, a value of 0.33. Subsequent work has found it to depend on the solar geometry and the optical properties of water. The method used here was put forward by Morel and Gentili (1991), stated to be accurate within 1.5 % for solar zenith angles smaller than 70°. It relates f to the proportion of backscattering due to water molecules ($\eta_b = b_{bw}/b_b$) and the cosine of the solar zenith angle (μ_s) as follows:

$$f = 0.6279 - 0.2227\eta_b - 0.0513\eta_b^2 + (-0.3119 + 0.2465\eta_b)\mu_s \tag{3.33}$$

Assuming a pigment concentration of 0.3 mg m^{-3}, representative η_b values are approximately 0.32 at 550 nm, 0.20 at 660 nm, 0.09 at 870 nm and 0.01 at 1.6 μm. The resultant variation of f is depicted in Figure 3.10, revealing a weak dependence on wavelength and a stronger dependence on solar angles. From this it can be seen that according to this parametrisation the constant value of 0.33 for f would in most cases be an underestimate.

The above equations allow computation of R_w at varying representative pigment concentrations, a selection of which are displayed in Table 3.4. The strong dependence on both λ and C should be noted.

C, mg m^{-3}	R_w at 550 nm	R_w at 660 nm	R_w at 870 nm	R_w at 1.6 μm
0.1	1.18×10^{-2}	1.30×10^{-3}	6.52×10^{-5}	3.03×10^{-7}
1.0	2.49×10^{-2}	3.04×10^{-3}	1.80×10^{-4}	9.49×10^{-7}
10.0	4.80×10^{-2}	5.16×10^{-3}	4.53×10^{-4}	2.70×10^{-6}

Table 3.4: Water body reflectance R_w at AATSR visible wavelengths and pigment concentrations C. The solar zenith angle was taken to be 40°, and the CDOM absorption coefficient a_{CDOM} 0.005 m^{-1} at 550 nm.

3.4.2 Magnitude of contribution

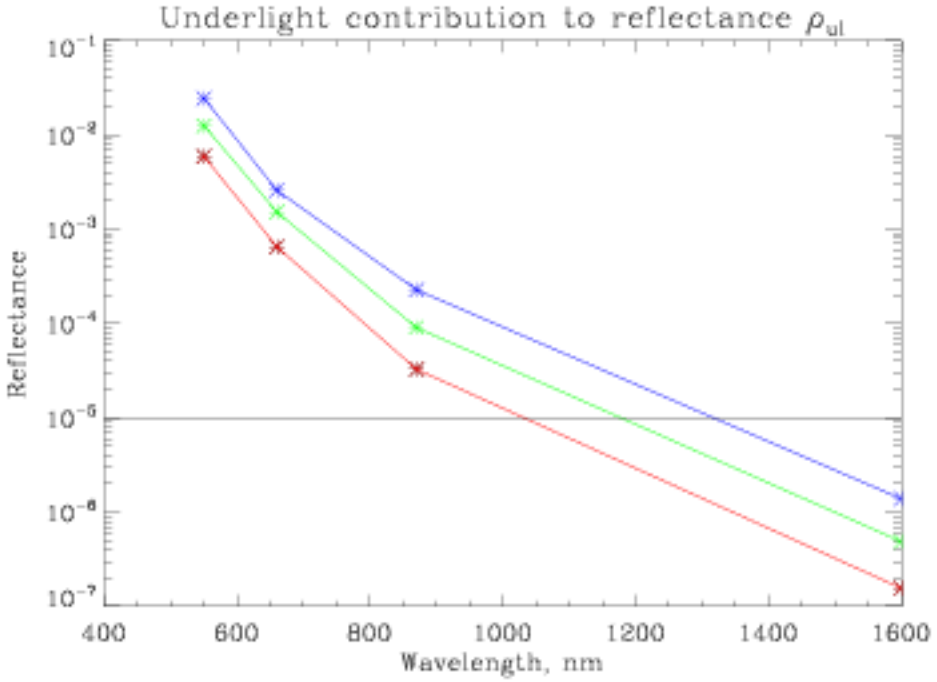

Figure 3.11: Variation of ρ_{ul} with wavelength, with data points at the AATSR visible channel positions. The satellite and solar zenith angles are taken as 10°. Red points are for a pigment concentration C of 0.1 mg m^{-3}, green for 1 mg m^{-3} and blue for 10 mg m^{-3}. the CDOM absorption coefficient a_{CDOM} was taken to be 0.005 m^{-1} at 550 nm. The line across ρ_{ul} =10^{-5} shows the minimum value of surface reflectance used in the AATSR retrieval.

The contribution ρ_{ul} is dominated by the shape of R_w, and is shown in Figure 3.11. At the shorter wavelengths it is of the order of 10^{-2}-10^{-3}, and is likely equal to or larger than other contributions to R_{SBD}. Hence knowledge of C is essential to accurately judge the magnitude of the contribution to the total reflectance. As it shows a stronger wavelength-dependence than ρ_{wc} and ρ_{gl}, the spectral shape of R_{SBD} will be largely determined by ρ_{ul} outside of the sun-glint region.

At longer wavelengths the contribution is low: for the 1.6 μm channel it is calculated to be lower than the minimum increment used in the AATSR retrieval at even the highest chlorophyll concentrations normally observed.

3.4.3 Uncertainties

The major uncertainties associated with ρ_{ul} are errors arising from poor characterisation of pigment and CDOM distributions and scattering. The small size of the underlight term at the longer two wavelengths means that errors in ρ_{ul} will only have minor impacts on the modelled reflectance.

- The relationship between $\bar{b}_{\mathrm{b}}$ and C was developed for Case I waters (according to the definitions of Morel and Prieur, 1977) and so may not accurately characterise scattering in Case II waters (where pigment and scattering particles do not covary in the same way). This may cause the algorithm to perform less well over Case II waters. Case II waters are largely coastal and as the inhomogeneity of coastal regions presents other problems for aerosol retrieval no adjustment for Case II waters has been developed at this stage. This would, however, be a desirable goal for future work.

- The assumption of zero absorption by pigments at the longer wavelengths is not considered to be a problem given that ρ_{ul} is very small at these wavelengths.

- Errors arising from use of monthly means for chlorophyll and CDOM values. The GlobColour chlorophyll products have a stated accuracy of 31 %. CDOM errors are not given by Barrot et al. (2006). Further errors arise due to variations on shorter timescales than a month.

- There may be regional biases from using 0.014 as a global CDOM spectral slope S, as Roesler et al. (1989) found values from 0.011 to 0.018 in different parts of the world.

- Neglecting CDOM and detritus absorption at wavelengths longer than 550 nm has a negligible effect on results due to the differing wavelength-dependences of absorption by pure water, chlorophyll and CDOM.

3.5 Typical values, patterns, and errors

3.5.1 The BRDF

Figures 3.12 and 3.13 show example nadir-view and forward-view BRDFs generated using this scheme at 550 nm and 1600 nm. Sun-glint is visible in the centre of the nadir-view swath and toward the north of the forward-view swath. These glint patterns persist for different orbits as the

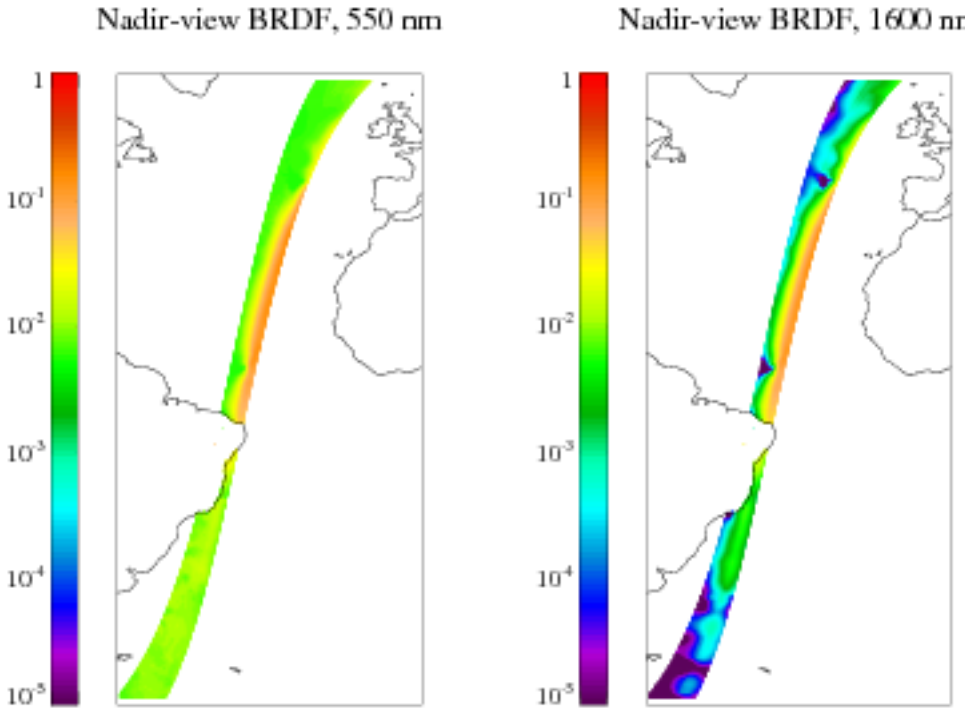

Figure 3.12: Nadir-view sea BRDF at 550 nm (left) and 1600 nm (right). Data for an AATSR swath on September 6 2004.

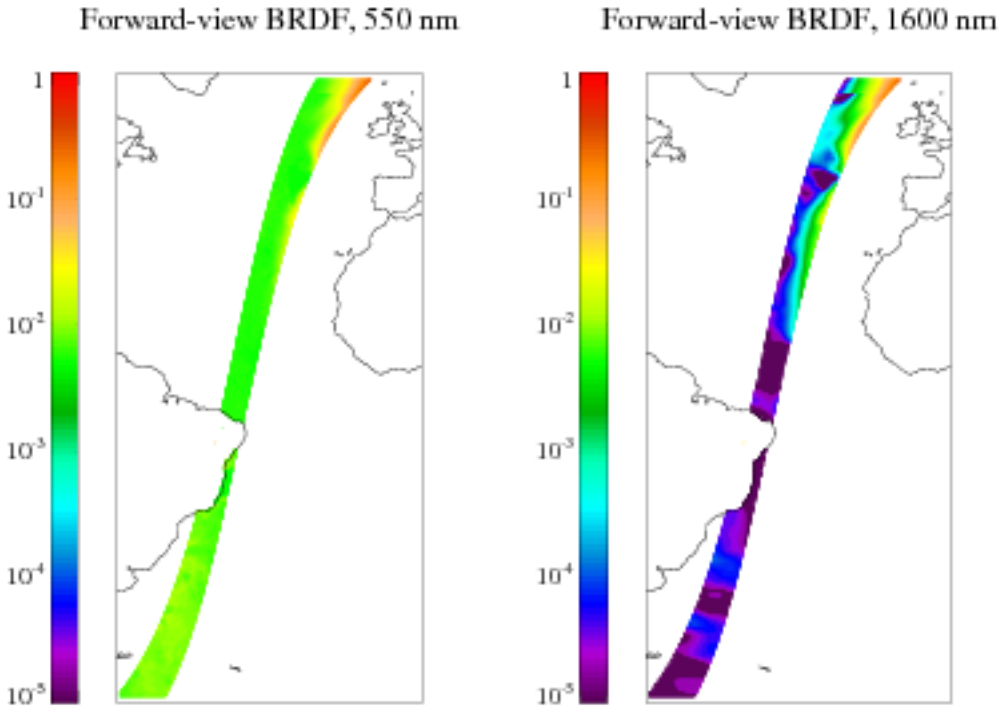

Figure 3.13: Forward-view sea BRDF at 550 nm (left) and 1600 nm (right). Data for an AATSR swath on September 6 2004.

satellite geometry remains the same. Perturbations to this glint shape arise due to variations in the wind speed and from increased reflectance due to oceanic whitecaps or underlight, particularly at 550 nm. BRDF at 660 nm and 870 nm takes values in between these two wavelengths; at 550 nm all contributions to the BRDF are important, while at 1600 nm the shape is glint-dominated. The exact location of the glint region varies seasonally.

The BRDF can take values from around 10^{-5} to 1 dependent on the wavelength and location with respect to the sun-glint region. The highest values are observed for shortest wavelengths; near the glint region the BRDF can be nearly spectrally flat while far away there can be orders of magnitude difference. The BRDF for the forward and nadir views is often quite different. For example, for $\theta_s = 50^\circ$, $\phi_r = 45^\circ$, $\phi_s = 45^\circ$, $\phi_w = 20^\circ$, $w = 10$ m s^{-1}, $C = 1$ mg m^{-3} and $a_{\mathrm{CDOM},550} = 0.005$ m^{-1}, with $\theta_v = 10^\circ$ the BRDF is 0.0183 at 550 nm, 0.00634 at 660 nm, 0.00457 at 870 nm and 0.00406 at 1600 nm. In contrast with a forward-view geometry of $\theta_v = 55^\circ$ the BRDF is much lower and more spectrally variable with 0.0131 at 550 nm, 0.00193 at 660 nm, 0.00033 at 870 nm and 0.00006 at 1600 nm.

3.5.2 Black-sky albedo

The black-sky albedo, shown at 550 nm and 1600 nm in Figure 3.14 for the same swath as the previous BRDF example, is almost identical for both of AATSR's viewing geometries (whose solar angles differ by under 1° for any given pixel). Values increase as the solar zenith angle increases; as with the shape of the BRDF, perturbations to the basic shape arise due to the wind and pigment distribution. The same scale is used as in Figures 3.12 and 3.13 to illustrate the comparative variability of the reflectance and albedo. Again, values at 660 nm and 870 nm are intermediate between these two wavelengths.

The black-sky typically takes value around 10^{-2} (where the sun is high) to 10^{-1} (where the sun is low), and is slightly larger at shorter wavelengths.

3.5.3 White-sky albedo

Figure 3.15 shows the *a priori* white-sky albedo generated for this overpass. Being independent of geometry it is the same for both instrument viewings and variability over the globe is determined by wind and pigment distributions: the glint shape of the BRDF is 'averaged out' by the integration,

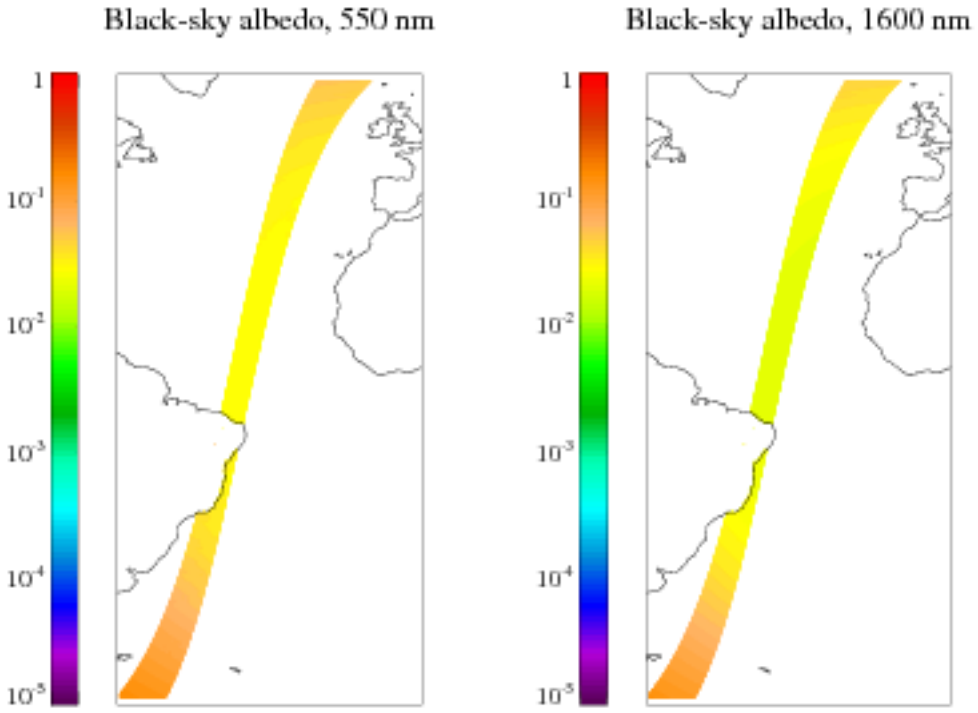

Figure 3.14: Sea black-sky albedo at 550 nm (left) and 1600 nm (right). Data for an AATSR swath on September 6 2004.

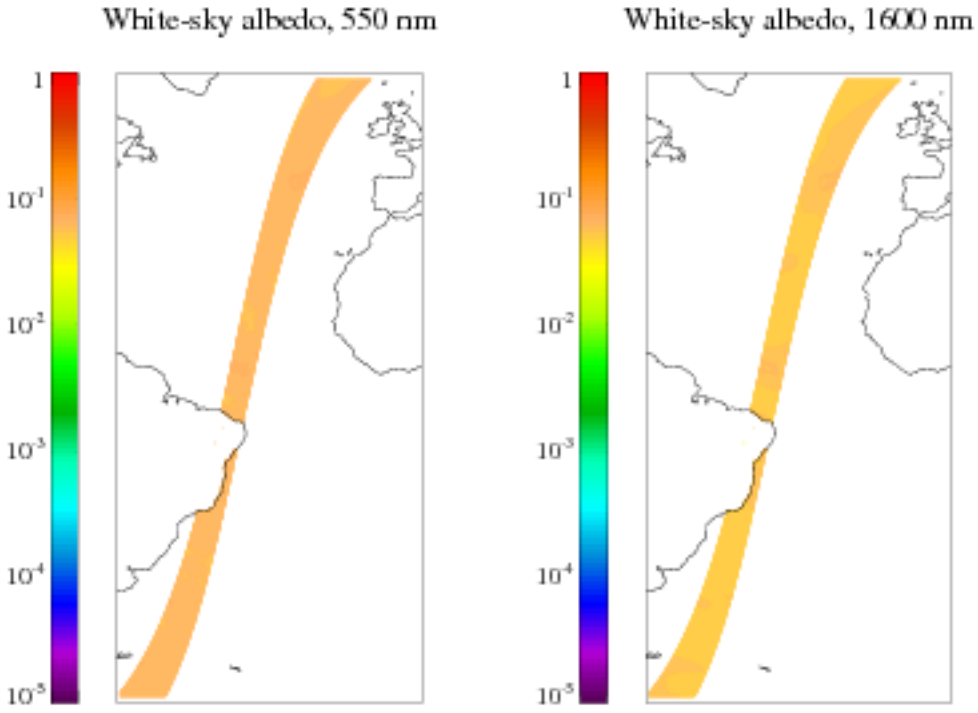

Figure 3.15: Sea white-sky (Lambertian) albedo at 550 nm (left) and 1600 nm (right). Data for an AATSR swath on September 6 2004.

which contributes around 0.05 to the albedo. The white-sky albedo is spectrally flatter and less variable than the BRDF or black-sky albedo. Typical values are in the range 0.05-0.08 at all wavelengths, with the shortest wavelengths being brightest. The same scale is again used as in Figures 3.12 and 3.13, to illustrate this comparative lack of variability.

3.5.4 *A priori* error and correlations

Uncertaintities in the different terms comprising the sea surface reflectance, such as pigment concentrations or foam reflectance, were detailed in the preceding sections. It is important to assign a reasonable error to the *a priori* albedo generated: too tight and the retrieval will be unduly constrained by an imperfect model, but too loose and some of the information on the state is effectively thrown away. Appropriate uncertainties have been determined in the following way:

- Generation of 1,000 random sets of typical ocean and viewing states (for example differing wind speeds, chlorophyll concentrations and geometries).

- For each ocean state, generation of R_{SBD} and integration for R_{SLB} and R_{SLW} for an ensemble (50 members each) of random perturbations to the uncertain model parameters (such as the foam reflectance factor or sea slope characteristics from Cox and Munk, 1954b). The magnitude of the perturbations is determined by the stated uncertainty on the model parameter as previously described in the text.

- Calculation of the ensemble mean reflectance (or albedo) and its standard deviation for each ocean state, for each of R_{SBD}, R_{SLB} and R_{SLW}. The ratio of the standard deviation to the mean is a measure of the sensitivity of the state to errors in the model parameters: a larger ratio means that the calculated reflectance or albedo is more sensitive to errors in the model parameters.

- Calculation of the mean of these sensitivity ratios for each wavelength, over all the ocean states. For R_{SLW}, multiplying these means by the *a priori* values generated for a pixel and squaring will give the diagonal elements of the *a priori* covariance matrix $\mathbf{S}_x$ corresponding to the four *a priori* R_{SLW} for that pixel. For R_{SBD} and R_{SLB}, the analagous values are used to estimate forward model error, discussed in Section 2.6.2.

This process leads to values of the mean proportional sensitivity to parameter error for R_{SLW} of 0.20 to 2 decimal places at all wavelengths. The equality between wavelengths is an indicator that the

dominating terms and, namely the integrated glint contribution, are similar for each. In the retrieval, multiplying these sensitivities by the *a priori* albedo gives the *a priori* error estimate for a pixel. Repetition of calculations with a larger ensemble of random perturbations gives the same results to 2 significant figures.

Analagous values for the sensitivity of the BRDF for nadir-view geometries are 1.23 at 550 nm, 1.29 at 660 nm, 1.34 at 870 nm and 1.35 at 1600 nm and for forward view geometries 1.13 at 550 nm, 1.14 at 660 nm, 1.15 at 870 nm and 1.15 at 1600 nm. For the black-sky albedo, sensitivities are 0.21 at 550 nm, 0.19 at 660 nm, 0.20 at 870 nm and 0.20 at 1600 nm. The higher sensitivities for R_{SBD} as compared to R_{SLB} or R_{SLW} arise due to its strong geometric dependence. These values were used to estimate the sensitivity of the forward model to errors in the ratios $R_{\mathrm{SLW}} : R_{\mathrm{SBD}}$ and $R_{\mathrm{SLW}} : R_{\mathrm{SLB}}$ as discussed in Section 2.6.2.

The collection of ocean states can also be used to calculate the off-diagonal elements of the covariance matrix. The value of matrix element *i,j* is the product of the uncertainty on *i*, the uncertainty on *j* and the correlation between the parameters *i* and *j*. Using the previously-generated set of white-sky albedos leads to the correlation matrix shown in Table 3.5; again, replication leads to agreement to two significant figures. A flag may be set in the retrieval to use these values (and if not set, the albedo in each channel is assumed independent). The high correlations are indicative of the fact that the same factors (i.e. high wind and pigment levels) lead to a brighter surface at all wavelengths.

Wavelength λ	550 nm	660 nm	870 nm	1600 nm
550 nm	1.00	0.47	0.37	0.36
660 nm	0.47	1.00	0.88	0.87
870 nm	0.37	0.88	1.00	0.94
1600 nm	0.36	0.87	0.94	1.00

Table 3.5: Correlation matrix for *a priori* white-sky albedo at AATSR visible channel wavelengths over ocean, calculated from a range of representative ocean states.

Chapter 4

Land surface reflectance

Accurate modelling of the surface reflectance is particularly crucial over the land, as its increased brightness as compared to the ocean results in a higher contribution to the satellite signal. As the land surface is both temporally and spatially highly variable, an auxiliary dataset is employed to set the *a priori*.

This chapter describes the use of the MOD43B1 data product detailed by Wanner et al. (1997) to obtain the MODIS surface BRDF, black-sky and white-sky albedo $R_{\mathrm{SBD,M}}$, $R_{\mathrm{SLB,M}}$ and $R_{\mathrm{SLW,M}}$ (the subscripted M indicates MODIS reflectance or albedo), which are converted to a reflectance and albedos suitable for AATSR for use in the retrieval. Error estimates and correlations between the white-sky albedo at the four wavelengths of interest are also discussed.

It should be emphasised that the MODIS BRDF retrieval algorithm is independent of the MODIS aerosol retrieval algorithm. Therefore errors in the MODIS aerosol product will not lead to errors in the BRDF (which could manifest as biases in the AATSR retrieval).

4.1 The MODIS BRDF product and forward model

4.1.1 The kernel-driven model

The MODIS BRDF product consists of a set of weights derived by fitting observed reflectances to a semi-empirical kernel-driven model of surface reflectance. The product allows use of the Ambrals (Algorithm for MODIS bidirectional reflectance anisotropy of the land surface) forward model (Wanner et al., 1997) to calculate the BRDF $R_{\mathrm{SBD,M}}$ for a given viewing geometry, given these weights

and a set of kernels (detailed by Wanner et al., 1995) which describe reflectance from various kinds of surface:

$$R_{\mathrm{SBD,M}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) = f_{\mathrm{iso}}k_{\mathrm{iso}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) + f_{\mathrm{vol}}k_{\mathrm{vol}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) + f_{\mathrm{geo}}k_{\mathrm{geo}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) \tag{4.1}$$

As before the following angular notation is adopted:

- Solar zenith angle θ_{s} (measured from vertical) and azimuth angle ϕ_{s} (measured from north).
- Viewing zenith angle θ_{v} (measured from vertical) and azimuth angle ϕ_{v} (measured from north).
- Relative azimuth angle, $\phi_{\mathrm{r}} = \phi_{\mathrm{s}} - \phi_{\mathrm{v}}$.

The three kernels k_{iso}, k_{vol} and k_{geo} parametrise isotropic scattering, volume scattering (from dense vegetation) and geometric scattering (shadowing effects between canopies, such as large isolated trees) as functions of the viewing geometry. The parameters f_{iso}, f_{vol} and f_{geo} are the respective weights of these kernels, characteristic for surfaces of different types, and so differ for each pixel. As the first kernel represents isotropic scattering it has no geometric dependence and is equal to 1 for all geometries. The Ross-Thick kernel k_{RT} is used for k_{vol} and the Li-Sparse kernel k_{LS} is used for k_{geo}. The forms and derivations of both of these kernels and others are given by Wanner et al. (1995). Hence Equation 4.1 can be reduced to the following form:

$$R_{\mathrm{SBD,M}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) = f_{\mathrm{iso}} + f_{\mathrm{vol}}k_{\mathrm{RT}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) + f_{\mathrm{geo}}k_{\mathrm{LS}}(\theta_{\mathrm{s}},\theta_{\mathrm{v}},\phi_{\mathrm{r}}) \tag{4.2}$$

In this retrieval scheme $R_{\mathrm{SBD,M}}$ is obtained using the above equations using the parameters provided in MOD43B1 with the appropriate kernels, and integrated to provide the black-sky and white-sky albedos for a MODIS-like sensor using a precalculated LUT of kernel integrals and an expression provided by Lucht et al. (2000) respectively. To account for the differing shapes of the instruments' visible channels, they are then adjusted to give the AATSR equivalents for use in the retrieval.

4.1.2 MOD43B1 output data

Values for the weighting parameters f are obtained by inversion of atmospherically-corrected reflectances from the MODIS instrument aboard Terra according to the algorithm of Lucht et al. (2000).

Collection 4 of the data are currently being used. Parameters are obtained for MODIS visible channels (4 of which are centred near AATSR channels) and 3 broad bands on a 16-day repeat cycle. This repeat cycle allows for sampling of multiple angular parts of the BRDF by different satellite overpasses, under the assumption that the surface properties do not show great variability during this time period. It also ameliorates the effects of cloud cover on sampling.

The weighting parameters are supplied on a 1 km×1 km sinusoidal grid. The data are supplied in granules, with the globe divided into 460 granules (of 1200 km×1200 km each); to reduce file size, each granule is down-sampled to approximately 2.5 km×2.5 km resolution and the broadband datasets removed before processing. Data are available for the whole of the AATSR mission.

Broadband white-sky albedo derived from the Collection 4 BRDF model parameters has been validated (Jin et al., 2003a, Jin et al., 2003b) and found to generally meet the accuracy requirement of 0.02 in comparison to field measurements.

For surface data from 2007 onwards, the MOD43B1 product is unavailable and MCD43B1, which uses combined observations from MODIS-Aqua and MODIS-Terra, is used instead. The methodology associated with MCD43B1 is otherwise identical to that described here. When reprocessing and validation of the new dataset are completed, it will be adopted instead of MOD43B1 (provided it is no worse). MCD43B1 is part of the newest Collection 5 of the MODIS product suite.

4.2 Adaptation to AATSR channels

MODIS and AATSR are different sensors and each channel measures not a single wavelength but a narrow band. As a result, the albedo reported by MODIS is an average albedo over the band, weighted by the instrument's response function. Figure 4.1 shows that the visible channel response functions for these instruments do not match, so if the true spectral albedo varies across each band then the albedo each instrument would see for the same underlying surface would differ. This effect would be the most pronounced for the 660 nm and 1.6 μm channels, as these have the least similar response functions between the two instruments. It is important that the magnitude of the error use of MODIS data introduces be quantified, and, if possible, the MODIS reflectances corrected for.

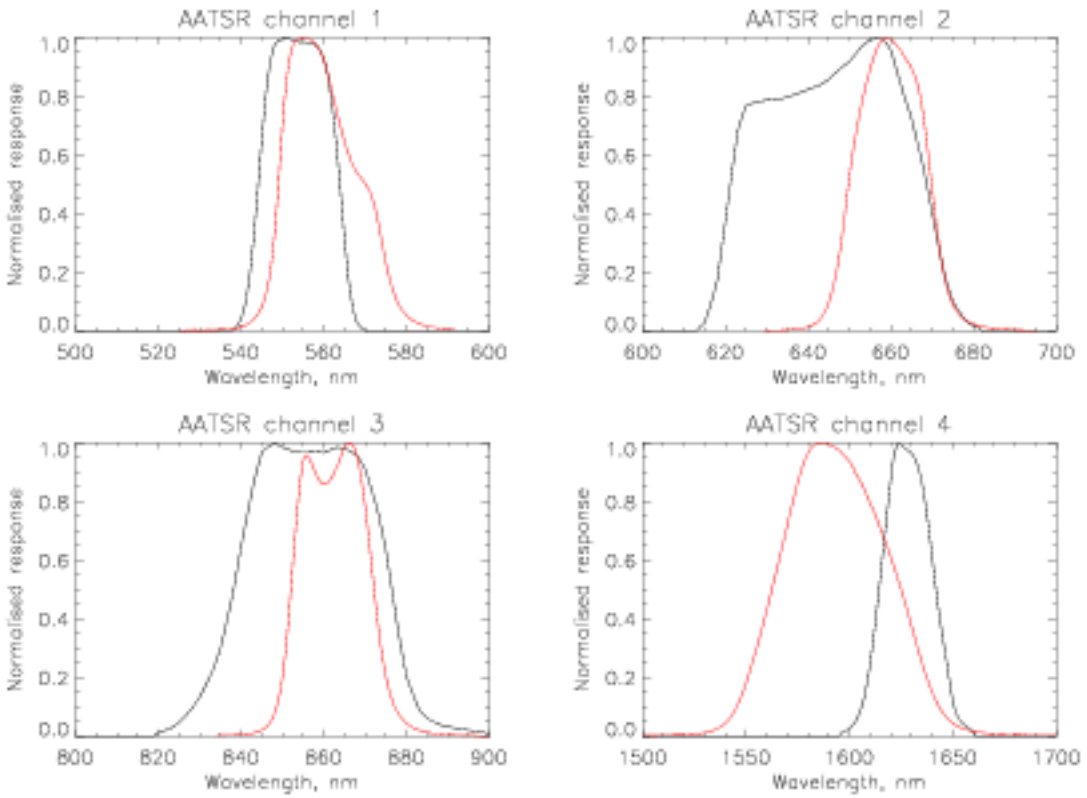

Figure 4.1: MODIS (black) and AATSR (red) normalised visible channel response functions. MODIS channels closest to corresponding AATSR channels are shown.

4.2.1 High-resolution spectral databases

The ASTER (Hook, 1999) and USGS (Clark et al., 2003) spectral libraries provide high-resolution data on the reflectance of a large variety of natural and man-made substances from laboratory measurements. A selection of 47 spectra from the ASTER library and 100 from the USGS library, chosen as representatives of surfaces that might be found in the real world (such as various types of vegetation, sand and soil, or man-made substances such as concrete), can be used to quantify the impact of the visible channel shapes shown in Figure 4.1 on the MODIS-reported white-sky albedo as compared to AATSR, and suggest routes to minimise that impact. This requires the following assumptions:

- The spectra chosen are representative of the spectra of surfaces that make up the real world.

- The difference found between integration of the spectra over the instruments' visible channel response functions can be related to differences between MODIS and AATSR white-sky albedos (as the spectra lack directional information), and any correction applied to the white-sky albedo can then also be applied to $R_{\mathrm{SBD,M}}$ and $R_{\mathrm{SLB,M}}$ to obtain R_{SBD} and R_{SLB}.

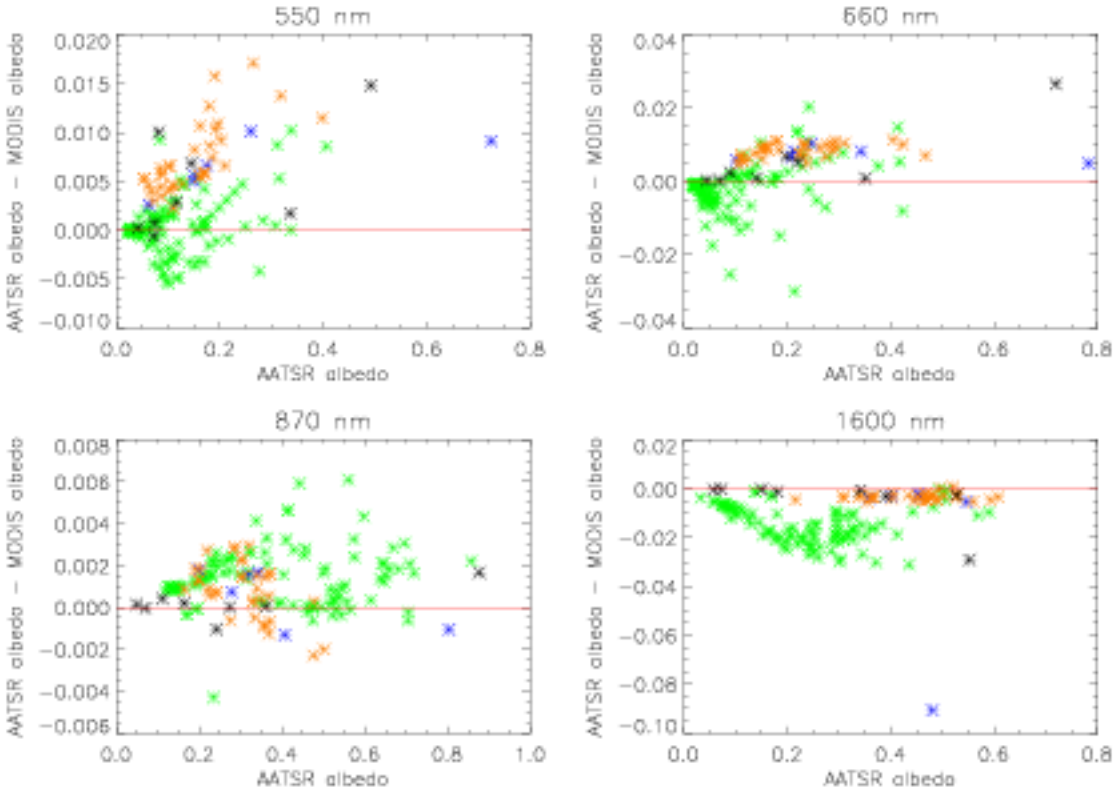

Figure 4.2: Differences between albedo reported by MODIS and AATSR as calculated for the 147 spectra selected from the ASTER/USGS libraries, for bands corresponding to AATSR visible channels, as a function of albedo calculated for AATSR. Green spectra are vegetation, blue sand, orange soil and black man-made substances.

The spectra were chosen with the first assumption in mind; the effect of surface inhomogeneity in a satellite pixel is discussed in Section 4.2.3. The second assumption requires that the satellite-observed reflectance is proportional to the surface's reflectance spectrum. For a Lambertian surface this should hold, but it is possible that for a non-Lambertian surface the difference between $R_{\mathrm{SBD,M}}$ and R_{SBD} will not be consistent over all viewing geometries. Unfortunately there is no way at present to test this; as the difference between the MODIS and AATSR visible channels is small to begin with, however, this is considered a minimal additional error.

These spectra can be grouped according to the type of surface that they represent: 35 soils, 6 sands, 8 man-made substances and 104 assorted types of vegetation. Spectra corresponding to snow or ice were not used as MOD43B1 (Collection 4) does not provide BRDF parameters over snow-covered surfaces. Convolving these spectra with the instruments' visible channel response functions leads to a set of 147 pairs of simulated 'observed' albedos. The differences between albedo reported from MODIS and AATSR calculated for these spectra is given as a function of the AATSR 'observed'

albedo in Figure 4.2.

The largest departures, as predicted from Figure 4.1, are for 660 nm and 1.6 μm, with a clear bias particularly visible at 1.6 μm. The mean absolute differences in Figure 4.2 are 0.0034 at 550 nm, 0.0058 at 660 nm, 0.0014 at 870 nm and 0.012 at 1.6 μm. Given the white-sky albedo derived from the MODIS product being accurate to within 0.02, these differences can be significant. Making some correction to account for these differering response functions is therefore desirable.

4.2.2 Singular value decomposition of spectrally-derived albedo

Theory of singular value decomposition

Singular value decomposition (SVD) is a statistical tool that may be applied to describe the variability in a set of data; in this case it is applied to allow the prediction of an AATSR-observed albedo for a spectrum given that observed by MODIS. The SVD of an m-by-n matrix M is defined as follows:

$$\mathbf{M} = \mathbf{U}\Sigma\mathbf{V}^* \tag{4.3}$$

This states that M may be expressed as a product of three matrices. U is an m-by-m unitary matrix and $\mathbf{V}^*$ an n-by-n unitary matrix. The matrix Σ is an m-by-n diagonal matrix with the nonzero elements in descending order. The transpose of $\mathbf{V}^*$, V, is the matrix of interest here as it describes the patterns of the variability in M. U contains a set of basis vectors which can be used to describe M. The columns of U and V are known as singular vectors. The elements of Σ, known as the singular values, describe how important each singular vector v in V is to the variability in M: the contribution of a singular vector in V to the total variability is its corresponding singular value as a proportion of the total value of the singular values. Hence the first singular vector contains the most information on the variability. Standard computational routines are available to compute the SVD in most programming languages.

The matrix M in this context consists of the simulated albedos reported by MODIS and AATSR for each spectrum; hence, it has one column for each spectrum used (147) and one row for each instrument channel. Results are presented using just the 4 AATSR channels and their equivalent MODIS bands (for a total of 8 rows), as well as three additional MODIS bands centred at 460 nm, 1.2 μm and 2.2 μm (for a total of 11 rows). In this discussion n_{M} refers to the number of MODIS

bands used, and n_{A} the number of AATSR bands (always 4).

SVD fitting and $\bar{R}_{\mathrm{SLW}}$ calculation

Element detailing visible channel	Singular vector							
	1	2	3	4	5	6	7	8
MODIS 550 nm	-0.191	0.150	-0.380	0.563	-0.0418	0.376	-0.296	-0.498
MODIS 660 nm	-0.215	0.376	-0.346	-0.355	0.0362	-0.650	-0.238	-0.294
MODIS 870 nm	-0.500	-0.477	-0.0985	-0.105	0.0471	-0.0412	0.593	-0.383
MODIS 1.6 μm	-0.415	0.292	0.468	0.109	-0.712	-0.0590	0.0381	0.00538
AATSR 550 nm	-0.195	0.168	-0.381	0.513	0.0569	-0.294	0.340	0.567
AATSR 660 nm	-0.217	0.404	-0.335	-0.500	-0.0418	0.583	0.188	0.227
AATSR 870 nm	-0.501	-0.479	-0.0930	-0.108	-0.0140	0.0408	-0.593	0.382
AATSR 1.6 μm	-0.399	0.322	0.491	0.0972	0.695	0.0587	-0.0351	-0.00822
Singular value	9.580	2.362	1.550	0.366	0.0638	0.0260	0.0115	0.008

Table 4.1: Values of the singular vectors obtained in the analysis of spectra, for the case of $n_{\mathrm{M}} = 4$ and $n_{\mathrm{A}} = 4$. Singular vectors are numbered in order of their significance (with 1 being most significant). The values of the elements are given to 3 significant figures here, but in calculations floating-point precision is used. The final row gives the corresponding singular value for each singular vector.

After computing the SVD as described above, the singular vectors V may be obtained. These may be used to make an estimation $\bar{R}_{\mathrm{SLW}}$ of the AATSR albedo given the MODIS albedo $R_{\mathrm{SLW,M}}$ at each wavelength, for any spectrum contained within M, by the following method:

- A simultaneous linear fit of the first n_{M} elements of each singular vector v_i (elements referred to as $v_{i,\mathrm{M}}$) to the n_{M} MODIS albedos $R_{\mathrm{SLW,M}}$ in question is performed. These elements of v describe the variability in the MODIS channels. This yields a series of n_{M} coefficients c_i, such that the MODIS albedo (at all n_{M} wavelengths) can be calculated exactly as the following:

$$R_{\mathrm{SLW,M}(1,2,\ldots,n_{\mathrm{M}})} = \sum_{i=1}^{n_{\mathrm{M}}} c_i v_{i,\mathrm{M}} \tag{4.4}$$

- These coefficients are used to estimate the n_{A} AATSR albedos using the same summation, but with the last n_{A} elements of the singular vectors (referred to as $v_{i,\mathrm{A}}$) instead of the first n_{M}. These describe the variability in the AATSR channels, and lead to the following prediction of the value of the AATSR albedo at all n_{A} wavelengths:

$$\bar{R}_{\mathrm{SLW}(1,2,\ldots,n_{\mathrm{A}})} = \sum_{i=1}^{n_{\mathrm{M}}} c_i v_{i,\mathrm{A}} \tag{4.5}$$

The more singular vectors used, the greater the proportion of the variability in the spectra will be captured by them and so the more accurate the prediction $\bar{R}_{\mathrm{SLW}}$ of R_{SLW}. A maximum of n_{M} singular vectors may be used; any more and the fit of $R_{\mathrm{SLW,M}}$ described above will have non-unique solutions.

In practice, from 4 to 7 singular vectors will be used with the same number of input MODIS albedos to provide 4 to 7 coefficients respectively to fit the AATSR-reported albedo at 4 wavelengths. If all singular vectors were used (8 for $n_{\mathrm{M}} = 4$ or 11 for $n_{\mathrm{M}} = 7$) then M would be reconstructed perfectly, but the values of the AATSR albedo would be needed in the input to fit the data and so there would be no predictive power. Hence there is a trade-off between the number of values which can be predicted and the accuracy of the predictions which can be made.

In this case if the SVD is performed using the 4 MODIS channels corresponding most closely to AATSR wavelengths, the first 4 singular vectors capture 99.2 % of the variability in the spectra. If all 7 MODIS wavelengths are used then this increases to 99.8 %. These values suggest that in both cases the fitting process described above should reconstruct the AATSR albedos well. The singular vectors obtained for $n_M = 4$ are tabulated in Table 4.1, together with their singular values; the similarity between elements corresponding to equivalent channels is an indication of the fact that the channels are generally observing similar albedos.

Quality of SVD fit

The value of the fit and estimation process can be quantified by comparing the error in the AATSR albedo that would result from using the MODIS albedo as an estimate, $R_{\mathrm{SLW}} - R_{\mathrm{SLW,M}}$, to the error resulting from using the SVD-predicted estimate, $R_{\mathrm{SLW}} - \bar{R}_{\mathrm{SLW}}$. This first quantity is referred to as the expected difference e_{e}, and the second the AATSR reconstruction error e_{r}: if $|e_{\mathrm{r}}| < |e_{\mathrm{e}}|$ then the SVD fit has enabled an improved estimate of R_{SLW}.

Plots of e_{e} against e_{r} are shown for the cases of $n_M = 4$ and $n_M = 7$ in Figure 4.3. In both cases, and for all classes of surface, the mean value of $|e_{\mathrm{r}}|$ is smaller than the mean of $|e_{\mathrm{e}}|$. The impact of using 7 singular vectors rather than 4 is shown directly in Figure 4.4. For $n_{\mathrm{M}} = 4$, the mean value of $|e_{\mathrm{r}}|$ is 0.0011 at 550 nm, 0.002 at 660 nm, 0.00097 at 870 nm and 0.0048 at 1.6 μm. For $n_{\mathrm{M}} = 7$, these values are 0.00087 at 550 nm, 0.0016 at 660 nm, 0.00060 at 870 nm and 0.0038 at 1.6 μm. These are all well below the 0.02 uncertainty quoted on $R_{\mathrm{SLW,M}}$ although individual spectra may

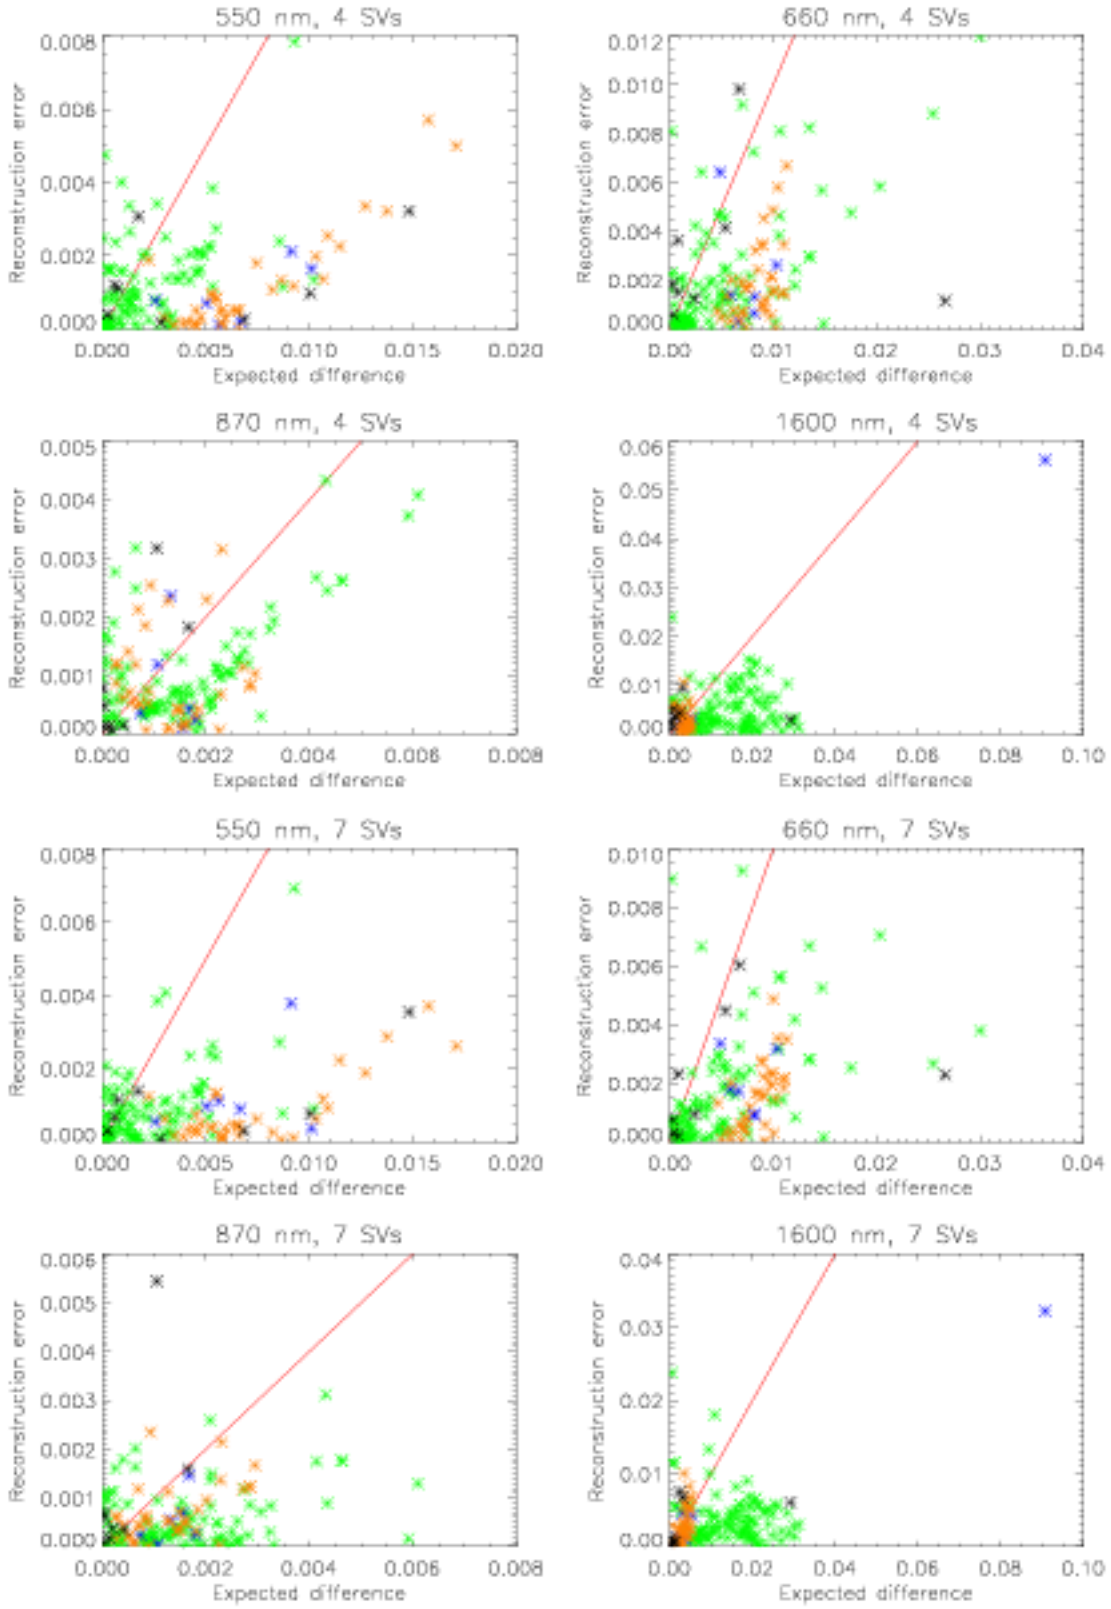

Figure 4.3: Expected difference in albedo e_e against reconstruction error ϵ_r for the 147 spectra used in the analysis, at AATSR visible channel wavelengths. Presented in units of absolute albedo, using 4 MODIS channels and singular vectors (top 4 graphs) and 7 MODIS channels and singular vectors (bottom 4 graphs). Green spectra are vegetation, blue sand, orange soil and black man-made substances. The $1:1$ line is in red.

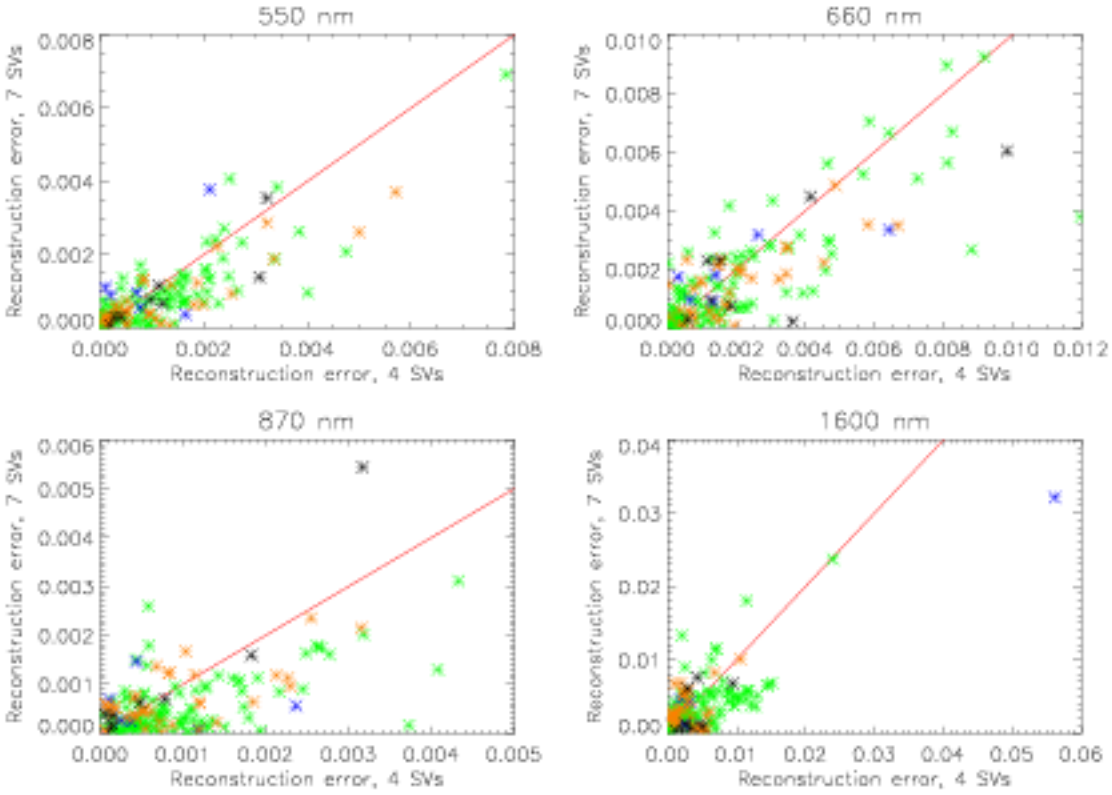

Figure 4.4: Reconstruction error e_r for the case of using 7 MODIS bands' singular vectors as opposed to using only 4, at AATSR visible channel wavelengths. Presented in units of absolute albedo. Green spectra are vegetation, blue sand, orange soil and black man-made substances. The $1:1$ line is in red.

have reconstruction errors larger than this value.

Figure 4.5 shows the reconstruction errors for both cases as a function of the true AATSR-observed albedo R_{SLW}. A comparison with Figure 4.2 shows that the use of singular vectors gives reconstruction errors with smaller overall biases (particularly at 1.6 μm) and less clustering of similar spectral types than direct use of the MODIS data.

In summary, the use of SVD techniques generally allows an improved prediction of the AATSR white-sky albedo for a surface as compared to using the white-sky albedo derived from MOD43B1 as-is. The fitting process results in estimates of R_{SLW} generally within the uncertainty on $R_{\mathrm{SLW,M}}$, with the largest discrepancies at 660 nm and 1.6 μm.

4.2.3 Composite spectra

As real surfaces are unlikely to conform exactly to those spectra present in the libraries used, it is important to quantify how well the singular vectors in Table 4.1 can predict R_{SLW} for surfaces consisting of a composite of these spectra. Generation of an ensemble of a million composite spectra,

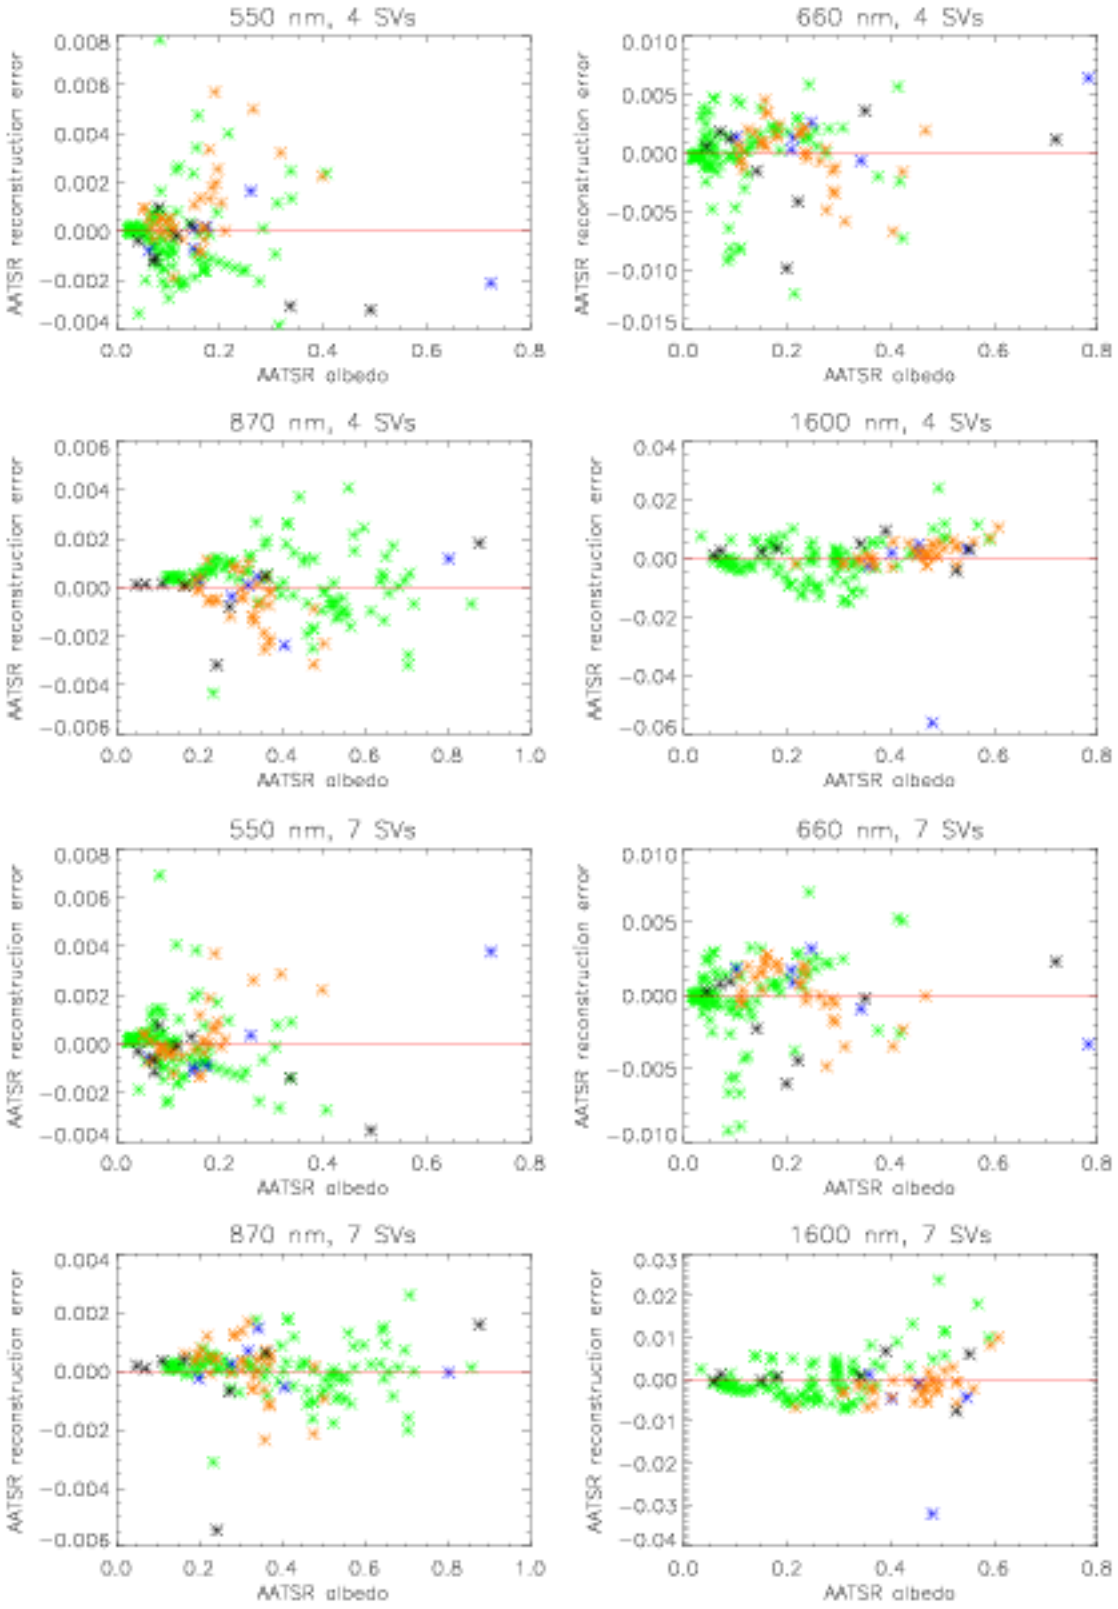

Figure 4.5: AATSR-observed albedo against reconstruction error for the 147 spectra used in the analysis, at AATSR visible channel wavelengths. Presented in units of albedo, using 4 MODIS channels and singular vectors (top 4 graphs) and 7 MODIS channels and singular vectors (bottom 4 graphs). Green spectra are vegetation, blue sand, orange soil and black man-made substances.

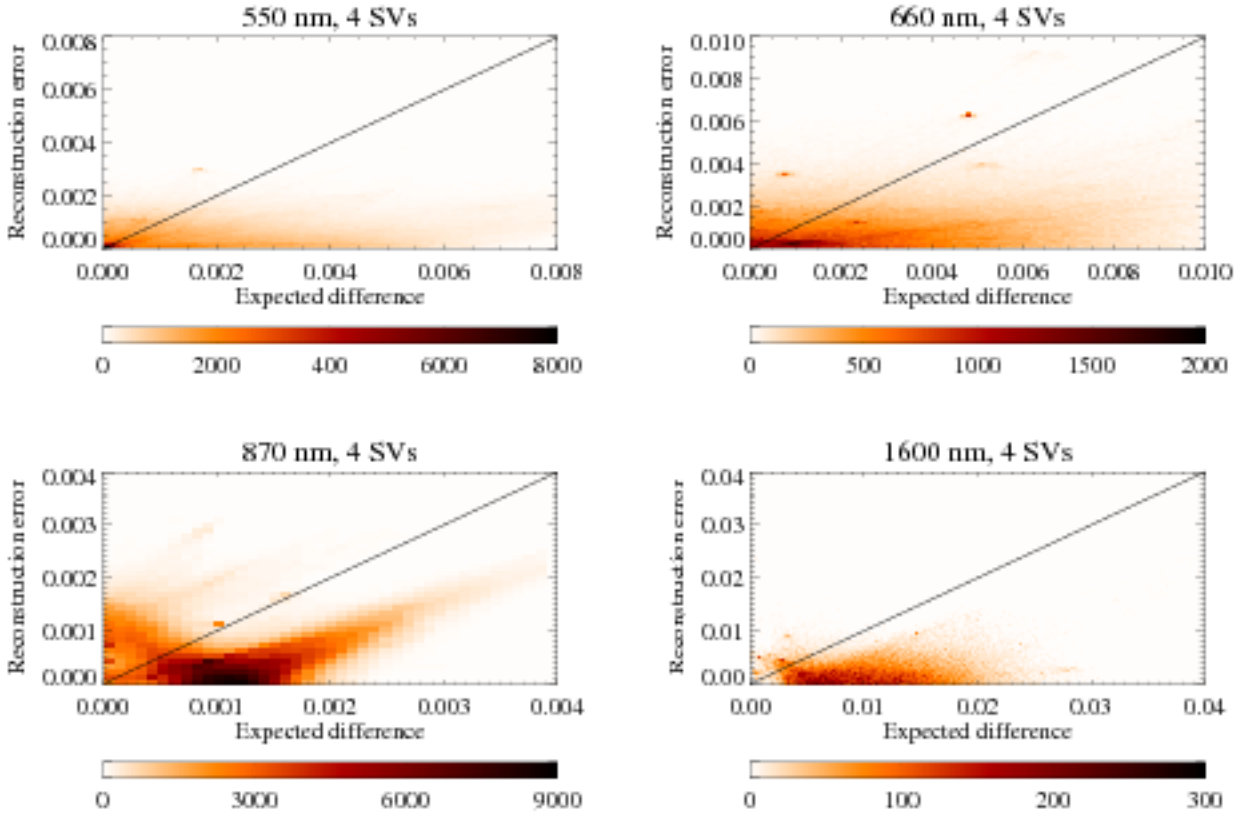

Figure 4.6: Scatter density plots of expected difference in albedo e_a against reconstruction error e_r for the ensemble of composite spectra used in the analysis, at AATSR visible channel wavelengths. Presented in units of absolute albedo, using 4 MODIS channels and singular vectors. The bin size for density calculations is 0.0001. The 1:1 line is in black.

each calculated from a linear combination of between 2 and 4 randomly-weighted and randomly-picked spectra, yields an ensemble of a million matched sets of R_{SLW} and $R_{\mathrm{SLW,M}}$. From there, distributions of e_a and e_r are calculated after making the SVD fit; a million ensemble members ensures statistics robust to 2 significant figures.

Except at 1.6 μm, the mean absolute expected differences for the ensemble of composite spectra are smaller than the corresponding values for the group of 147 single spectra. This suggests that to some extent the impact of the differences between MODIS and AATSR is cancelled out by the different materials in the composite spectra. The mean absolute values are 0.0031 at 550 nm, 0.0041 at 660 nm, 0.0012 at 870 nm and 0.012 at 1.6 μm.

The SVD technique proves effective in the case of composite spectra; for $n_M = 4$ the mean absolute e_r are 0.00086 at 550 nm, 0.0015 at 660 nm, 0.00066 at 870 nm and 0.0034 at 1.6 μm. A slight improvement is again found by using all 7 MODIS singular vectors, giving values of 0.00065

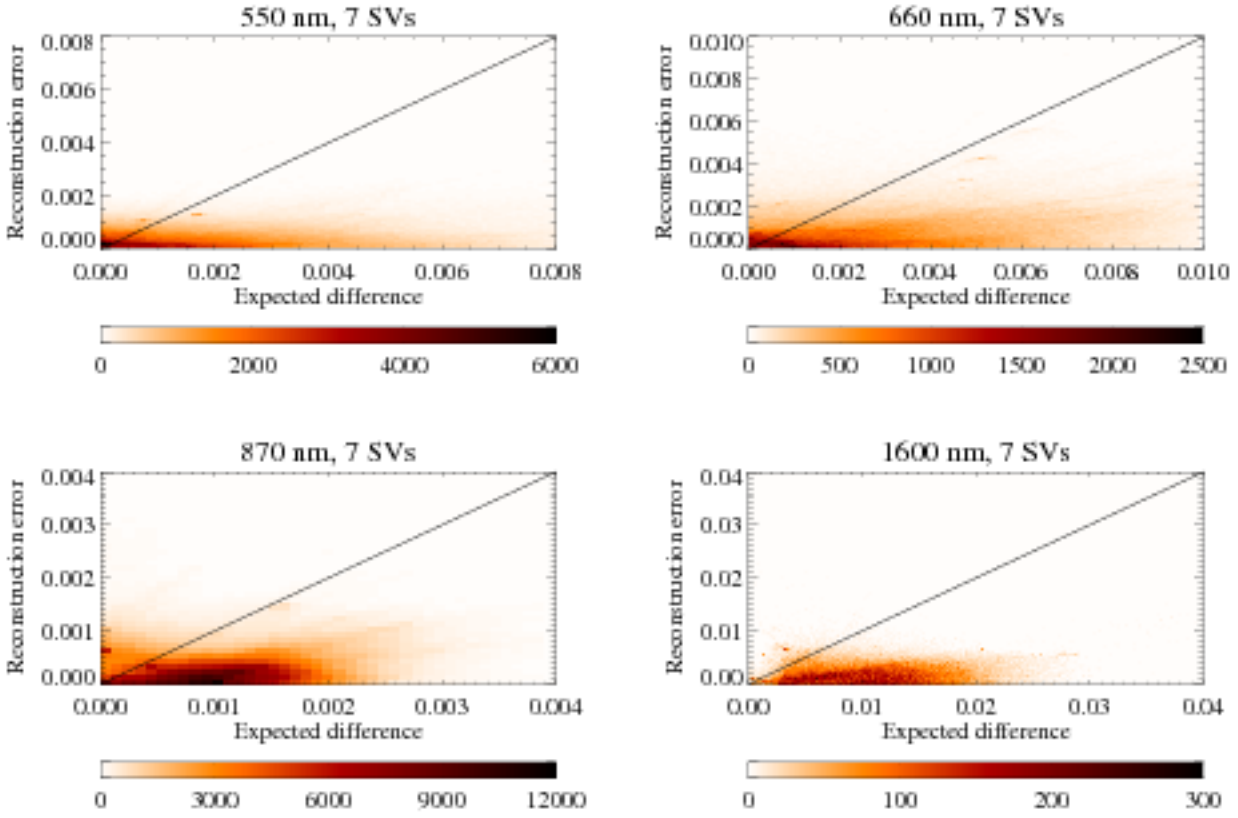

Figure 4.7: Scatter density plots of expected difference in albedo e_e against reconstruction error e_r for the ensemble of composite spectra used in the analysis, at AATSR visible channel wavelengths. Presented in units of absolute albedo, using 7 MODIS channels and singular vectors. The bin size for density calculations is 0.0001. The 1:1 line is in black.

at 550 nm, 0.0012 at 660 nm, 0.00046 at 870 nm and 0.0027 at 1.6 μm.

Figures 4.6 and 4.7 show that e_r is generally independent of e_e. This suggests that the SVD technique performs equally well for all surface types. Additionally, the reconstruction error is smaller than the expected difference in the majority of cases, and always lower than the $R_{\mathrm{SLW,M}}$ accuracy threshold of 0.02. The reconstruction error is also independent of R_{SLW} (not shown in plots), with small correlation coefficients $|r| < 0.1$ for all wavelengths for $n_{\mathrm{M}} = 4$ or 7. Given these independencies it is useful to consider the frequency distributions of the reconstruction errors themselves. These are shown in Figure 4.8; for both $n_{\mathrm{M}} = 4$ and $n_{\mathrm{M}} = 7$ the distributions drop off rapidly as compared to the distribution of expected differences. Additionally, the distributions for $n_{\mathrm{M}} = 4$ and $n_{\mathrm{M}} = 7$ have very similar forms and close means.

Similar results are obtained if the number of spectra in each composite spectrum is changed (such as allowing 2-6 spectra rather than 2-4); the higher the maximum number of spectra allowed,

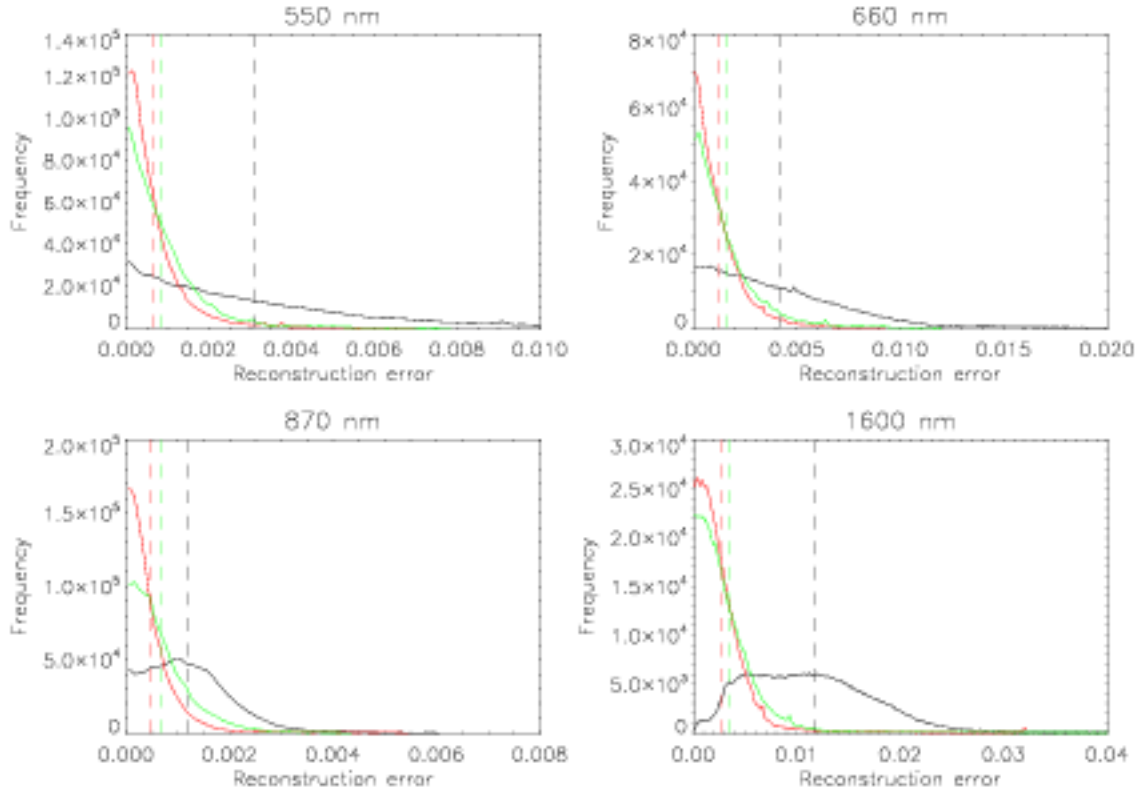

Figure 4.8: Frequency distributions of reconstruction error for the ensemble of composite spectra used in the analysis, at AATSR visible channel wavelengths. Presented in units of absolute albedo. Results in black indicate the error arising from using the MODIS albedo unadjusted. Green indicates the use of 4 singular vectors and MODIS wavelengths, and red the use of 7. Vertical dashed lines show the means of the distributions. The bin size for frequency calculations is 0.0001.

the smaller the expected differences and reconstruction errors. The ratio of the mean of e_r to the mean of e_a remains similar.

4.2.4 Application to the retrieval

In the preprocessing of the aerosol retrieval $R_{\mathrm{SBD,M}}$, $R_{\mathrm{SLB,M}}$ and $R_{\mathrm{SLW,M}}$ are obtained from the MOD43B1 product for each pixel. The singular vectors obtained from the 147 spectra are used to calculate $\bar{R}_{\mathrm{SLW}}$, which is used as the *a priori* value of R_{SLW} in the retrieval. The ratio $\bar{R}_{\mathrm{SLW}} : R_{\mathrm{SLW,M}}$ is calculated for each pixel and then applied as a scaling factor to $R_{\mathrm{SBD,M}}$ and $R_{\mathrm{SLB,M}}$ to obtain the AATSR BRDF and black-sky albedo. This preserves the ratio between the three reflectances as obtained from MOD43B1.

Albedo is calculated at full instrumental resolution; as the MODIS grid is downsampled to approximately 2.5×2.5 km resolution some MODIS data may be replicated in the AATSR superpixel.

The unique values (typically up to around 20) within each superpixel are then averaged to give the superpixel mean R_{SBD}, R_{SLB} and R_{SLW} used in the retrieval. The singular vectors for the case $n_{\mathrm{M}} = 4$ are used as the improvement gained by using 7 singular vectors instead is small when compared to the 0.02 accuracy of the MODIS white-sky albedo. Only requiring the BRDF parameters for 4 wavelengths as opposed to the 7 amounts to a substantial saving in space required to store the data.

If for any wavelength the ratio $\bar{R}_{\mathrm{SLW}} : R_{\mathrm{SLW,M}}$ is larger than 1.5 or smaller than 0.5 in a superpixel, then it is deemed a poor fit of the singular vectors and no adjustment is made. This constraint helps to prevent unphysical estimates of R_{SBD}, R_{SLB} and R_{SLW}.

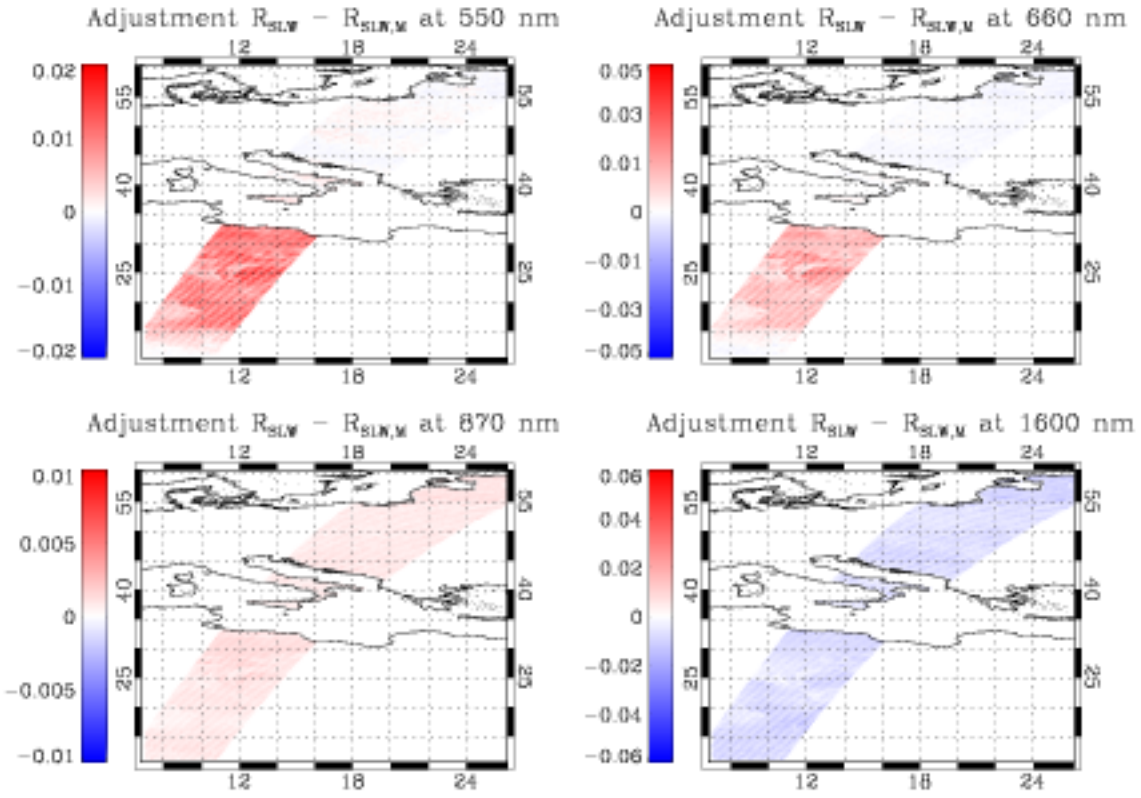

Figure 4.9: Adjustment $R_{\mathrm{SLW}} - R_{\mathrm{SLW,M}}$ made to the MODIS white-sky albedo at AATSR visible channel wavelengths. Presented in units of absolute albedo, using 4 MODIS channels and singular vectors. The scene covers continental Europe and parts of Northern Africa for an AATSR swath on September 5, 2004.

An example of this in practice is shown in Figure 4.9. Well-defined areas of differing land surface cover result in adjustments made to $R_{\mathrm{SLW,M}}$ of differing signs and magnitudes, with a strong agreement between the sign of the adjustment and the sign of the expected AATSR-MODIS difference for spectra in different categories (such as vegetation, sand and soil) shown in Figure 4.2.

Histograms of the magnitude of the adjustment for this scene are presented in Figure 4.10; the wavelength-dependence of the magnitude of the adjustment (smaller at 550 nm and 870 nm, larger

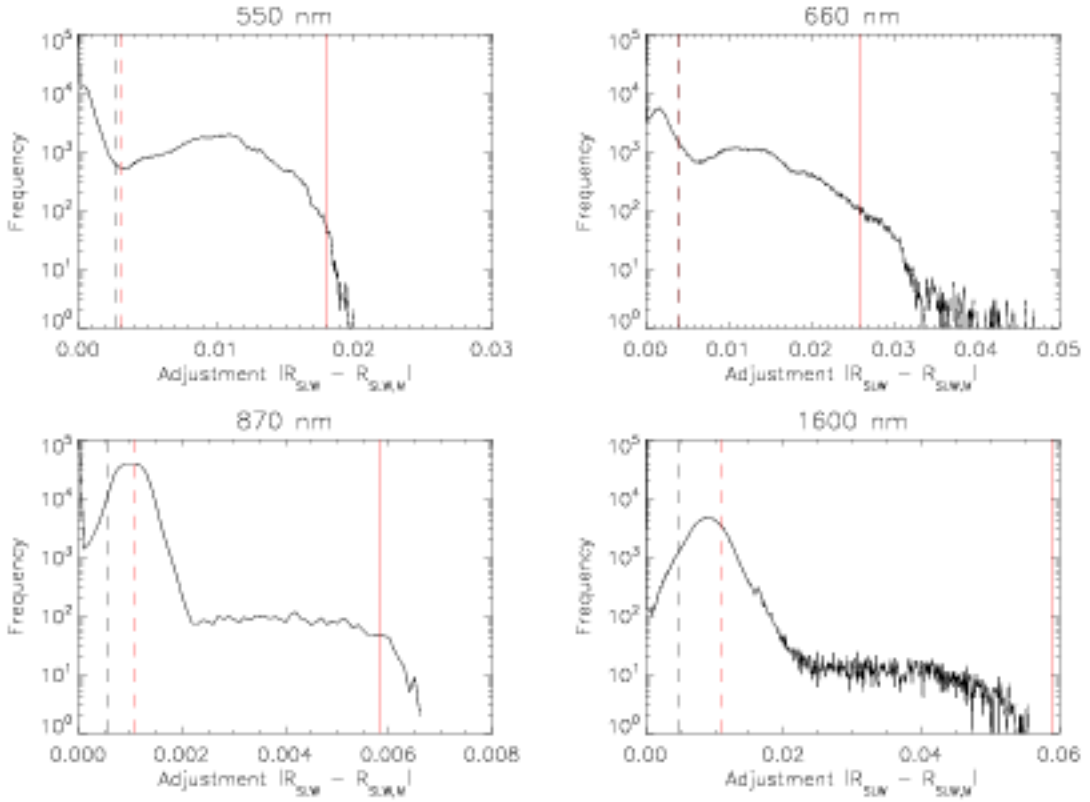

Figure 4.10: Frequency histograms of the magnitude of the adjustment $|R_{\mathrm{SLW}} - \tilde{R}_{\mathrm{SLW,M}}|$ made to the MODIS white-sky albedo at AATSR visible channel wavelengths for the scene shown in Figure 4.9. Presented in units of absolute albedo, using 4 MODIS channels and singular vectors. The bin size for density calculations is 0.0001. The dashed black line indicates the mean of the distribution. The dashed red line indicates the mean adjustment of the ensemble of a million composite spectra discussed in Section 4.2.3, and the solid red line the maximum adjustment of the million-member ensemble.

at 660 nm and 1.6 μm) is also consistent with the single-spectrum results. Additionally, the means of the distribution of adjustments for this scene and in the ensemble of composite spectra from Section 4.2.3 are similar to each other: they would not be expected to match exactly, as the mean of the scene depends upon the composition of the land surfaces in it, but the fact that they are similar suggests that the spectra provide a reasonable representation of land surfaces.

Together these figures provide confidence that the method, which was shown in earlier sections to provide results valid for individual and composite spectra, also yields physically reasonable estimates of the AATSR white-sky albedo when applied to real MODIS BRDF data.

4.3 Uncertainties and correlations

The retrieval algorithm requires knowledge of the uncertainty on *a priori* parameters and the correlations between them. As with the sea surface, the value of element index i, j in this covariance matrix $\mathbf{S}_x$ is the product of the uncertainties on parameters i and j and the correlation between them.

4.3.1 Inter-wavelength albedo correlations

Generating R_{SLW} at full instrumental resolution (as opposed to superpixel resolution) enables the calculation of correlation coefficients $r_{(i,j)}$ for each pair of wavelengths for each superpixel from the unique MODIS albedos inside the superpixel. If fewer than 10 unique MODIS values are found in a superpixel then the correlations are set to zero. This is because, with a small number of data points, the derived correlation may not be reliable. Only the unique values are used to determine correlations as repeat values would over-emphasise the strength of any existing correlation.

In the case of the retrieval being performed at full AATSR resolution, the assumed correlations become zero.

4.3.2 *A priori* errors for white-sky albedo

Neither the stated errors of the MODIS white-sky albedo nor the absolute AATSR albedo reconstruction errors e_{r} depend upon the value of the white-sky albedo in question. Therefore the *a priori* uncertainty on this retrieval parameter is treated as being the sum of these two sources of error. The *a priori* covariance matrix element *(i,j)* is the product of the uncertainty on variable i, the uncertainty on j and the correlation between them. For the land surface, this is expressed by:

$$\mathbf{S}_{x(i,j)} = e_{\mathrm{r}(i)} e_{\mathrm{r}(j)} \delta_{(i,j)} + S_{R_{\mathrm{SLW,M(I)}}} S_{R_{\mathrm{SLW,M(}}} r_{i,j} \tag{4.6}$$

In the above the SVD reconstruction errors e_{r} for 4 singular vectors and the 147 single spectra are used, namely 0.001 at 550 nm, 0.002 at 660 nm, 0.001 at 870 nm and 0.004 at 1.6 μm. $\delta_{(i,j)}$ is the Kronecker delta; reconstruction error was found to be uncorrelated between wavelengths so the SVD error term vanishes for the off-diagonal elements of $\mathbf{S}_x$. $S_{R_{\mathrm{SLW,M(I)}}}$ is the MODIS albedo uncertainty (0.02). The final term, $r_{i,j}$, is the correlation between the superpixel $R_{\mathrm{SLW,M}}$ at the two wavelengths, as described above.

Analagously to decreasing the measurement error in $\mathbf{S}_y$ resulting from superpixelling, decreasing the *a priori* uncertainty on R_{SLW} by a factor of the square root of the number of unique MODIS values in the superpixel has been attempted. This approach was abandoned as it was found to substantially increase the retrieval cost, and additionally led to a poorer match of retrieved aerosol optical depth with AERONET data (in terms of the correlation coefficient and root mean square difference between the datasets). This suggests that this was too strong a constraint and there are other sources of error which should be taken into account, such as temporal variability of albedo.

4.3.3 Temporal variability of albedo

The MOD43B1 product provides the mean albedo from observations over a 16-day period. As the AATSR retrieval is for an instantaneous observation, temporal variability of the surface over this period may result in the instantaneous true albedo being significantly different from the 16-day mean. Unfortunately it is difficult to infer the variability of surface albedo within the period from the data itself, which is why the retrieval uses the quoted uncertainty of 0.02.

One method to probe the variability is to look at the variabilty of surface albedo between consecutive 16-day periods. Figure 4.11 shows a time series of the MODIS white-sky albedo during 2002 for a vegetated region in Southern France, as well as the change in albedo between each 16-day period. From this it can be seen that, while the variability between each period is often smaller than the quoted MODIS uncertainty of 0.02, in many cases (particularly at longer wavelengths) it is not. Furthermore, the variability itself is not constant with time.

Figure 4.12 shows an analagous series of plots for an arid region on the Moroccan-Algerian border. Variability here is generally smaller. Together these plots illustrate that the temporal variability of surface albedo is complicated and a factor of surface type as well as local meteorology. Consequently there is no simple way to account for the temporal variability of albedo in the *a priori* error. One possibility would be to estimate the variability on a pixel-by-pixel basis based on the data for the preceding and succeeding 16-day periods, but this has not yet been attempted as it involves a greatly increased computational overhead.

Another aspect is the variability of albedo within a day. Although difficult to quantify, this is not anticipated to cause significant error—the BRDF model should be able to capture changes in surface reflectance resulting from differing solar geometries, and as AATSR and the MODIS platforms orbit

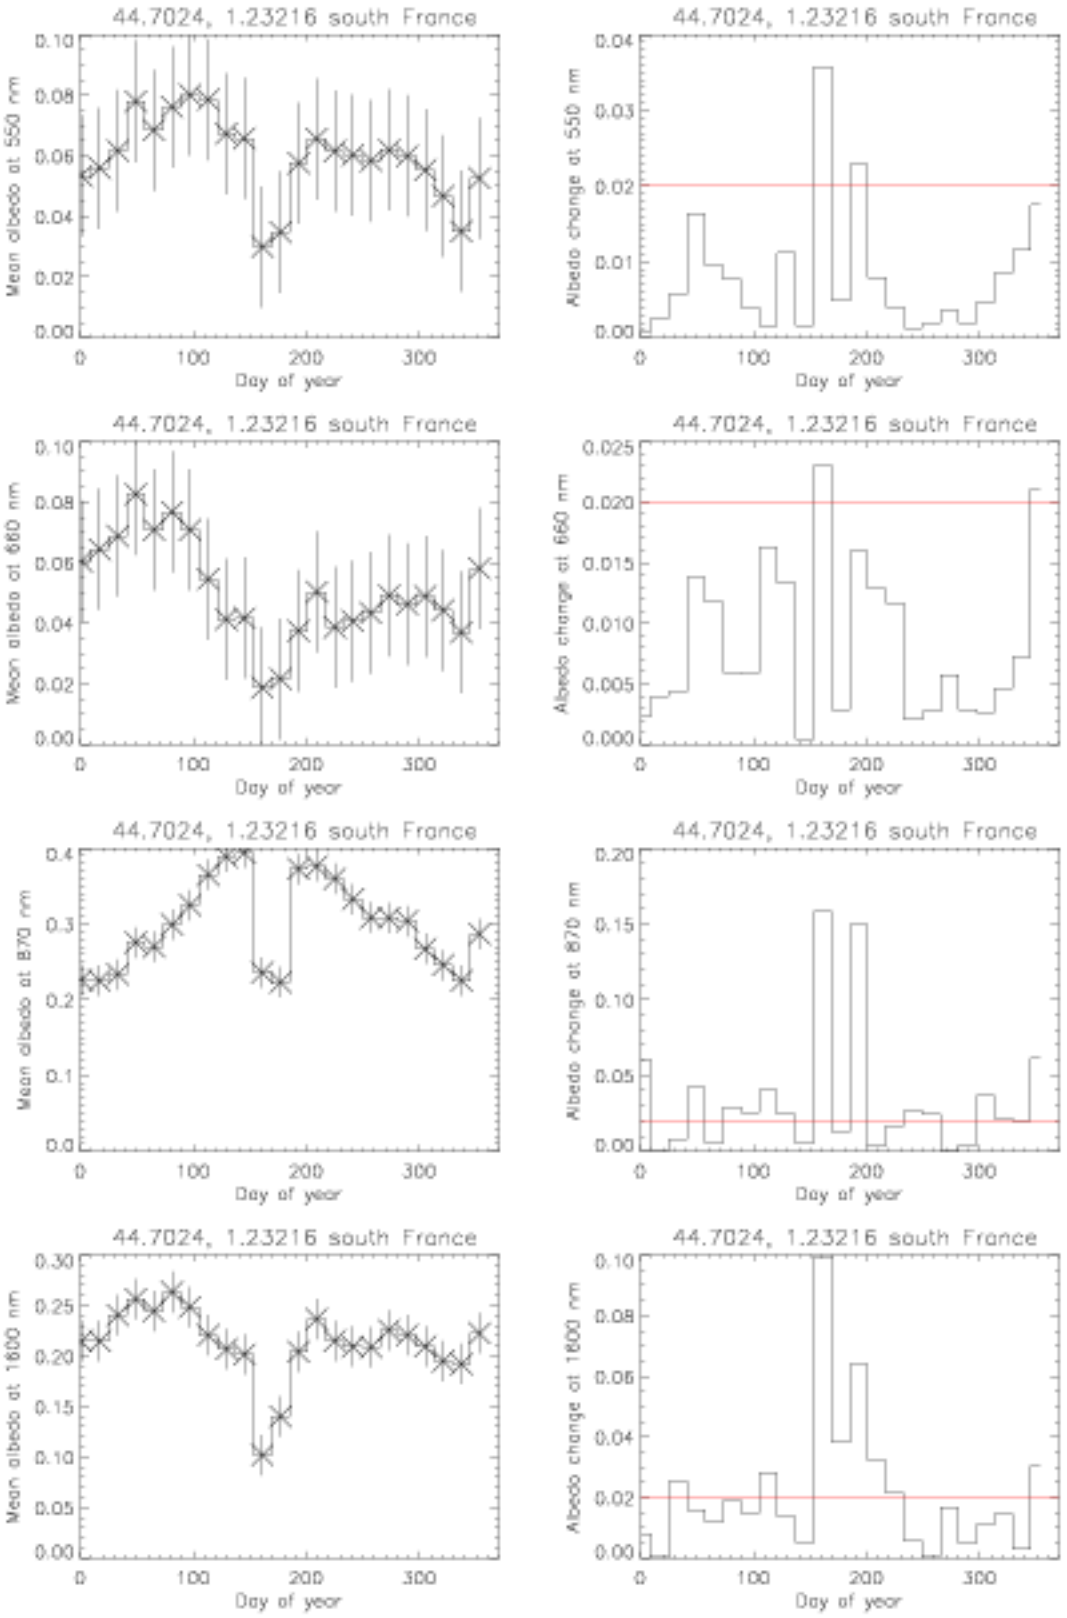

Figure 4.11: Surface albedo for a 10 km× 10 km superpixel centred on 44.7024° N, 1.23216° E in Southern France. The left-hand side shows the mean MODIS R_{SLW} for the superpixel for each 16-day period during 2002, for all 4 bands near AATSR channels. Error bars are the quoted MODIS uncertainty of 0.02. The right-hand side shows the difference between the albedo for each period and the albedo of the previous 16-day period (with the first point being the difference from the year's final 16-day period). The horizontal red line shows the MODIS uncertainty of 0.02 on the 16-day mean albedo.

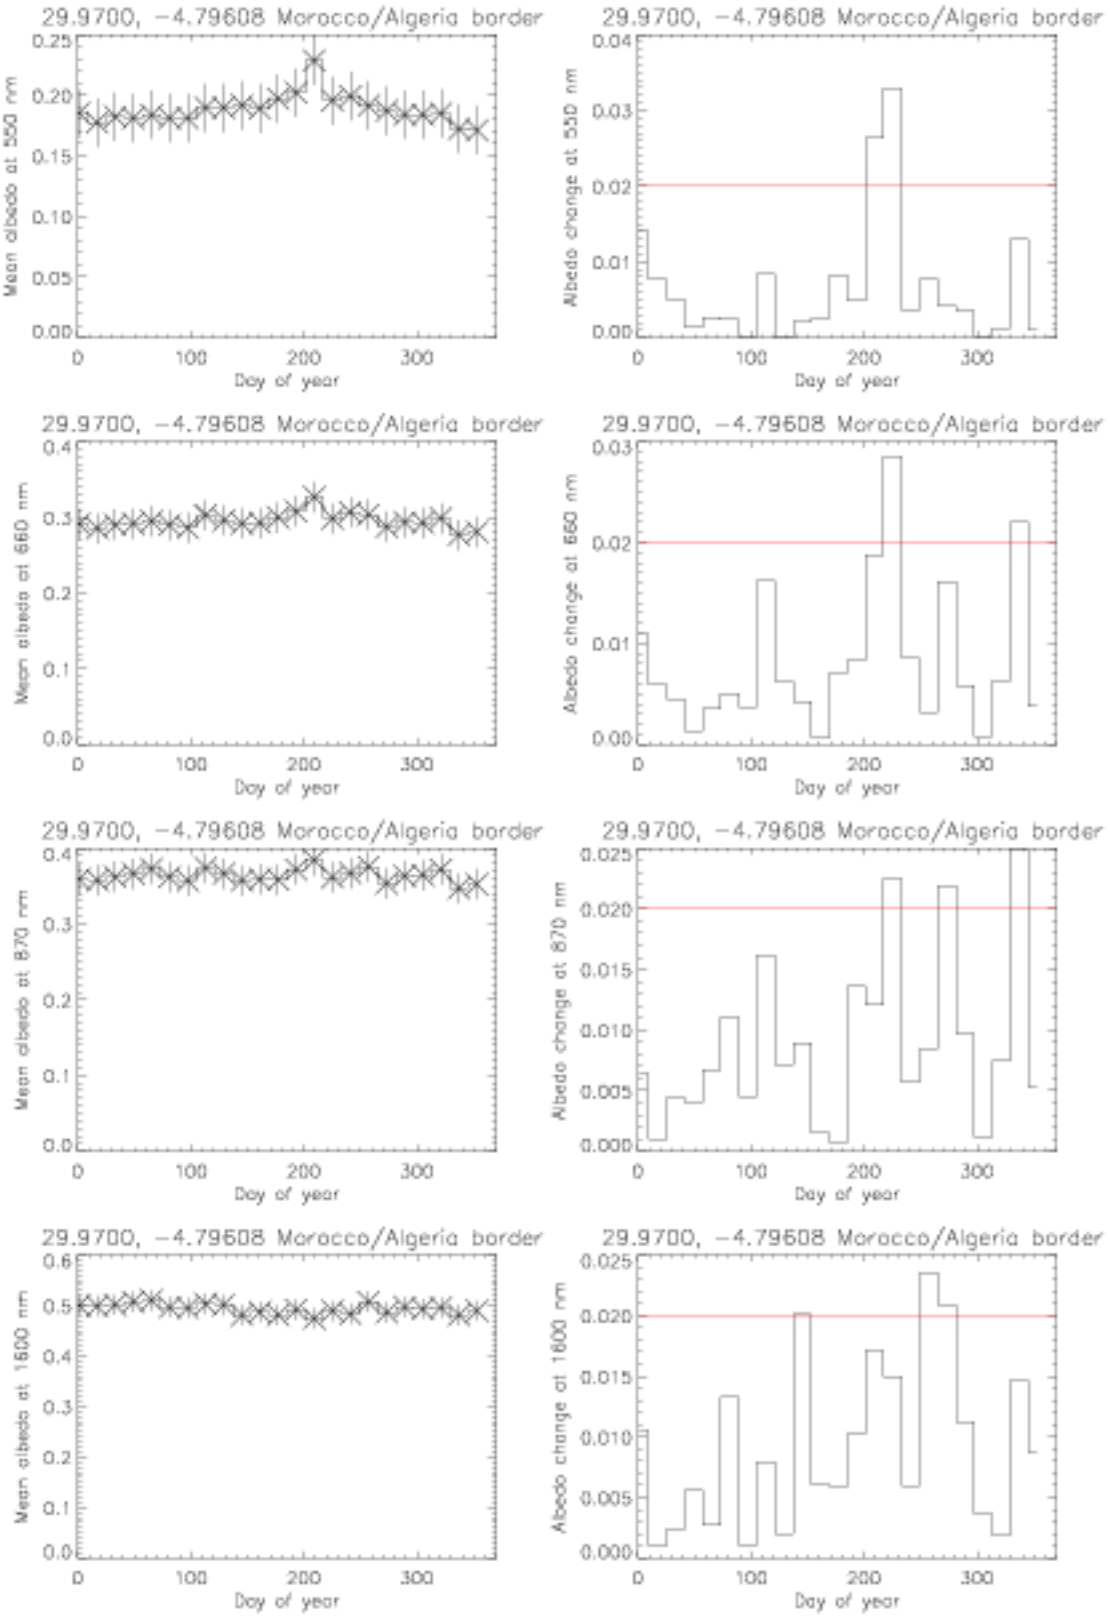

Figure 4.12: Surface albedo for a 10 km × 10 km superpixel centred on 29.9700° N, 4.79608° W near the border between Morocco and Algeria. The left-hand side shows the mean MODIS R_{SLW} for the superpixel for each 16-day period during 2002, for all 4 bands near AATSR channels. Error bars are the quoted MODIS uncertainty of 0.02. The right-hand side shows the difference between the albedo for each period and the albedo of the previous 16-day period (with the first point being the difference from the year's final 16-day period). The horizontal red line shows the MODIS uncertainty of 0.02 on the 16-day mean albedo.

at similar times of day the white-sky albedos provided by MODIS are derived at a similar local time to the AATSR overpass.

In summary, neglecting to account for temporal variability of albedo means that there is scope for further refinement of the treatment of *a priori* albedo errors. However, as in many cases the variability of surface albedo between consecutive 16-day periods is smaller than the quoted uncertainty on the MODIS data used, this is not suspected to be a significant source of error. The implications of this are that the *a priori* component of cost at the retrieval solution may be larger than expected over surfaces showing high temporal variability, and data about where this occurs could be a useful first step into quantifying where temporal variability may be important over the globe.

Chapter 5

Results and validation of the new algorithm

The development and application of a new algorithm should be complemented by a study into how well it performs. This validation gives confidence that the results are believable and can provide an indication into where it performs well and where it is poor. Furthermore, a validation exercise can provide information on what parts of the scheme contain scope for further improvement, and perhaps suggest avenues for improvement in other datasets.

The validation exercises presented here are both internal and external. The internal validation of the algorithm examines the statistics of the output data. External validation consists of a comparison with coincident AERONET ground-based measurements in different parts of the world, and a global comparison of monthly mean data with aerosol optical depth derived from MISR data. Statistics and AERONET comparisons with the previous nadir-only version of this algorithm are also discussed, to illustrate the impact of the improvements made to the new algorithm.

In all cases, AATSR data from the month of September 2004 were used. This month was chosen because all relevant instrumental and ancillary data needed for the retrieval, as well as the AERONET and MISR data, were available.

5.1 Intercomparison with AERONET data

5.1.1 Methodology

Validation of retrieval results against an independent set of measurements is an important method of determining the reliability of the derived data products. To this end, AATSR data for the month of

September 2004 was processed in the regions of ground-based AERONET stations, resulting in 58 cloud-free coincidences between AATSR and AERONET. Dates, times and station locations of these coincidences are presented in Appendix B. AERONET is a standard dataset for validation of satellite retrievals of aerosol optical depth, and here the validation method of Ichoku et al. (2002) is adopted for use. In all cases, Level 2.0 AERONET data (cloud-screened and quality assured) is used.

A direct comparison of AERONET and satellite retrieved data is difficult; the ground-based station is a point measurement, whereas the satellite retrieval pixel is much larger in size (here, 10 km $\times$ 10 km). Hence the ground measurement may not be representative of the region as a whole. Additionally there will be some time delay between the ground measurement and satellite overpass, during which conditions may change. Ichoku et al. (2002), as well as this analysis, attempt to account for these issues by assuming that the variability in space and time of the two datasets is correlated through wind. Additionally, AERONET stations thought to be unrepresentative of the larger region they are in (Kinne, 2007) are removed from the analysis.

The validation then proceeds via the following method:

- The pixel containing the AERONET station is identified.

- The retrieved τ_{550} and τ_{870} within 20 km of this pixel (after Ichoku et al., 2002) are taken, with the number of good retrievals, their mean and standard deviation noted. Quality control is applied to the AATSR retrieved data, as described later. The standard deviation is used as opposed to a retrieval error as frequently variability in the AATSR measurements will be larger than the retrieval error (typically small and fairly continuous), suggesting a change in the aerosol properties over the box. Thus using just the retrieval error would overestimate the extent to which the AATSR data would be expected to agree with AERONET; using the standard deviation in this case includes both the variability resulting from the random error on the AATSR retrieval and the true aerosol variability.

- AERONET measurements for the given ground station within 30 minutes of the satellite overpass are extracted from the Level 2.0 files. Again following Ichoku et al. (2002), this time is used assuming an aerosol transport rate of 50 km h^{-1}, as was found to be typical in that work. Where τ_{550} and τ_{870} are not directly recorded for a station, they are estimated from optical depths at other wavelengths and the Ångstrøm exponents provided in the Level 2.0 data.

Again, the number of good retrievals, their mean and standard deviation are recorded.

The quality control applied to the AATSR data are as follows:

- The retrieval cost, J, must be under 10. This ensures poor-quality fits are discounted.
- The number of iterations must be between 2 and 25, to ensure that the retrieval converged and did not sit on the first guess value.
- The retrieved τ_{550} must be larger than 0.01 and smaller than 5, as retrieved states lying exactly on the limit are not reliable. Similarly, r_e must be between 0.01 and 5 μm; although the retrieval permits r_e up to 10 μm, large radii are generally evidence of cloud contamination.
- The retrieved 550 nm surface albedo must be greater than 0 and less than 0.2. This is because the retrieval is less reliable over bright surfaces, and a bright albedo can also be an indication of cloud contamination. Note that no such constraints are placed upon retrieved albedo at the other wavelengths, as the land surface is brighter and more variable at longer wavelengths so this would not be an efficient test.
- The fraction of cloudy pixels identified within the superpixel during the preprocessing stage of the retrieval (Section 2.5) must be less than 0.5, to minimise residual cloud contamination and errors arising from, for example, cloud adjacency effects.

The choice of an appropriate cost threshold to cut the data is an important one; too low and good retrievals are missed, leading to a low number of comparisons, but too high and poor quality data included will decrease the overall quality of the comparison.

With 8 measurements and 5 constrained retrieved parameters, the expected mean retrieval cost is approximately 1.6. Assuming residuals on measured data and *a priori* variables to have Gaussian distributions and their errors to be set appropriately, it follows that most (over 99 %) good retrieved pixels should have costs less than 3 times this mean (approximately 5). The probability density function (PDF) for cloud-contaminated pixels should have a somewhat higher mean cost. Therefore, it would be expected that there would be little change in the statistics of the comparison for cost thresholds set between 5 and near typical values for poorly-retrieved cloud-contaminated pixels.

The change in the statistics of the comparison for the continental aerosol class is shown in Figure 5.1. The general trends are as expected; many matches are obtained by a cost threshold of 5, but

the number of matches continues to increase until a threshold of approximately 15. The fact that the best-fit statistics exhibit only a small change over this range is encouraging, as it suggests that the extra matches obtained by the increase in cost threshold are largely well-retrieved points. As a result a cost threshold of 10 has been chosen, as it provides a high number of matches but is low enough to weed out most poorly-retrieved pixels.

The AERONET comparison was carried out for τ_{550} and τ_{870} for each of the five aerosol classes in use. Additionally, a comparison using a simple attempt at speciation (by considering the lowest-cost aerosol model for each pixel) is presented.

No comparison is presented over the sea. AERONET stations are all located on the land; there are some coastal stations, but these are naturally inhomogeneous locations where the retrieval is not expected to perform as well. Additionally, the sea surface model is designed for open water and changes in, for example, sea bottom depth and sediments or debris in the water which occur on small scales may result in it not providing a representative model of the sea reflectance. It would be a useful exercise to see how well the retrieval does in such situations, but for this test month of data there were only 3 coincidences with coastal stations available, which is insufficient information from which to draw conclusions. Aerosol optical depth over the ocean is, however, compared with MISR data in Section 5.4.

A note on AERONET instrumentation

Central wavelength	Bandwidth
0.340 μm $\pm$1.5 nm	2 nm $\pm$0.5 nm
0.380 μm $\pm$1.5 nm	2 nm $\pm$0.5 nm
0.440 μm $\pm$2 nm	10 nm $\pm$2 nm
0.500 μm $\pm$2 nm	10 nm $\pm$2 nm
0.675 μm $\pm$2 nm	10 nm $\pm$2 nm
0.870 μm $\pm$2 nm	10 nm $\pm$2 nm
0.937 μm $\pm$2 nm	10 nm $\pm$2 nm
1.020 μm $\pm$2 nm	10 nm $\pm$2 nm
1.640 μm $\pm$4 nm	25 nm $\pm$5 nm

Table 5.1: AERONET Sun-photometer channel locations and bandwidths.

The standard instrumentation used at AERONET sites is the CIMEL CE-318 Sun-photometer, which has nominal channel centres and bandwidths as given in Table 5.1 (CIMEL, 2008). Each

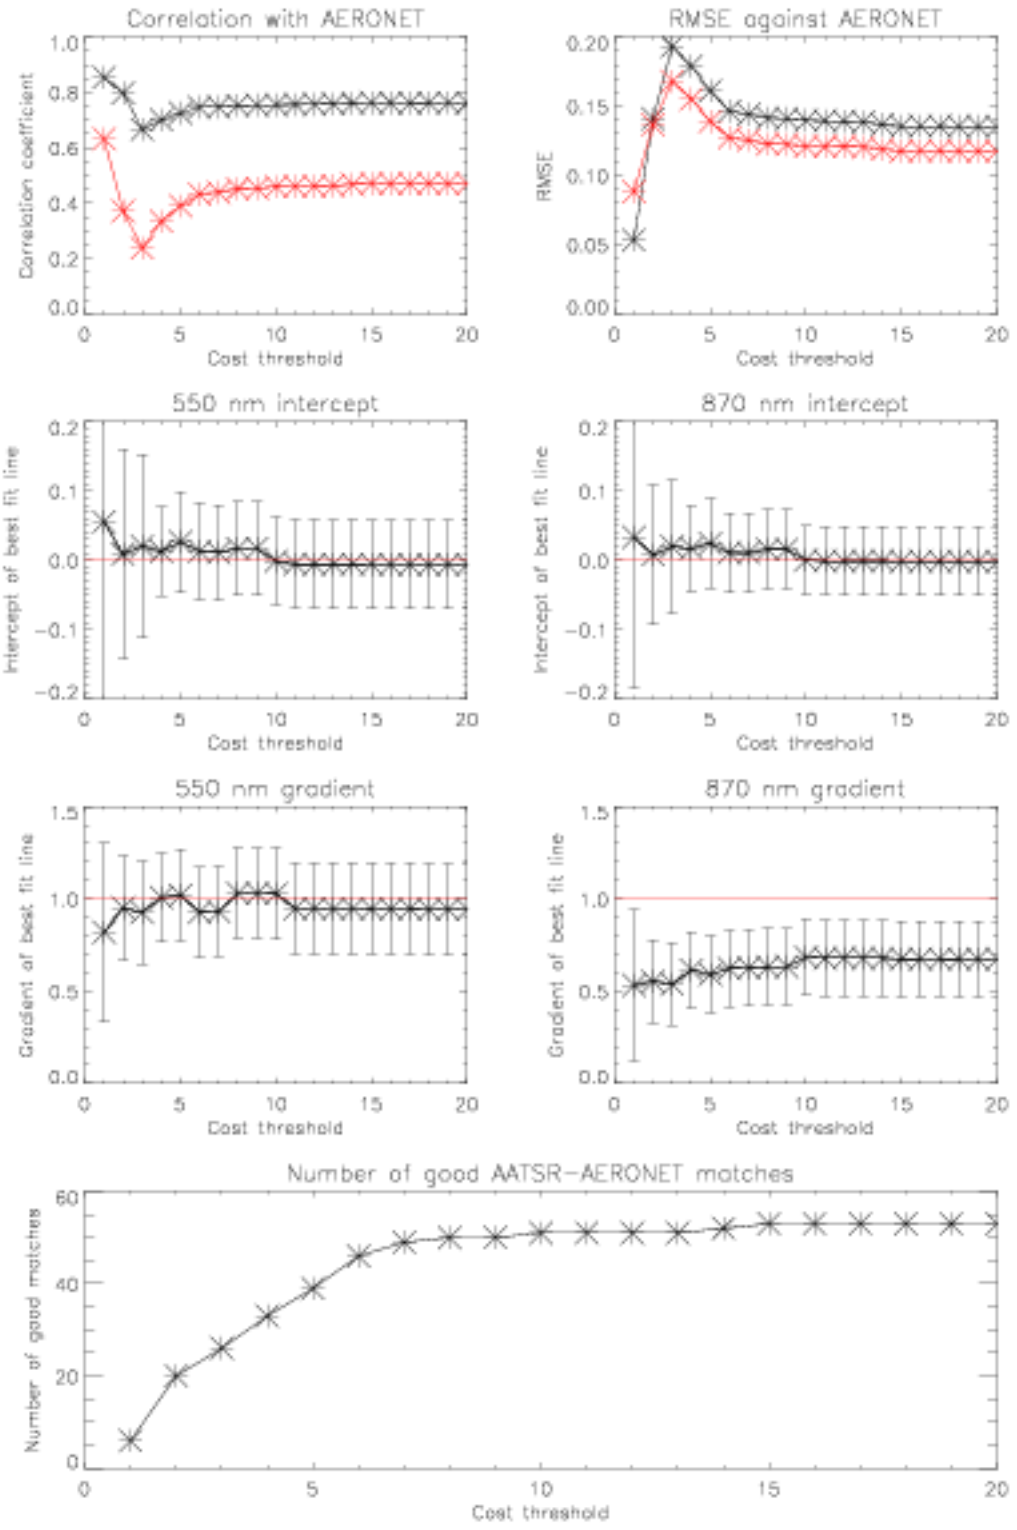

Figure 5.1: Statistics of AATSR-AERONET optical depth comparison for September 2004 for different retrieval cost thresholds. Top-left: Correlation coefficient between the two datasets at 550 nm (black) and 870 nm (red). Top-right: RMSE of the two datasets at 550 nm (black) and 870 nm (red). Second row: Intercept of the line of best fit of the data, at 550 nm (left) and 870 nm (right). Error bars show the uncertainty on the intercept. The red line indicates an intercept of zero. Third row: Gradient of the line of best fit of the data, at 550 nm (left) and 870 nm (right). Error bars show the uncertainty on the gradient. The red line indicates a gradient of zero. Bottom: The number of good coincidences between AATSR and AERONET during the month. All results for the continental aerosol model.

instrument is unique, however, leading to slight differences in band centres and widths. These bands are narrower than those of AATSR (Table 1.1).

5.1.2 550 nm optical depth validation

Scatter plots comparing dual-view AATSR and AERONET retrievals of τ_{550} are shown in Figure 5.2; as the majority of AERONET stations in the analysis are continental land stations, it is suspected that the most appropriate aerosol class to use is the continental class (top-left part of the figure). This is also found to be the class that provides the best match to the AERONET data.

In all cases, aside from the urban aerosol model and lowest-cost speciation attempt, the correlation between these independent aerosol datasets is high, with best fit lines which, within uncertainty, pass through the origin. This is an important result as it demonstrates the consistency between satellite-based and ground-based measurements of aerosol. Such agreement at AERONET sites is a piece of evidence that the AATSR retrieval could be used to obtain reliable estimates of aerosol properties at locations where there is no corresponding ground-based measurement.

Several further comments can be made about the individual graphs in Figure 5.2:

- The best match (in terms of smallest intercept and gradient closest to 1 of line of best fit) was for the continental aerosol model, suggesting that if the true aerosol type for a scene is unknown then use of the continental model as a first guess is the best idea for the retrieval.

- The continental model comparison has some outlying points with a positive AATSR bias; one of these is a coastal station, in which case it is possible that continental is not the most appropriate aerosol type for this pixel. The four largest outliers also have a large standard deviation on the AATSR estimate, suggesting perhaps some residual cloud contamination in some superpixels or natural aerosol inhomogeneity meaning that the assumptions of Ichoku et al. (2002) are not valid for these points. Alternatively it is possible that these stations are either unrepresentative of the wider area or the assumed *a priori* uncertainty on the MODIS surface reflectance product is not very good for these pixels.

- Both the desert and maritime classes are more strongly correlated than the continental class with AERONET data ($r > 0.85$), but show larger biases (gradients of 1.24 and 0.91). As desert aerosol can contain nonspherical particles, it is possible that phase function errors introduced

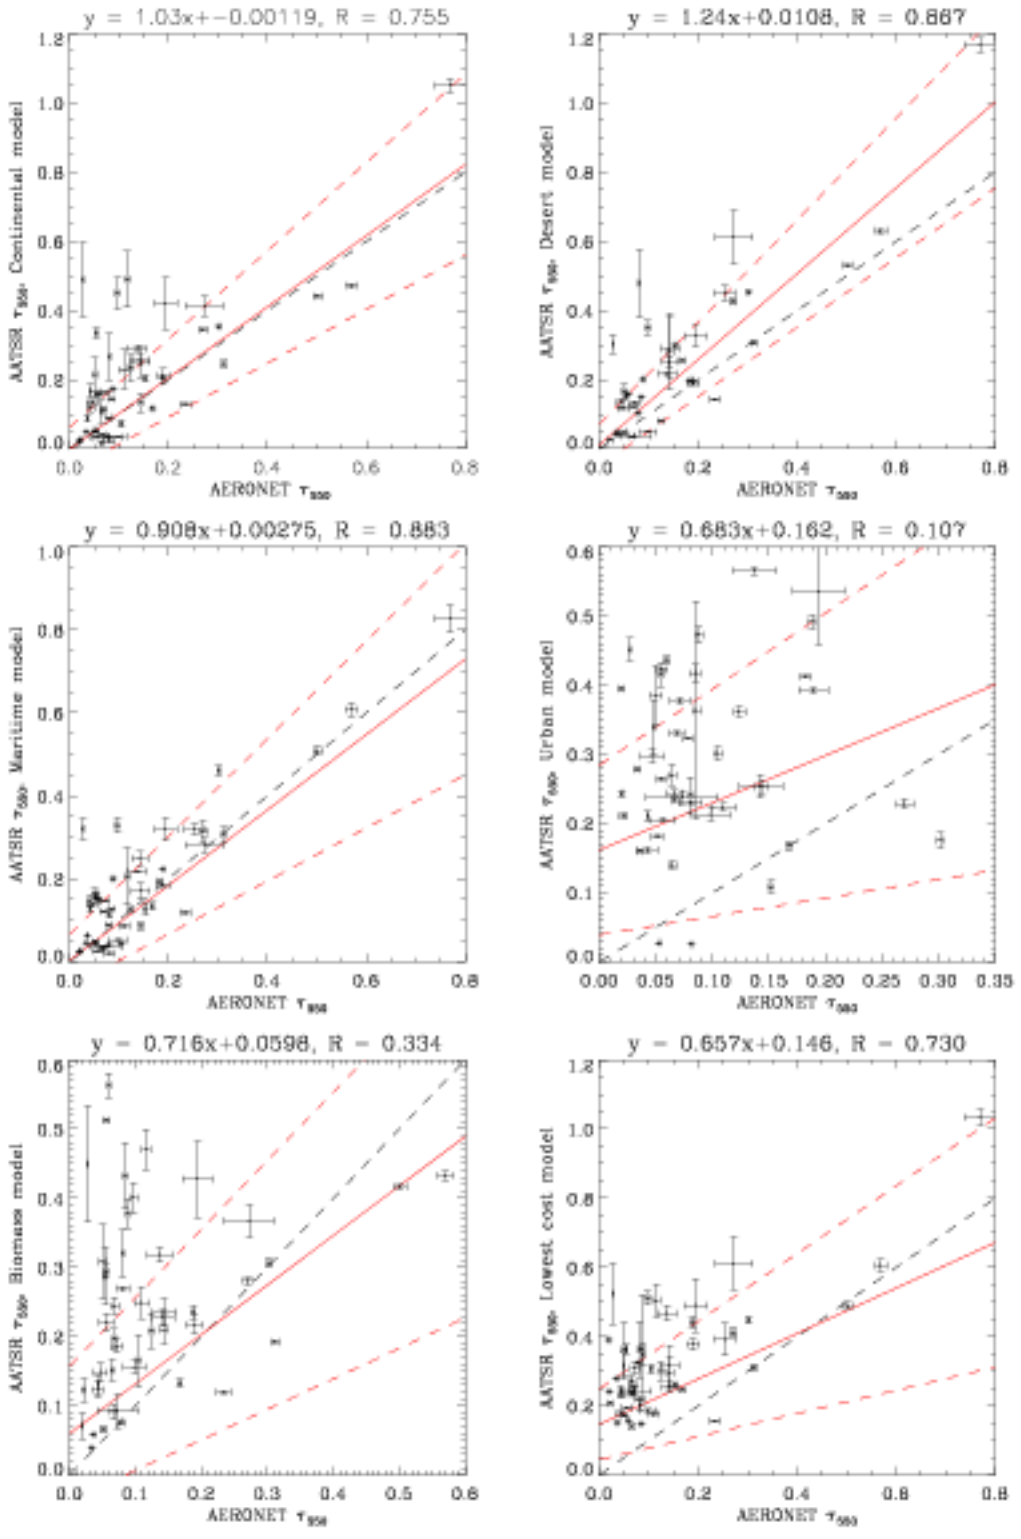

Figure 5.2: Comparison of spatio-temporally averaged dual-view AATSR and AERONET aerosol optical depths at 550 nm during September 2004, assuming different aerosol models for the AATSR retrieval. Top-left: Continental aerosol model. Top-right: Desert aerosol model. Centre-left: Maritime aerosol model. Centre-right: Urban aerosol model. Bottom-left: Biomass burning aerosol model. Bottom-right: Comparison using the lowest-cose aerosol model for each pixel. The black dotted line is the 1:1 line, the red solid line the best fit and red dotted lines its upper and lower limits. Slope and intercept coefficients of the line of best fit are given above the plots, along with the correlation coefficient. Error bars on the data are the standard deviations of the spatio-temporally averaged measurements.

by using Mie theory's spherical approximation lead to systematic errors in the phase function (and hence optical depth) at scattering angles sampled by AATSR. For maritime aerosol, similar errors could arise if the relative humidity (assumed 50 % in LUT generation) is incorrect.

- For the urban aerosol model the agreement with AERONET is almost non-existent. This is not unexpected—the aerosol present at most of the AERONET stations is unlikely to be similar to the urban aerosol model, which is highly absorbing (in contrast to the mostly scattering continental, desert and maritime classes).

- Attempting aerosol speciation by lowest-cost methods gives a poorer result than assuming that the aerosol conforms to the continental model; the gradient of the line of best-fit is poor (0.66) and there is a substantial offset (0.15). This suggests that AATSR does not have sufficient information to discriminate between aerosol types, and the retrieval is able to fit the measured reflectances within observation error for a variety of aerosol models. This is not necessarily a problem with the retrieval scheme but a limitation of the instrument itself. Therefore, in a comparative study it is useful to have some idea of the type of aerosol that is being observed.

5.1.3 870 nm optical depth validation

Figure 5.3 presents analagous results to Figure 5.2 but for τ_{870}. This is not a quantity directly retrieved in the ORAC scheme, but instead calculated from τ_{550}, r_e and the ratio between β_e at 550 nm and 870 nm in a separate LUT for the aerosol model chosen. Together with validation of τ_{550}, this provides a test of the validity of the aerosol models used.

The same general comments as made for validation of τ_{550} can be made for validation of τ_{870}, although the level of agreement with AERONET is poorer than for τ_{550}. This is an indication that the retrieval is not always able to partition correctly the spectral dependence of the measurements between effects of the particle size and the surface's spectral shape. It may also indicate errors in the assumed optical properties or size distributions. Nevertheless there is clearly some correspondence between the two datasets.

Improvement of retrieved τ_{870} may be aided by eventual adoption of the infrared channels into the retrieval scheme as, particularly in the case of absorbing aerosols, particle size has a large impact on the signal in the infrared for large particles.

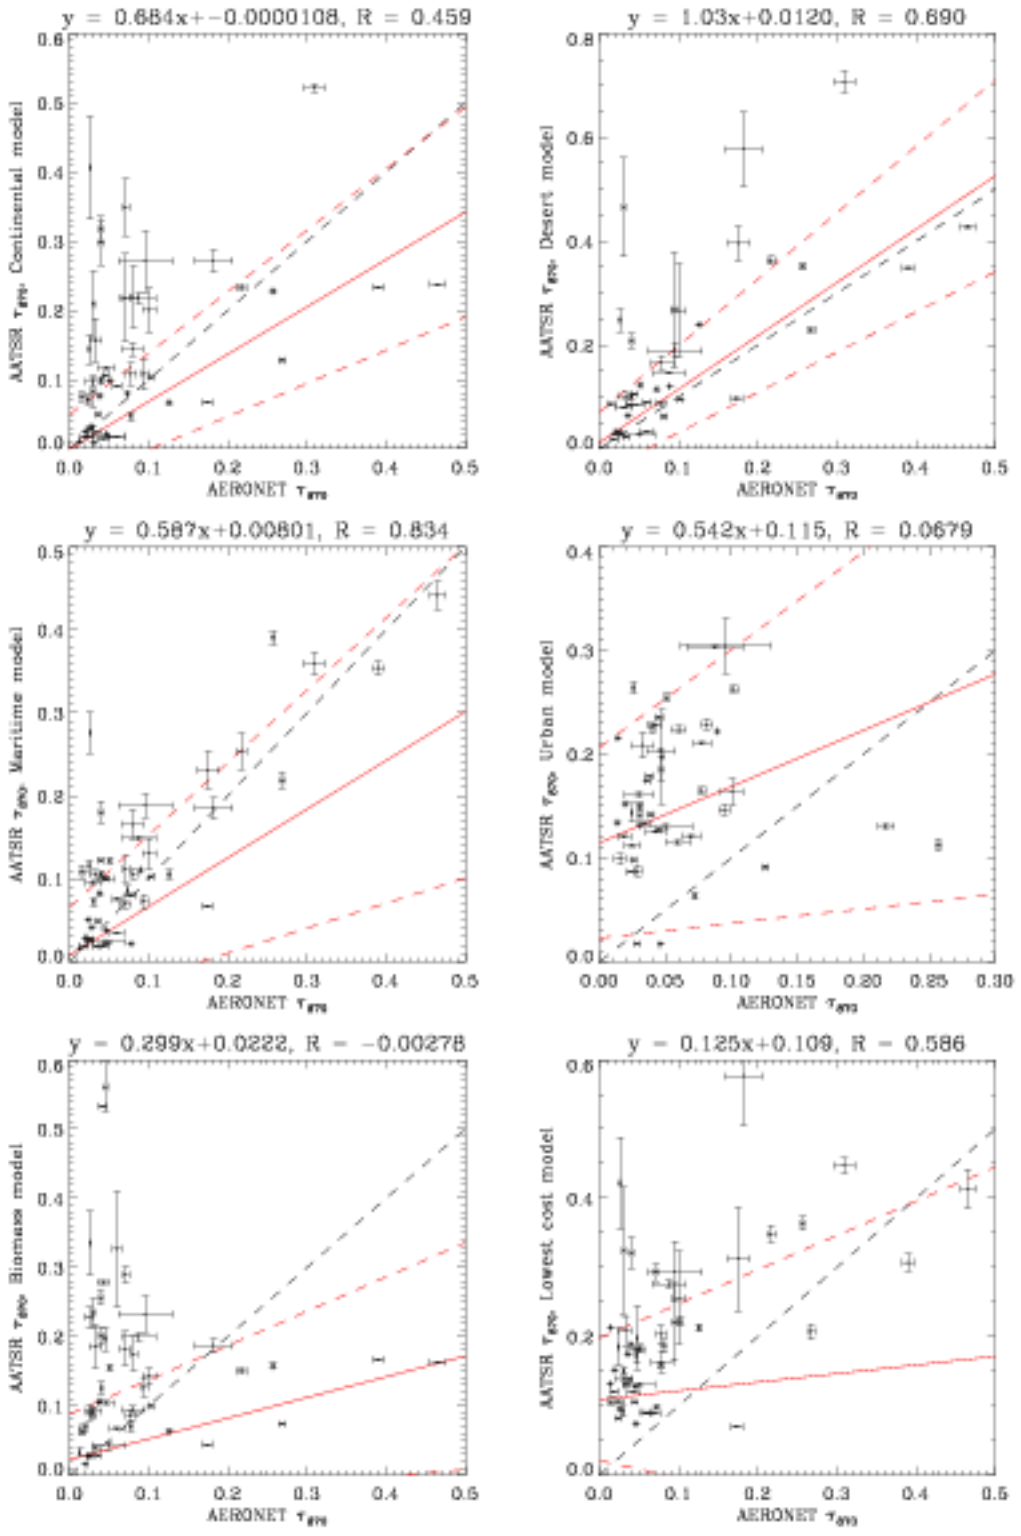

Figure 5.3: Comparison of spatio-temporally averaged dual-view AATSR and AERONET aerosol optical depths at 870 nm during September 2004, assuming different aerosol models for the AATSR retrieval. Top-left: Continental aerosol model. Top-right: Desert aerosol model. Centre-left: Maritime aerosol model. Centre-right: Urban aerosol model. Bottom-left: Biomass burning aerosol model. Bottom-right: Comparison using the lowest-cose aerosol model for each pixel. The black dotted line is the 1:1 line, the red solid line the best fit and red dotted lines its upper and lower limits. Slope and intercept coefficients of the line of best fit are given above the plots, along with the correlation coefficient. Error bars on the data are the standard deviations of the spatio-temporally averaged measurements.

5.1.4 Improving aerosol speciation: class-based cost thresholds

Aerosol class	Maximum cost over land	Maximum cost over sea
Continental clean	10.0	0.5
Desert dust	0.5	0.5
Maritime	0.5	10.0
Urban	0.5	0.5
Biomass burning	1.0	1.0

Table 5.2: Cost thresholds for each class, as used in the aerosol speciation described in the text.

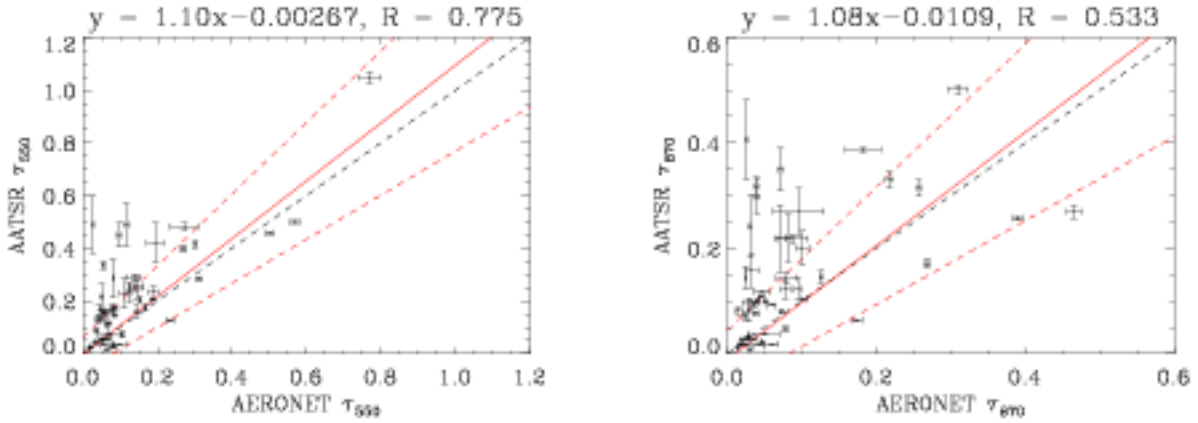

Figure 5.4: Comparison of spatio-temporally averaged dual-view AATSR and AERONET aerosol optical depths during September 2004, assuming aerosol speciated by the limited lowest cost method, using the thresholds in Table 5.2. Top: Comparison at 550 nm. Bottom: Comparison at 870 nm. The black dotted line is the 1:1 line, the red solid line the best fit and red dotted lines its upper and lower limits. Slope and intercept coefficients of the line of best fit are given above the plots, along with the correlation coefficient. Error bars on the data are the standard deviations of the spatio-temporally averaged measurements.

Section 2.7 discussed the limits of lowest-cost aerosol speciation, and showed that particularly for low aerosol loadings multiple aerosol models may fit the data well. The next simplest approach to aerosol speciation is to adopt different thresholds of cost for different aerosol classes when deciding whether a retrieved value is 'good' or not. This approach has been explored to allow for a slightly more intelligent attempt at speciation. Different methods were experimented with to see which gave an improved match between the two datasets; the coefficients in Table 5.2 were found to be suitable thresholds to use in this limited lowest cost speciation.

When comparing costs to determine the best model to pick, the lowest-cost aerosol model is chosen, so long as it has a cost lower than the threshold value given in Table 5.2. With these values, this effectively means that the continental class is favoured over the land and maritime over the sea

(as typical background aerosol types), unless the data strongly suggests that one of the other classes provides a very good fit. The threshold for biomass aerosol is higher than that for the other atypical aerosol classes as this aerosol type shows behaviour more distinguishable than continental, desert or maritime aerosol and so it is less likely that it will be misidentified (see also Figure 2.9). Although urban aerosol is also distinct from the more scattering classes, it is thought unlikely to dominate areas the size of retrieval pixels, and more relaxed thresholds for it resulted in markedly poorer comparisons with AERONET.

This is a conservative approach to speciation, meaning aerosol may be wrongly classified as continental or maritime. However for an aerosol event extending over several pixels it is likely that some of the aerosol will be flagged appropriately, giving an indication of, for example, a dust event. Conversely the conservative thresholds minimise the chances of a dust event being flagged in the data when none occurred in reality. It should be emphasised again that this method is intended to provide an indication of the dominant aerosol type, and any other information on likely aerosol type should be taken into account for any use of this data.

As Figures 5.2 and 5.3 show, the fact that similar optical depth comparisons are obtained for each aerosol class (aside from urban) means that the misclassification of a dust or biomass event as continental or maritime is still likely to result in a reasonable optical depth, particularly at 550 nm. The fact that different aerosol classes give similar 550 nm optical depths in turn means it is difficult to find a set of thresholds which do make a noticeable difference, as many of the changes in the comparison are within their error bars.

Application of the cost thresholds shown in Table 5.2 to the data for September 2004 leads to the comparison shown in Figure 5.4. The resulting dataset consists mostly of continental aerosol, with small amounts of the other classes. The correlation and RMSE of the data are both slightly improved as compared to the matches with the continental class and particularly the full lowest-cost speciated data, as is the gradient of the best-fit line at 870 nm. The gradient at 550 nm is slightly further away from 1 than for the continental class, though still matches within error bars.

In summary, the use of different cost thresholds for different aerosol classes provides a slightly improved match of the data to AERONET measurements as compared to lowest-cost speciation with a uniform cost threshold. Addition of further constraints to the retrieval, in the form of more tightly-constrained measurements or *a priori* data, may help in the future to differentiate between the aerosol

classes more strongly (as an inappropriate aerosol model may be affected more strongly by these constraints than the most appropriate one). It is surmised that a future addition of the infrared measurements into the retrieval could improve speciation due to the differing behaviour of the aerosol models in the infrared. Carboni et al. (2008) have found that measurements at 11 μm and 12 μm are sensitive to aerosols with radii of approximately 1 μm or larger.

5.1.5 General comments on AERONET validation

All aerosol retrievals from visible/near-infrared satellite imagers suffer from similar difficulties, which lead to imperfect comparisons with each other and ground-based measurements such as AERONET. These include the following:

- Radiances from cloud are generally stronger than radiances from aerosol, meaning undetected residual cloud in a superpixel will lead to (generally positive) biases in retrieved aerosol optical depth. Cloud masks and the cost function can ameliorate this problem, but there is a compromise between lax cloud flagging (and lots of residual cloud contamination) and over-zealous cloud flagging (which may wrongly eliminate strong aerosol events, leading to a negative bias in optical depth). As noted by Koren et al. (2007) and others, there is a 'twilight zone' between clouds and aerosols, with modified aerosol properties and enhanced radiance in this region. In this case, the correlation between the AATSR-AERONET difference and the AATSR cloud fraction was 0.33 at 550 nm and 0.43 at 870 nm, suggesting some possible cloud contamination (assuming unflagged cloud is likely to occur near flagged cloud).

- Assuming cloud flagging to be perfect, errors may also arise due to adjacency effects whereby there is an enhancement of radiance near cloud, from scattering into the field of view on one side, or a decrease, from shadowing on the other side. This is not currently accounted for by the retrieval forward model, and could also explain elevated optical depth retrieved from increased radiance near clouds.

- As is shown by the scatter plots presented, the retrieved aerosol properties are strongly dependent upon the assumed composition of the aerosol (refractive index, size distribution and particle shape). Satellite measurements generally provide insufficient information to retrieve these; speciation by methods such as lowest cost, as shown above, cannot always accurately

pick from the best aerosol model provided. Errors in the assumed aerosol model manifest as errors in the retrieved properties.

- Over bright land and sun-glint regions in the sea, the majority of the TOA reflectance comes from the surface. The aerosol contributes only a small amount, meaning that it can be difficult to attribute accurately radiance to surface or atmosphere; small errors in the land surface reflectance can cause larger errors in retrieved aerosol properties. This explains, for example, the difficulty of performing retrievals over sand or ice as opposed to the dark ocean.

- Different instruments do not have identical measurement wavelengths or filter shapes. The standard AERONET photometer, for example, lacks a 550 nm measurement (Table 5.1) and so errors are introduced by interpolating optical depth at other wavelengths to 550 nm.

- Although the spatio-temporal averaging method of Ichoku et al. (2002) is designed to minimise the effects of representiveness error, it is still unrealistic to expect perfect agreement between the datasets within their errors.

Finally it should be emphasised that, although as an independent data source it is useful to validate against AERONET measurements, AERONET itself is not necessarily a perfect representation of the true state of the atmosphere. Therefore while useful to check results are reasonable, excessive tuning of results to improve the match against AERONET (or any other data source) is undesirable and may even be counterproductive. As an example, rather than SST derived from the (A)ATSR sensors being validated against buoys, in some cases now the buoys are calibrated against the satellites (Donlon et al., 2008). Of equal importance to external validation should be the investigation of retrieval statistics, as these can be diagonostic of where the retrieval algorithm is doing well and where it is doing poorly.

5.2 Retrieval statistics

As well as validation of the retrieval products with independent data sources such as AERONET, it is an important exercise to examine the statistics of the retrieval. Unexpected patterns in the data can be indicative of either a problem with the input data or the retrieval process itself. Figures 5.5, 5.6

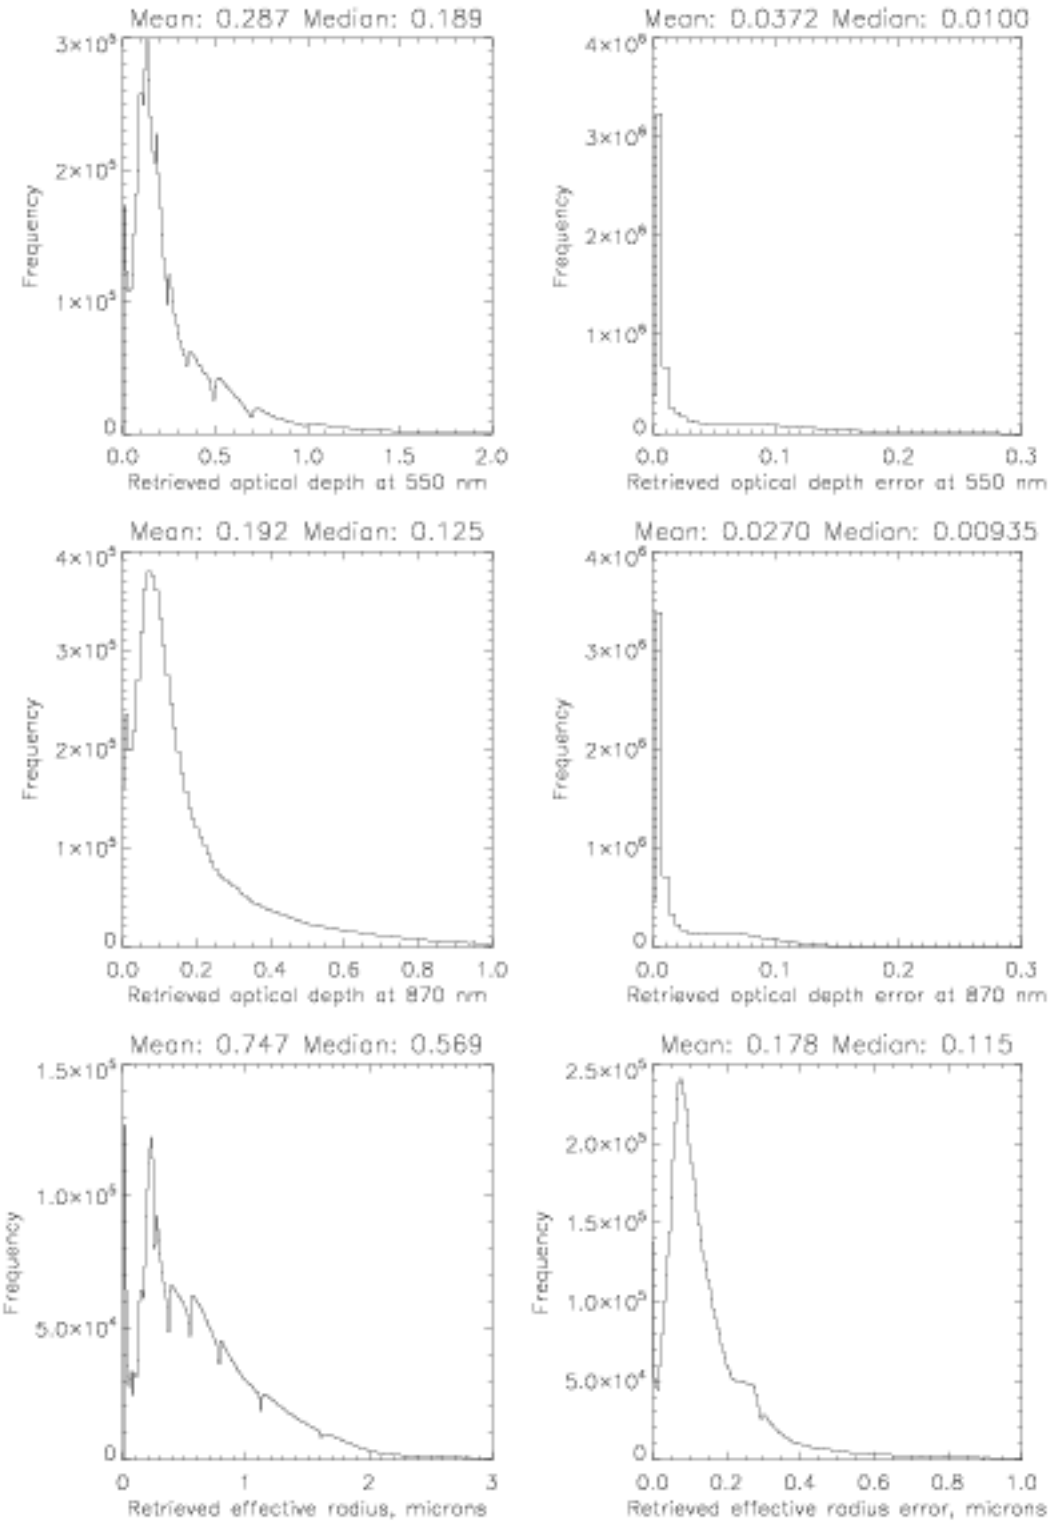

Figure 5.5: Histograms of retrieved aerosol properties for September 2004, using the continental aerosol model. Results presented for retrievals that converged with a cost lower than 20. The left-hand sides show the retrieved values and the right-hand sides their errors. The mean and median values are indicated above each histogram. From top to bottom: optical depth at 550 nm, optical depth at 870 nm and aerosol effective radius.

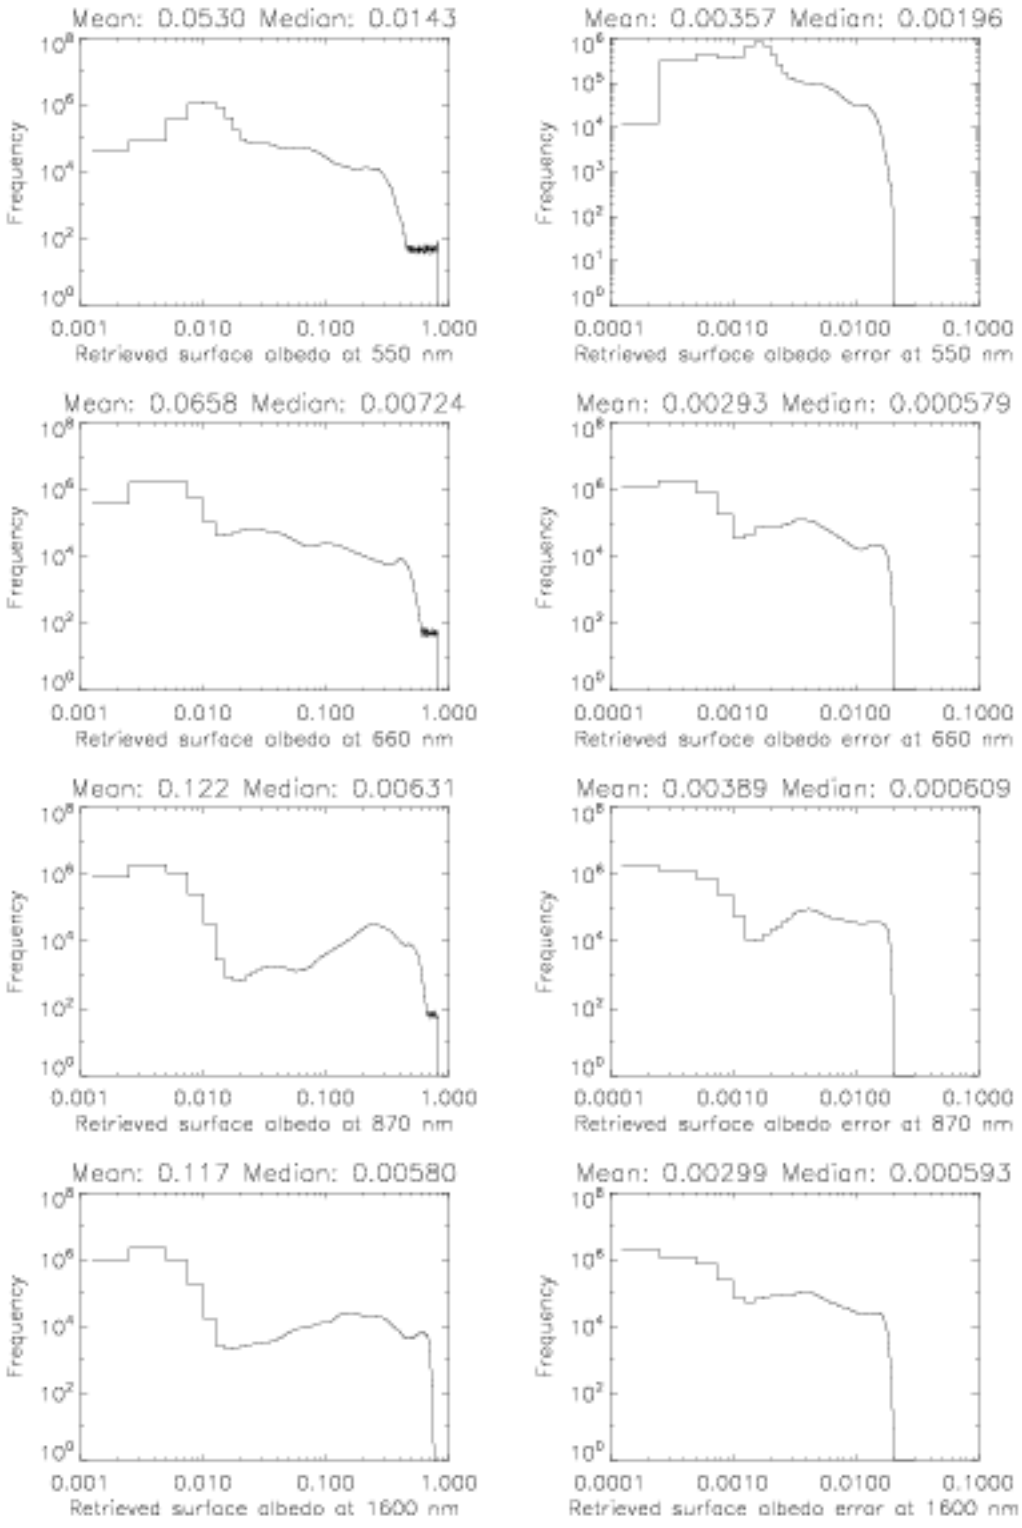

Figure 5.6: Histograms of retrieved surface properties for September 2004, using the continental aerosol model. Results presented for retrievals that converged with a cost lower than 20. The left-hand sides show the retrieved values and the right-hand sides their errors. The mean and median values are indicated above each histogram. From top to bottom: surface albedo at 550 nm, surface albedo at 660 nm, surface albedo at 870 nm and surface albedo at 1.6 μm. Axes are logarithmic as otherwise the histogram detail is dominated by the numerous dark sea pixels.

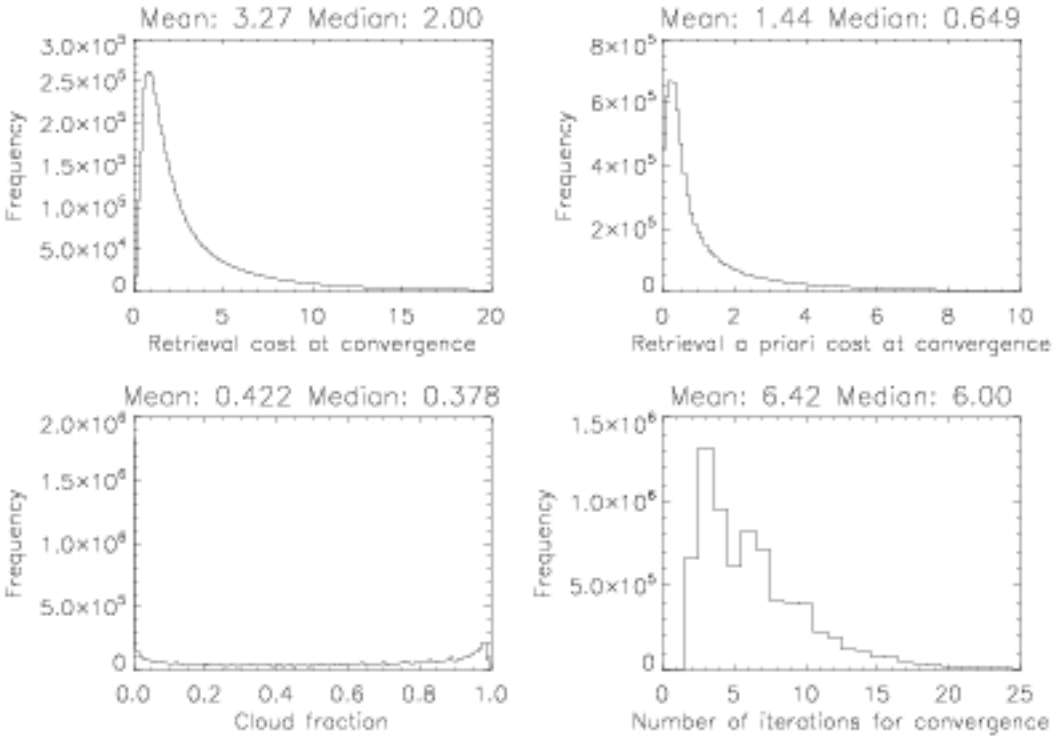

Figure 5.7: Histograms of retrieval quality statistics for September 2004, using the continental aerosol model. Results presented for retrievals that converged with a cost lower than 20. The mean and median values are indicated above each histogram. Top-left is the total retrieval cost, top-right the *a priori* component, bottom-left the superpixel cloud fraction, and bottom-right the number of iterations for convergence. Note that the cloud fraction is defined as the proportion of cloudy pixels within a superpixel for only those pixels where an aerosol retrieval was performed, and hence is not a measure of the true cloud cover.

and 5.7 present histograms of retrieval statistics for aerosol properties, surface properties and quality control statistics respectively for retrievals during September 2004.

Results are presented for retrievals using the continental aerosol class model. The only quality control applied to the data is that the retrieval converged and had a cost less than 20; this removes most remaining cloud contamination and pixels where there was a problem (e.g. missing channels or missing land surface reflectance data). Similar results (omitted for brevity) are obtained using the other aerosol models or the speciated dataset. The continental model is assumed here to allow a direct comparison with the previous nadir-view retrieval; their different setups would mean that the speciated datasets would not be directly comparable.

5.2.1 Retrieved parameters

Figures 5.5 and 5.6 show the frequency distributions of retrieved quantities; aside from τ_{550} and r_e the distributions are smoothly varying, with no tendency for the retrieval to stick at particular values.

The sudden dips in the τ_{550} and r_s distributions occur as at LUT point values there is a discontinuity in the forward model gradient (as LUTs are linearly interpolated) which can hamper the iteration. As the convergence criteria depend on the steepness of the cost function, which will be steeper on the side of the LUT point with the steeper forward model gradient, this leads to a slight clustering of retrieved data on this side of the LUT value.

This introduces some small-scale biases into the retrieval, although such errors will likely be encompassed by the retrieval's error estimate for that state parameter. One method to solve this problem would be to move to higher-order LUT interpolation, to provide a smoothly-varying forward model gradient. A finer grid would also decrease the size of these artefacts (by spreading them out). Both methods carry an increased computational overhead and are discussed later as possible future developments to the algorithm.

The distributions for surface albedo are multimodal because different surface types (such as ocean, dense dark vegetation and bright vegetation or rock) have their own characteristic brightnesses.

In all cases the mean value of the error frequency distribution is significantly smaller than the mean of the retrieved value frequency distribution. Multiple peaks in the error distributions can generally be linked to whether the pixel is a land or sea pixel. Over the dark sea surface, it is possible to assign the contributions of surface and atmosphere to TOA reflectance with more certainty, resulting in a lower error estimate on the retrieved parameter.

5.2.2 Quality control statistics

Cost

The mean retrieval cost, as indicated by Figure 5.7, is 3.27 with a measurement component of 1.83 and an *a priori* component of 1.44. As the cost is divided by the number of channels on output, with 8 measurements and 6 constrained parameters (of which 5 are strong constraints; the optical depth constraint is weak) an expected mean cost would be 1.63 (13/8) of which measurement cost would account for 1.0 (8/8) and *a priori* cost 0.63 (5/8).

The relative contributions of measurement and *a priori* cost are approximately as expected, suggesting that the correct weight is being given to measurement and *a priori* by the retrieval scheme. However, as the mean total cost is larger than expected this indicates that the forward model is not ac-

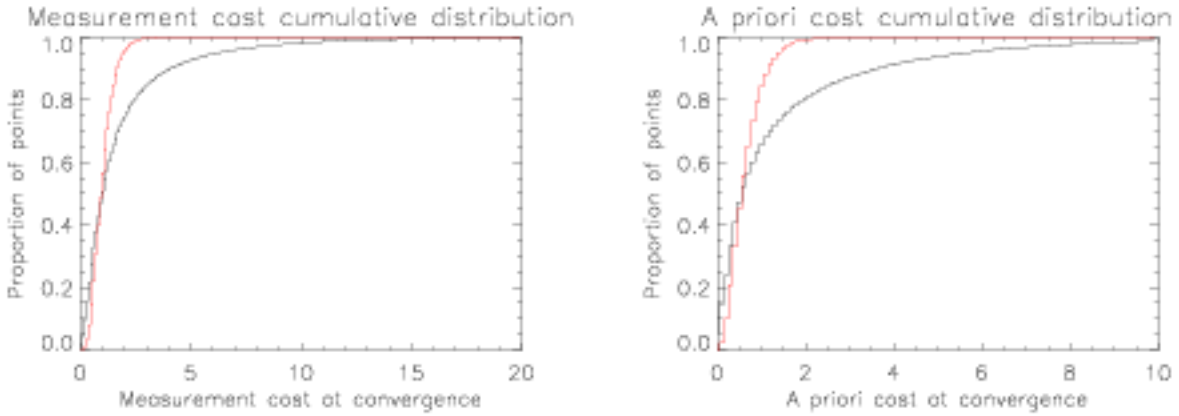

Figure 5.8: Left: Cumulative function of the measurement component of retrieval cost. The x axis shows the cost, with the y axis the proportion of pixels converging with that cost or smaller. The black line is for retrieval data using the continental clean aerosol class throughout September 2004, and the red line is a theoretical curve for a chi-squared distribution with 8 degrees of freedom (remembering that the cost function is divided by 8 on output) Right: The same, except for the *a priori* contribution to cost, and a chi-square distribution with 5 degrees of freedom.

counting for all the factors affecting the TOA reflectance. Such factors may include imperfect aerosol or surface models, improper constraints on the measurements or *a priori* data, or the approximations made of aerosol vertical distribution or sphericity. These histograms also show that by a cost of 10 (the threshold adopted for AERONET comparisons) most of the converged pixels are accepted.

Cost distributions can be used to check whether the retrieval results are consistent with the measurements and *a priori* data. Residuals should follow a Gaussian distribution, meaning the retrieval costs would have a χ^2 distribution, which is tested for with the test of the same name. For this purpose it is not appropriate to consider the distribution of total cost as the retrieved state depends on the measurements, so the measurement and *a priori* components of cost are not independent.

Following Rodgers (2000), separate χ^2 tests are required to check for consistency of the retrieval with measurement and *a priori* data. The covariance matrices for these tests are not $\mathbf{S}_y$ and $\mathbf{S}_x$ respectively, as the fit of the measurements is constrained by the *a priori* (and the converse for the fit to the *a priori* data). Instead, for the measurements the covariance $\mathbf{S}_{\delta\hat{y}}$ of the residuals $\hat{\mathbf{y}} - \mathbf{y}$ is required, which can be shown to be:

$$\mathbf{S}_{\delta\hat{y}} = \mathbf{S}_y(\mathbf{K}_x\mathbf{S}_x\mathbf{K}_x^T + \mathbf{S}_y)^{-1}\mathbf{S}_y \tag{5.1}$$

From Equation 5.1 it follows that $\mathbf{S}_{\delta\hat{y}}$ will have smaller values than $\mathbf{S}_y$. This is intuitively true

as the with more state variables $\hat{\mathbf{y}} - \mathbf{y}$ should become smaller as the retrieval is better able to fit the measurements. The term $\mathbf{K}_x\mathbf{S}_x\mathbf{K}_x^T$ maps the *a priori* uncertainty into measurement space; again, as the *a priori* becomes less well-defined or the weighting functions get steeper the fit of $\mathbf{y}$ should improve. Rodgers (2000) also suggests the approximation of using $\mathbf{S}_y$ for $\mathbf{S}_{\delta\hat{y}}$, noting that it will provide a 'loose' estimate of the goodness-of-fit. Calculation of $\mathbf{S}_{\delta\hat{y}}$ explicitly for each pixel carries a large computational overhead. It has been found that, for the majority of cases, performing the χ^2 test with $\mathbf{S}_y$ instead of $\mathbf{S}_{\delta\hat{y}}$ yields similar results (with χ^2 using $\mathbf{S}_{\delta\hat{y}}$ about 10 % higher). Therefore $\mathbf{S}_{\delta\hat{y}}$ is approximated by $\mathbf{S}_y$ here.

The resulting distribution is compared to a theoretical χ^2 distribution with 8 degrees of freedom in Figure 5.8. The distributions are similar, suggesting the retrieval performs well. An overabundance of pixels with a lower-than-expected cost may indicate regions of overestimates of measurement error. Conversely the tail of the observed distribution is larger than that of the theoretical, indicating errors where the forward model is not adequately fitting the data. Part of this region will be due to the aerosol model being inappropriate (and repetition with the speciated product leads to an improved match with the theoretical distribution in this region); other reasons for the discrepancy would be factors not accounted for in the retrieval (such as aerosol vertical distribution or surface elevation) or convergence criteria that are too loose.

Despite these differences the steepest parts of the observed and theoretical distributions occur for similar values of measurement cost. For a χ^2 distribution with 8 degrees of freedom (and normalised by dividing χ^2 by 8), approximately 99 % of pixels would be expected to show a χ^2 statistic of 3 or less. For the observations that figure is around 80 %, suggesting the retrieval is generally performing well and consistent with the measurements.

A similar matrix $\mathbf{S}_x$ is formed to check for consistency with the *a priori*. Analogously to Equation 5.1, this maps the uncertainty on the measurements into state space and results in a matrix slightly tighter (in the sense of having values smaller in magnitude) than the *a priori* covariance matrix $\mathbf{S}_x$:

$$\mathbf{S}_{\hat{x}} = \mathbf{S}_x\mathbf{K}_x^T(\mathbf{K}_x\mathbf{S}_x\mathbf{K}_x^T + \mathbf{S}_y)^{-1}\mathbf{K}_x\mathbf{S}_x \tag{5.2}$$

Again the resulting matrix can be approximated by $\mathbf{S}_x$, and the resulting χ^2 distribution has approximately 5 degrees of freedom (really slightly more than 5 but fewer than 6, as τ_{550} is weakly constrained over the range of the state limits). The approximation of $\mathbf{S}_{\hat{x}}$ by $\mathbf{S}_x$ has again been found

to have small impact using simulated data, and the resulting distribution is shown in Figure 5.8; similar comments as made for measurement cost apply.

Additional checks are defined by Rodgers (2000) to check for consistency of the measured signal and *a priori*, and the retrieved signal and *a priori*. As these are similar in principle to those discussed here, and generally speaking all checks should give the same answer if the retrieval has correctly converged, they are omitted for brevity.

In summary, the retrieval is in general consistent with measurements and *a priori* data. This gives confidence that it is performing correctly. Departures from theoretical distributions indicate factors unaccounted for by the retrieval forward model.

Cloud fraction

The histogram of superpixel cloud fractions shows that most of the retrievals were under largely cloud-free conditions. As these histograms show only retrievals which successfully converged, this is not unexpected: fully cloudy superpixels are rejected in the retrieval preprocessing, and many other superpixels with a high cloud fraction may not converge if there was nearby unflagged cloud.

Number of iterations

The mean number of iterations for convergence is 6.42. Almost all iterations require 2 or more steps, indicating that they move away from the *a priori* values, and most have converged within 15 iterations. The fact that very few pixels require more suggests that the cutoff of a maximum of 25 steps before the pixel is abandoned is not having a significant impact on coverage.

5.2.3 Residuals

Examination of the residuals on both measurement and *a priori* data is an important task: it reveals how well the forward model and retrieval method are able to fit the data, and gives an indication of any biases in the scheme, as well as providing an indication as to whether the assumed error budget is appropriate or not.

All the residuals discussed in this section are given as normalised residuals; the normalised residual is simply obtained by dividing the residual by its assumed error (from $\mathbf{S}_y$ or $\mathbf{S}_x$ for measurements and *a priori* data respectively). The result of this is that the distributions of residuals would for a

perfect retrieval be Gaussian with a mean of 0 and standard deviation of 1. Approximately 68 % of the distribution should fall between ±1 and over 99 % within ±3.

Results are presented for the continental aerosol model for all pixels which converged successfully with a cost smaller than 20 and state variables not lying exactly on their boundaries as specified by the retrieval. As such it represents a 'worst case' of successful retrievals, as this model may not be globally applicable. Similar results, with slightly different biases, are obtained using other aerosol classes or speciated products.

Residuals on measurements

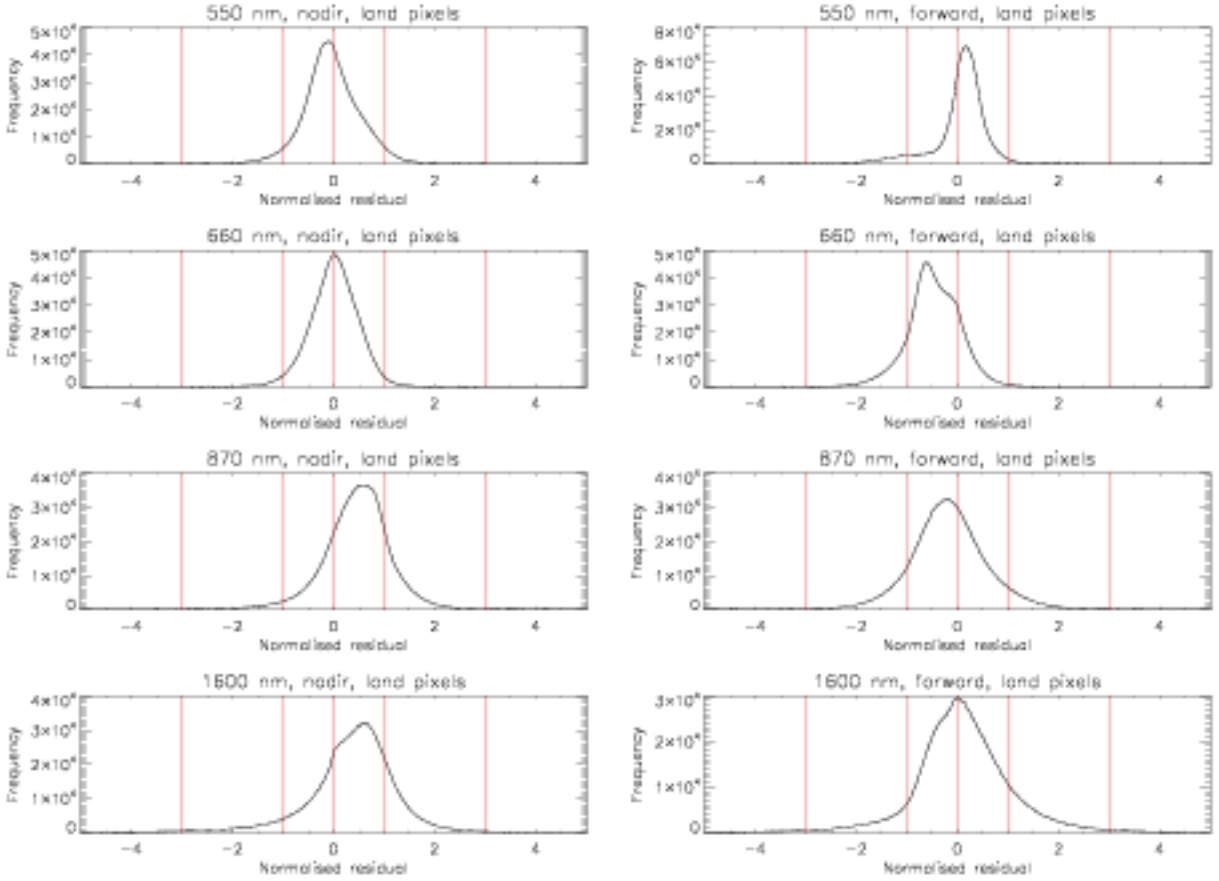

Figure 5.9: Histograms of residuals on measured TOA radiance for September 2004, normalised to the total uncertainty (forward model and measurement) on the value. The left-hand side shows nadir-view measurements and the right-hand side forward-view measurements. From top to bottom the graphs show the modelled-predicted residual at 550 nm, 660 nm, 870 nm and 1.6 μm. Red lines indicate zero, ±1 and ±3. Data for successfully-converged land retrievals using the continental aerosol model.

Distributions of normalised residuals on measured data are presented for land pixels in Figure 5.9 and sea pixels in Figure 5.10. Positive values indicate the forward model predicting a larger TOA

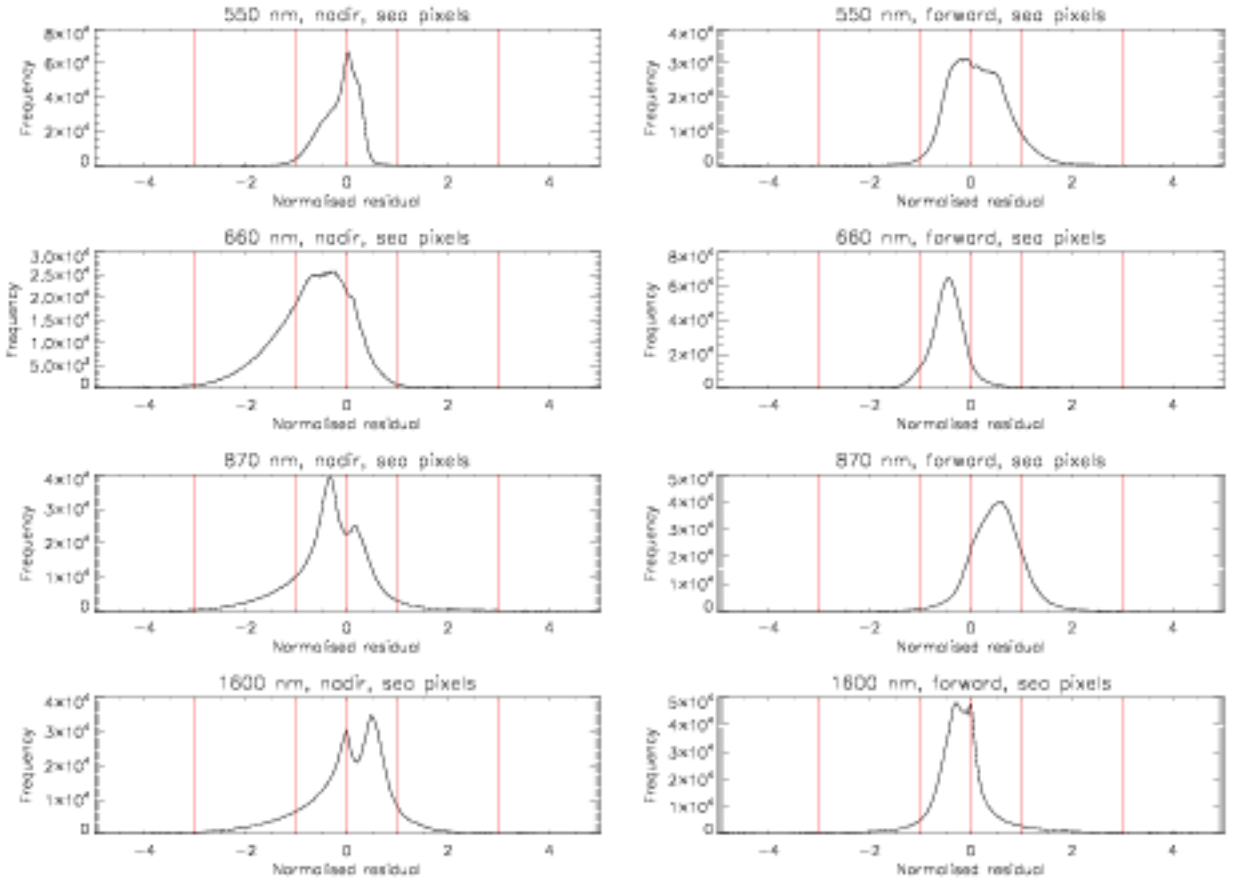

Figure 5.10: Histograms of residuals on measured TOA radiance for September 2004, normalised to the total uncertainty (forward model and measurement) on the value. The left-hand side shows nadir-view measurements and the right-hand side forward-view measurements. From top to bottom the graphs show the modelled-predicted residual at 550 nm, 660 nm, 870 nm and 1.6 μm. Red lines indicate zero, ± 1 and ± 3. Data for successfully-converged sea retrievals using the continental aerosol model.

radiance than observed, and vice-versa. Over land, distributions are generally monomodal; slight biases in the mean are found for different wavelengths and views; these may be partially explained by the aerosol model not being able to fit the data perfectly or errors associated with an improper land surface reflectance. This may be associated with the humps in some of the distributions; these tend to arise from data clustered in particular regions of the world, rather than being uniformly spread, which supports this hypothesis of regional issues. There is some hint that the assumed errors at 550 nm and 660 nm are slightly loose, as the distribution of residuals is narrower than expected.

Residuals for sea pixels as shown in Figure 5.9 also show evidence of generally being narrower distributions than expected, again suggesting the errors in $\mathbf{S}_y$ could be slightly tightened. Multiple peaks in some distributions arise due to the differing behaviour of the water in shallow water (such as coastal regions) and open ocean: the surface reflectance model is designed for open water, so may

not provide an accurate BRDF in turbid shallow waters.

In summary, these distributions suggest that the retrieval scheme is generally able to fit the measurements well. There are hints that the assumed error budget may be too lax in some cases; smaller values of S_g should lead to smaller error estimates on retrieved quantities, and may aid the lowest-cost speciation method as fitting errors from an improper aerosol model would be amplified. Distributions are generally monomodal. Departures from a Gaussian shape can be explained by a combination of an improper aerosol model being used, and errors in the surface reflectance.

Residuals on *a priori* data

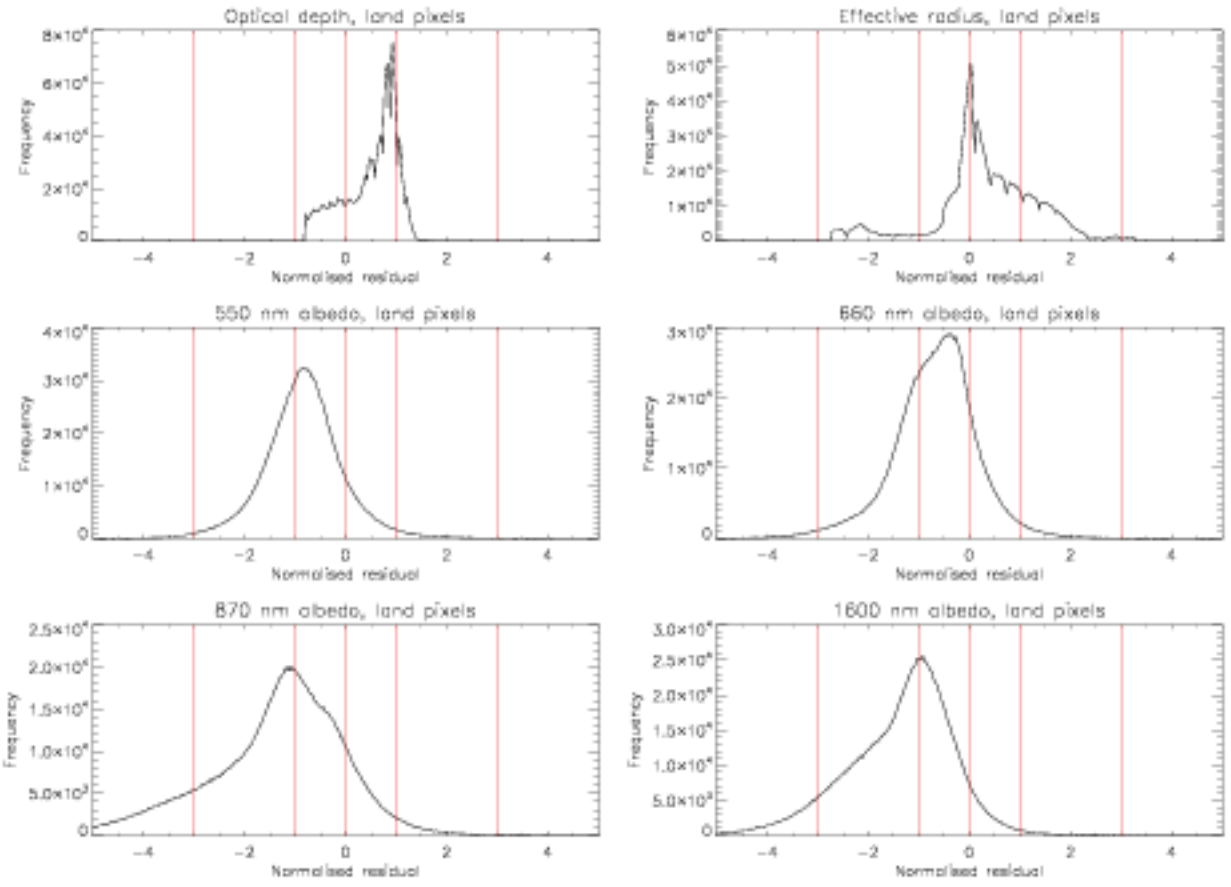

Figure 5.11: Histograms of residuals on *a priori* data for September 2004, normalised to the *a priori* uncertainty on the value. The top row shows retrieved-*a priori* τ_{550} (left) and r_e (red). The middle row shows retrieved-*a priori* $R_{\mathrm{SLW},550}$ (left) and $R_{\mathrm{SLW},660}$ (right). The bottom row shows retrieved-*a priori* $R_{\mathrm{SLW},870}$ (left) and $R_{\mathrm{SLW},1600}$ (right). Red lines indicate zero, ± 1 and ± 3. Data for successfully-converged land retrievals using the continental aerosol model.

The *a priori* data used by the algorithm acts as a constraint to stop the retrieval of unphysical values. As such, looking at the *a priori* residuals is important, as a retrieval overconstrained by

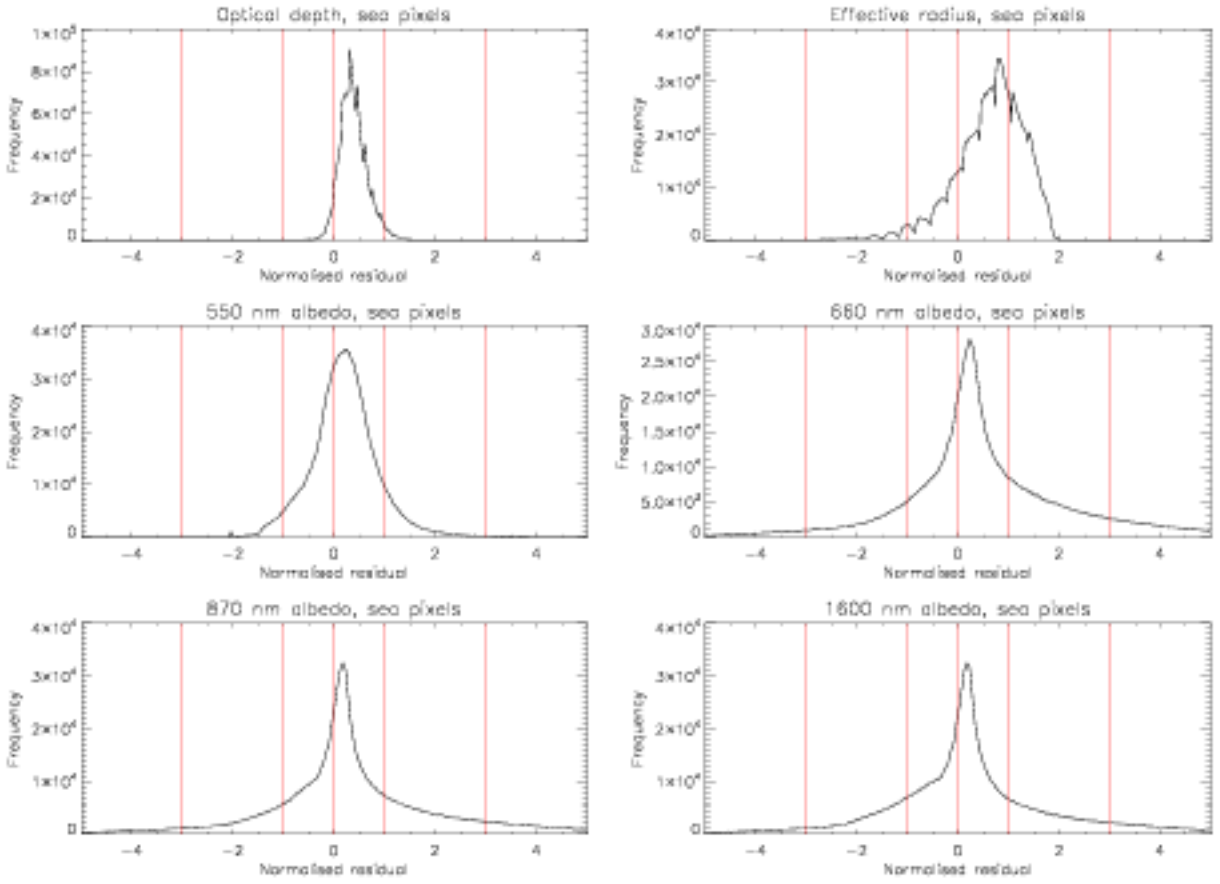

Figure 5.12: Histograms of residuals on *a priori* data for September 2004, normalised to the *a priori* uncertainty on the value. The top row shows retrieved-*a priori* τ_{550} (left) and r_e (red). The middle row shows retrieved-*a priori* $R_{\mathrm{SLW},550}$ (left) and $R_{\mathrm{SLW},660}$ (right). The bottom row shows retrieved-*a priori* $R_{\mathrm{SLW},870}$ (left) and $R_{\mathrm{SLW},1600}$ (right). Red lines indicate zero, ± 1 and ± 3. Data for successfully-converged sea retrievals using the continental aerosol model.

a priori data will not give sufficient weight to the measurements. Figures 5.11 and 5.12 present the distributions of normalised residuals on *a priori* values for land and sea pixels respectively; a positive value indicates the retrieved state being higher than the *a priori*, and vice-versa.

Over both land and see the histograms for τ_{550} are significantly thinner than normalised Gaussians; this is because, as the aerosol optical depth is the primary product of interest, the *a priori* constraint on it is intentionally weak so it will be fully described by the observations. Hence the *a priori* uncertainty on this parameter is large and its resultant residuals (and contribution to cost) are small. Structure in these distributions arises from the slight preference of the retrieval for state values lying on LUT points, as mentioned previously. Sharp cutoffs at the edge represent the state limits imposed on the retrieval ($0.01 < \tau_{550} < 5$).

Similar structure is observed for effective radius residual histograms, although as the effective

radius is more tightly constrained for the various aerosol classes the distribution is wider.

For the retrieved surface albedo, the distributions look quite different for land and sea pixels:

- Over land, although the distributions have appropriate widths, there is a persistent bias whereby the retrieved surface albedo is generally lower than the *a priori* at all wavelengths. This could indicate the MODIS product systematically undercorrecting for aerosol, a problem with the BRDF model or a problem with the aerosol models used in this scheme. This bias is also seen in histograms for other aerosol classes, and may result in the AATSR retrieval returning a higher aerosol optical depth than others to compensate.

- Over sea, distributions also generally have appropriate widths. This indicates that the real sea shows variability similar to that predicted by the model (as this was the method used to estimate appropriate uncertainties). The distributions have their peaks close to zero, although there is a small positive bias at all wavelengths. This may stem from a factor such as an underestimate of the reflectance of oceanic whitecaps (which, as mentioned in Section 3.2.3, is highly uncertain). Foam is a likely cause as it occurs globally and is dependent only on wind speed; the monomodality of residual distributions suggests a persistent factor. Alternatively the choice of aerosol model could be to blame; the bias is reduced (but not eliminated) if the maritime class is used instead of the continental. Together these results give confidence that the sea surface model developed in this work is appropriate for use.

These figures show that the *a priori* constraints are generally appropriate, as the distributions are close to Gaussian. The exception is τ_{550}, which is thinner because it is comparatively unconstrained within the state limits. The sea surface model provides *a priori* data in good agreement with the retrieved state, with the resulting distributions of residuals close to theoretical shape. The land surface albedo retrieved is generally lower than the *a priori*, although the reasons for this are unclear.

5.2.4 Averaging kernels

The averaging kernel matrix, A, provides information on how much of the retrieved state comes from the measurements and how much from the *a priori* value (Rodgers, 2000). It is derived from

the inversion operator $\mathbf{D}_y$ (see Equation 2.32):

$$\mathbf{A} = \mathbf{D}_y \mathbf{K}_x \tag{5.3}$$

Each row and column of $\mathbf{A}$ corresponds to an element of $\mathbf{x}$, with elements indicating how much a change in the true state leads to a change in the retrieved state. Hence the diagonal terms represent the sensitivity of the retrieval to the true state, such that in the perfect case the trace of the matrix will equal the number of elements in the state vector (i.e. the state is fully described by the measurements and the *a priori* information is not necessary). Off-diagonal elements, then, indicate how sensitive the retrieved state is to changes in the *other* true state variables; in the ideal case, the off-diagonal elements will be zero and the state vector elements will be independent of each other. From the definitions of $\mathbf{D}_y$ and $\mathbf{K}_x$ it follows that in the ideal case the averaging kernel $\mathbf{A}$ will be the identity matrix $\mathbf{I}$.

Due to the limits of the information provided by the satellite measurements, the aerosol retrieval averaging kernels are non-diagonal and generally have a trace significantly lower than 6 (the number of state vector elements). Three general cases of pixels with differing averaging kernels can be described:

1. Sea pixels. In such cases the trace of the averaging kernel is typically 5 or higher. Off-diagonal elements are small. The element corresponding to optical depth is typically close to 1, with the others being lower (as the effective radius and surface are more constrained by their *a priori* information).

2. Normal land pixels (forest, field). Over most land pixels the averaging kernel's trace is between 4 and 5. Effective radius and surface albedo kernel elements are again lower than those for optical depth. There are typically off-diagonal elements between elements relating to effective radius and surface albedo, as these both strongly influence the spectral dependence of the measured TOA reflectance.

3. Bright land pixels (desert, snow). In these cases the trace of the averaging kernel is often lower than 4. All elements relating to state vectors are noticeably less than 1, particularly effective radius, and the off-diagonal elements between effective radius and surface albedo are large.

These values are as expected, and indicate that the most well-retrieved parameter is τ_{550} (its averaging kernel diagonal element is generally closest to 1). Over the ocean particle size and surface albedo may be retrieved, but over brighter surfaces the *a priori* becomes more important. This indicates the importance of thorough investigation into *a priori* errors, as they have a non-negligible influence on the final solution.

Improved constraints on land surface albedo should result in off-diagonal elements of A between effective radius and surface albedo becoming smaller. This may also lead to an increase of the averaging kernel element describing effective radius over land, but it is likely that this will still contain a significant amount of information from the *a priori*. Off-diagonal elements between surface albedo at different wavelengths will likely remain significant because of the strong constraints on the surface albedo, particularly over land (see Chapter 4), but this is not a problem as the primary data products of interest are the aerosol information.

5.3 Improvement on validation of previous nadir-view retrieval

Prior to this work the AATSR ORAC retrieval algorithm was simpler, as follows:

- It used only the nadir-view channels (as opposed to both views).
- The surface reflectance was assumed Lambertian, with the sea albedo approximated by the glint contribution to the BRDF and the land albedo using the unadjusted MODIS albedo.
- As a result of only using the nadir-view measurements and having consequently less information, the spectral shape of the surface was fixed by the *a priori* spectral shape, and only the 550 nm albedo (rather than all 4) was retrieved. The *a priori* surface error was set to 0.01 for both sea and land pixels: a likely overestimation of the former, and underestimation of the latter.
- S_y was simplified and not optimal for aerosol retrieval (flat measurement error, no consideration of numerical precision, simplified model parameter error).

To illustrate the improvement in the aerosol retrieval arising from improvements in the aerosol algorithm, the month of September 2004 was additionally processed using the old nadir-view scheme.

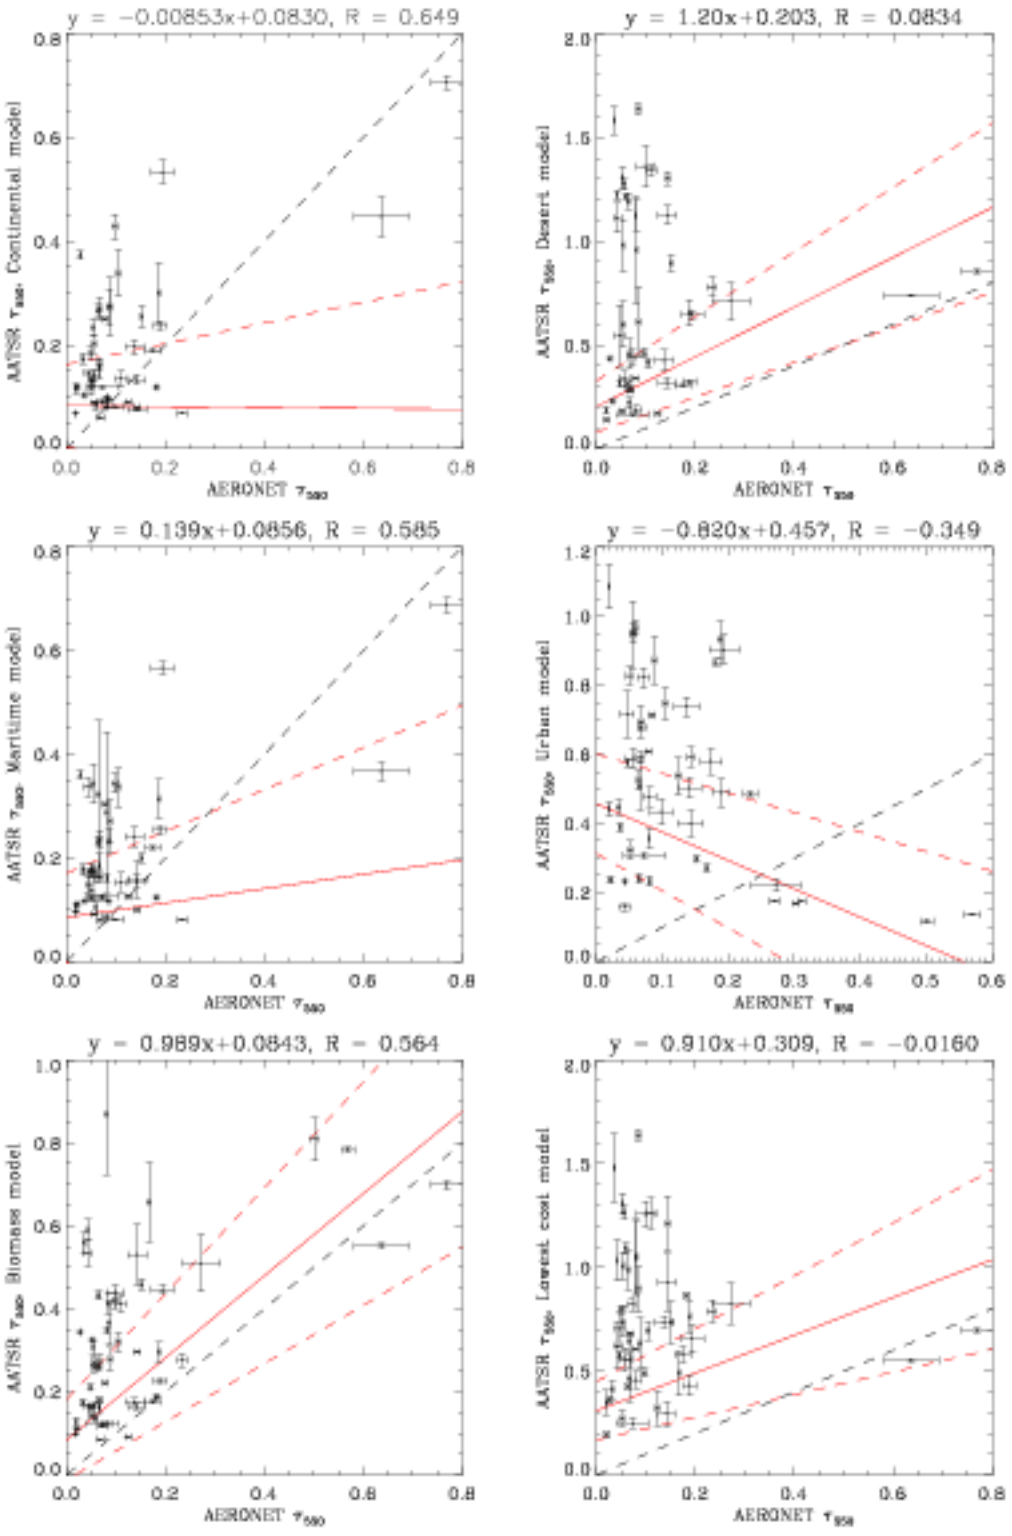

Figure 5.13: Comparison of spatio-temporally averaged nadir-view AATSR and AERONET aerosol optical depths at 550 nm during September 2004, assuming different aerosol models for the AATSR retrieval. Top-left: Continental aerosol model. Top-right: Desert aerosol model. Centre-left: Maritime aerosol model. Centre-right: Urban aerosol model. Bottom-left: Biomass burning aerosol model. Bottom-right: Comparison using the lowest-cost aerosol model for each pixel. The black dotted line is the 1:1 line, the red solid line the best fit and red dotted lines its upper and lower limits. Slope and intercept coefficients of the line of best fit are given above the plots, along with the correlation coefficient. Error bars on the data are the standard deviations of the spatio-temporally averaged measurements.

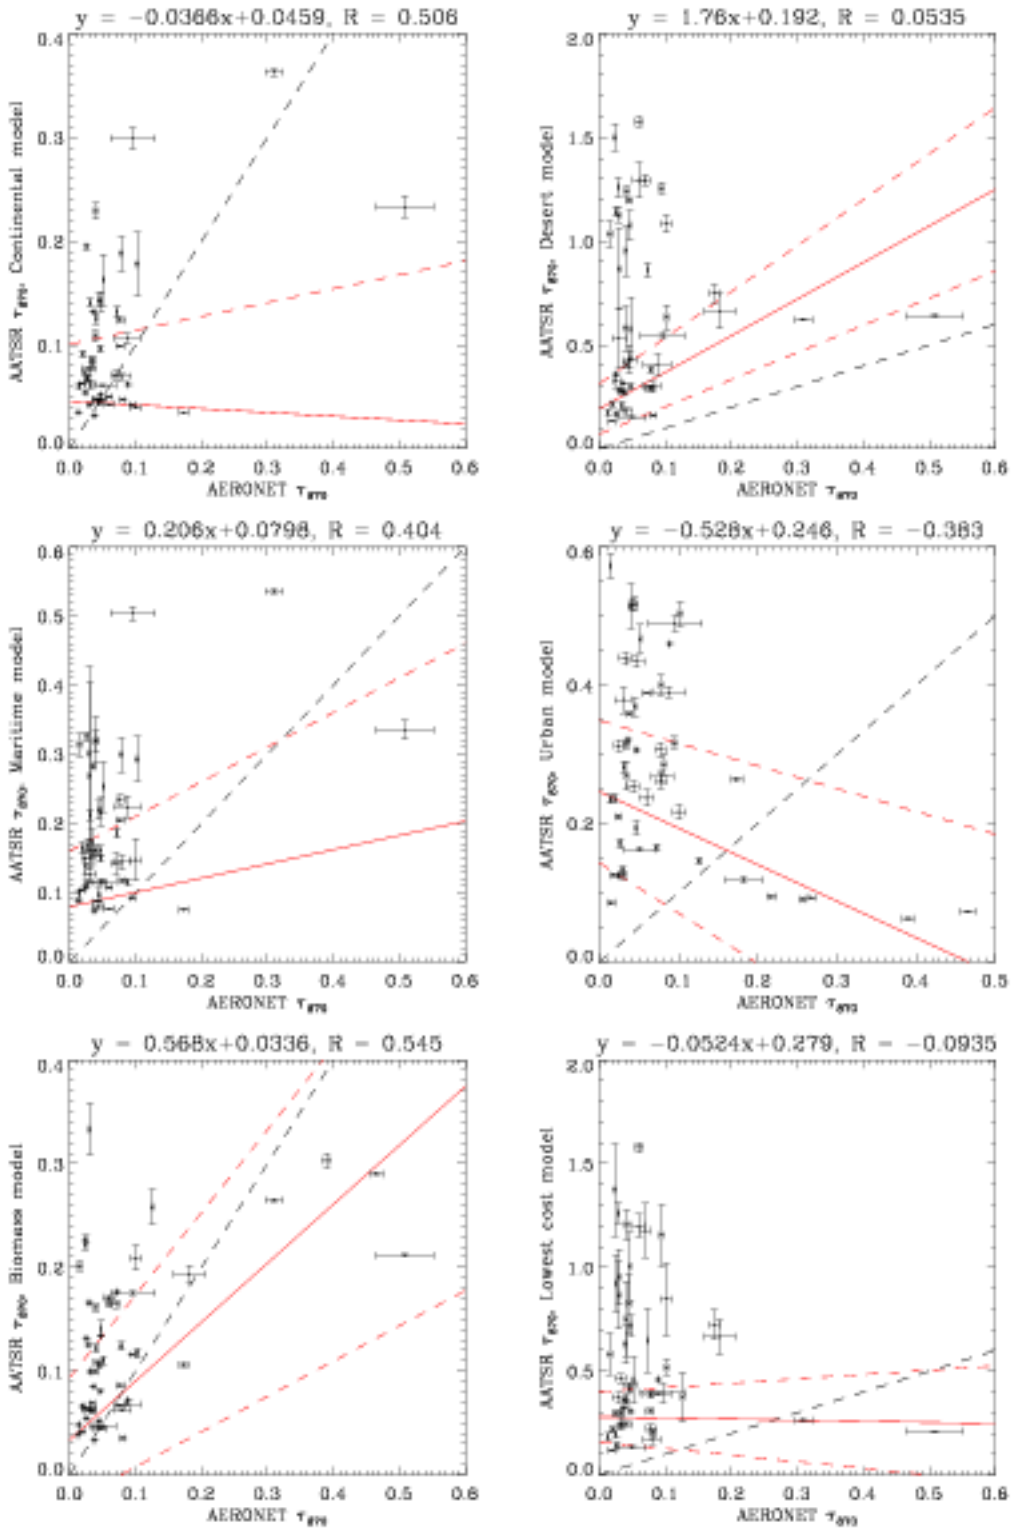

Figure 5.14: Comparison of spatio-temporally averaged nadir-view AATSR and AERONET aerosol optical depths at 870 nm during September 2004, assuming different aerosol models for the AATSR retrieval. Top-left: Continental aerosol model. Top-right: Desert aerosol model. Centre-left: Maritime aerosol model. Centre-right: Urban aerosol model. Bottom-left: Biomass burning aerosol model. Bottom-right: Comparison using the lowest-cost aerosol model for each pixel. The black dotted line is the 1:1 line, the red solid line the best fit and red dotted lines its upper and lower limits. Slope and intercept coefficients of the line of best fit are given above the plots, along with the correlation coefficient. Error bars on the data are the standard deviations of the spatio-temporally averaged measurements.

5.3.1 AERONET validation

AERONET validation of this previous nadir-view retrieval is shown for τ_{550} in Figure 5.13 and for τ_{870} in Figure 5.14. In comparison with Figures 5.2 and 5.3, these results show that particularly in the case of τ_{550} there is a great improvement in the level of agreement with AERONET in terms of the correlation between the two datasets and the gradient and intercept of the line of best fit between them. This clearly demonstrates the advantages that the dual-view technique of AATSR allows over other sensors, and represents a large step forward in the ability to quantify accurately the distribution of atmospheric aerosol from space with AATSR.

5.3.2 Retrieval Statistics

Retrieval quantity	Mean, dual-view	Mean, nadir-view	Median, dual-view	Median, nadir-view
τ_{550}	0.287 ±0.037	0.244 ±0.156	0.189 ±0.01	0.116 ±0.142
τ_{870}	0.192 ±0.027	0.137 ±0.090	0.125 ±0.009	0.059 ±0.078
r_e, microns	0.747 ±0.178	0.260 ±0.257	0.569 ±0.115	0.237 ±0.270
$R_{\mathrm{SLW},550}$	0.053 ±0.004	0.083 ±0.006	0.014 ±0.002	0.052 ±0.007
$R_{\mathrm{SLW},660}$	0.066 ±0.003	-	0.007 ±0.001	-
$R_{\mathrm{SLW},870}$	0.122 ±0.004	-	0.006 ±0.001	-
$R_{\mathrm{SLW},1600}$	0.117 ±0.003	-	0.006 ±0.001	-
Retrieval cost	3.27	1.36	2.00	0.74
Retrieval *a priori* cost	1.44	0.48	0.65	0.20
Superpixel cloud fraction	0.42	0.40	0.38	0.31
Number of iterations	6.42	5.15	6.00	3.00

Table 5.3: Comparative retrieval statistics for the dual-view algorithm presented here and previous nadir-view ORAC algorithm. Results presented for the month of September 2004 and the continental aerosol model. The uncertainties quoted are the mean or median retrieval uncertainty on each retrieved quantity over the month, as opposed to the standard error on the monthly mean value, to provide an indication of retrieval precision. Note that the cloud fraction is defined as the proportion of cloudy pixels within a superpixel for only those pixels where an aerosol retrieval was performed, and hence is not a measure of the true cloud cover.

A summary of the mean and median retrieval statistics from Figures 5.5, 5.6 and 5.7 for the dual-view retrieval are shown together with the equivalent values for the nadir-view retrieval in Table 5.3. Again results are presented for the month of September 2004 and the continental aerosol class; in the interest of brevity, the distributions themselves are not shown (although are discussed below).

These results show that the nadir-view retrieval estimates, on average, smaller values of τ_{550} (and hence τ_{870}) than the dual-view retrieval for this aerosol class. The fact that the distributions of optical

depth change when the forward view measurements are added into the retrieval is important, as it tells us that the forward-view channels, sampling a different part of the atmospheric phase function than the nadir-view channels, are providing new information. Section 5.1 shows that the dual-view retrieval provides a better match to ground-based aerosol optical depths; clearly, the use of the dual-view technique is an advantage to the retrieval scheme. The nadir-view retrieval also shows larger errors on retrieved optical depth than the dual-view (which is to be expected, given fewer channels are used) and suffers from the same issue of periodical 'sticking' at values close to LUT points.

The nadir-view effective radius generally remains close to the *a priori* value, with its error estimate taking a similar value. This shows that by itself use of the nadir viewing geometry does not provide enough information to adequately retrieve r_e. The distribution of errors is bimodal, with smaller errors over the sea and errors of 100% or larger over land. The dual-view results in a significant improvement in the retrieved effective radius: it moves away from the *a priori* value, and the error estimates are smaller.

Surface albedo at 550 nm is retrieved with a similar distribution of values and errors for both the nadir-view and dual-view algorithms. This may be explained because, having a strong effect on the TOA reflectance, it is a simple quality for the retrieval to fit and so the addition of further measurements to the retrieval then results in improved characterisation of aerosol properties. The error arising in the nadir view retrieval from not having enough measurements to retrieve surface albedo at all 4 wavelengths simultaneously may explain the poor comparison of retrieved nadir-view optical depth to AERONET, shown in Figures 5.13 and 5.14.

Both total and *a priori* cost take lower values on average for the nadir-view retrieval than the dual-view retrieval. This is again expected behaviour: as discussed in Section 5.2 the expected mean cost for a dual-view retrieval is 1.63, while for the nadir-view retrieval with 4 measurements and 2 constrained state variables (again, τ_{550} is not tightly constrained) leads to an expected cost of 1.50 (6/4). The distribution of retrieval costs is similar, with the nadir-view retrieval converging on average slightly faster. This may be related to the cost convergence criteria: with fewer measurements to fit, the cost function gradient is likely shallower for the nadir-view retrieval than the dual-view, meaning the point to stop iterating can be reached sooner.

Finally, the mean cloud fraction shows a slightly higher value for the dual-view retrieval than the nadir-view, which can be explained as the dual-view retrieval checks for cloud in both views and so

is more likely to flag a pixel as cloudy. Although it decreases the potential coverage of the retrieval, it also acts as a secondary check to avoid cloud-contamination.

In Table 5.3 medians, as well as means, are presented. In most cases for both dual-view and nadir-view retrievals the distributions are skewed with the median value smaller than the mean. In the case of aerosol optical depth, this may be an indication that stronger aerosol events are comparatively uncommon (i.e. the background aerosol optical depth is low). For all state parameters this may also be because sea pixels outnumber the land, and so a median value is likely to correspond to a sea state (hence low median surface albedos in particular). For quality control statistics such as cost this is an indication that the majority of retrievals perform better than the mean would suggest, and a few outliers where the forward model fits poorly inflate the statistic.

5.4 Monthly mean data and comparison with MISR

5.4.1 Methodology of MISR comparison

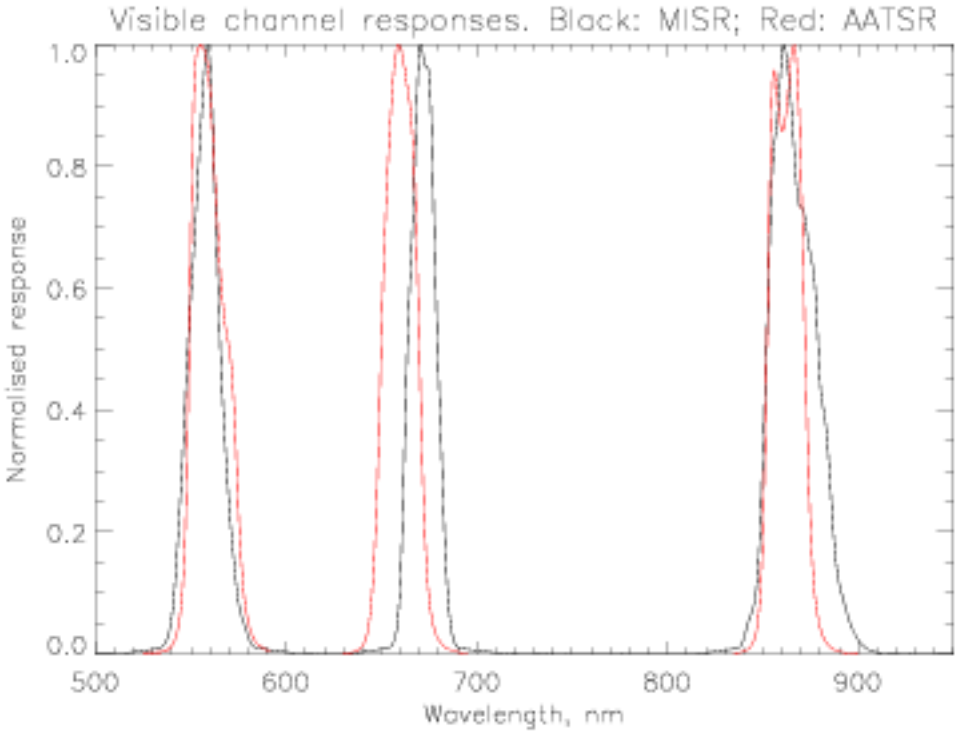

Figure 5.15: MISR (black) and AATSR (red) normalised visible channel response functions, for spectral regions where both instruments measure.

AERONET provides a long time-series of aerosol measurements at distinct locations; a major

advantage of satellites is their greater global coverage, so as an additional validation exercise the AATSR τ_{550} has been compared with MISR data. MISR was chosen as, like AATSR, it is a multi-angle sensor on a sun-synchronous polar orbit and so is able to provide global aerosol data over both land and sea. Monthly mean aerosol optical depth is compared rather than daily optical depth; the satellites orbit at different times, so there is no instantaneous overlap, and the overlap of daily orbital swaths is poor.

Figure 5.15 compares the spectral responses of MISR and AATSR for spectral regions where both instruments measure. There is a good correspondence at 550 nm and 870 nm but measurements at 660 nm are less directly comparable. MISR has an additional band near 440 nm but lacks the AATSR 1.6 μm and longer-wavelength channels.

The MISR dataset is on a $1^\circ \times 1^\circ$ grid. Therefore the AATSR orbits must be processed to provide an equivalent dataset, as follows:

1. All AATSR orbits are processed on an approximately 10 km$\times$10 km sinusoidal grid, and the limited lowest cost speciation method described in Section 5.1.4 applied.

2. Additional quality control constraints (cloud fraction less than 0.5, number of iterations greater than 1 and state variables not lying on state limits) are applied to the retrieved τ_{550} to filter which pixels to consider in the comparison.

3. The globe is divided into $1^\circ \times 1^\circ$ bins, matching the MISR data. For each bin, all pixels throughout the month lying within are averaged to provide a mean τ_{550} for that grid cell. The standard deviation, number of good pixels, and proportion of pixels corresponding to each aerosol class are also noted.

5.4.2 Intercomparison of optical depth

The resulting monthly mean maps of τ_{550} for MISR and AATSR are plotted (for points where both instruments have data) in Figure 5.16. Visual inspection reveals similar features in both maps, such as elevated optical depth in biomass burning (South America and west-central Africa), desert dust (China, Sahara desert and Atlantic Ocean off the coast of Africa) and industrial (south-east Asia) regions. In contrast aerosol optical depths over open ocean are generally low.

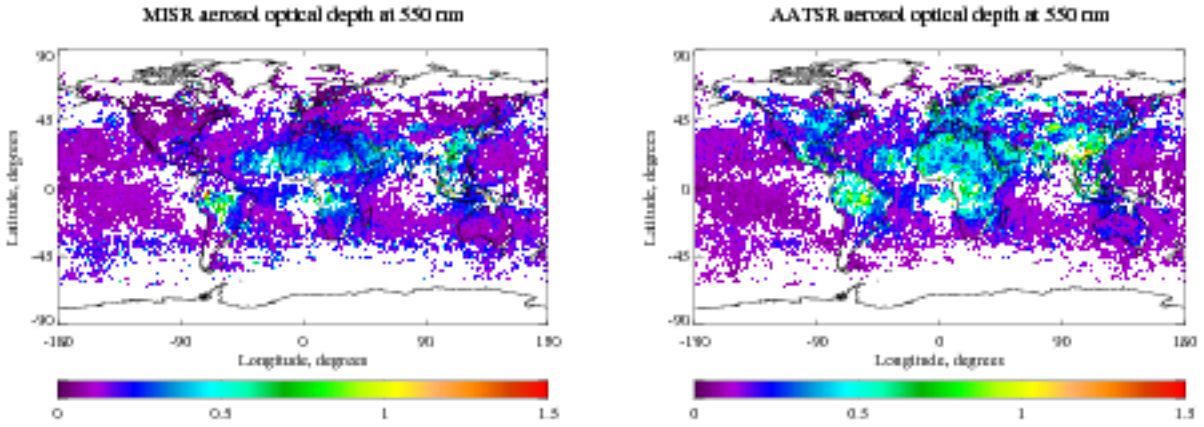

Figure 5.16: Monthly mean aerosol optical depth at 550 nm from MISR (left) and AATSR (right) on a $1^\circ \times 1^\circ$ grid. Data for the month of September 2004.

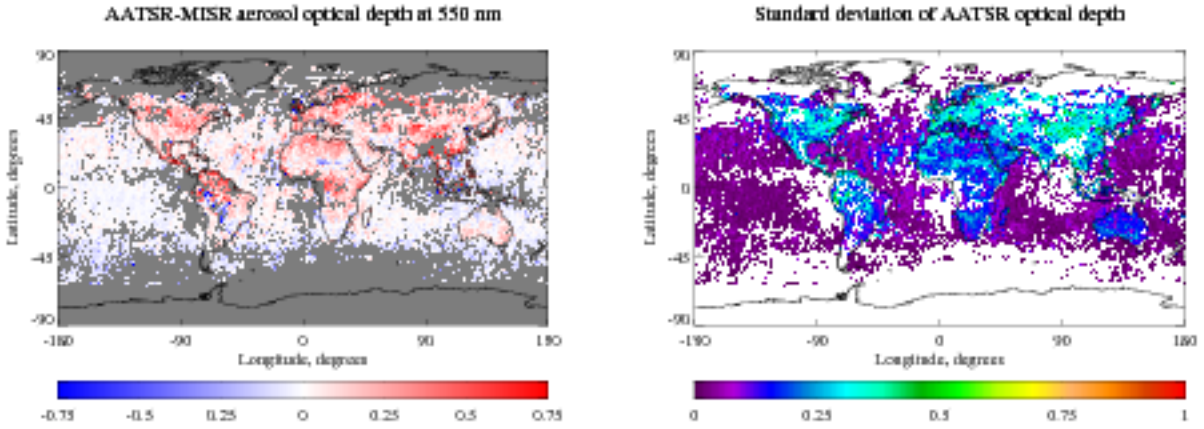

Figure 5.17: Left: Difference (AATSR-MISR) in monthly mean τ_{550}. Missing data are indicated with a grey fill value. Right: Standard deviation of AATSR optical depth. Data for the month of September 2004, on a $1^\circ \times 1^\circ$ grid.

Figure 5.17 shows the difference (AATSR-MISR) in monthly mean optical depth reported from the two satellites; over sea there are regions of high and low bias as compared to MISR, while over land AATSR generally has a high bias. There are several plausible reasons for the differences between the optical depth from the two sensors:

- AATSR generally retrieves lower-than-*a priori* surface albedo over land (Figure 5.11); with a duller surface, the AATSR retrieval may be compensating with an increased aerosol optical depth. If the MODIS BRDF product is correct, this would explain the discrepancy. The MISR data does not make use of the MODIS BRDF product.

- MISR lacks infrared channels and so its capabilities to detect cloud are poorer than those of AATSR. Hence a conservative cloud mask is applied to the data, which may lead to the removal

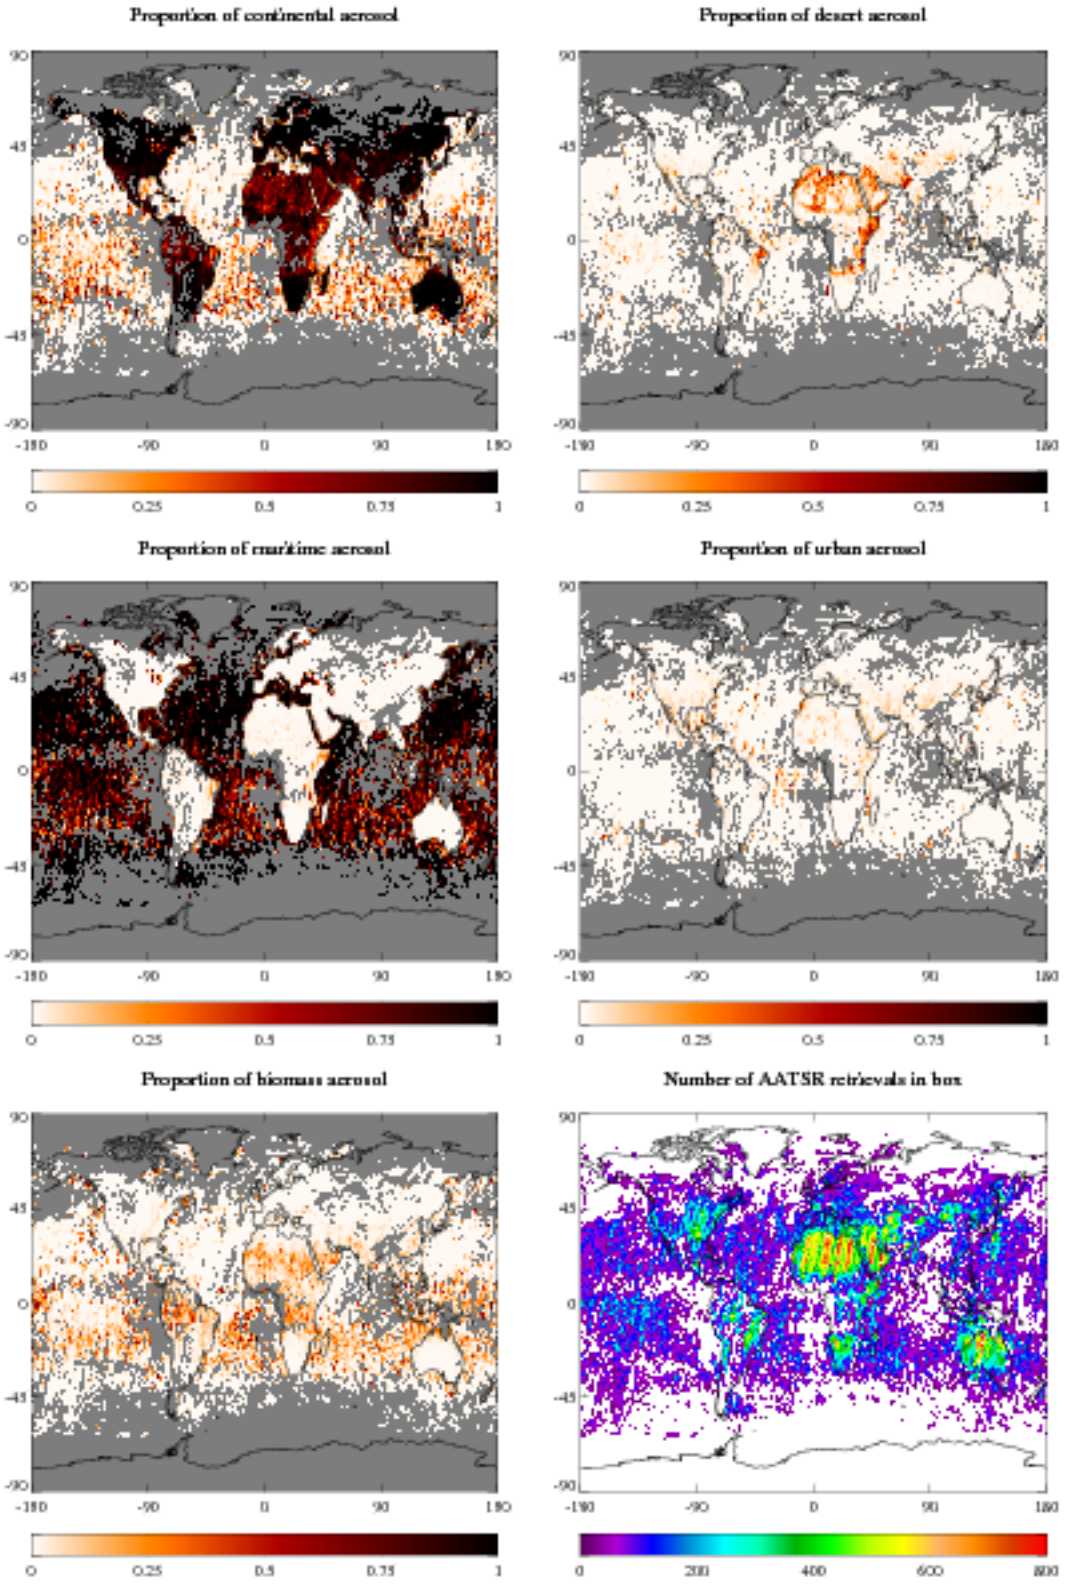

Figure 5.18: Speciation of aerosol retrieved by AATSR for the month of September 2004. The top row shows the proportion of retrieved pixels with each $1^\circ \times 1^\circ$ grid box assigned the continental aerosol class (left) and the desert aerosol class (right). The middle row shows the proportion of maritime (left) and urban (right). The bottom-left plot shows the proportion of pixels classified as biomass burning aerosol, and the bottom-right plot the number of good data points going into each grid box. For the speciation maps, missing data are indicated with a grey fill value.

of stronger aerosol events and a resulting low bias in the MISR aerosol field.

- The satellites' swaths are not coincident: MISR overpasses at 10:30 rather than 10:00 and has a 360 km swath. Additionally, the daily orbital paths do not track each other. Hence atmospheric conditions may differ when the satellites view a particular piece of the globe; Figure 5.17 shows a high standard deviation in AATSR τ_{550} in many regions where the AATSR-MISR difference is large. This is likely over, for example, Saharan dust regions where aerosol travels a long distance over a few days. The problem could be ameliorated by averaging over a longer timescale, which has not yet been done due to data access issues and computational time required.

- Assumed aerosol models for AATSR (shown in Figure 5.18) or MISR could be inappropriate, resulting in incorrect τ_{550}.

- The sampling of one or both of the satellites could be insufficient to provide a representative view of the aerosol within each grid box. Figure 5.18 also shows the number of AATSR data points within each box; this is highly variable, from approximately 600 over the Sahara, Australia and southern Africa (where skies are typically clear) to under 100 for much of the ocean (which is often cloudy). As AATSR has an orbital repeat cycle of 3 days, for some parts of the globe the monthly mean consists of data from a single orbit. This problem could also be reduced by averaging over a longer timescale.

Despite these difficulties, the correlation coefficient over land is 0.46 and sea 0.50, which are statistically significant. The fact that AATSR compares much more favourably to AERONET, which is an instantaneous comparison, than it does to the monthly-averaged MISR data, suggests that sampling issues can explain a large part of the discrepancy between the two datasets. The speciation of the AATSR data as shown in Figure 5.18, further discussed in Section 5.5, provides some confidence in the method used.

It should be emphasised again that perfect agreement between the two datasets is not to be expected so the differences are not necessarily indicative of problems with the retrieval. The instruments make different measurements in different locations at different times, and use different retrieval algorithms to obtain the optical depth. The correlations are significant, and both instruments report similar trends of τ_{550} over different regions of the globe.

5.5 Examination of other monthly mean AATSR data products

The monthly-mean datasets generated for the MISR comparison provide interesting information in themselves; although only τ_{550} is compared with MISR, as the retrieval is performed simultaneously for all parameters it is useful to examine them too.

Plots show only those regions of the globe where MISR optical depths are available over the month; as gaps are generally due to persistent cloud cover, there is little increase in coverage by adding in those grid cells where no MISR data are available but AATSR data are.

5.5.1 Speciation of data

Figure 5.18 shows the result of the speciation method described in Section 5.1.4 applied to the month of data (on an individual-pixel scale). The arising distributions of the various aerosol models are generally plausible, lending some confidence to this method.

Continental aerosol is predominant over land; continental flagged over the ocean may be aged biomass or desert aerosol (which, when mixed with the maritime aerosol, becomes more scattering). It is also likely that some is incorrectly flagged, due to incorrect optical properties for maritime aerosol meaning the model used is a poor fit at some geometries.

Desert dust is sometimes picked over the Sahara, though not over several other source regions. Additionally, it is absent in the plume off north-western Africa. These results show that the limited lowest cost speciation employed is likely too conservative—although the fact that some dust *is* flagged suggests that dust is present.

The maritime aerosol class is picked for most ocean pixels, and is generally not chosen over land. This is partially a result of the speciation method, but is also not surprising as the ocean is relatively homogeneous so the model would be expected to fit well globally.

Urban aerosol is rarely chosen and does not dominate anywhere. There is some indication that it is more prevalent over certain urban regions, such as parts of North America and southern Asia, although as it is also sometimes picked over open ocean or desert these may be artefacts.

Finally, biomass burning aerosol is found in the expected source regions of South America and southern Africa. Apparant biomass burning over the Sahara is likely an artefact of the speciation method; some also appears over southern Asia, which is plausible. Biomass aerosol present over

the sea is likely in some areas (off the coast of south-west Africa) but may be an artefact in others. The fact that biomass aerosol over the sea is often found in regions where continental aerosol is also picked may indicate ageing of the biomass aerosol; it may also indicate problems with the maritime phase function at certain scattering angles.

5.5.2 Effective radius

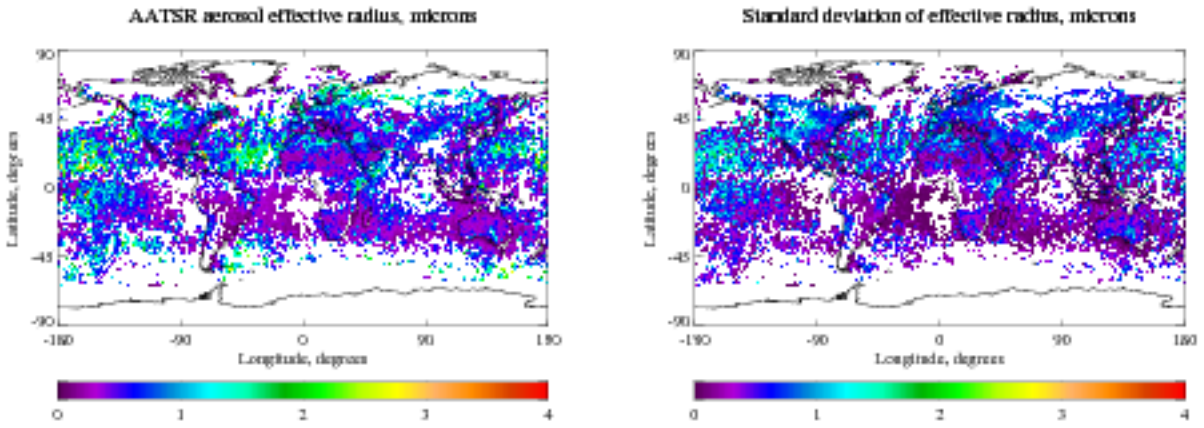

Figure 5.19: Left: Monthly mean aerosol effective radius (microns) derived from AATSR data. Right: Standard deviation of effective radius. Data for the month of September 2004, on a $1^\circ \times 1^\circ$ grid.

Particularly over land, the retrieved effective radius contains a strong constraint from the *a priori* data, and so the quality of the retrieved effective radius field is dependent upon the quality of the aerosol speciation. The monthly mean effective radius for each grid cell and its standard deviation are shown in Figure 5.19.

There is a visible plume of aerosol with a large effective radius off the west coast of the Sahara desert, which corresponds to desert dust aerosol being blown over the ocean. Figure 5.18 does not flag dust over the ocean, but the fact that dust-like particle sizes are being retrieved shows that, over the sea, information about the particle size can be retrieved despite possibly the wrong aerosol model being chosen. The lower effective radius over the Sahara itself shows that, for bright surfaces, most of the retrieved state comes from the *a priori* and so correct speciation is important.

Another identifiable feature is small particle sizes off the coast of south-west Africa, which corresponds to biomass burning particles blowing off the coast. This is coincident with the speciation picking out biomass burning aerosol. Other features in the effective radius field are difficult to interpret, although as τ_{870} (calculated from τ_{550} and r_e) shows reasonable agreement with AERONET

over land there is some confidence in retrieval results.

High standard deviations associated with larger effective radii may indicate heterogeneity within each grid cell, but may also arise from poor speciation (if several classes with differing *a priori* effective radii are flagged within each cell). Evidence would point towards heterogeneity within the regions, as they are often seen in Figure 5.18 to consist largely of just one aerosol type.

5.5.3 Surface albedo

The retrieved surface albedo, shown in Figure 5.20, looks much as would be expected. The sea surface is comparatively dark and the land brightness generally increases with wavelength.

The sea surface has a low standard deviation, suggesting spatial and temporal homogeneity. The land surface shows more variability, which is evidence of both spatial and temporal inhomogeneity. This is particularly pronounced at 870 nm, which is strongly affected by the presence of vegetation. These large standard deviations indicate why a dataset such as the MODIS BRDF used in this retrieval is used instead of assumed regional values: the land surface is too variable to adopt a simpler approach.

Retrieved land surface albedo has not been validated in this work. Simultaneous retrieval of all state parameters means that, to a certain extent, a good validation of aerosol properties means that the surface albedos described must be reasonable for the retrieval to arrive at that solution. The fact that similar models to that developed in Chapter 3 are widely used in the literature and retrieved and *a priori* data match as expected given uncertainties (Figure 5.12) adds further confidence. Validation of the land surface albedo against MODIS albedo is unsuitable as the latter is the retrieval *a priori*, but Figure 5.11 similarly indicates reasonable agreement between *a priori* and retrieved states.

It should be emphasised, however, that these comments do not constitute a validation of the retrieved surface albedo; rather, they suggest that the values are physically reasonable. No extensive validation was carried out due to time constraints and the difficulty of finding an appropriate dataset to validate against. Hence any use of this AATSR-derived surface albedo should take careful note of the quality control statistics as an indicator for where the data should be most reliable.

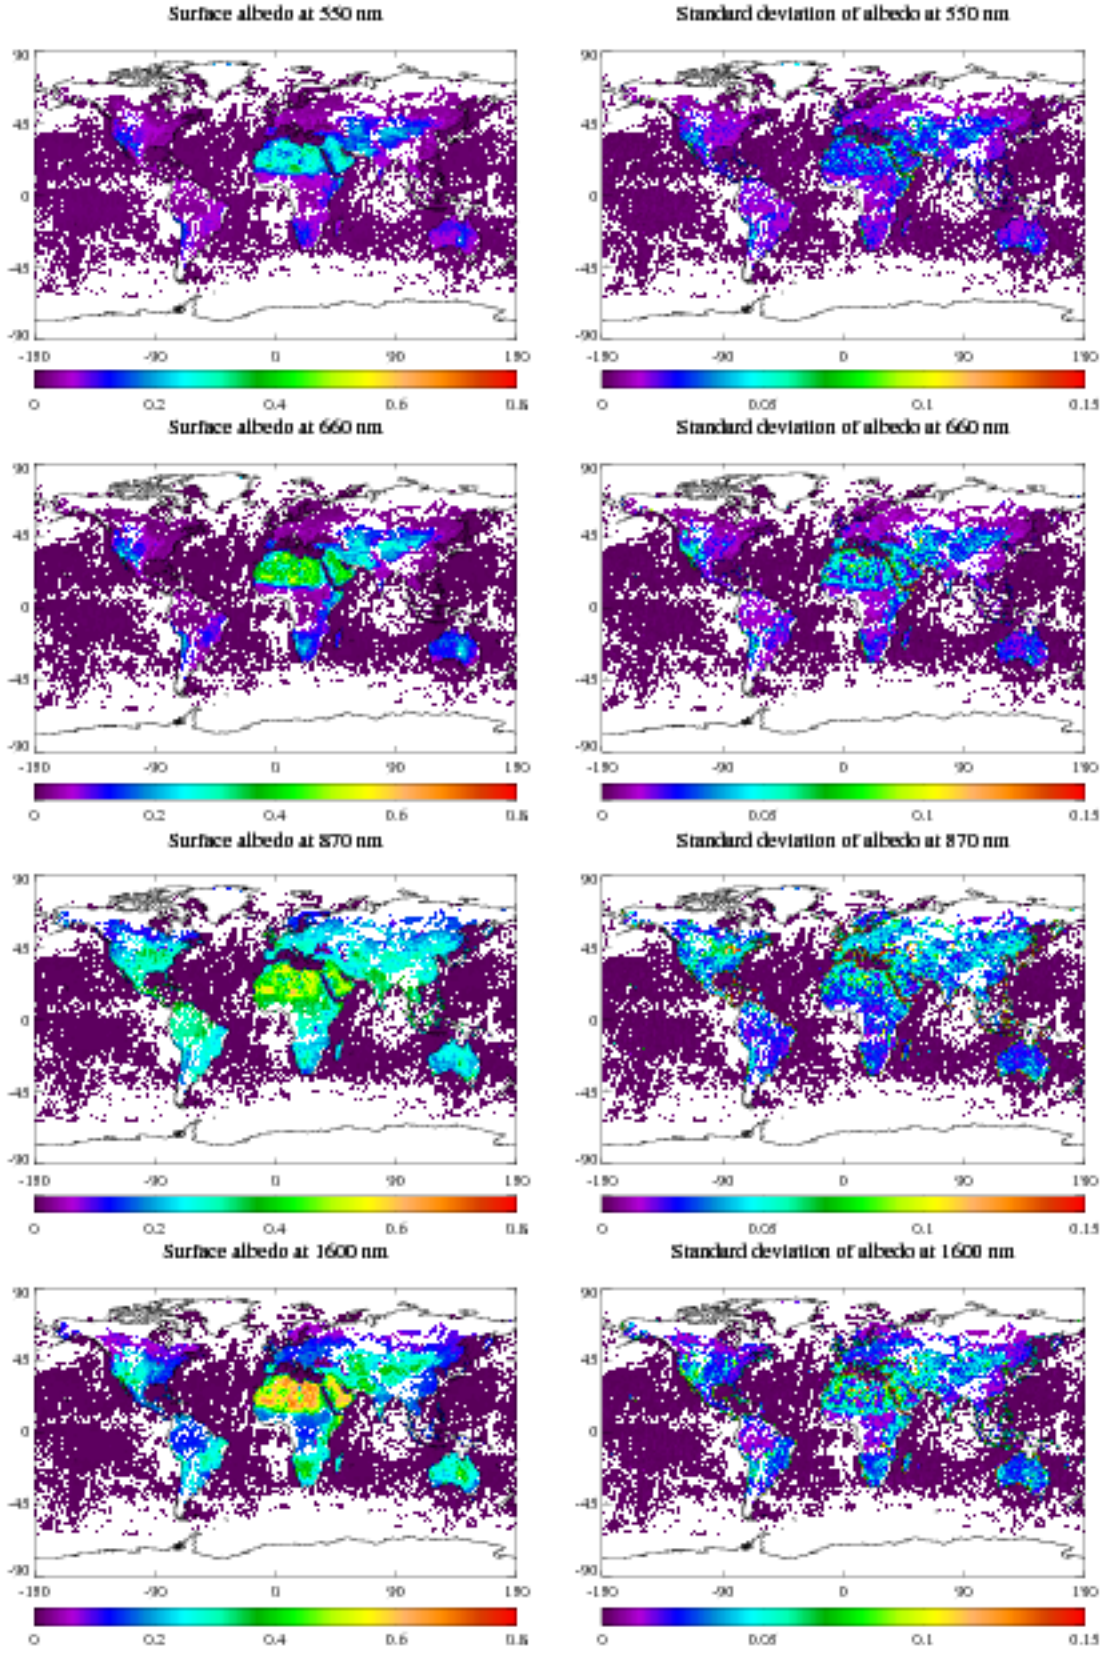

Figure 5.20: Monthly mean surface white-sky albedo derived from AATSR data on a $1^\circ \times 1^\circ$ grid. Data for the month of September 2004. The left-hand plots show the albedo and the right-hand the standard deviation within each grid cell. From top to bottom, data for 550 nm, 660 nm, 870 nm and 1.6 μm.

5.6 Summary of validation exercises

The validation exercises carried out have shown that the algorithm is able to retrieve aerosol properties well. Validation against independent AERONET data reveals a good match of the instantaneous observed τ_{550}; agreement for τ_{870} is poorer. This suggests some scope for further refinement of aerosol models and speciation. Lowest-cost speciation cannot accurately pick the most appropriate aerosol type in all cases but applying different quality control thresholds for the lowest-cost speciation improves matters. The dual-view results show a significantly better match to AERONET data than the previous nadir-view retrieval results.

An intercomparison with MISR τ_{550} reveals similar features in both datasets, although the agreement is poorer, due in part to the low sampling over a month.

Together validation of τ_{550} and τ_{870} describe the performance of the retrieved aerosol data products (τ_{550}, r_e and the aerosol models used). Surface albedo has not yet been validated. This is due both to time constraints and the difficulty of finding an appropriate datset to validate against. However, it provides physically reasonable fields and so should be suitable for analysis, if proper consideration is given to the retrieval quality control. Additionally, as the surface reflectance provides a large contribution to the measured signal over land the good validation of aerosol products suggests the surface characteristics are likely to be well-retrieved.

Examination of the statistics of the retrieval indicate that it performs adequately. Evidence suggests that the constraints on measurments and *a priori* data are appropriate; retrieval costs are higher than predicted, which indicates that the forward model does not model all aspects of atmospheric and surface variability. This is to be expected as certain parameters, such as aerosol optical properties and vertical distribution, are fixed in the retrieval.

Again, the dual-view retrieval provides generally smaller estimated uncertainties on retrieved state variables. This is due to a combination of the fact that 8 instead of 4 measurements are used, and the improved error analysis going in to $\mathbf{S}_y$ and $\mathbf{S}_x$.

Several areas in which further improvement is possible are, however, raised by this exercise:

- The retrieval shows a preference for τ_{550} and r_e slightly to one particular side of each LUT point. This introduces a small bias into results, and arises due to discontinuities in the gradient of the (linearly-interpolated) forward model at LUT points, meaning weighting functions are steeper on one side of the LUT point. This could be fixed by higher-order LUT interpolation,

providing a continuous forward model gradient.

- Surface albedo retrieved over land is generally lower than the *a priori* value from MODIS data. It is not certain whether this is due to an underestimate in the MODIS data, some factor due to the different instruments' measurements or overpasses, or an issue with the AATSR retrieval algorithm. New versions of MODIS data, and research into the temporal variability of surface reflectance, may shed light on this matter.

- Cost-based speciation cannot distinguish between aerosol types for low aerosol loading unless tuned. Addition of infrared channels to the retrieval, while a complicated problem in itself due to the additional forward model required, will increase the contrast between the different aerosol classes and improve the speciation. However in the case of small particles or low aerosol loadings, the improvement will be slight.

- The retrieved effective radius over land shows low averaging kernel values indicating a significant contribution from the *a priori*. Effective radius retrieval (and hence τ_{870} calculation) will be improved by addition of infrared channels for high aerosol loadings or large particles, due to their impact on infrared radiance. Again the impact of this improvement will be small for small particles or low aerosol loadings.

- Consideration of nonspherical scattering is likely to improve the aerosol phase functions used, particularly for the desert dust class.

Chapter 6

Applications of retrieval output data

In this chapter various applications of the retrieval's output data products are explored. Firstly, the ability of the retrieved surface albedo to provide atmospherically-corrected false-colour images of the Earth's surface is described. Secondly, the long lifetime of the AATSR mission is put to use to investigate differences in aerosol and surface properties on a global scale between August 2002 and August 2007. Thirdly, a region of South America focussing on the Amazon is examined for the whole of this period, and retrieval results interpreted in the context of the area.

The studies presented in this chapter are not intended to be exhaustive. They are applications of an exploratory nature and show the diverse uses to which the data products generated by the retrieval scheme may be put.

6.1 False-colour images of the Earth's surface

The retrieved surface albedo datasets provide information on what the surface of the Earth would look like if the atmosphere was not 'in the way'. The albedo may be processed to provide a false-colour image of the Earth, for a visualisation of the different types of surface cover in the world. This requires the selection of different wavelengths for each of the red, green and blue colours. In this regard AATSR is at a disadvantage when compared to sensors like MODIS and MERIS, as with only four wavelengths in the visible/near-infrared region, and particularly no measurement at blue wavelengths (around 440 nm), the colours of the resulting false-colour image will be distorted.

In this analysis, the 670 nm channel is taken as the 'red' measurement, 870 nm as 'green' and 550 nm as 'blue'. The choice of red channel is appropriate as it lies in the red region of the spectrum.

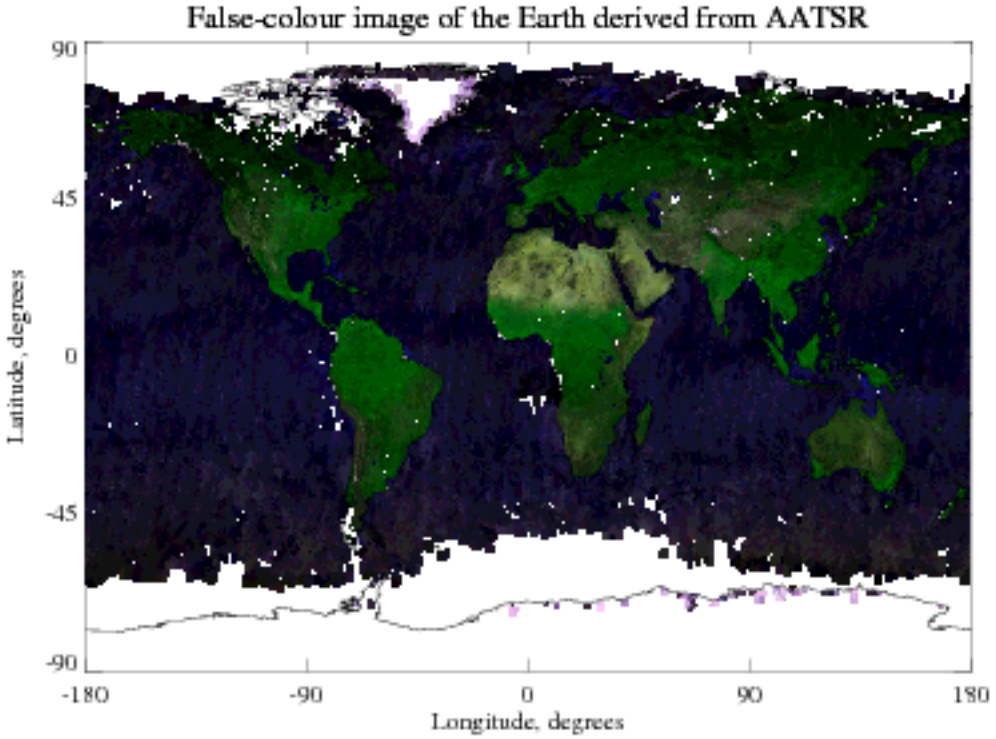

Figure 6.1: False-colour image of the Earth's surface, obtained from AATSR visible-channel surface albedo for September 2004.

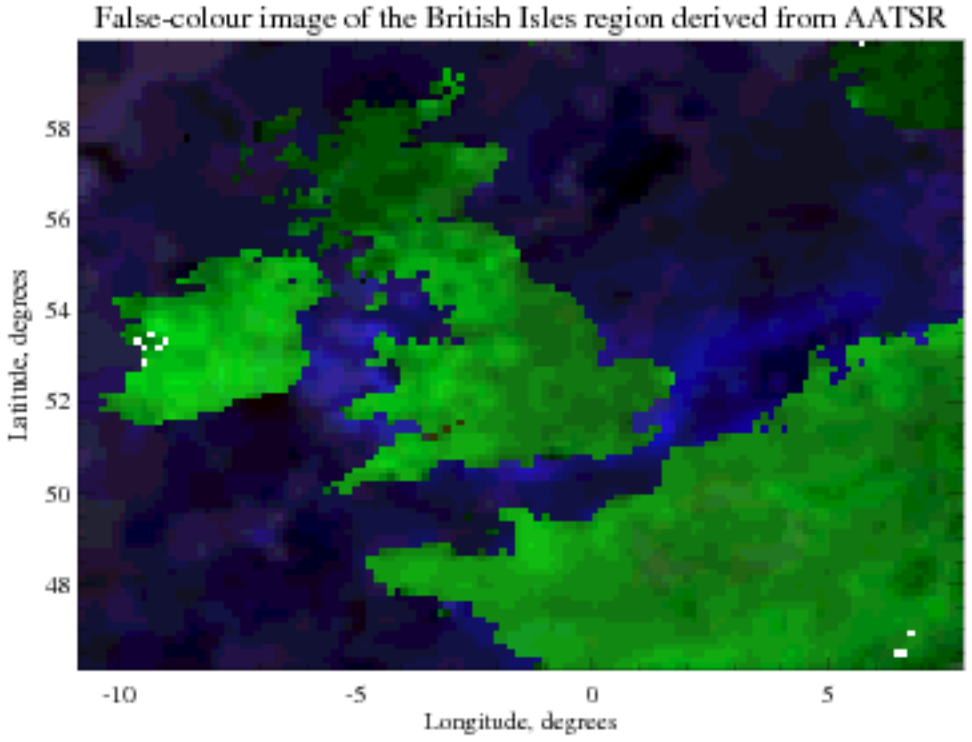

Figure 6.2: False-colour image of the region around the British Isles, obtained from AATSR visible-channel surface albedo for September 2004.

Although the 550 nm channel is in the green region of the spectrum, it is chosen instead for the blue RGB intensity as it is the closest of channels to blue wavelengths. 870 nm is an appropriate proxy for the green intensity as vegetation is bright at 870 nm as well as green wavelengths. The surface albedos at the three wavelengths are then normalised and scaled by the number of levels in the RGB colour-table of the plotting software (8-bit for IDL 7.0 file output) to produce the false-colour image.

This method has been applied to the surface albedo dataset obtained for September 2004, as described in Chapter 5. Small gaps in the data have been filled by nearest-neighbour sampling (separately for land and water pixels), and the resulting fields smoothed with a low-pass filter (with the final value of albedo at each point being composed of 4/7 of the pixel's albedo and 3/7 the average of its surrounding pixels). Additionally, the brightness of the sea surface has been enhanced 15-fold to allow variations in ocean colour to be seen by eye. The resulting global map of the surface is shown in Figure 6.1, with Figure 6.2 showing the region around the British Isles in more detail.

The fact that the resulting map intuitively resembles the surface of the world, without major distortions due to aerosol or cloud contamination in the surface reflectance, lends confidence to the retrieval results. Different types of land cover are all distinct within the image. Over the ocean the variations in water colour due to differing plankton populations and the presence of turbid matter are also visible.

Closer examination of coastal waters in Figure 6.2 shows strong signatures of differing water colour (in the Irish Sea, Severn Estuary, English Channel and the region north and east of the Thames estuary in particular) as compared to non-coastal regions of the North Sea. As the blue component of the RGB value used in the false-colour image comes from AATSR's 550 nm channel, in reality these very blue waters would appear very green, and this is consistent with turbid, shallow waters (around the coasts) and dirty waters from busy shipping lanes (e.g. the Irish Sea). Although AATSR is not a dedicated ocean-colour sensor and the surface reflectance model described in Chapter 3 was optimised for open Case I waters rather than turbid Case II, the retrieval is still able to distinguish these different regions of water to a degree.

The necessary distortion of the true colours is clear in Figure 6.1, where snow-covered land appears pale purple rather than white and very dark waters in the region of 50° S and further poleward appear brown, as well as the greenish tinge to bare land. The sparse sampling of AATSR over the month is also evident by the persistence of gaps in typically cloudy or snow-covered environments,

and the presence of features resembling orbital tracks in some of the data over the sea.

The utility of these false-colour images depends on their intended purpose; for display, it is suggested that using more than 1 month of data and attempting some correction for AATSR's measurement wavelengths would be desirable. To monitor regional land-use or ocean colour changes over time, performing a retrieval at higher resolution would allow observation of smaller-scale features such as urban areas in more detail.

6.2 Comparison of the world in August 2002 and August 2007

6.2.1 Methodology

As a well-calibrated sensor (Smith, 2008, Smith et al., 2008) with a long time-series (on satellite timescales) of observations, and the potential to combine with the ATSR-2 record, the AATSR sensor is suited to monitor global change on an interannual timescale. In this section retrieval results from the months of August 2002 and August 2007 are compared. August 2002 is the earliest-available full month of data, and August 2007 the most recently-available August. Although slightly more recent months of level 1 data are available, it is preferable that the two months chosen should be the same, to mitigate the influence of seasonality on the comparison.

Examination of the two months and their differences provides information on both the variability of aerosol and surface conditions within a month, and also on longer-term changes. Although two points are insufficient to draw conclusions about global trends, picking one early-mission month and one late-mission increases the likelihood that any significant trends will show up in the data. Processing of sufficient data for consideration of global long-term aerosol trends is out of the scope of this work and carries a significant computational overhead; it is being undertaken by the GlobAerosol project (Thomas et al., 2006), which involves ORAC with ATSR-2 and AATSR.

Both months were processed with the algorithm on the standard 10×10 km sinusoidal grid. As no MODIS Collection 4 surface data were available for 2007, for consistency both months were processed with the newer Collection 5 dataset. The retrieval output data were speciated as described in Section 5.1.4 and, due to a poor level of overlap on an individual-pixel scale between the two months (due to missing data from cloud cover), the data were averaged onto a 2.5° grid after applying the quality control measures described in Sections 5.1 and 5.4.

To minimise the impacts of poor sampling on the comparison, grid cells containing fewer than 20 pixels in either month are excluded from the analysis.

6.2.2 Difference in aerosol properties

Mapping global differences

The monthly mean retrieved τ_{550}, its standard error (standard deviation divided by the square root of the number of retrievals in the grid cell) and the difference between the two months is shown in Figure 6.3. These reveal generally similar patterns of aerosol optical depth and variability for both months (as well the September 2004 results previously shown in Figure 5.16). The similarity is an indication that many global aerosol sources are generally persistent (i.e. on a 5-year timescale, the locations of deserts, forests, cities and oceans do not change much).

The right-hand plots provide a measure of the variability of the aerosol within a month; again, these are reasonably consistent between the two months, indicating that the spatio-temporal variability of aerosol is similar from one August to the next. The relatively high variability of the land as compared to the sea is also visible; this is expected, as the sea is largely homogeneous, although the standard errors for both land and sea are frequently of the order of 0.01 or smaller. The smoothness of the fields suggests that the majority of the variability comes from real spatio-temporal variability in the aerosol, rather than noise on the individual retrievals.

Examination of the differences in optical depth between the two months reveals clear regions across several 2.5° grid cells where the optical depth is persistently higher or lower in one month. The bottom-right part of Figure 6.3 shows that this difference between the two months is significant over most parts of the world. For this calculation the difference in τ_{550} is divided by the root sum of squares of both months' standard errors; this gives a measure equivalent to the z-score for a point in a Gaussian distribution with mean zero (no optical depth change) and variance equal to the sum of the standard errors. If this z-score exceeds that of the 90% significance threshold ($z = 1.645$) then the difference is termed 'significant' in this discussion. In this context, it indicates that the probability of observing such a difference in τ_{550} betwen the two months if the true τ_{550} has *not* changed is less than 10 %.

The significance of the differences over most grid cells—even over the ocean, where changes in optical depth may be very small—makes it difficult to state with confidence any possible long-term

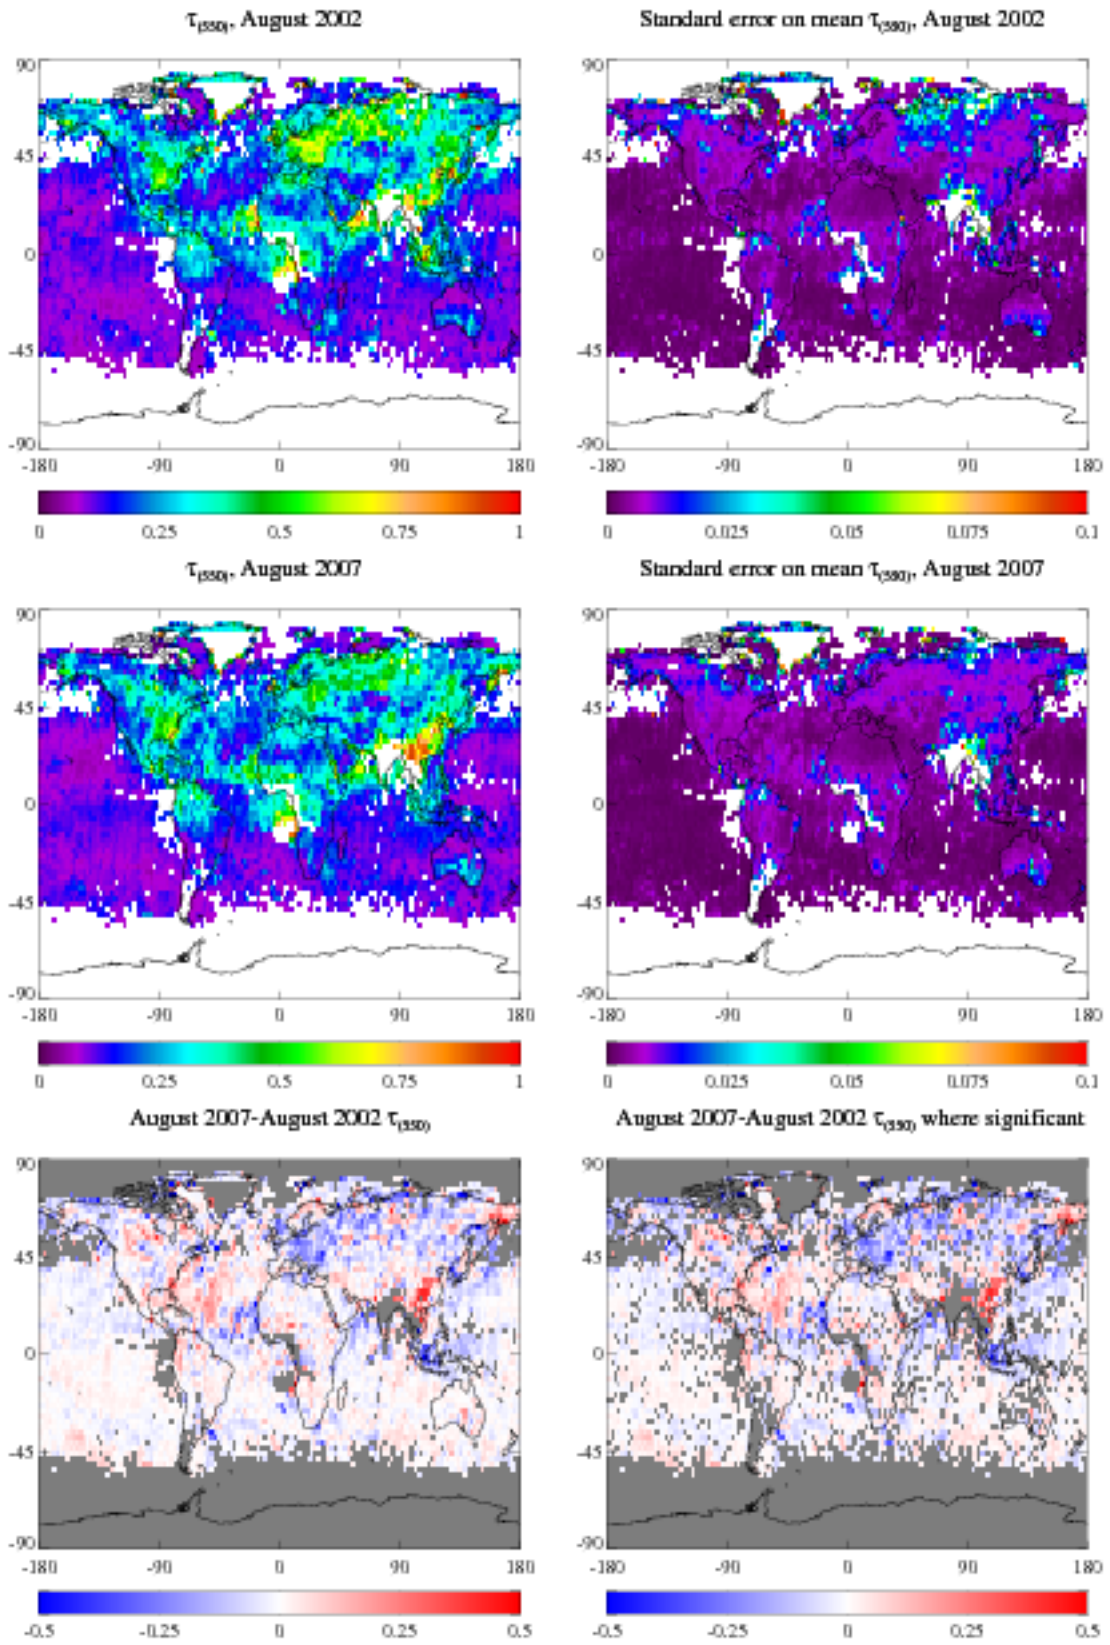

Figure 6.3: Retrieved τ_{550} in August 2002 and 2007 on a 2.5° grid. Top-left shows the mean optical depth for August 2002, and top-right the standard error on the mean for each grid cell. The middle row is the same as the top, except for August 2007. The bottom-left shows the 2007-2002 difference in aerosol optical depth, with the bottom-right the same except only plotted where significant at the 90 % confidence level. In the difference plots, missing data are indicated by a grey fill value.

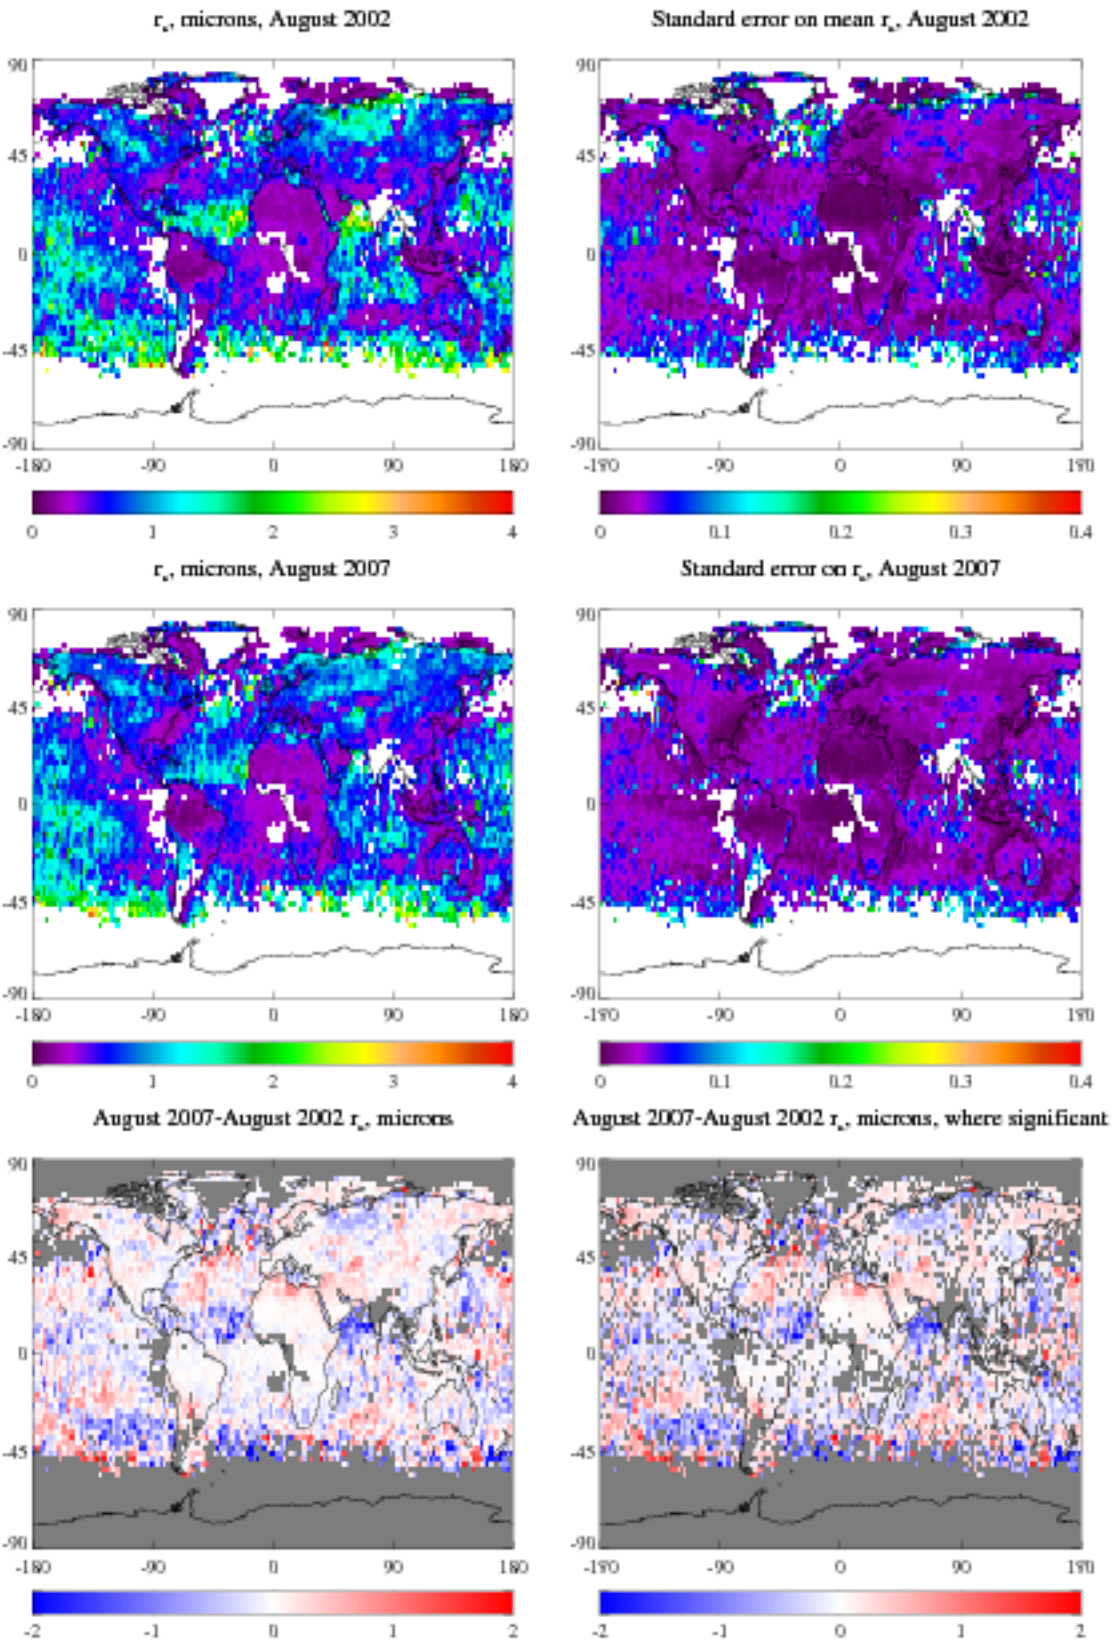

Figure 6.4: Retrieved r_e in August 2002 and 2007 on a 2.5° grid. Top-left shows the mean effective radius for August 2002, and top-right the standard error on the mean for each grid cell. The middle row is the same as the top, except for August 2007. The bottom-left shows the 2007-2002 difference in effective radius, with the bottom-right the same except only plotted where significant at the 90 % confident level. In the difference plot, missing data are indicated by a grey fill value.

trends in the data, unless the 5-year impact of such a trend is large. It is likely that many of the regions where a small difference is observed correspond to a change in some factor such as local humidity or aerosol height which is not captured by the retrieval and so leads to a bias in results in one month, or a difference in meteorological conditions between the two years causing, for example, faster wet removel of aerosol. In this respect it is a further indication that more data are needed to separate meteorological effects from long-term changes. The comparison performed is still useful to shed light on the magnitude of such natural interannual variability.

Similar comments regarding significance may be made for the change in effective radius (shown in Figure 6.4) with the caveat that, as effective radius can change significantly depending on the assumed aerosol model, and speciation is imperfect and unreliable for low aerosol loadings, extra care must be taken with the interpretation of results. The combination of Figures 6.3 and 6.4 highlights various regions around the world where the differences can be explained.

- In general, over most land regions the change in aerosol properties is small and likely explicable in terms of changes in meteorological conditions at the time of sampling.

- The change in both τ_{550} and r_e is significant over the Atlantic off the western Sahara and the Indian Ocean off the Arabian peninsula, with more aerosol and larger particles in August 2002 (and the reverse further out into the North Atlantic). This is likely due to differences between the two months in the wind, which blows dust off the Sahara. The r_e retrieval over the Sahara itself is poor due to the bright surface, explaining the apparent land-sea discontinuity.

- Over much of the rest of the ocean there are significant small differences, which may be due to changes in aerosol loading arising from changes in plankton populations or wind speed. This could be linked to the fact that August 2002 was during an El Niño episode and August 2007 was not (NOAA, 2008), and El Niño events impact upon sea surface temperature and wind speed in the tropics (Schneider, 1996a). Differences in the wind may cause the change in particle size observed in the storm track region around 45° S.

- Parts of Europe and China show a significant decrease and increase in τ_{550} (with small increases and decreases in particle size) respectively. These may indicate cleaning-up of eastern Europe and increased industrial activity in China; it is suggested that these would be good places to generate long time-series of data to look for trends.

- Regions at high Northern latitudes show occasional significant differences in one or both aerosol parameters, which may be due to differences in fire activity or snow cover between the two months.

One possibility for apparant changes in aerosol loading is the algorithm retrieving a higher optical depth but lower surface albedo when the true state has not changed, and vice-versa. This is not thought to be a problem, as the correlation between the change in optical depth and change in surface albedo is small (0.08 at 550 nm and similar for other wavelengths). Correlations between other changes are small and generally similar to correlations in the retrieved values themselves (0.17 for τ_{550} and r_e and 0.08 for r_e and $R_{SLW(550)}$), suggesting that this compensation is not a problem for any of the parameters.

Latitudinal and frequency distributions of aerosol properties

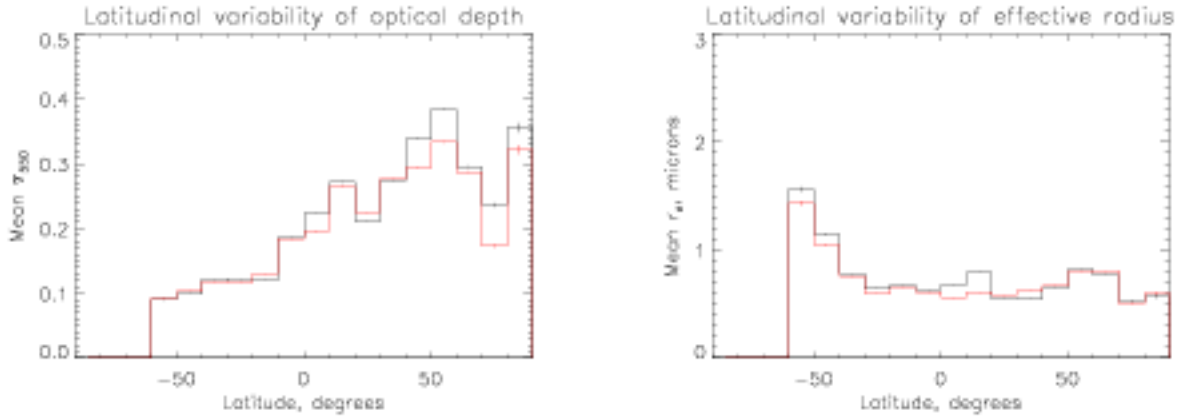

Figure 6.5: Latitudinal mean aerosol optical depth (left) and effective radius (right) for the two months. Results for August 2002 are shown in black, and those for August 2007 are in red. Error bars indicate the uncertainty on each latitude band (the standard deviation of pixels in that band divided by the square root of the number of pixels averaged in it).

There is additional information about the state of the atmosphere in the latitudinal distributions of the retrieved data. In this analysis distributions are constructed from results on the full sinusoidal grid used in the retrieval, for sinusoidal grid cells where both months have successful retrievals. This ensures that the geographic coverage of the distributions is the same between the two datasets.

Figure 6.5 shows 10° latitudinally-averaged aerosol optical depth and effective radius over the months, with Figure 6.6 detailing the number of retrievals in each latitude bin. Note that these plots should not be directly used as a reference for the true mean aerosol properties across each latitude

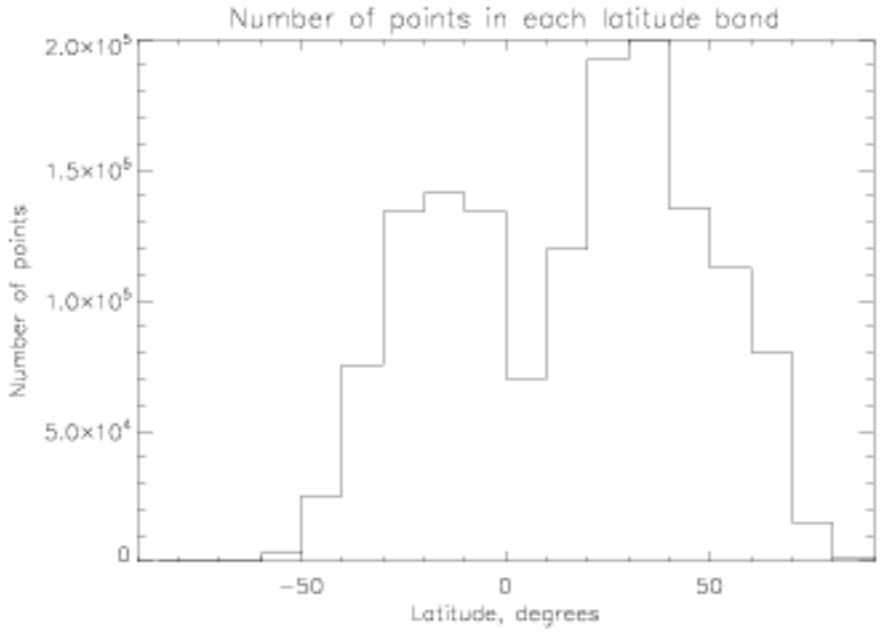

Figure 6.6: Histogram of the number of data points in each latitude band in this analysis.

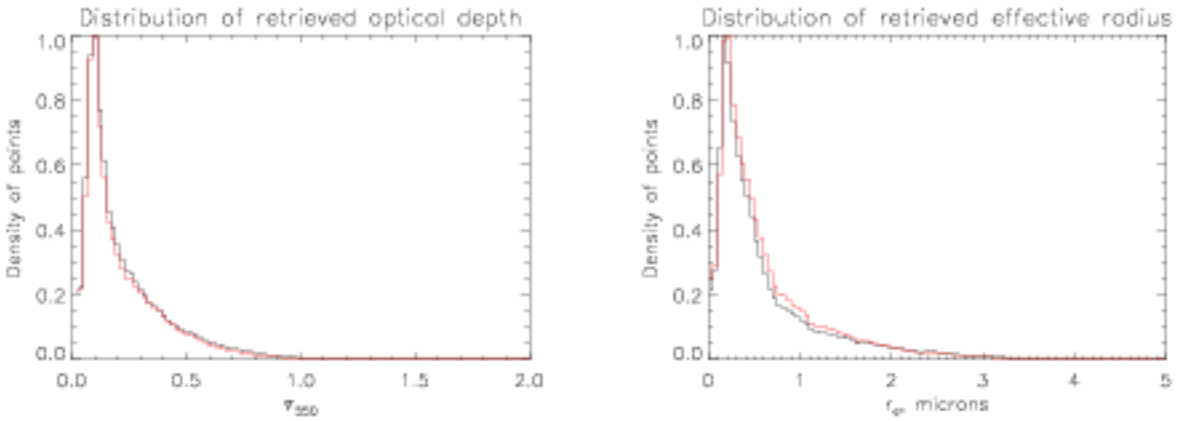

Figure 6.7: Frequency density plot of retrieved aerosol optical depth (left, bin size 0.02) and effective radius (right, bin size 0.05 microns) for the two months. Results for August 2002 are shown in black, and those for August 2007 are in red. Distribution normalised to the most populated bin.

band, as the coverage of the data is incomplete. Coverage at high latitudes is poor or absent; over the southern polar region this is due to the sun being very low in the sky or beneath the horizon at AATSR overpass times.

Overall, the latitudinal trends in both years and both aerosol properties are similar. This is an indication that they share similar global aerosol source regions. There is a suggestion of decreased aerosol optical depth in the Northern Hemisphere and decreased particle size in the Southern Hemisphere and particularly from 0°-20° N in 2007. As the bulk of the world's population lives in the

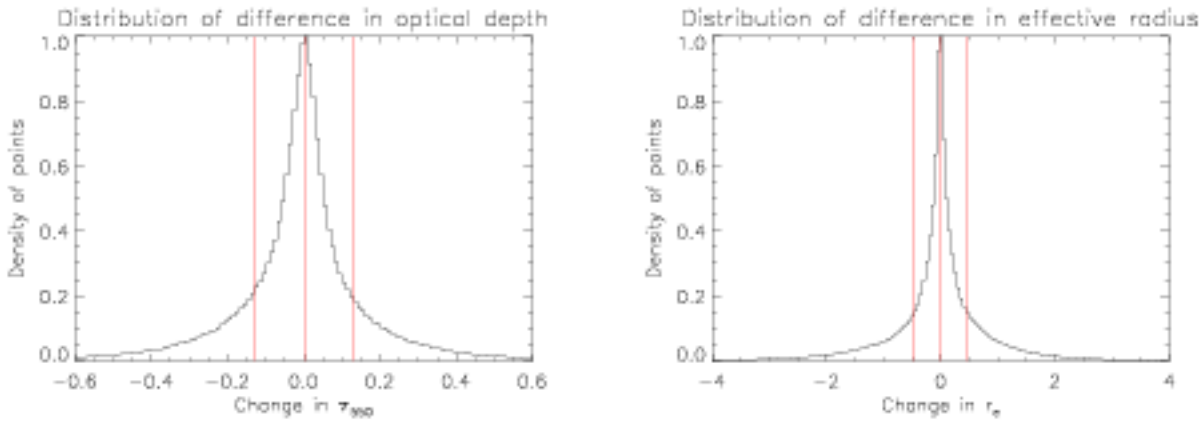

Figure 6.8: Frequency density plot of the differences in retrieved aerosol optical depth (left, bin size 0.01) and effective radius (right, bin size 0.05 microns) between the two months. The red lines indicate zero difference and $\pm$ the root sum of squares mean retrieval error on each parameter from the two months. Distribution normalised to the most populated bin.

Northern Hemisphere, the decrease in optical depth could be due to a decrease in industrial emissions resulting from improved technology. Reasons for the change in particle size are unclear, although the shifts in tropical regions could be due to different patterns of transport of Saharan dust between the two months. Shifts may also be due to errors in the aerosol speciation for both months.

The distributions of retrieved aerosol properties as shown in Figure 6.7 are similar for the two months. August 2007 has a slightly thinner distribution of τ_{550} than August 2002; this is linked to the lower latitudinal-mean τ_{550} shown in Figure 6.5. The retrieved size distribution also shows that 2007 has comparatively more aerosol with $0.2\ \mu\text{m} < r_e < 2\ \mu\text{m}$; the reasons for this are unclear, although it could, for example, indicate 2007 had less small pollutant aerosol or more large mineral dust.

Distributions of the differences between the two months are given in Figure 6.8, along with estimates of the mean uncertainty on each retrieved parameter. These estimates are the root sum of squares of the mean 1σ retrieval error on τ_{550} (or r_e, as appropriate) for each month. Differences smaller than this value can likely be ascribed to noise on the retrieved state. This is a simplistic measure of the significance of any difference as the retrieval error is different for each pixel, although it does provide a quick measure of how many of the differences are likely to be due to a real change in the atmosphere. Both distributions peak close to 0 and have the majority of their area within the red bars, suggesting that aerosol in most regions of the globe was similar in August 2007 to how it was in August 2002. In this respect, both τ_{550} and r_e have a similar proportion of their area within the bars, suggesting similar interannual variability (within the limits of retrieval precision).

6.2.3 Difference in surface properties

Mapping global differences

Rather than provide figures analagous to Figure 6.3 for the retrieved surface albedo, changes in the normalised different vegetation index (NDVI) are considered instead. The NDVI has long been known and used as a semi-quantitative measure of the density of vegetation via chlorophyll abundance (for example Tucker, 1979, or Myneni et al., 1995 for a more mathematical treatment). Hence, consideration of the change in NDVI between the two months can be a useful indicator of where land use has changed which is more immediately obvious than a direct comparison of the surface albedo. The index is defined as follows:

$$\mathrm{NDVI} = \frac{R_{\mathrm{nIR}} - R_{\mathrm{red}}}{R_{\mathrm{nIR}} + R_{\mathrm{red}}} \tag{6.1}$$

The values R_{nIR} and R_{red} are the observed reflectance in the near-infrared and red regions of the spectrum respectivety. It follows that the ratio can take values between -1 and 1. Plants scatter light at near-infrared wavelengths and absorb that in the red for use in photosynthesis; hence, dense vegetation has an NDVI approaching 1, with the NDVI of sparse vegetation being lower (Holben and Kimes, 1986). The NDVI of bare soil is typically small and positive (0.3 or less), while that of water can be small and negative (as the brightness of water, as apparant from Chapter 3, decreases with wavelength).

Here, R_{nIR} is taken as $R_{\mathrm{SLW}(870)}$ and R_{red} as $R_{\mathrm{SLW}(660)}$. An advantage of AATSR over application to older satellites such as AVHRR (Holben and Kimes, 1986), is that by using the retrieved surface albedo and not the measured TOA reflectance AATSR is able to provide accurate atmospherically-corrected NDVI values while older sensors cannot as they lack sufficient information to do so. The precision σ_{NDVI} of the AATSR-derived NDVI is typically 0.03-0.06, and may be determined on a pixel-by-pixel basis from the proportional uncertainties $\sigma_{R_{\mathrm{SLW}}}$ on each R_{SLW}:

$$\sigma_{\mathrm{NDVI}} = \mathrm{NDVI}\sqrt{\left(\frac{\sigma^2_{R_{\mathrm{SLW}(660)}}}{R^2_{\mathrm{SLW}(660)}} + \frac{\sigma^2_{R_{\mathrm{SLW}(870)}}}{R^2_{\mathrm{SLW}(870)}}\right)} \tag{6.2}$$

NDVI mapped over the land for both months is shown in Figure 6.9. Regions of dense vegetation are found throughout much of the tropics regions as well as agricultural regions (largely in the Northern Hemisphere). In contrast, arid and dry regions have a lower NDVI. Before consideration is

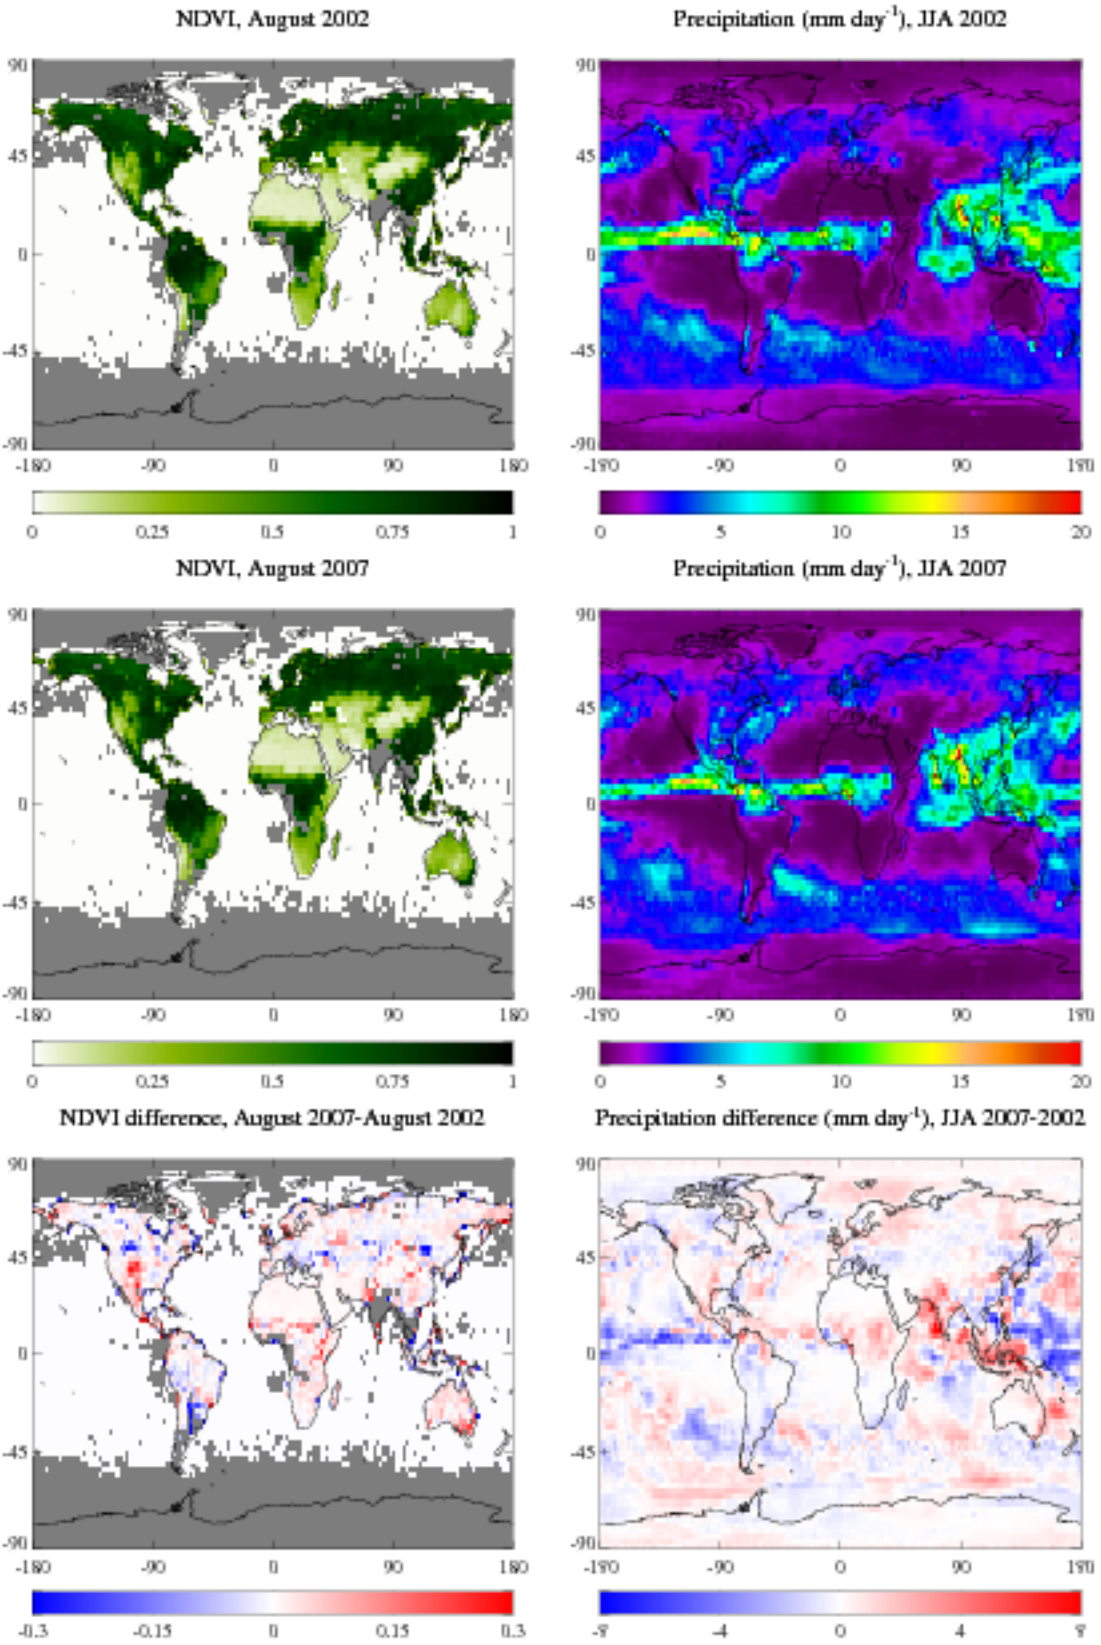

Figure 6.9: NDVI and precipitation in August 2002 and 2007. Top-left shows the mean NDVI for August 2002 on a 2.5° grid. Top-right is a seasonal mean precipitation map for JJA 2002 derived from GPCP data (Adler et al., 2003) on a 2.5° grid. The middle row is the same as the top, except for 2007, and the bottom row shows the difference in NDVI (left) and precipitation (right) between the two months. In the NDVI plots, missing data are indicated by a grey fill value.

given to differences in the NDVI of the two months, the following points should be mentioned:

- As the retrieval is constrained by the MODIS BRDF product, derived products are not completely independent datasets (although averaging kernels for surface albedo, as well as optical depth, are often close to 1, indicating the majority of the retrieved state comes from the measurements). Therefore changes in the quality of the MODIS data could manifest as small changes in the NDVI. The newest MODIS data were used in both months to minimise this potential source of error.

- The NDVI is a semi-quantitative measure of the chlorophyll abundance in a region (Myneni et al., 1995) but changes in NDVI cannot be easily directly mapped to a corresponding change in plant number, type or size.

- Looking on a scale of 2.5°, small changes in land cover may not manifest in detectable changes in NDVI.

Also shown in Figure 6.9 are global precipitation maps at 2.5° resolution derived from Global Precipitation Climatology Project (GPCP) data (Adler et al., 2003). This provides monthly satellite-based estimates of precipitation. Here the June, July and August precipitation estimates have been used to generate a seasonal mean precipitation map and difference between the two years. Over many regions of the world's land—notably the North, Central and parts of the South Americas, central Africa, Europe and much of Asia—there is a strong link between increased precipitation in 2007 and increased NDVI in 2007, and vice-versa. This can be explained by more rainfall in preceding months leading to increased growth of plants, or a comparative drought leading to less plant matter. This suggests that a large part of the difference in NDVI is due to different meteorological conditions in the two years, as opposed to land-use change (although, as mentioned, small-scale land-use change is likely to be missed using this technique).

It is possible that the strong positive precipitation difference over Indonesia and parts of Papua New Guinea may have led to increased wet removal of aerosol in these regions, explaining the observed decrease in aerosol optical depth there (Figure 6.3). There are few obvious regions of an anti-correspondence between precipitation change and optical depth change.

Aside from some regions of central Asia, there is little NDVI change to suggest land cover change over most of the world. Therefore it is reasonable to conclude that, on a 2.5° scale, land cover in

August 2002 and August 2007 was generally the same.

Latitudinal and frequency distributions of surface properties

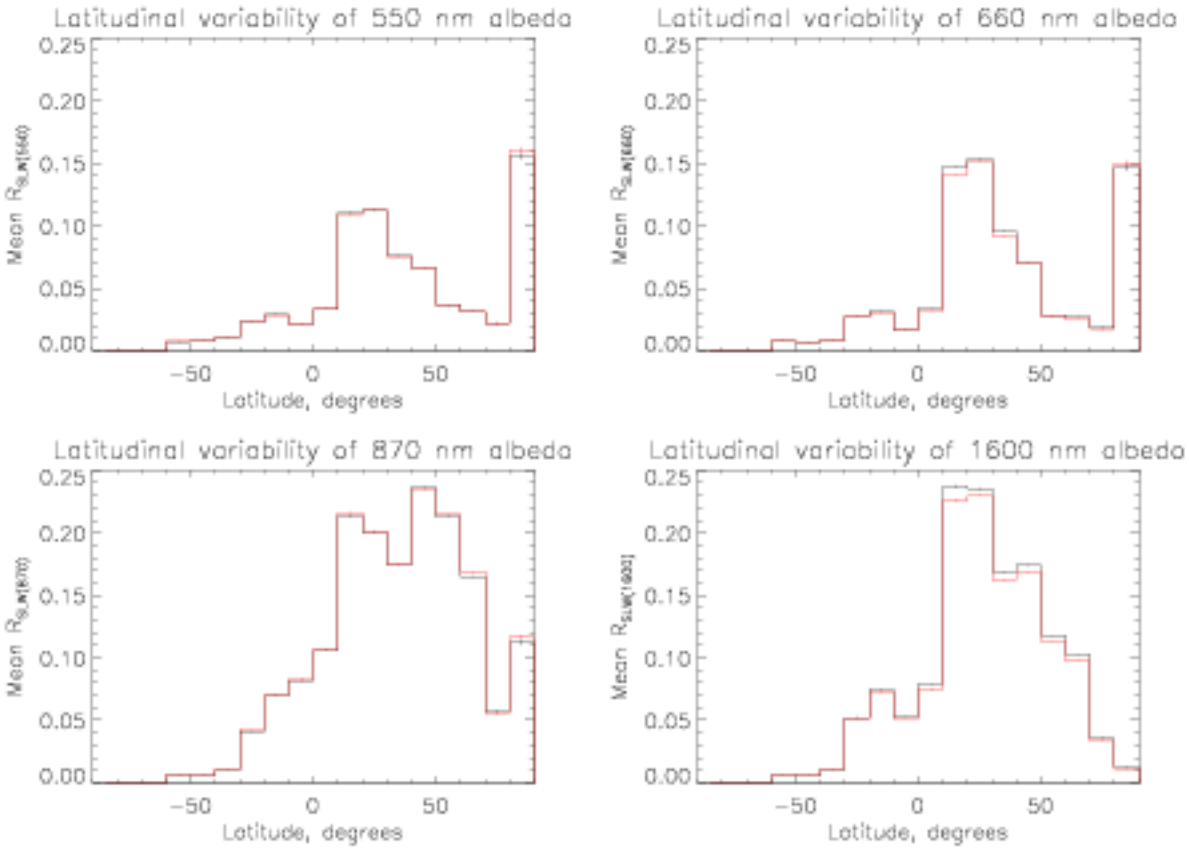

Figure 6.10: Latitudinal mean white-sky surface albedo at the 4 AATSR visible-channel wavelengths. Results for August 2002 are shown in black, and those for August 2007 are in red. Error bars indicate the uncertainty on each latitude band (the standard deviation of pixels in that band divided by the square root of the number of pixels averaged in it).

The latitudinal mean surface albedo (again, for points in each 10° latitude band where both months contained successful retrievals) at each wavelength is shown in Figure 6.10 and there is a very strong similarity between the months at all wavelengths. The largest differences are found at 1600 nm, which is sensitive to moisture, so it is possible that the differences here are linked to precipitation changes. The Northern Hemisphere shows larger changes than the Southern Hemisphere; as previously mentioned, the Northern Hemisphere has a higher population so these differences may be linked to land-use change. Frequency distributions of the retrieved albedo, shown in Figure 6.11, reveal a similarly high correspondence. Finally, the distributions of the differences in Figure 6.12 are, again, largely within retrieval error; the larger negative tail at 1600 nm is a possible indicator of

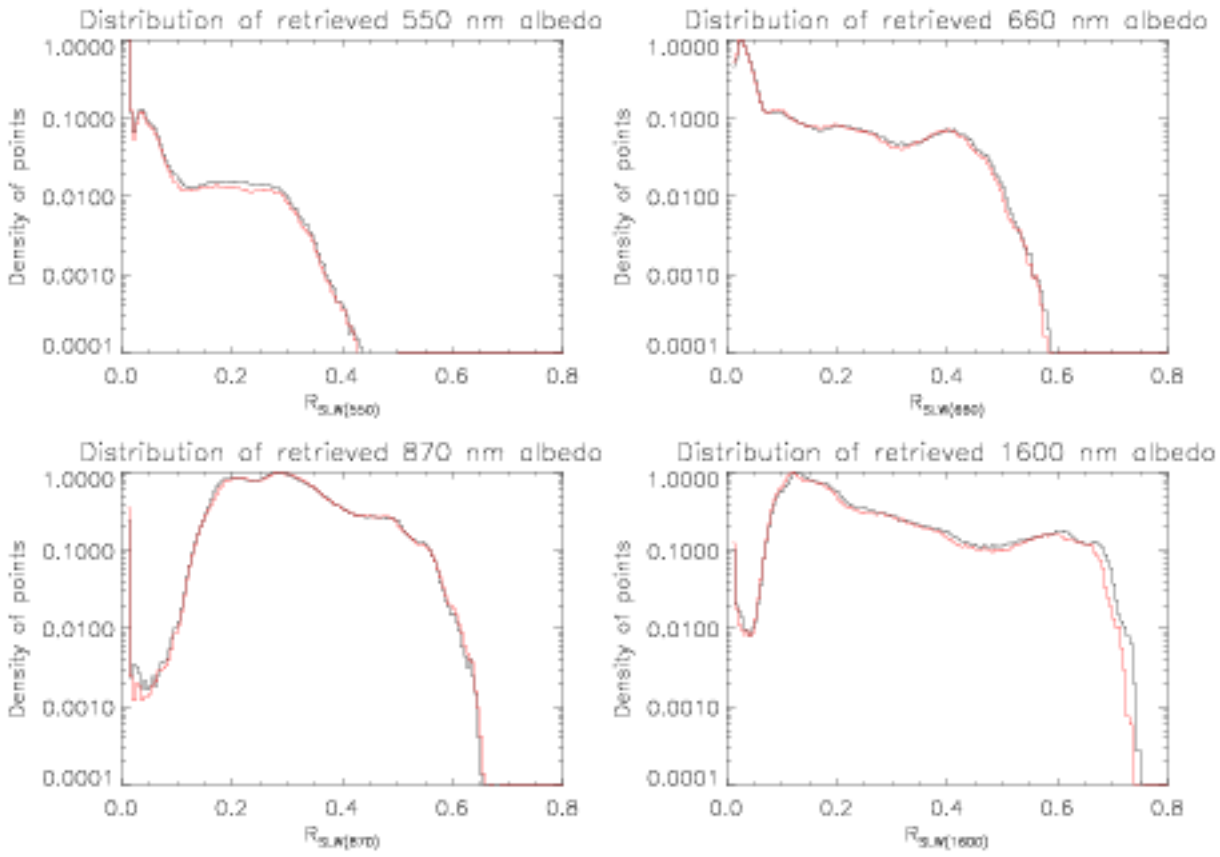

Figure 6.11: Frequency density plot of retrieved white-sky surface albedo at the 4 AATSR visible-channel wavelengths (bin size 0.005) for the two months. Results for August 2002 are shown in black, and those for August 2007 are in red. Distribution normalised to the most populated bin.

precipitation differences.

These results suggest that the surface was not significantly different in August 2002 and August 2007. The surface albedo shows less variability over this period than the aerosol.

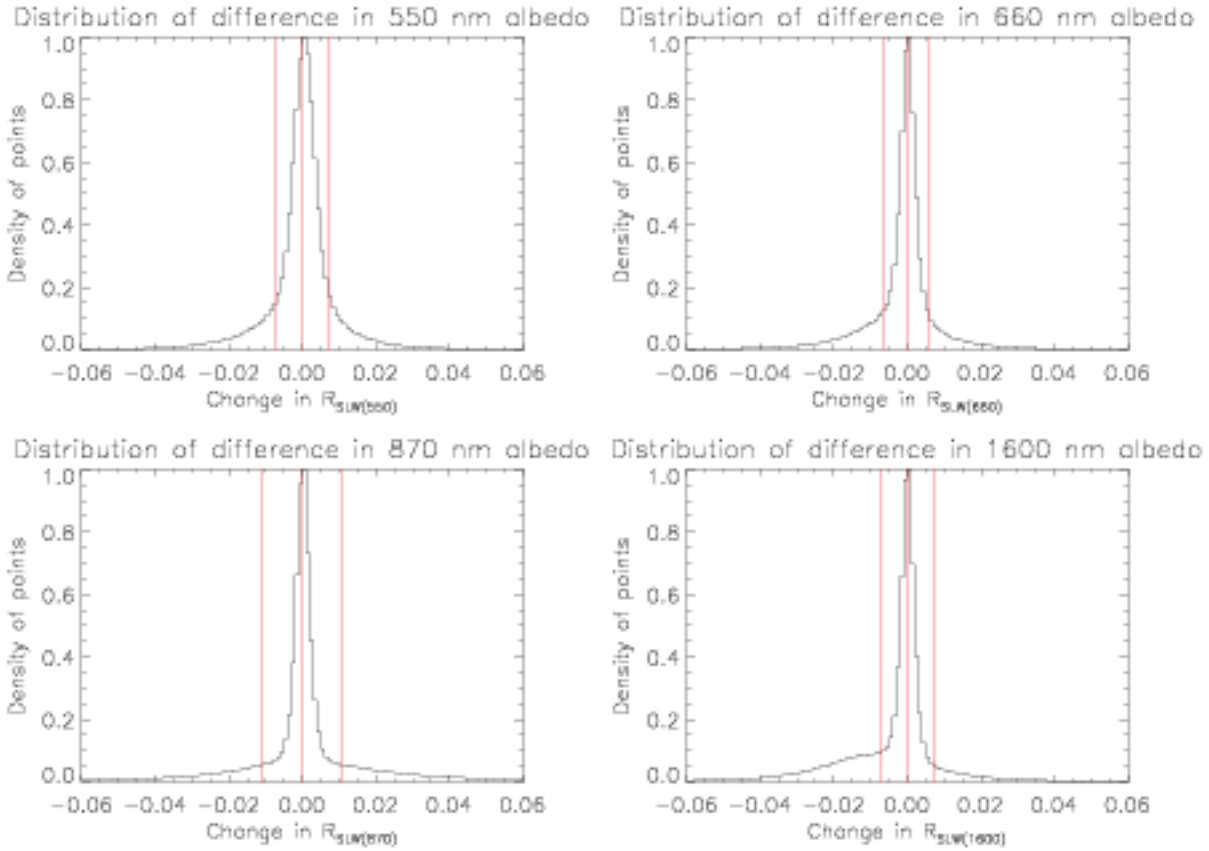

Figure 6.12: Frequency density plot of the differences in retrieved white-sky surface albedo at the 4 AATSR visible-channel wavelengths (bin size 0.001) between the two months. The red lines indicate zero difference and ± the root sum of squares mean retrieval error on each parameter from the two months. Distribution normalised to the most populated bin.

6.3 Aerosol and surface properties in the Amazon region

6.3.1 Area of study

The final component of this chapter is the analysis of a long time-series of data in the Amazon region. AATSR data from August 2002 until August 2007 inclusive have been processed between the limits of 20° S and 10° N, and 90° W to 30° W. The dual-view algorithm, speciation and quality control as discussed previously were applied to the output data. A false-colour image derived from the whole dataset illustrates this region of study in Figure 6.13. Many different types of terrain are encompassed by this region, including:

- The Amazon rainforest (largely located in Brazil; near the centre of the map, along and around the Amazon river basin).

- The Cerrado savannah (to the south-east of the Amazon rainforest).

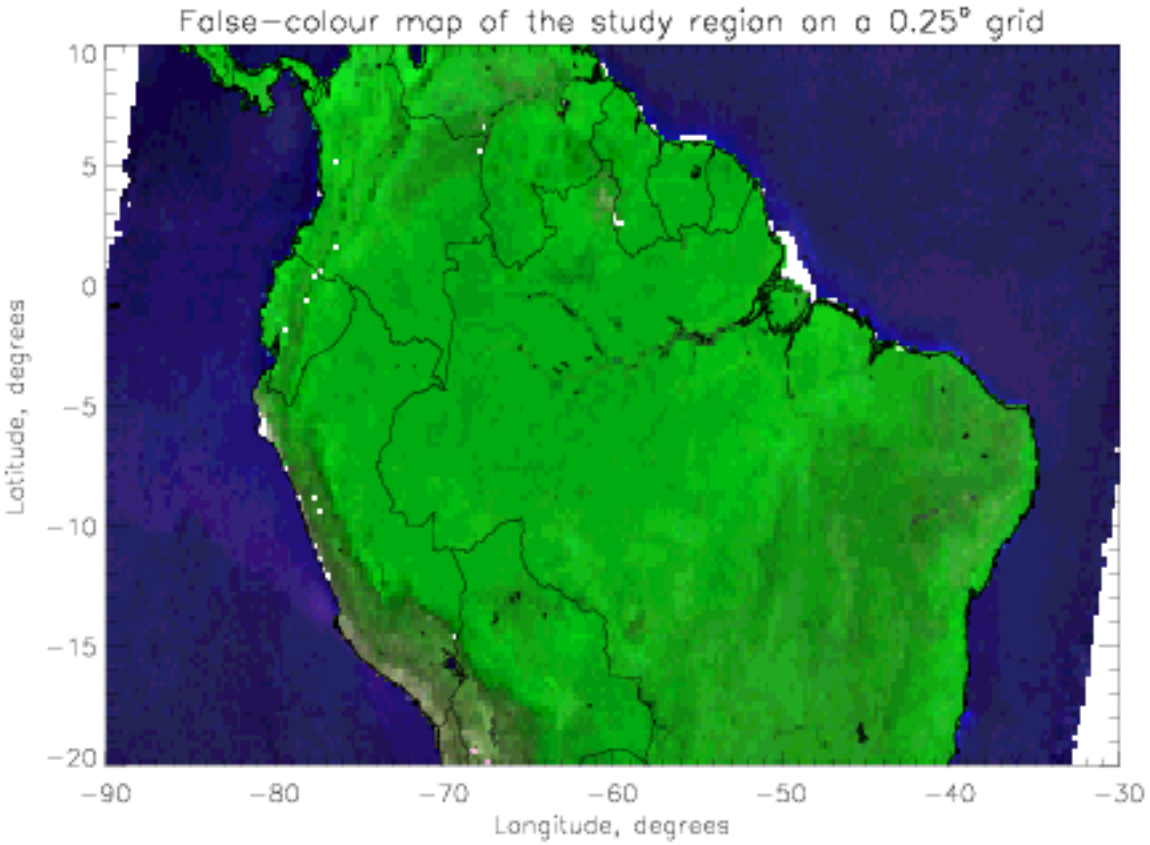

Figure 6.13: False-colour image of the Amazon region on a 0.25° grid derived from the processed 5 years of AATSR data. The brightness of the sea has been enhanced 15-fold.

- Part of the Andes mountain range (along the western coast of the continent).
- Agricultural areas (largely in the south-east of the region and along the north-east coast).
- Urban areas (largely along the coast, although there are some cities inland).
- The arid Sertão region (in the easternmost part of Brazil).
- Two oceans; the Atlantic (to the east) and Pacific (to the west).

Despite comprising 5 years of data, Figure 6.13 contains gaps. The gaps in the south-east and north-west corners arise as orbits were only selected if they contained a certain minimum amount of data in the study region. Reasons for the other gaps are discussed later.

As it is located around the equator, the study region exhibits little seasonal variation of temperature. Much of the region exhibits comparatively high rainfall year-round. It does, however, have comparatively 'wet' and 'dry' seasons, arising from variations in the position of the intertropical convergence zone (ITCZ). This is a region of the atmosphere of low pressure caused by convective ascent

of warm, wet air, which forms a belt across the equator of almost continual rainfall (Gordon et al., 1998). This frequent and persistent occurrence of precipitation substantially decreases the sampling of satellite aerosol retrievals over this region, due to the high cloud cover. Over the bulk of the study region the 'wet season' is December-May and 'dry season' June-September. This may be broken down on a more regional basis after Ronchail et al. (2002):

- The region around and north of the equator is wet year-round; the rainfall maximum occurs between April and August, although there is no strong difference between the wet and dry seasons.

- The western Amazon region is similarly wet year-round, with the most rain between January and May.

- The southern Amazon has a wet season largely October-January, slightly out-of-sync with the bulk of the area of study, due to advection of cold air along the Andes.

- Between the mouth of the Amazon and the city of Manaus further inland (3.1° S, 60.0° W) the wet season is December-May.

- The Sertão is largely outside the typical range of the ITCZ and is dry year-round, although still exhibits comparative wet and dry seasons in line with the rest of the region (Glantz, 1996).

Interannual variability of climate in this region is linked to the El Niño-Southern Oscillation (ENSO) phenomenon (Schneider, 1996a, Glantz, 1996, Ronchail et al., 2002, Nagai et al., 2007), with increased rainfall during El Niño episodes. El Niño episodes are defined by positive sea surface temperature anomalies greater than 0.5 K persisting for five or more months in the Pacific off the coast of Peru (Gordon et al., 1998), and La Niña by anomalies of the opposite sign. These episodes are accompanied by correspoding decreases and increases respectively in the atmospheric pressure and changes in the magnitude of the Walker circulation, which drives the westward 'trade winds' in the tropics (Schneider, 1996b). Episodes arise due to changing paths of ocean currents in this region and can occur on a semi-annual timescale, persisting for up to a year or more at a time.

Recent weak El Niño episodes occurred in 2002-2003, 2004-2005 and 2006-2007 (NOAA, 2008), beginning typically in the late (Northern Hemisphere) summer and lasting until the following spring.

The most recent strong El Niño was from 1997-1998 (Nagai et al., 2007). There were no La Niña episodes during the study period, although one began in late summer 2007 (NOAA, 2008).

6.3.2 Seasonal and interannual variability

Figures 6.14 to 6.20 summarise various datasets derived from application of the retrieval algorithm. Each one is a seasonal mean plot on a 0.25° grid. All quality control parameters were as previously used, except the maximum permissable cloud fraction was increased to 0.75, to increase coverage. The seasons are defined as December-January-February (DJF), March-April-May (MAM), June-July-August (JJA) and September-October-November (SON), and are arranged such that viewing left-right across a row gives the seasonal cycle over the year and top-bottom down a column the interannual cycle of that season. The exception is the first plot, which contains data for only August 2002.

The 550 nm aerosol optical depth (Figure 6.14) exhibits clear seasonal cycles which appear consistent from year to year. The different cycles observed over the region, together with the spatial variability of 1600 nm albedo (Figure 6.19), dovetail well with a classification into distinct zones over which to consider trends in the assorted retrieval products. The discussion of dust and biomass proportions which follows assumes the speciation to be reasonable. As it is not reliable for low aerosol loadings, these conclusions may be incorrect, but are the best that can be drawn at the present time.

The grid structure visible in Figure 6.20 is a result of the interplay of the retrieval's grid (approximately 10 km) and the 0.25° grid the seasonal data is averaged to. As the two are not directly scaled versions of each other, the sampling over the 0.25° grid is not quite regular. As can be seen in the other fields, this does not affect the distributions of the seasonal mean retrieved data.

The Amazon rainforest

The extent of the Amazon rainforest is well defined by the area with 1600 nm albedo between approximately 0.08 and 0.18 (Figure 6.19) in all seasons. The forest basin extends outward from the path of the Amazon river and its tributaries. It is bounded on the east by the Cerrado savannah, the west by the Andes, and the north by the Andes and Atlantic Ocean. The majority of this territory lies within Brazil. The elevated region of the Guyana Shield to the north is excluded from this zone.

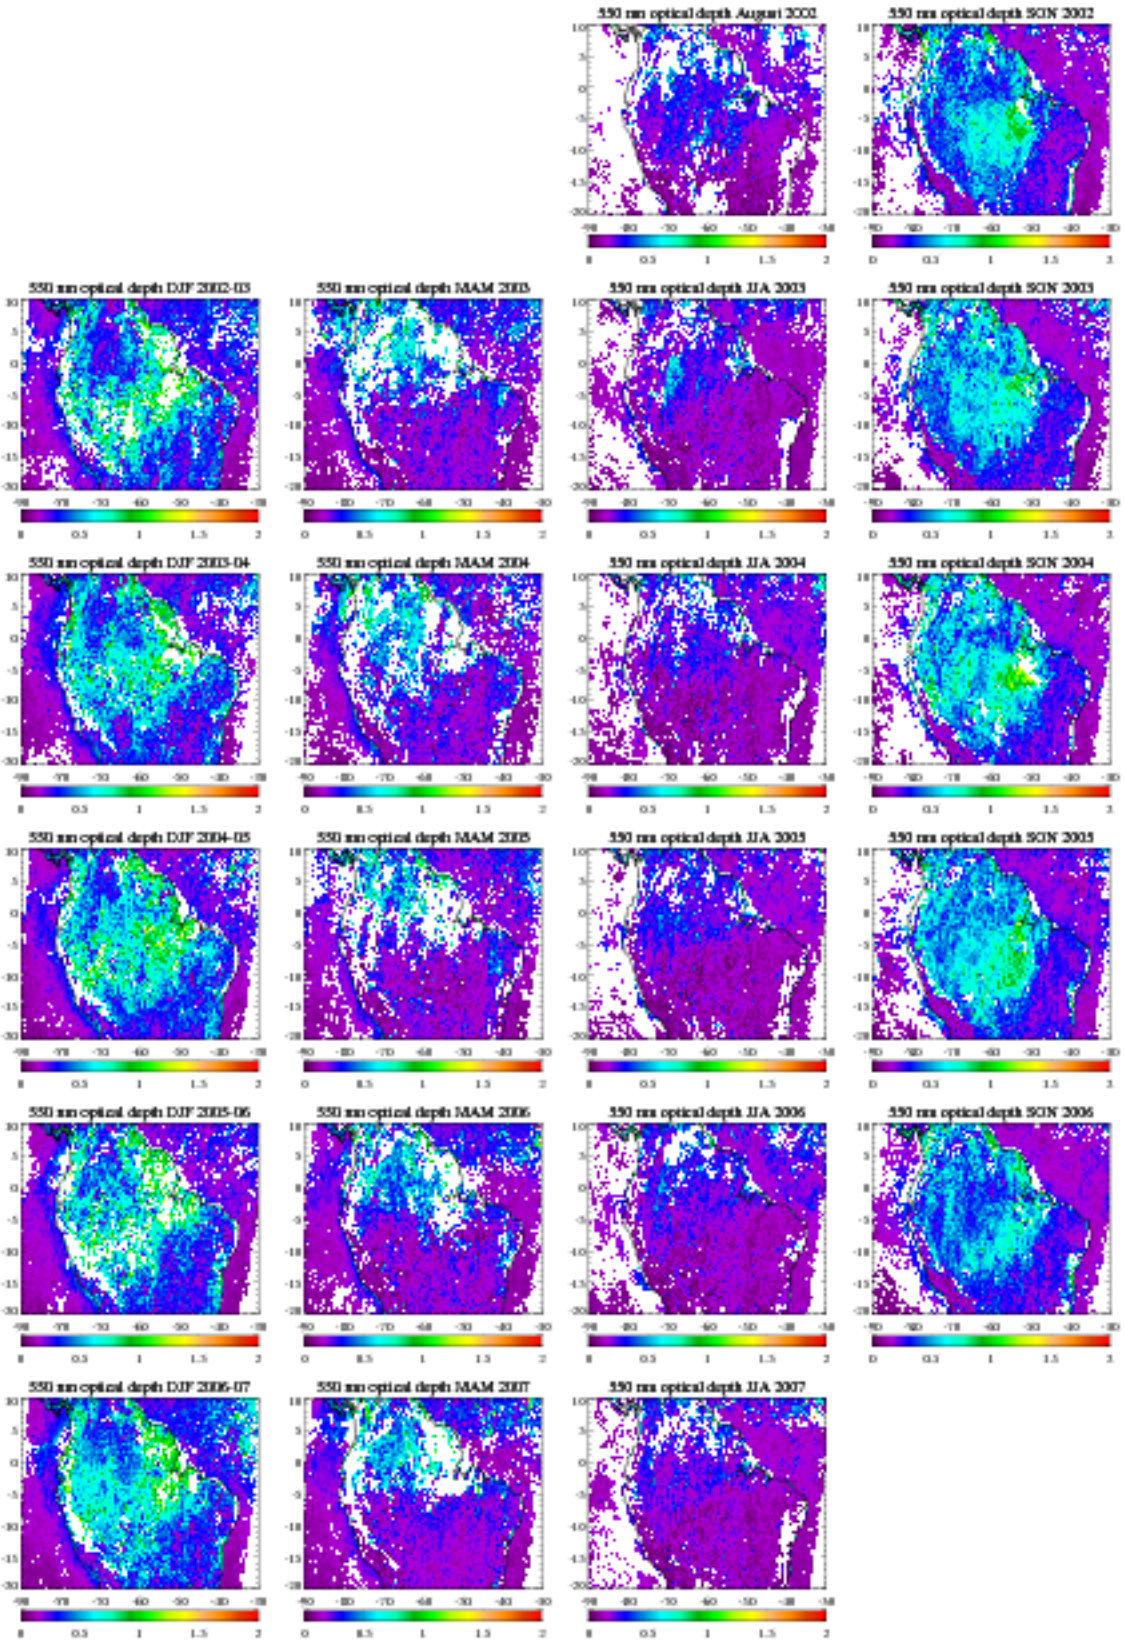

Figure 6.14: Seasonal mean 550 nm aerosol optical depth in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only.

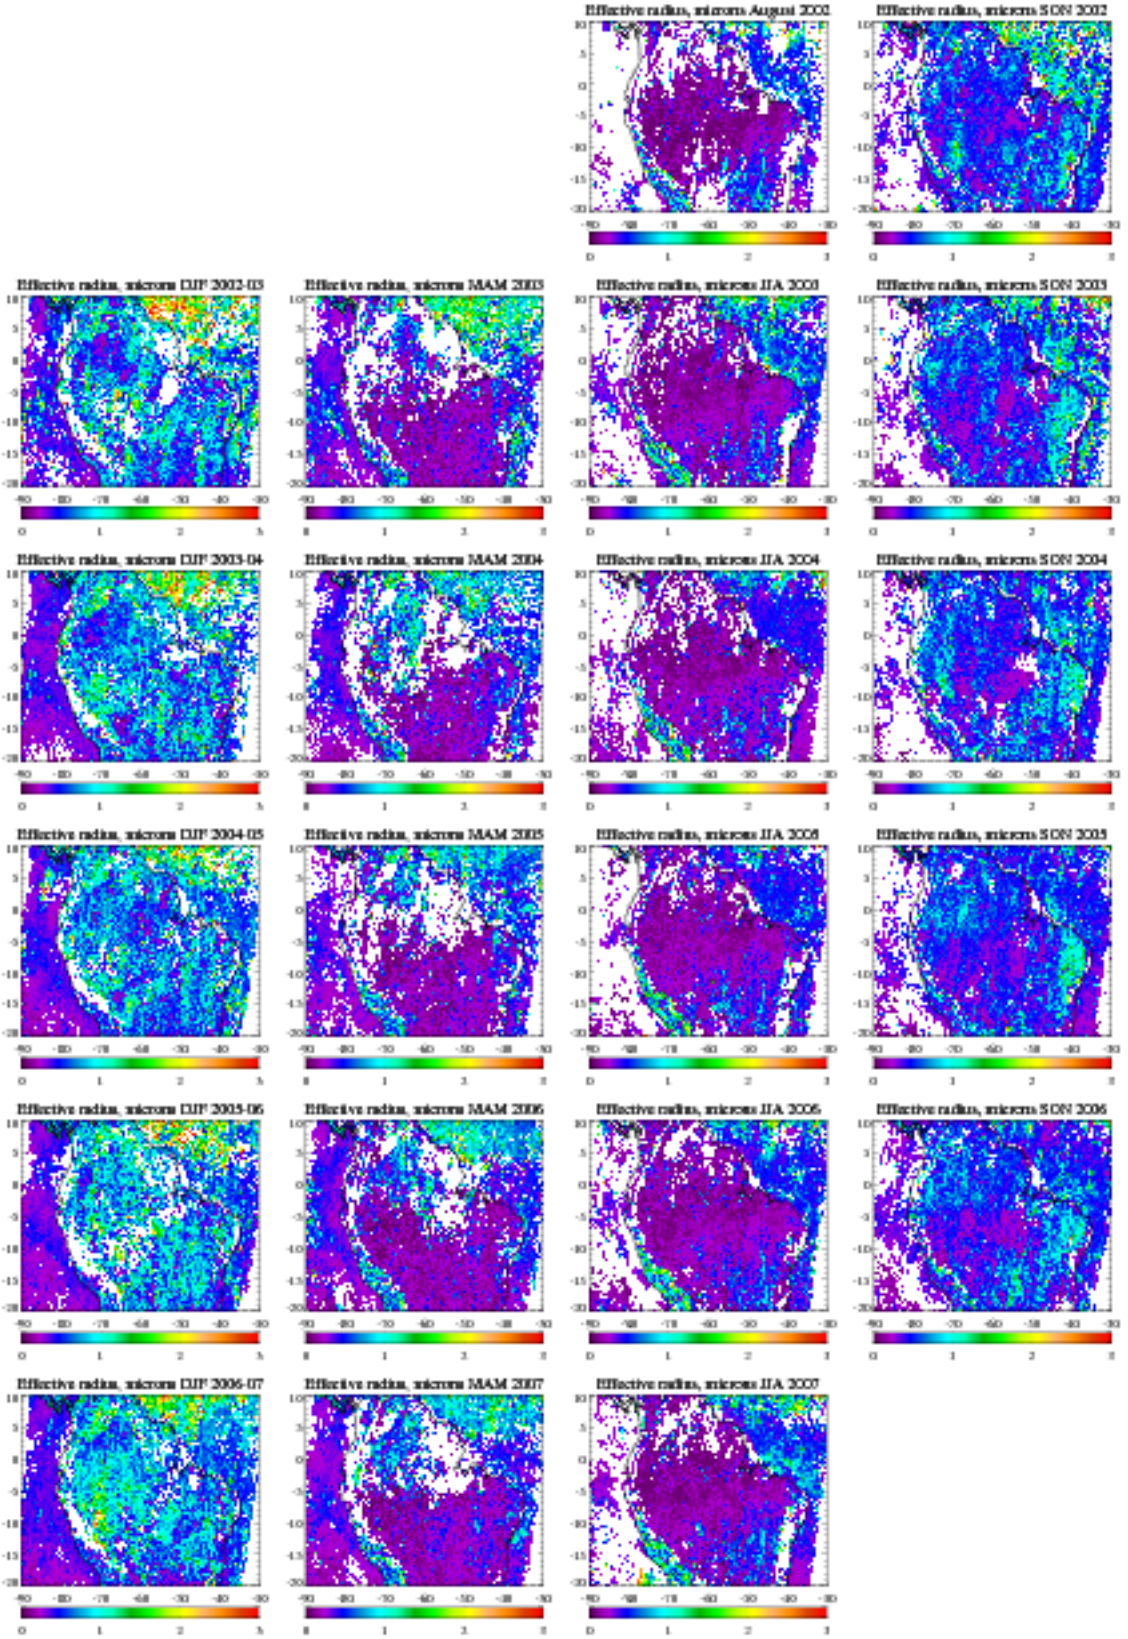

Figure 6.15: Seasonal mean aerosol effective radius in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only.

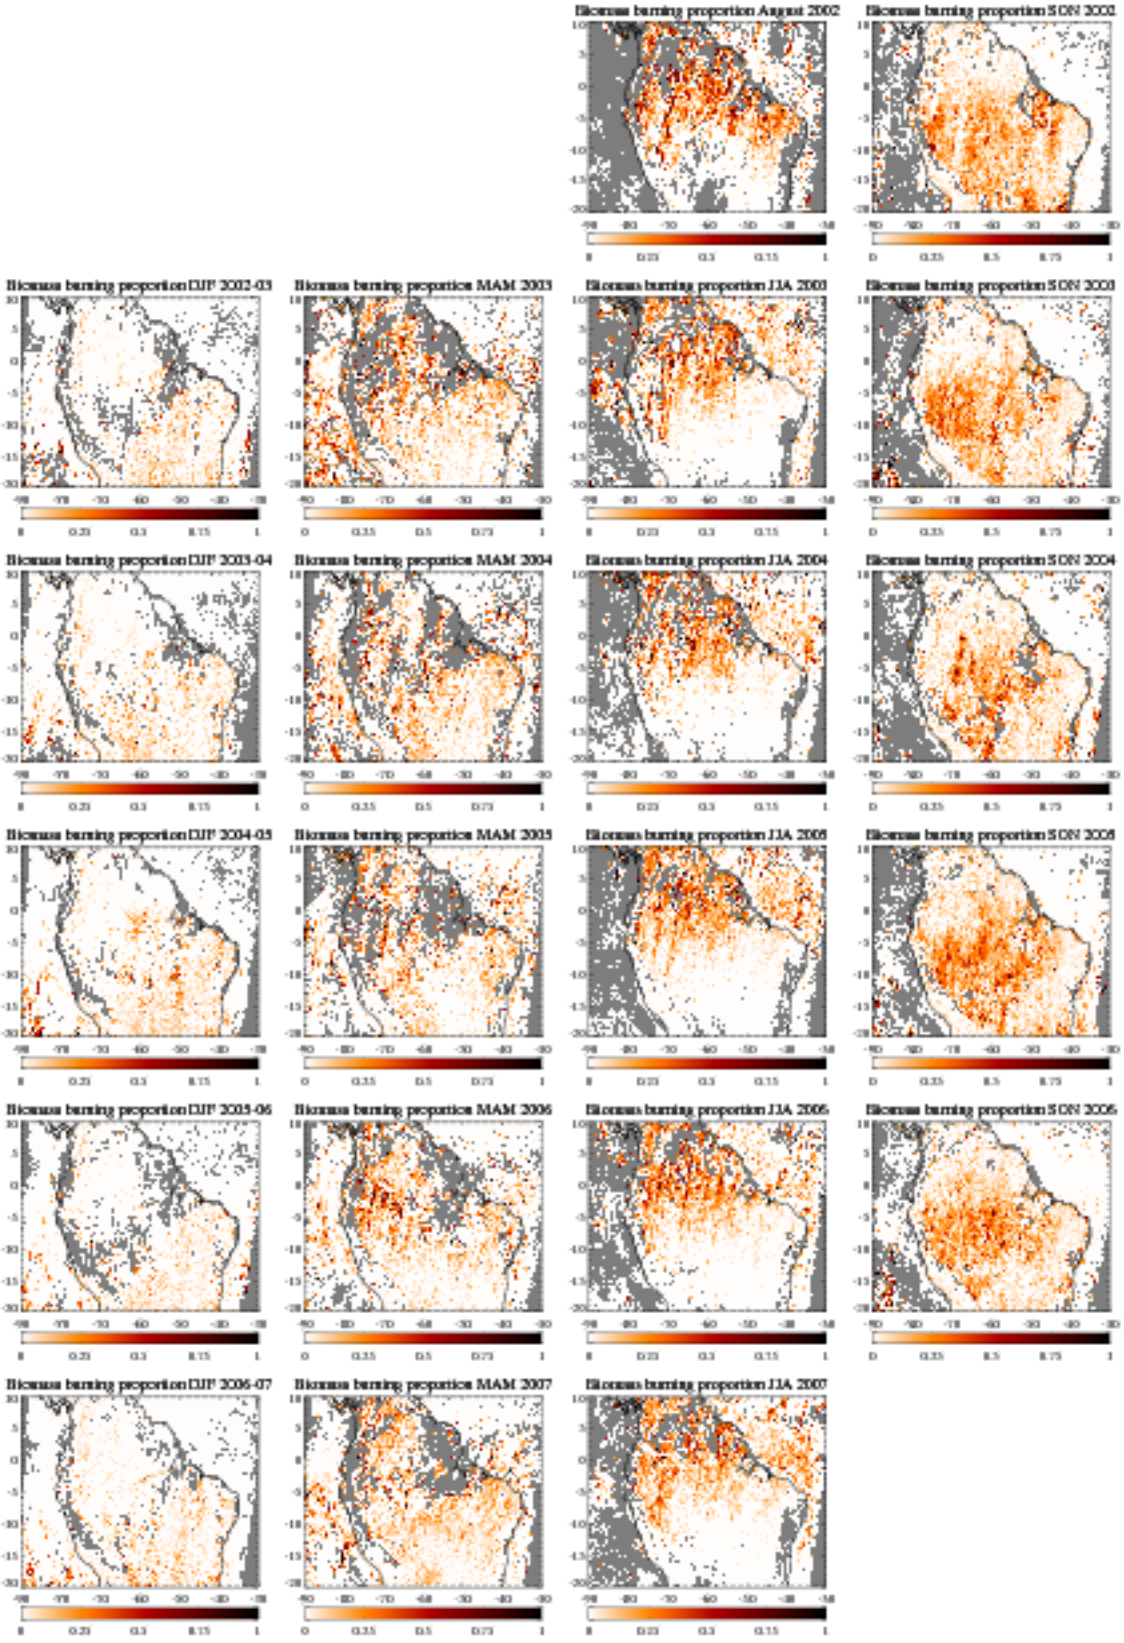

Figure 6.16: Seasonal proportion of aerosol retrievals best matching the biomass burning aerosol class in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only. Missing data are indicated in grey.

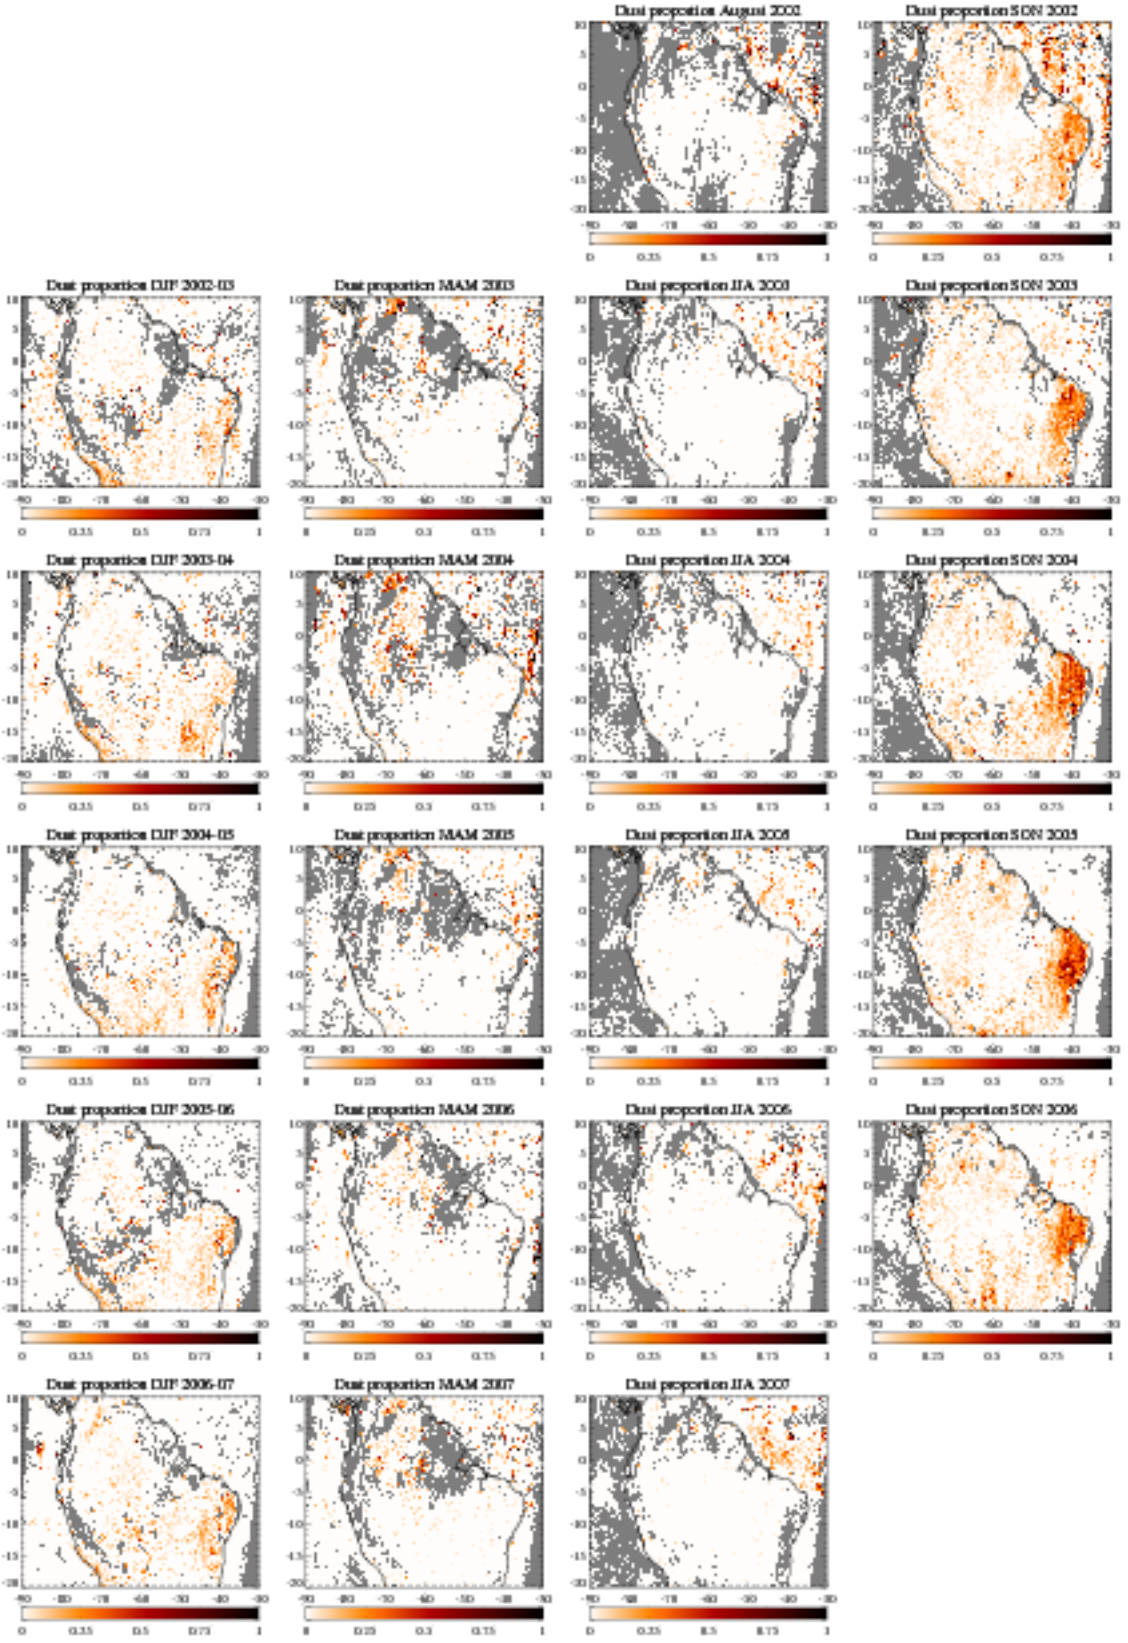

Figure 6.17: Seasonal proportion of aerosol retrievals best matching the dust aerosol class in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only. Missing data are indicated in grey.

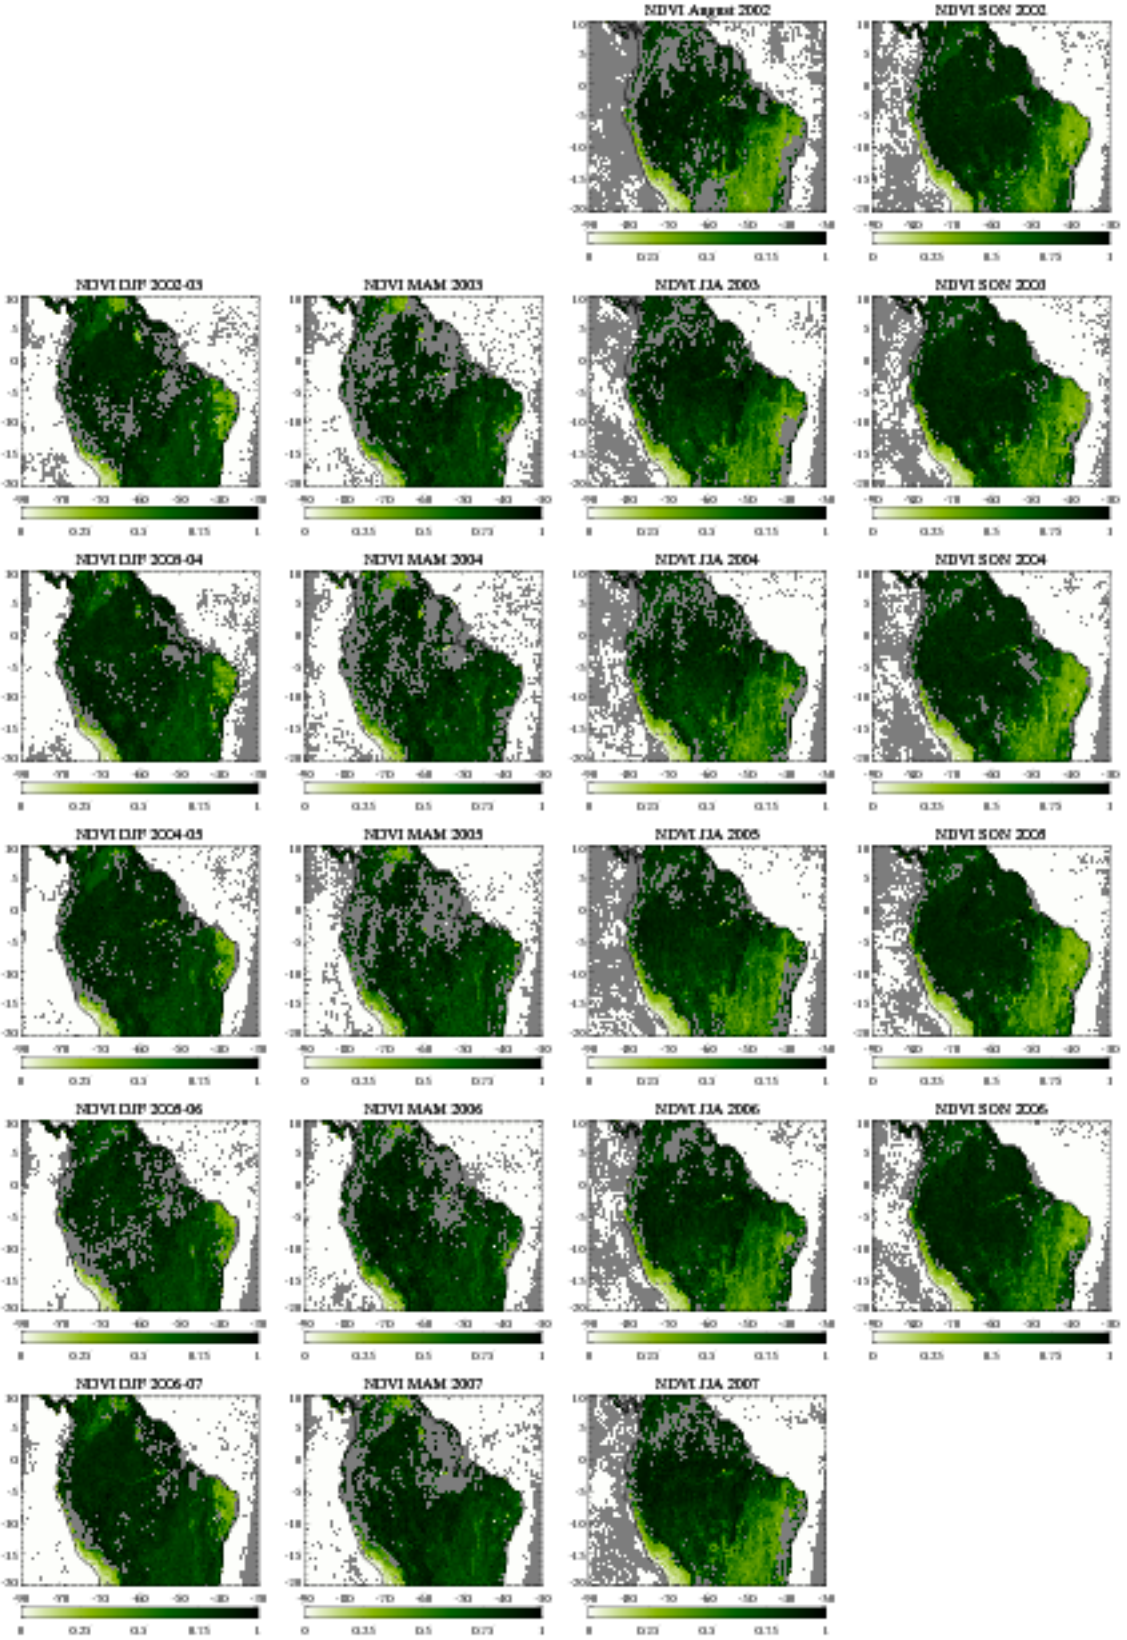

Figure 6.18: Seasonal mean NDVI in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only. Missing data are indicated in grey.

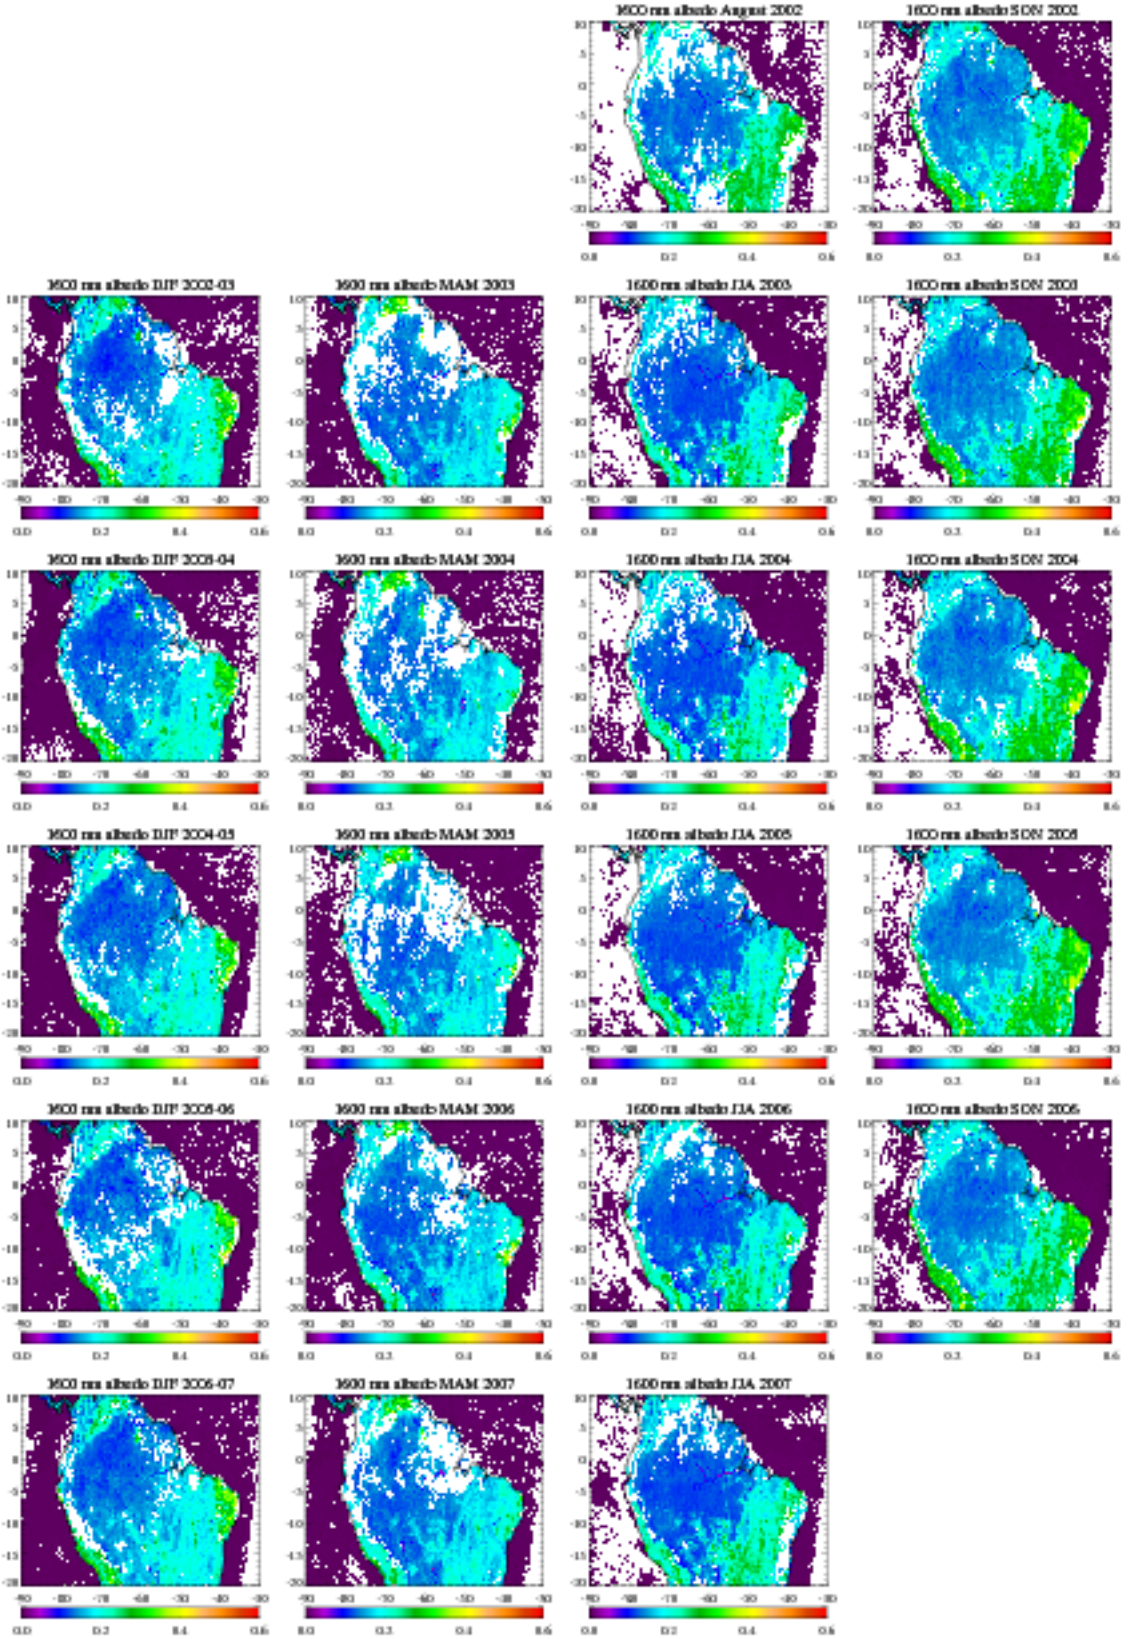

Figure 6.19: Seasonal mean 1600 nm surface albedo in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only.

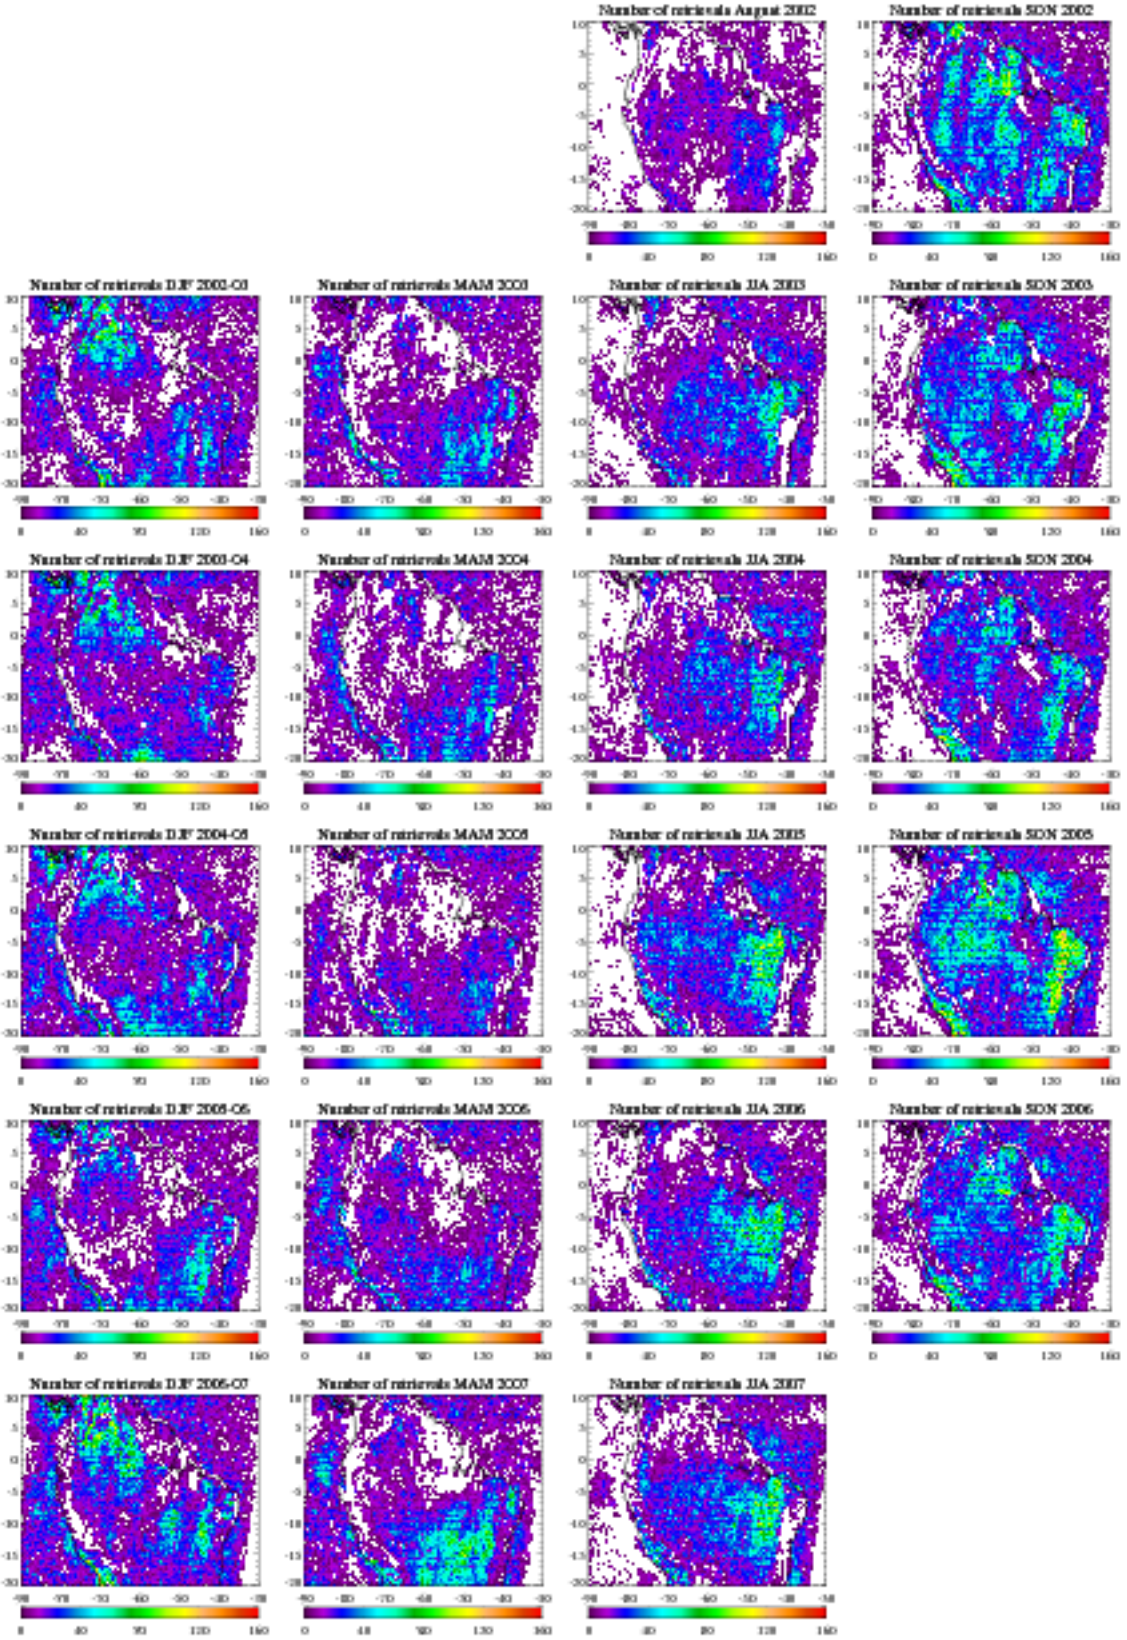

Figure 6.20: Seasonal number of successful retrievals in the study region, on a 0.25° grid. Plots are from left to right Northern Hemisphere winter (DJF), spring (MAM), summer (JJA) and autumn (SON) and from top to bottom the years 2002-2007, aside from the first plot in the series which is August 2002 only.

Throughout the year, the Amazon shows higher optical depths than the rest of the study region, with peak values of 1.5 and above. Aerosol optical depth shows a strong seasonality, with the highest optical depths in DFJ and SON further south than in MAM and JJA. The lowest overall optical depths are observed in JJA, with typical values lower than 0.5. For the majority of the Amazon, the wet season is DJF and MAM with cloud cover most limiting the sampling during this period. The strong north-south divide and cycling of optical depth requires some explanation.

- In DJF τ_{550} reaches its peak values and most of the aerosol is assigned the continental type. The NDVI is also very high (0.75 or more). These are signatures of productive, dense vegetation. Effective radius is variable, but is generally around 1 μm. There is a small amount of biomass burning aerosol in the southern Amazon; the wet season there ends in January (Ronchail et al., 2002), suggesting the onset of fires afterwards.

- During MAM the northern Amazon remains similar, while τ_{550} over the central and southern Amazon decreases sharply. The wet season in the southern Amazon has ended and there is an increase in the proportion of biomass burning aerosol, suggesting forest fires. Smaller particles are found at this time, closer to the *a priori* size for the biomass burning class, though as τ_{550} has decreased there is less information about r_e in the AATSR signal.

- By JJA the wet season has ended in the northern Amazon, although on the equator there is no true seasonality (Ronchail et al., 2002). Optical depths are highest from 0°-5° S. Biomass burning peaks here in the northern and central region; as the forest is drier, this is again likely due to a combination of natural and man-made forest fires.

- In SON optical depth again increases to levels similar to those of DJF, although the highest optical depths are more tightly concentrated to the south-west of the mouth of the Amazon river than in DJF. This period is characterised by high optical depths and frequent biomass burning aerosol, particularly in the central and southern Amazon. Deforestation occurs most often during the dry season (Kirby et al., 2006), and regions of strong biomass burning in Figure 6.16 are located in regions where deforestation has both historically (Tucker et al., 1984) and recently (Alencar, 2005, Souza et al., 2005, Câmara et al., 2006) taken place. Deforestation of the Amazon and NDVI change is dealt with in more detail in Section 6.3.7.

Finally, Figure 6.17 shows dust is rarely flagged in the Amazon. This is consistent with it being a moist, densely-vegetated environment with little bare soil for the wind to lift.

The Sertão

The Sertão is the arid land in the east of the region of study, well-defined by the area of eastern Brazil with a 1600 nm surface albedo around 0.3-0.4 during DJF (Figure 6.19). It comprises parts of the Brazilian states of Bahia, Ceará, Rio Grande do Norte, Paraíba, Pernambuco, Alagoas and Sergipe.

The seasonal cycles exhibited in this region are reasonably consistent year-to-year.

- During DJF, τ_{550} is typically 0.3-0.7 and r_e 0.5-1 μm. A small amount of dust aerosol is observed, with the rest continental. The NDVI is low (0.2-0.4), suggesting sparse vegetation. This is the early phase of the comparatively wet season (Glantz, 1996), suggesting some of the aerosol is wind-blown dust from the arid ground. This bare soil is indicated by a high 1600 nm albedo. Biomass burning also has a presence, implying burning of dry vegetation.

- By MAM the wet season is underway, and there is a corresponding decrease in τ_{550} (values normally lower than 0.5) and r_e (often 0.1-0.5 μm) as compared to DJF, with the proportion flagged as dust almost zero. This is accompanied by a large increase in the NDVI (values around 0.5-0.7). The interpretation is that following the wet season the soil is moist and vegetative cover increases, decreasing the exposure of the soil to the wind, resulting in less dust aerosol.

- JJA shows a further decrease of optical depth to approximately 0.05-0.2, and slight decrease of effective radius. Aside from August 2002, neither dust nor biomass aerosol is common. The resemblance of August 2002 results to SON seasonality suggests that August in this region resembles SON, and the JJA season shows a transition between MAM and SON profiles. This is consistent with the wet season ending at the beginning of this time, and vegetation gradually drying out into August (Glantz, 1996). Consequently NDVI decreases to 0.3-0.5; reaching lower levels than MAM but higher than DJF.

- In SON aerosol and surface properties look like a more extreme version of DJF. This is the second half of the dry season, and the arid nature of the terrain is indicated by the low NDVI (0.2-0.4), bright 1600 nm albedo (0.3-0.4) and frequent occurrence of dust aerosol (with dust

fractions frequently above 0.25 and sometimes approaching 1, and an increase of τ_{550} and r_e as compared to the previous season).

The Cerrado savannah and coastal lands

The remainder of the lowland in the region of study is either the Cerrado savannah or a mixture of urban and agricultural land on the coast of the Atlantic. It is bordered on the north and west by the Amazon and the east by the Sertão. This area is located within Brazil, primarily the staes of Goiás, Bahia, Piauí and parts of Maranhão. Bahia state contains Brasilia, Brazil's capital.

The seasonality of aerosol properties in this area shows general similarities with that of the southern Amazon and the Sertão near those regions. This suggests transport of aerosol from these regions. There are marked NDVI shifts between the wet seasons (NDVI 0.4-0.6) and dry (NDVI 0.2-0.5), with some dust and biomass burning occurring late in the dry season when the terrain is arid. Optical depths are higher in the late dry season and early wet (SON and DJF, τ_{550} 0.3-0.8) than other times (MAM and JJA τ_{550} 0.1-0.25). Effective radius shows similar seasonal shifts, suggesting a change in the nature of the aerosol throughout the year.

The Atlantic coast has a slightly elevated optical depth as compared to the region, although cloud cover means coverage is sparse. This is a region of relatively dense population, implying an urban origin for the aerosol.

The Andes and Guiana Shield

The Andes and Guiana Shield are mountainous regions along the west of the area of study (Peru, Ecuador and Colombia) and through Colombia and Venezuela respectively.

The Andes show a fairly constant low optical depth (0.3 or less) over all years and seasons. Effective radius is unreliable as the optical depth is low, and the vertical and horizontal inhomogeneity of the terrain mean retrieval quality is likely to be low as these factors are not accounted for by the forward model. The same factors may affect the accuracy of the MODIS surface reflectance data used in this region. This is particularly evident in the frequent line of missing data traced along the east of the Andes, which may be in the shadow of the mountains when observed by AATSR. Cloud cover also reduces sampling.

Biomass burning aerosol is rarely observed over the Andes, as it would have to rise to a high

altitude to reach them. As mentioned, however, speciation at low optical depths is particularly unreliable. Dust is often found near the Peruvian coast (at the northern end of the Atacama desert), showing a similar seasonality to dust in the Cerrado and Sertão, implying similar causes of arid, vegetation-sparse land. The NDVI is low (less than 0.25) and 1600 nm albedo high (0.2 or more), suggesting bare ground.

The Guiana Shield show characteristics in between those of the Amazon and the Andes, corresponding to their comparative elevation and ground cover.

The Pacific Ocean

This is the body of water lying to the west of the region of study.

Aerosol optical depth is fairly low and constant across the Pacific for all years and seasons. Typical τ_{550} values are 0.2 or less for the open ocean, with a region of enhanced optical depth (0.2-0.5) near the coast. Maxima are found in the vicinity of major cities (Lima in Peru, Guayaquil and others in Ecuador, Cali in Colombia and Panama City in Panama) suggesting an industrial origin. Aerosol effective radius is generally 0.5 μm or smaller, although as optical depths are low r_e and aerosol speciation are highly uncertain. Dust and biomass aerosol are only rarely flagged, aside from in MAM when biomass burning aerosol is more common. This suggests transport of biomass burning aerosol west across the Amazon basin during this time; the direction is consistent with the usual westward flow of the prevailing winds in this region (Walker circulation, Schneider, 1996b). The precise origin of this aerosol is uncertain, as it would have had to cross the Andes to get there from the Amazon basin, so it is likely that some is incorrectly flagged or else the source regions are coastal.

Determination of any seasonal cycle is difficult in this area due to the high cloud cover limiting observations, particularly during JJA and SON. This is due to the movements of the ITCZ and orographic clouds forming from air moving westward over the Andes (Schneider, 1996a). There is no clear difference between conditions during El Niño episodes and at other times, although again poor coverage limits the conclusions which can be drawn.

The Atlantic Ocean

This is the body of water lying to the east of the region of study. Like the Pacific, the Atlantic Ocean exhibits low optical depths with enhancements around the coast. However, the extent of the coastal enhancement is much thinner than over the Pacific coast. Additionally, it can be observed more reliably because cloud cover is less of a problem.

The southern portion of the Atlantic in this region shows similar optical depths and effective radii to the Pacific. However, the northern part often contains elevated optical depths of 0.3-1 and effective radii up to 3 μm. These sizes are typical values for dust particles (Hess et al., 1998). The proportion of aerosol flagged as dust in this region is small (Figure 6.17), but the retrieval has difficulty distinguishing between the dust and maritime aerosol classes (Figure 2.9). As the prevailing winds are easterly (Schneider, 1996b), this aerosol is likely from the Sahara rather than blown off the arid Sertão region. This interpretation is consistent with the fact that optical depths off the coast of the Sertão are *not* elevated, and that long-range transport of Saharan dust to this region has been observed by a variety of instruments and techniques, such as Formenti et al. (2001), Kaufman et al. (2005), Koren et al. (2006) and Liu et al. (2008). The enhancement of optical depth and effective radius is greatest in DJF and MAM, fairly consistently over all years, suggesting seasonality in the deposition of Saharan dust. Seasonality in dust optical depth and single scatter albedo over the Atlantic has been noted from satellites by Kaufman et al. (2005) and more recently with airborne measurements by McConnell et al. (2008).

Biomass burning aerosol is also sometimes flagged in this region during MAM and JJA of every year. This coincides with high levels of it also being flagged over the land, so it is possible that some of this is transported from there. As the prevailing winds are westward, however, it is likely that some of it is transported from central Africa. This transport has been observed with other satellite instruments during these seasons, such as by Liu et al. (2008) for a case during August 2006. It is also possible that some of the aerosol is incorrectly attributed as biomass burning.

The region as a whole

Seasonal profiles derived for different years show a strong similarity. This suggests fairly low interannual variability in this region, and gives confidence in the results of the AATSR retrieval and calibration of the instrument, as for example the *a priori* information on surface reflectance is dif-

ferent for each year. This consistency also suggests that the results presented here can be taken as a climatology of aerosol and surface properties in the region. If desired the data could be averaged over the years and aggregated monthly to provide a profile of a 'typical year'. These plots have not been included here for brevity.

There is little evidence of cloud-contamination in the retrieval results, as fields are smoothly-varying without obvious discontinuities near missing data. This shows that the combination of the pre-retrieval cloud mask and post-retrieval quality control have been successful at filtering out contaminating cloud. A few regions have frequent or constant missing data, and the reasons for this are as follows:

- The mouth of the Amazon, opening on to the Atlantic. This is due to quality control as opposed to cloud cover. The silty water here is poorly described by the sea surface reflectance model used, suggesting that improvements to account for turbid waters would improve coverage in this situation.

- As mentioned previously, shadowing and high elevation as found in the region of the Andes mountains are not accounted for in the forward model, leading to retrievals which more frequently fail to converge than in other regions. A topographic correction could be developed to improve the quality of the retrieval over elevated land.

- Large portions of the study area, such as the Andes, the Pacific to the west of them, and the coast of Bahia state, are seasonally affected by almost perpetual cloud cover. Unfortunately nothing can be done to alleviate this with AATSR.

- Persistent missing MODIS data for a pixel (due to, for example, persistent cloud cover or failure of the BRDF model to describe the surface) would lead to corresponding missing AATSR retrievals. These gaps would require further generations of MODIS data, with increased coverage, to fill.

No significant differences were observed between periods within El Niño episodes and those without. However, none of the El Niño episodes during the period were strong, and there were no La Niña episodes. Therefore nothing can be said about the impact of strong ENSO episodes on the region.

6.3.3 Seasonality of NDVI in the Amazon and response to the 2005 drought

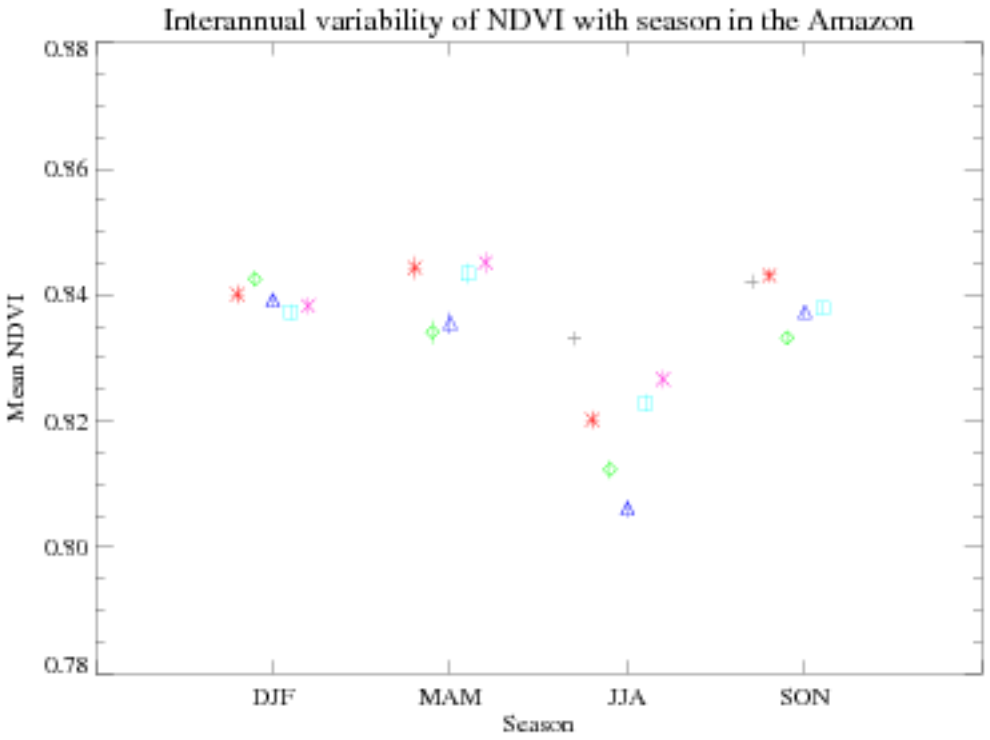

Figure 6.21: Seasonal mean NDVI for the Amazon basin. Error bars indicate the standard deviation of each season's NDVI, divided by the square root of the number of 0.25° grid cells containing data. Grid cells were only considered where data from every season from JJA 2002 to JJA 2007 were available, and the region fell inside the Amazon rainforest. Black crosses indicate August-November 2002, red stars Dec 2002-Nov 2003, green diamonds Dec 2003-Nov 2004, blue triangles Dec 2004-Nov 2005, cyan squares Dec 2005-Nov 2006 and purple crosses Dec 2006-Aug 2007. For each season the years are arranged in chronological order.

In 2005 the Amazon basin suffered a drought, peaking from July to September, in which large areas received 1-2 standard deviations less precipitation than normal (Aragão et al., 2007). The preceding season had a lesser negative precipitation anomaly. While this might be expected to have a negative impact on plant growth, Saleska et al. (2007) found instead an overall increase in the greenness of the Amazon as derived from MODIS data during the July-September period. This was explained via sunlight, rather than water availability, being a limiting factor for plant growth over large parts of the Amazon, and the decreased cloud cover during the drought permitting enhanced photosynthetic activity.

Examination of Figures 6.18 shows that the NDVI of the region as a whole did not show stark differences between late 2005 and the same seasons of different years. Furthermore, a generally higher count of successful retrievals is observed in Figure 6.20 for JJA and SON 2005, which is consistent with decreased cloud cover during this period.

To further investigate the interannual variability of NDVI in the Amazon basin, and any response in NDVI to the drought, the seasonal mean NDVI of the Amazon basin is presented in Figure 6.21. In this diagram, the mean NDVI for grid cells lying in the Amazon basin was found; in a sense, each point in the graph represents the central region of a plot in Figure 6.18. To ensure that changing coverage did not bias the results, grid cells were only considered if they contained data for every season. This left a total of 1,805 grid cells per season. The results show that the NDVI of the Amazon is generally high (0.8-0.85), indicating dense green vegetation cover. There is a clear contrast between JJA and the other seasons, with the former showing more variability between years. However, SON does not appear less green than the wet season. This behavour has also been observed by Huete et al. (2006), again suggesting that sunlight may be the limiting factor for growth in parts of the Amazon as there is sufficient moisture.

In contrast to the findings of Saleska et al. (2007), the NDVI during JJA 2005 was the lowest during the 2002-2007 period. This suggests an overall decrease in greenness coinciding with the peak of the drought, although the difference is small on an absolute scale. Interannual NDVI anomalies relating to ENSO in the Amazon can be of the order of 0.02 (Nagai et al., 2007), so while the JJA 2005 NDVI was low in this respect it was not abnormally so. As the data from SON 2005 onwards shows, if there was a negative NDVI response to the drought, the vegetation recovered quickly as the NDVI after that date does not appear different from non-drought years. Although these results do not show an active greening in SON, the lack of a difference between it and the DJF/MAM wet season is similar to the conclusions of Huete et al. (2006).

It is possible that sampling explains some of the discrepancy between the two sets of results: as this analysis only considered grid cells where data is available for every season, it is biased in favour of less-cloudy regions of the Amazon, where water may be more of a limiting factor to plant growth. The discrepancy may also be due in part to the differing vegetation indices used by Huete et al. (2006) and Saleska et al. (2007) (the enhanced vegetation index, EVI) and the NDVI used here.

6.3.4 Comparison of retrieved optical depths against AERONET

The study region contains several AERONET that are suggested by Kinne (2007) to be representative of their region on satellite spatial scales, namely Abracos Hill, Alta Floresta, Arica, Belterra, Cuiaba-Miranda, Rio Branco and Santa Cruz. A comparison of the retrieved optical depths at 550 nm and

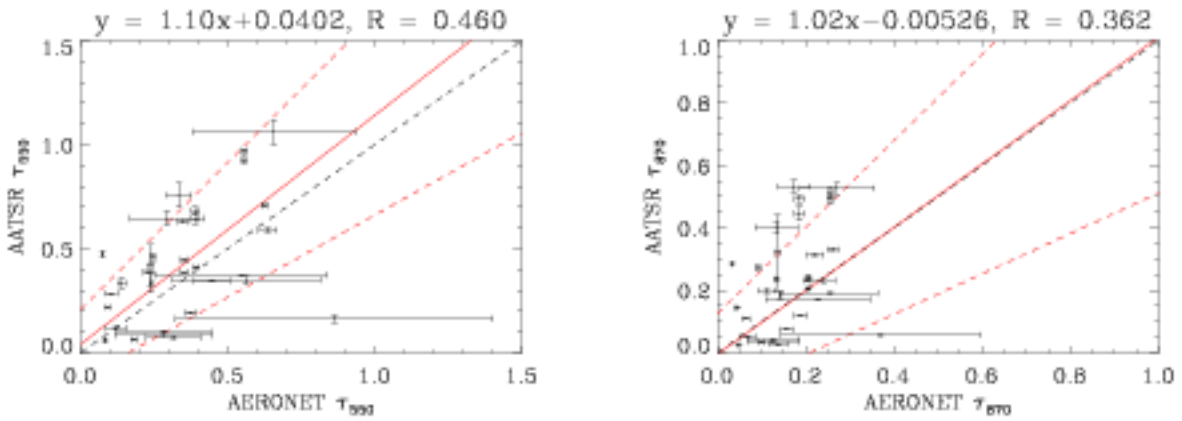

Figure 6.22: Comparison between AERONET and AATSR τ_{550} (left) and τ_{870} (right). Coefficients of fit are given above the plots.

870 nm against measurements from these sites is presented in Figure 6.22, following the method previously used in Section 5.1. The following points can be made about this comparison:

- The coefficients of the line of best-fit between the two stations are good (gradient close to 1, intercept close to 0), albeit with a high uncertainty on them.
- The correlation between the two datasets, however, is poor (0.46 and 0.36 at 550 nm and 870 nm respectively). The AERONET data in particular show very large variability for many data points, indicating a considerable change in aerosol optical depth over a short time-span. This high temporal inhomogeneity may be the reason for a poor match with AATSR, as it means the assumptions made in the spatio-temporal averaging may not be valid as the satellite and Sun photometer are not measuring the same aerosol. This suggests that sometimes it is not appropriate to use AERONET as a comparison tool for satellite data.
- Despite having 5 years of AATSR data, there are fewer than 40 successful comparisons. This is an indicator of the persistent frequent cloud cover in this region of the world, which drastically reduces the sampling of both AATSR and AERONET.

6.3.5 Relationships between population density and aerosol optical depth

The retrieval data can be used to investigate the relationship between population density and aerosol optical depth. Population density data are available from the Gridded Population of the World (GPW) project (CIESIN, 2005) at 5-year intervals; data for 2005 on a 0.25° grid are mapped in Figure 6.23,

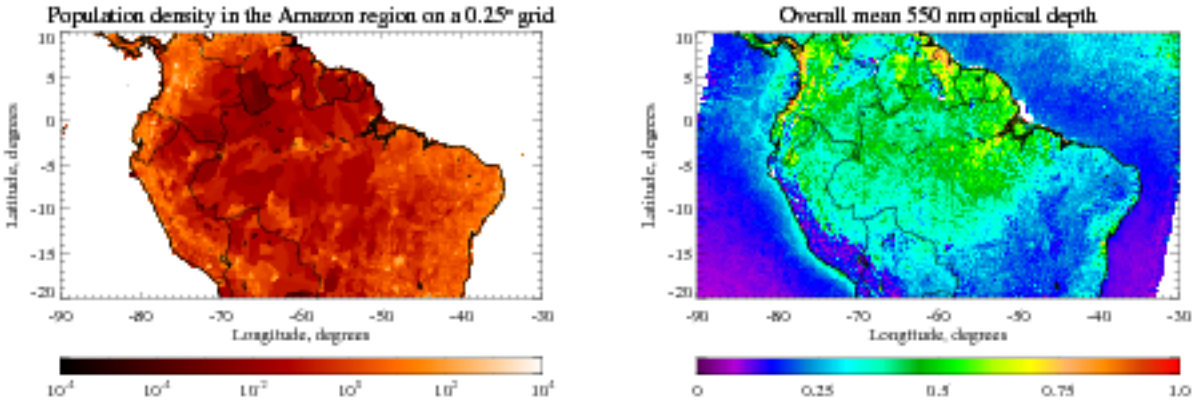

Figure 6.23: Left: Population density (km^{-2}) in the study region in 2005. Data from the Gridded Population of the World dataset version 3 (CIESIN, 2005). Right: Mean aerosol optical depth in the study region over the 5-year period. Both are on a 0.25° grid.

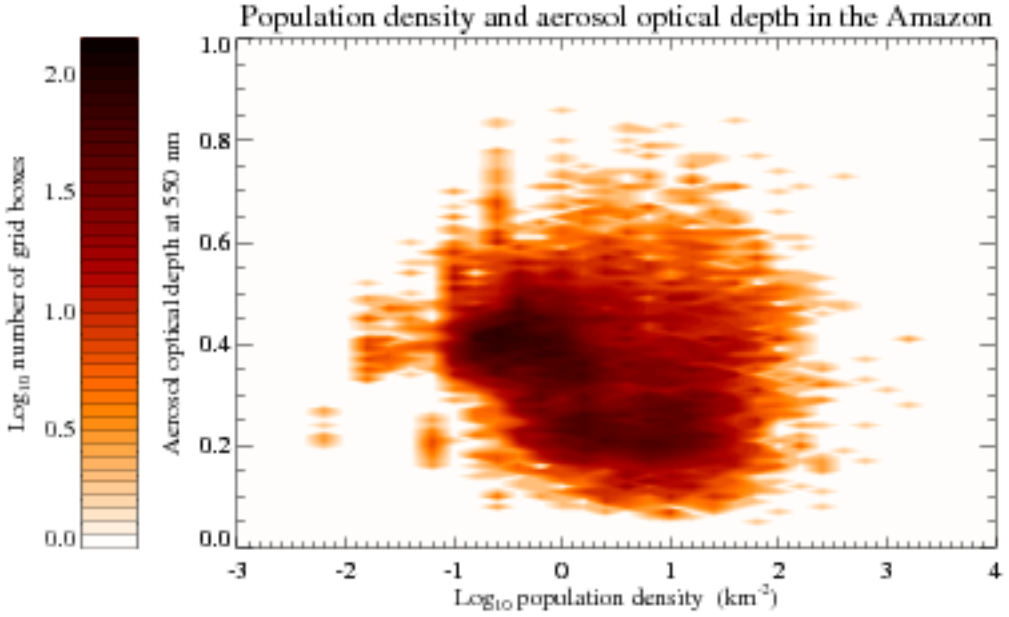

Figure 6.24: Density scatter plot of the base-10 logarithm of population density in the study region against mean aerosol optical depth (Figure 6.23), both on a 0.25° grid.

with a corresponding map of the mean aerosol optical depth at the same resolution. Averaging over all 5 years of data provides a seasonally-independent map of optical depth. Averaging on a 0.25° spatial scale as opposed to a finer resolution also helps to account for the fact that aerosol travels so elevated optical depths may be found away from source regions.

A scatter-density plot of the two datasets is shown in Figure 6.24 and reveals some clusters, as

opposed to a random relationship between the variables.

- A large number of pixels have a population density of between approximately 0.1 and 1 km^{-2} and a moderate mean optical depth of approximately 0.4. These correspond largely to productive forested areas.

- There is another cluster at slightly higher population densities (around 1 to 10 km^{-2}) with lower optical depths (0.25). These regions correspond generally to agricultural and non-forested vegetated regions and have a similar optical depth to many continental regions worldwide (see, for example, Figure 6.3).

- More densely-populated regions tend to have higher optical depths than this background value, although there is a large spread to the data. The highest optical depths tend to be encountered near the coasts of the study region, particularly the northern coast, which have moderate population densities (10 to 100 km^{-2}).

- Regardless of population density, the Andes mountains have a very low optical depth. This may be related to aforementioned inadequacies in modelling elevated regions.

The clusters provide evidence of links between population size and aerosol optical depth, although limited conclusions can be drawn from this analysis. More data are desirable as the majority of the study region has population densities on the lower end of the scale, meaning any signal from high populations may be masked. Additionally, a simple mapping between the two variables neglects the very different lifestyles (and hence probably resultant aerosol emissions) of urban and rural populations. Despite these shortcomings, the fact that clusters are observed in the results suggests that a more extensive analysis would be fruitful.

6.3.6 Weekly trends in aerosol optical depth

Concept of the 'weekend effect'

Recent work by Xia et al. (2008) used AERONET and MODIS data to investigate the weekly cycle of aerosol optical depth in various regions globally. The work found a larger weekly cycle in aerosol optical depth at urban locations as compared to rural locations. Additionally, in the USA, central Europe, south America, South Korea and the Middle East aerosol optical depths were seen to be

lower than the weekly mean around the cultural weekend (Sunday and Monday for the USA and central Europe). Decreased industrial activity at the weekends was presented as a possible reason for this, and the lack of a decrease on Saturday being due to it being preferred to Sunday for recreational trips. The lag between aerosol emission and removal may also explain the lack of a visible decrease on Saturday, and the continuance of the effect at the start of the working week on Monday. Diurnal cycles of aerosol optical depth have also been observed at urban AERONET sites (Smirnov et al., 2002) along with weekly cycles of tropospheric NO_2 (Beirle et al., 2003), providing further evidence for the impact of human activity on the atmosphere on such temporal and spatial scales. Bäumer et al. (2008) perfomed a similar study to Xia et al. (2008), using AERONET, for Europe alone, finding a similar weekly cycle.

A 'weekend effect' was formally defined in Xia et al. (2008) by taking firstly the average percentage departure from the weekly mean optical depth in a region during Wednesday, Thursday and Friday, and secondly the average departure during Saturday, Sunday and Monday. The weekend effect is the difference between these two 3-day averages, such that a larger weekend effect means greater variability of aerosol between week and weekend days. The study did not consider optical depth on Tuesdays while calculating the weekend effect. Positive values imply higher optical depths during the week, and vice-versa.

AERONET was used to examine the weekly cycle of individual sites, and MODIS of global regions as a whole. A similar analysis carried out here using AATSR.

Analysis with AATSR data: individual sites

The region of study contains 17 cities with populations of 1 million or greater (INE, 2008, IBGE, 2008, DANE, 2008, INEC, 2008, DEC, 2008, INEI, 2008). Daily mean optical depth was obtained from averaging all retrieval pixels lying within 0.15° of the centres of these cities (corresponding to a box approximately 30 km across, of typically 9-25 sinusoidal superpixels) over the 5 years of data. Choosing a box of data rather than a single pixel per city and averaging over the whole time-series helps to counter the low sampling of AATSR as compared to AERONET.

Of these 17 cities, for all but Lima the AATSR dataset contained good retrievals for each day of the week. The daily mean optical depths, together with their standard error (the standard deviation of the day's values divided by the square root of the number of points), are shown for each city

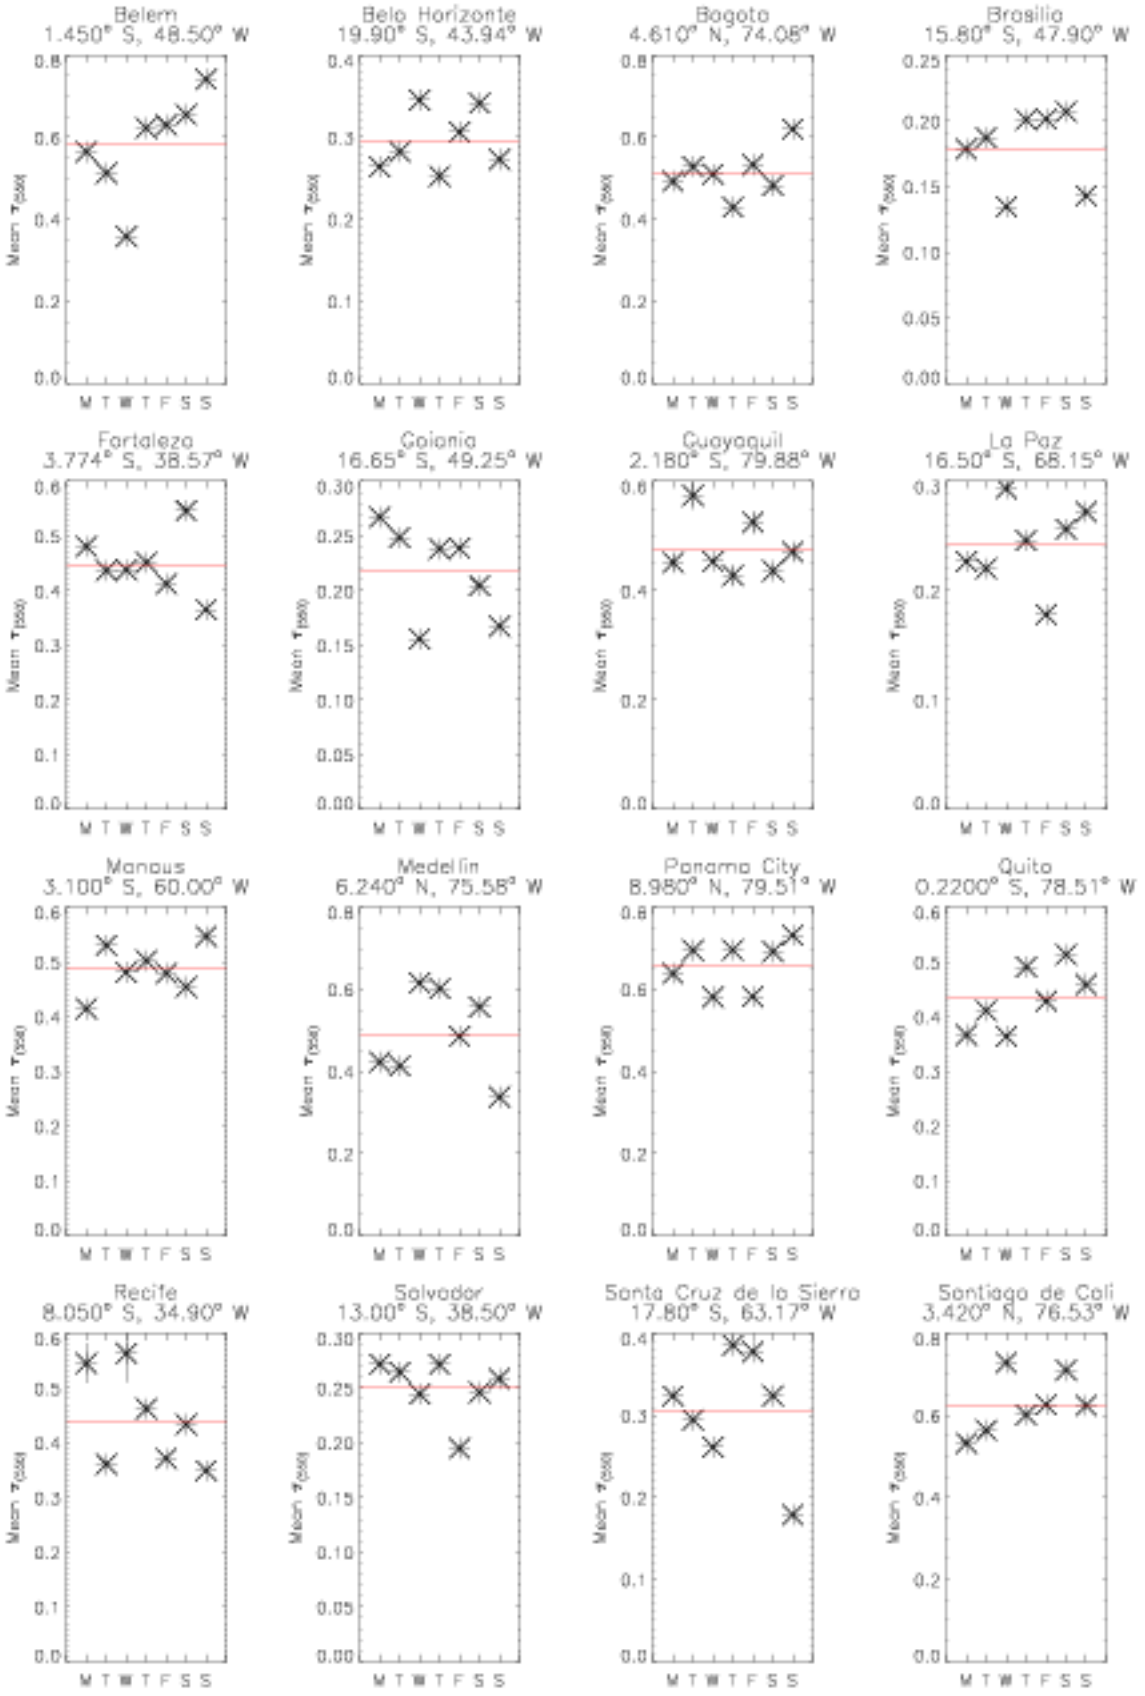

Figure 6.25: Daily mean τ_{550} of cities with a population of 1 million or more in the study region. Error bars indicate the uncertainty on the daily mean and the red line the weekly mean value. Labels on the x axis stand for the days of the week (from Monday to Sunday).

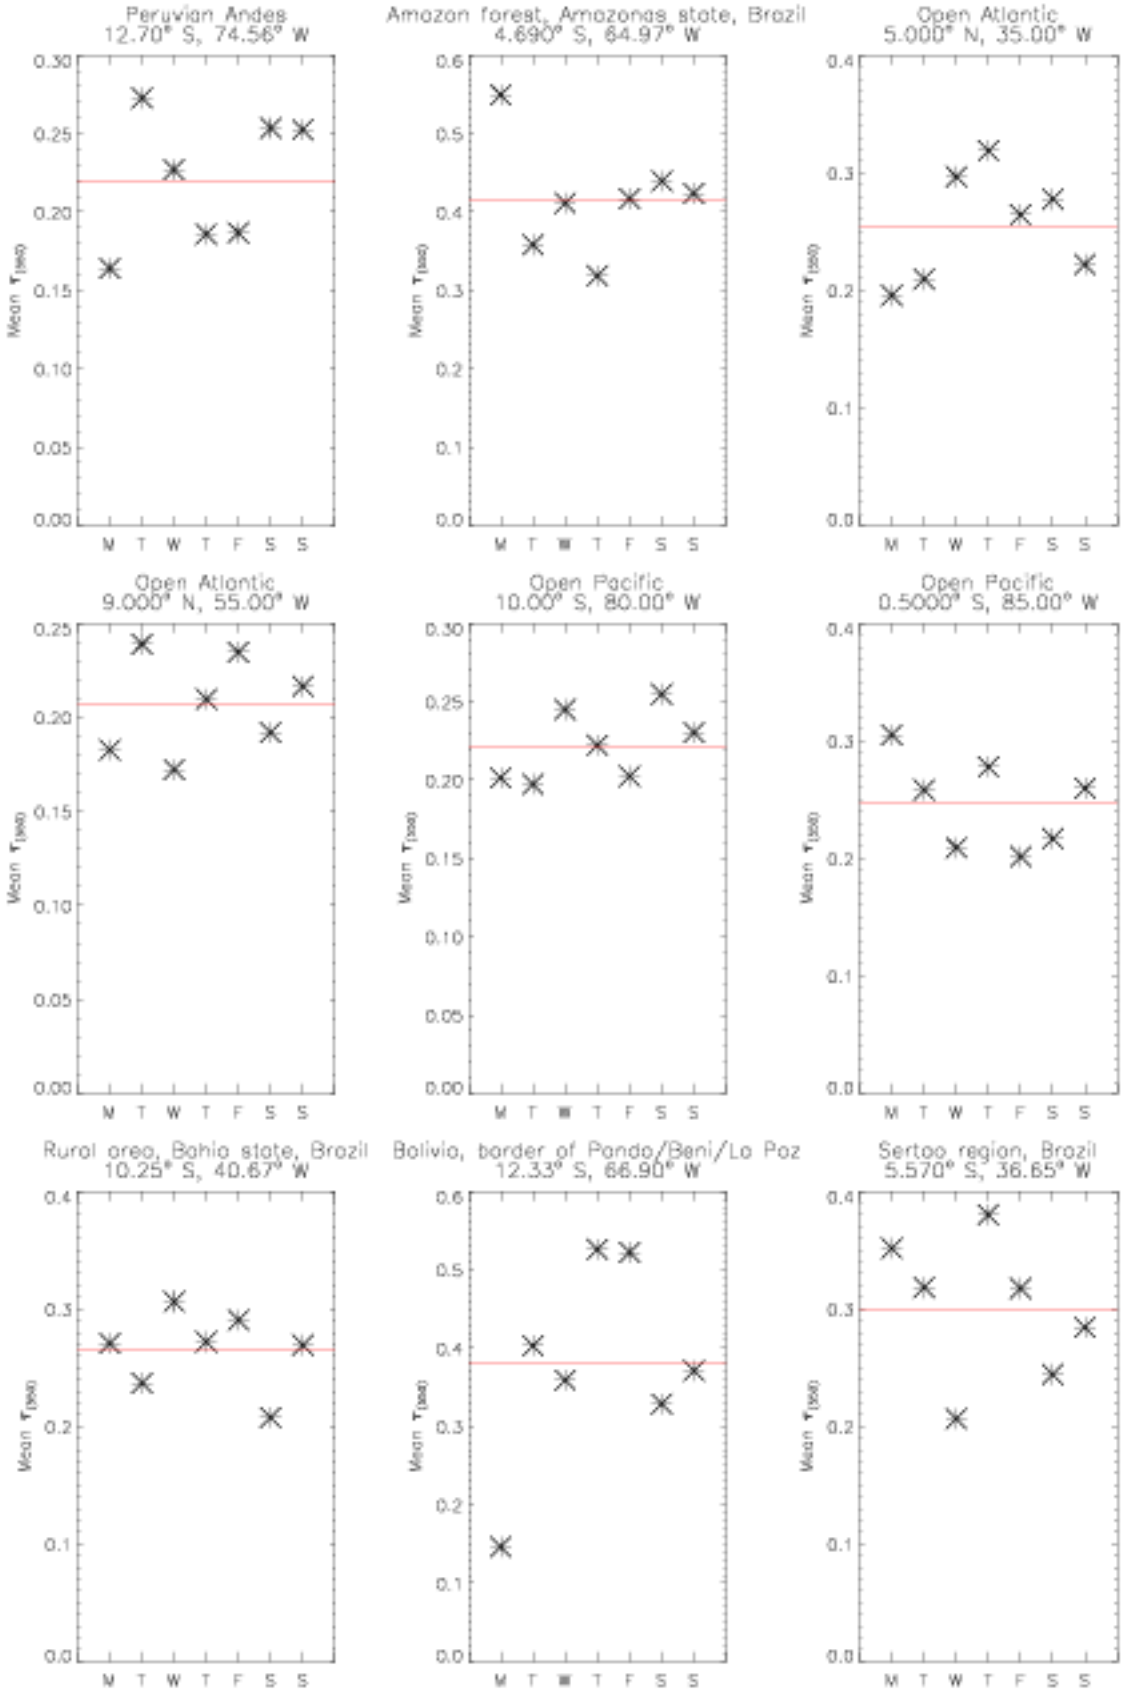

Figure 6.26: Daily mean τ_{550} of 9 rural regions in the study region. Error bars indicate the uncertainty on the daily mean and the red line the weekly mean value. Labels on the x axis stand for the days of the week (from Monday to Sunday).

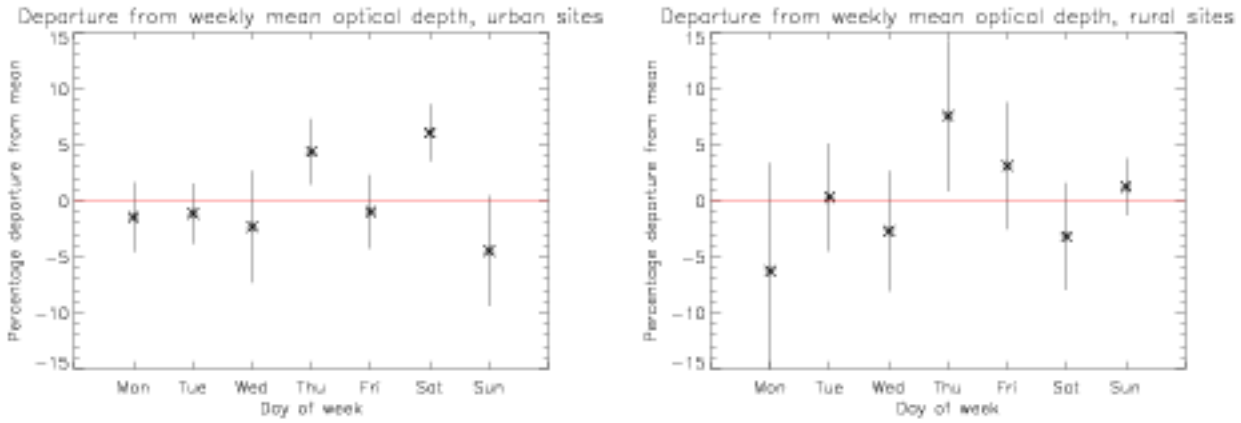

Figure 6.27: Left: Average percentage departure from weekly mean τ_{550} over all 16 cities for each day of the week. Right: The same, but for over all 9 remote sites. In both cases, error bars indicate the uncertainty on the daily departure and the red line zero.

in Figure 6.25. An analagous plot for 9 locations chosen for remoteness from cities is shown in Figure 6.26. The daily mean departure from weekly mean τ_{550}, with uncertainties defined in the same way, is shown for all sites in Figure 6.27.

It is apparent that each city has a unique weekly cycle, and in many cases the day-to-day variability is not significantly different from the weekly mean. Even upon averaging over all cities (Figure 6.27), the uncertainty on each day's departure from the weekly mean is of a similar size to the departure itself. The average weekend effect for these sites is $0.30\,\% \pm 2.89\,\%$. This is smaller than the typical values found using AERONET data by Xia et al. (2008) (which were of the order of 4 %) for cities in the USA and central Europe, implying for these cities there is not a significant change in aerosol between the week and weekend. Countries in this region are classified as developing economies (IMF, 2008) which may suggest that their industry does not shut off at the weekend to the extent that that in the USA and central Europe does. This is consistent with the lack of a strong cycle in Figure 6.27. It is also possible that domestic emissions are a large proportion of the anthropogenic signal in this region and do not show strong weekly variability.

Figure 6.26's rural locations show differing weekly cycles, such that after averaging them together the daily departure is zero within 1σ error bars for all days aside from Thursday. It would be expected for rural sites to lack a weekly cycle in τ_{550}, as the environment does not directly 'know' what day of the week it is. The average weekend effect over all sites is $5.39\,\% \pm 7.21\,\%$, which is zero within the 1σ uncertainty. The fact that the individual sites themselves *do* exhibit weekly cyclings leads to several possibilities, of which all are probably true to an extent:

- As aerosol is carried by wind from source regions, it is possible for it to pass over the remote areas and impact upon their weekly cycle.

- The sites (picked by analysis of satellite imagery for their apparent lack of human activity) are not as remote as assumed.

- The combination of high cloud cover and AATSR's lower temporal sampling as compared to AERONET means that the daily means can be based upon only a few different days of observations each. This may lead to skewed daily means for individual sites (which, as shown in Figure 6.27, average to nearly zero over all rural sites). This sampling argument also applies to cities in the region.

Weekend effects and their uncertainties are shown for all 25 sites in Table 6.1. In most cases the 1σ uncertainties are larger than the weekend effects themselves, implying, as suggested, either there is no overall weekly cycle of optical depth or AATSR's sampling is insufficient to characterise the weekly cycles on this scale.

In conclusion, the AATSR data do not show a weekend effect significantly different from zero overall for urban or rural sites in this region of South America. For industrialised regions, this may be an indication that industrial activity does not vary between the week and the weekend. For rural areas, this corresponds to the idea that the environment is not aware of which day of the week it is. In both cases the value, although smaller than its uncertainty, was positive which implies that the weekend may have lower optical depths than the week at these sites. However, this analysis is hampered by AATSR's low sampling, which results in large uncertainties when attempting an analysis on this scale. AERONET, as employed by Xia et al. (2008), is a better tool for this purpose as observations are available much more frequently.

Analysis with AATSR data: the region as a whole

The second component of Xia et al. (2008) was to use MODIS data to examine the magnitude of the weekend effect at 1° resolution on a subcontinental scale for the USA, Europe, China and India. In this section, AATSR data are used to perform a similar analysis for the region of South America examined. The daily percentage departure from weekly mean optical depth over this region is shown in Figure 6.28, from which a weekend effect of -1.37±0.30 % can be calculated. This indicates that

Site name	Latitude, degrees	Longitude, degrees	Weekend effect, % (0.15° circle)	Weekend effect, % (1° box)	Consistent at 90 % level?
Amazon forest, Brazil	-4.69	-64.97	-21.26±8.69	-8.50 ±3.29	Yes
Andes, Peru	-12.70	-74.56	*-10.64±20.12*	*-4.66 ±8.70*	Yes
Belém	-1.45	-48.50	-20.06±10.48	-28.46 ±3.48	Yes
Belo Horizonte	-19.90	-43.94	*2.91±16.50*	*-2.41 ±6.05*	Yes
Bogotá	4.61	-74.08	*-7.96±13.09*	5.39 ±3.25	Yes
Brasilia	-15.80	-47.90	*1.47±12.06*	*2.55 ±4.08*	Yes
Fortaleza	-3.77	-38.57	*-6.73±12.17*	*-0.53 ±4.85*	Yes
Goiânia	-16.65	-49.25	*-0.98±11.08*	*1.85 ±3.94*	Yes
Guayaquil	-2.18	-79.88	*3.35±19.46*	20.00±8.82	Yes
La Paz	-16.50	-68.15	*-5.11±15.51*	*-2.29 ±5.88*	Yes
Manaus	-3.10	-60.00	*3.27±8.98*	*4.21 ±3.01*	Yes
Medellín	6.24	-75.58	26.28±10.16	17.71±4.18	Yes
Open Atlantic	5.00	-35.00	24.17±8.55	7.69 ±3.33	Yes
Open Atlantic	9.00	-55.00	*4.09±6.14*	-5.63 ±2.29	Yes
Open Pacific	-10.00	-80.00	*-2.52±4.16*	*2.20 ±1.83*	Yes
Open Pacific	-0.50	-85.00	-12.56±6.47	-9.67 ±2.34	Yes
Panama City	8.98	-79.51	*-10.29±7.24*	*5.49 ±4.14*	No
Quito	-0.22	-78.51	*-4.23±9.02*	*-1.85 ±5.79*	Yes
Recife	-8.05	-34.90	*5.16±25.24*	10.85±6.86	Yes
Rural Bolivia	-12.33	-66.90	49.27±14.86	14.87±4.40	No
Rural Brazil	-10.25	-40.67	*15.32±12.73*	11.33±3.92	Yes
Salvador	-13.00	-38.50	*-8.65±12.56*	*-2.15 ±3.66*	Yes
Santa Cruz de la Sierra	-17.80	-63.17	21.56±10.94	16.90±3.87	Yes
Santiago de Cali	3.42	-76.53	*4.79±7.39*	15.57±4.32	Yes
Sertão region, Brazil	-5.57	-36.65	*2.66±8.73*	*-0.080±3.37*	Yes

Table 6.1: Weekend effects and their uncertainties (1σ) for the 16 urban and 9 rural sites observed at finer (0.15° radius circle) and coarser (1° box) resolution. Values in italics are not significantly different from zero at the 90 % confidence level.

the region has slightly higher optical depth over the weekend than during the week (although as the definition of the study region is arbitrary, the utility of this measure is limited). The day-to-day variability of the region is smaller than that reported by Xia et al. (2008) for the regions considered there (the USA, central Europe, eastern China and India); aside from the spike on Thursday, the weekly aerosol cycle is similar to that of India.

There is no obvious explanation for the elevated optical depths on Thursday as compared to other days of the week. As the daily percentage departures are all reasonably small (2 % or less), it is possible that the cycle in Figure 6.28 arises due to a few days of high or low aerosol loading skewing the daily mean values.

Mean aerosol optical depths over Wednesday, Thursday and Friday (the 'week') and Saturday,

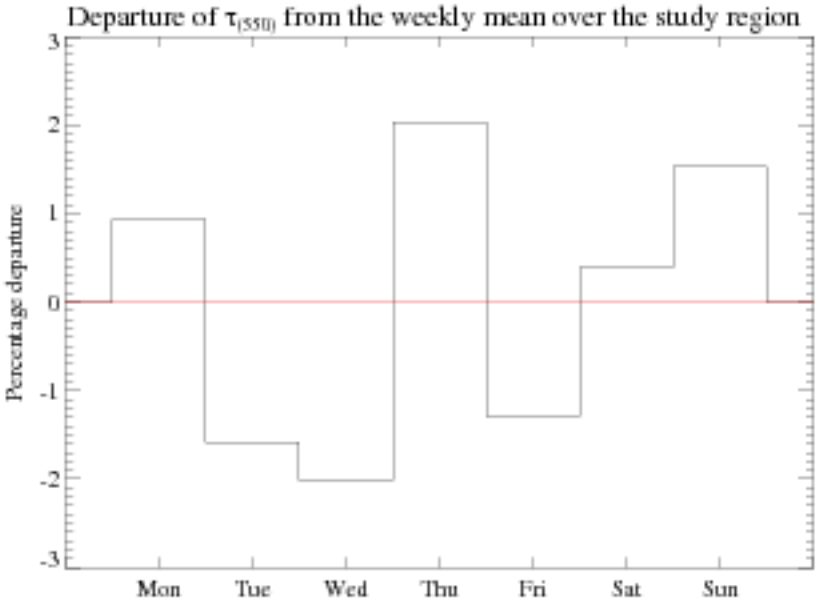

Figure 6.28: Percentage departure of daily mean aerosol optical depth from the weekly mean value for the study region.

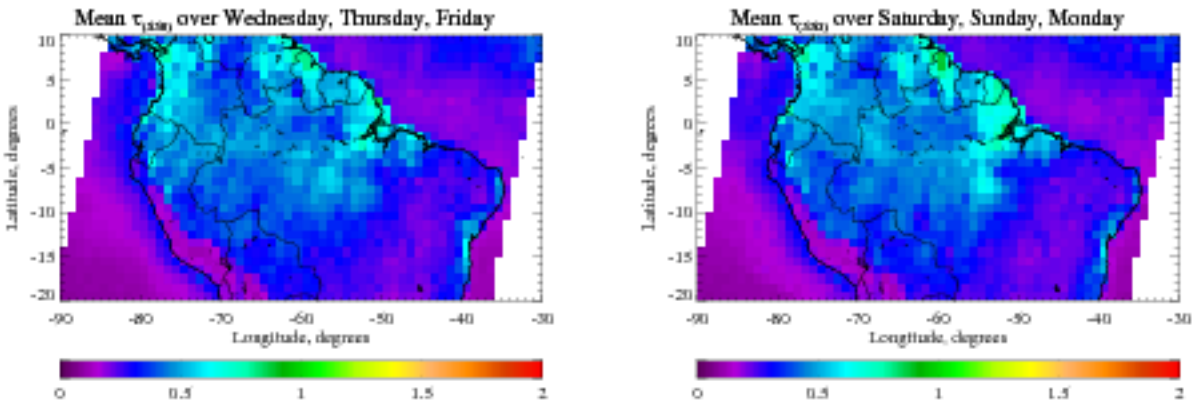

Figure 6.29: Left: Mean aerosol optical depth for Wednesday, Thursday and Friday (the 'week') for the study region. Right: The same, except for Saturday, Sunday and Monday (the 'weekend') Results on a 1° grid.

Sunday and Monday (the 'weekend') are shown in Figure 6.29. The difference between these periods, together with the calculated weekend effect, is given in Figure 6.30. Optical depths during the week and weekend show similar spatial patterns, and over much of the region the difference between week and weekend is small (0.05 or smaller). Differences approaching 0.3 are, however, observed and the patterns do not appear to be strongly linked to population density (Figure 6.23). If not real differences, it is possible that these large differences are the extreme outliers of a Gaussian distribution of small differences.

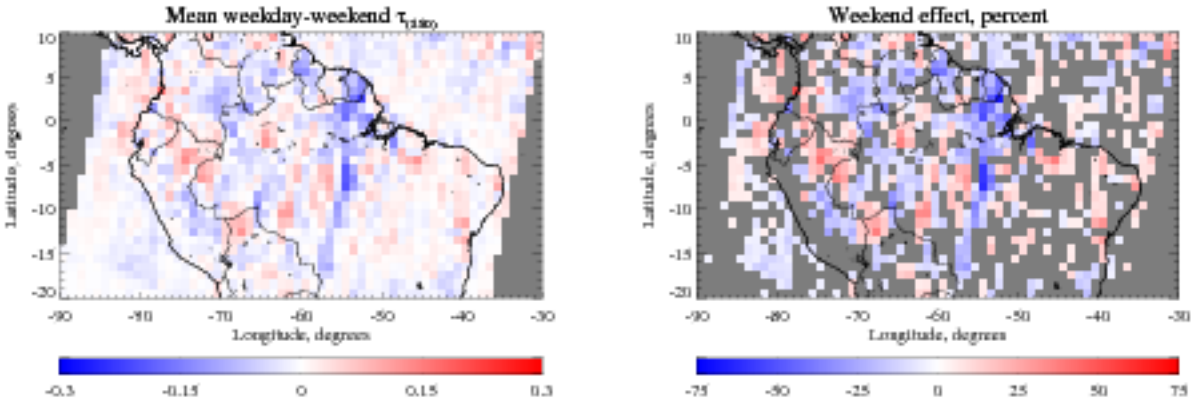

Figure 6.30: Left: Difference between 'weekday' and 'weekend' (Figure 6.29) aerosol optical depth for the study region. Right: The 'weekend effect' for the study region. Values are only shown for points significantly different from zero at the 90 % confidence level. Positive values imply a lower optical depth during the weekend. In both cases, the data are on a 1° grid and missing data are indicated by a grey fill value.

As the weekend effect is only shown where significantly different from zero (at the 90 % confidence level), coverage is reduced by approximately a half, with the majority of the removed pixels being areas with a small optical depth difference. This suggests an absence of a weekly aerosol cycle over a large portion of the study region. Much of these areas are also sea, mountainous terrain or agricultural land, which intuitively would not be expected to show a weekly aerosol cycle. Effects in the remaining regions of Figure 6.30 require further explanation.

- Various parts of the Atlantic and Pacific oceans show either positive or negative weekend effects, although generally weaker than over the land. As oceanic optical depths are low, these differences are small and it is possible that the weekend effects over some of these regions are a consequence of retrieval error and AATSR's sampling. There is no obvious explanation for a weekly cycle.

- Rural regions (low population density in Figure 6.23) frequently show negative weekend effects (higher optical depth at weekends). This is in agreement with the results obtained by Xia et al. (2008) for rural AERONET sites in the region, which exhibited this 'reverse weekend effect', although there is no obvious explanation for why this should be the case.

- Urban regions (high population density in Figure 6.23) frequently show positive weekend effects. This is in agreement with the AERONET measurements discussed in Xia et al. (2008) for Sao Paulo and Santiago (although that work does not mention other cities in this region)

and corresponds to the intuitive idea of reduced industrial activity at the weekend. The fact that a weekend effect can be distinguished at 1° spatial scales is an indication that the impacts of urban activity from a smaller source region can span these scales, and/or that urban areas can typically extend to this scale.

There are, however, points in Figure 6.30 where these explanations fail to explain the observed statistically-significant weekend effect. It is possible that some more complicated activity is at work producing the weekly cycles in these regions, and additionally AATSR's low sampling is likely to contribute to the uncertainty.

Comparison of weekend effect derived at both scales

This work has attempted to quantify the weekend effect over part of South America at local (0.15° circle) and coarser (1° box) spatial scales. At first glance, the methods appear to lead to different conclusions as the fine-scale approach found no significant weekend effect averaged over a selection of urban and rural sites, while the coarser-scale investigation found evidence of a positive weekend effect near many cities, as well as a reverse weekend effect in other areas.

The agreement between the two approaches can be examined by looking at the consistency of the weekend effect at the 25 sites where it was calculated at the finer scale. Estimates of the weekend effect made at both scales, together with their uncertainties, are presented in Table 6.1. For 23 of the 25 locations the estimated weekend effects agree at the 90 % confidence level. The two sites where they do not agree are as follows:

1. At Panama City, the first approach found a negative weekend effect while the second a positive one. Panama City is coastal so it is possible that poorer-quality retrievals, due to the inhomogeneity of coastal regions, lead to poor estimation of the weekend effect.

2. At the rural location in Bolivia, both methods reported positive weekend effects but of an approximately threefold difference in magnitude. Aerosol inhomogeneity or AATSR sampling may explain these results, although why this isolated area shows a positive weekend effect is unclear.

The general good agreement is due in part to the very large uncertainties on the weekend effect when calculated on fine spatial scales. This highlights the effectiveness of AERONET for small-scale

studies like this, where satellite sampling is inadequate. Only 7 of the 25 sites yielded significant weekend effects at the fine scale, as compared to 13 at the 1° scale, with all of the former set also present in the latter (non-italic values in Table 6.1). As weekend effect uncertainties are, in general, much lower at the 1° scale this demonstrates the utility of satellite data such as MODIS and AATSR for examining aerosol cycles at this level. The 1° data are therefore discussed in the remainder of this section.

At six of the 16 cities with populations of 1 million or more, a positive weekend effect was observed at the 90 % confidence level. This ranged in magnitude from 5.4 % in Bogotá to 20 % in Guayaquil, which is a similar magnitude to the value of 16 % found by Xia et al. (2008) for Sao Paolo and Santiago. Only one of the cities, Belém, showed a reverse negative effect of -28 %. The remaining nine did not have any significant difference between aerosol levels during the week and at the weekend. This suggests changes in industrial activity and lifestyle in some cities in the region, as found by Xia et al. (2008) for many other regions worldwide, but that in many cities any weekly cycle is too small to be detected.

Of the nine rural sites, four show a positive, two a negative and three no weekend effect. Small negative weekend effects were reported by Xia et al. (2008) for rural AERONET sites in the region; of the four with a positive effect here, two were located in the Atlantic Ocean (and one with a negative in the Pacific) and it is possible that aerosol in these oceanic regions is affected by weekly patterns of shipping or are spurious results arising from poor sampling.

6.3.7 Deforestation and regeneration in the Amazon basin

The Amazon rainforest has most of its area in Brazil (with the 'Legal Amazonia' region in that country being composed of the states of Acre, Amapá, Amazonas, Mato Grosso, Pará, Rondônia, Roraima, Tocantins, and part of Maranhão) and Peru. It also comprises a part of each of Bolivia, Colombia, Ecuador, French Guiana, Guyana, Suriname and Venezuela. It is a densely-forested region extending throughout the basin of the Amazon river. Beginning several decades ago large regions have been deforested, primarily to clear land for cattle-ranching or subsistence farming, with soybean farming becoming increasingly prevalent in recent years (Kirby et al., 2006). Particularly in the non-Brazilian Amazon, mining can also be a significant cause (Fonseca, 2004). The Amazon is an area of high biodiversity (see, for example, Lewinsohn and Inácia Prado, 2006), and also plays an

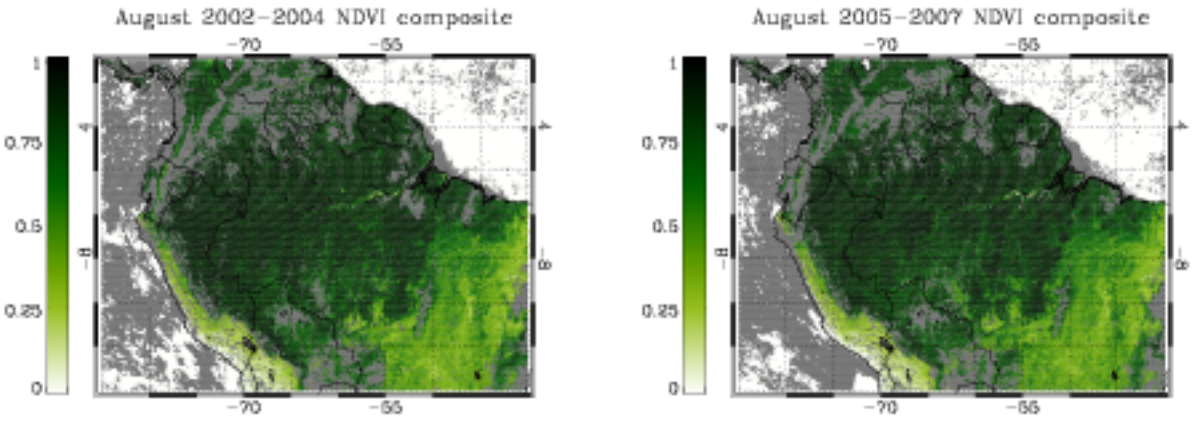

Figure 6.31: NDVI during August during the early part of the time-series spanned (left, composite from 2002-2004) and the later part of the time-series spanned (right, composite from 2005-2007). Results presented on the retrieval's approximately 0.1° sinusoidal grid. Missing data are indicated by a grey fill value.

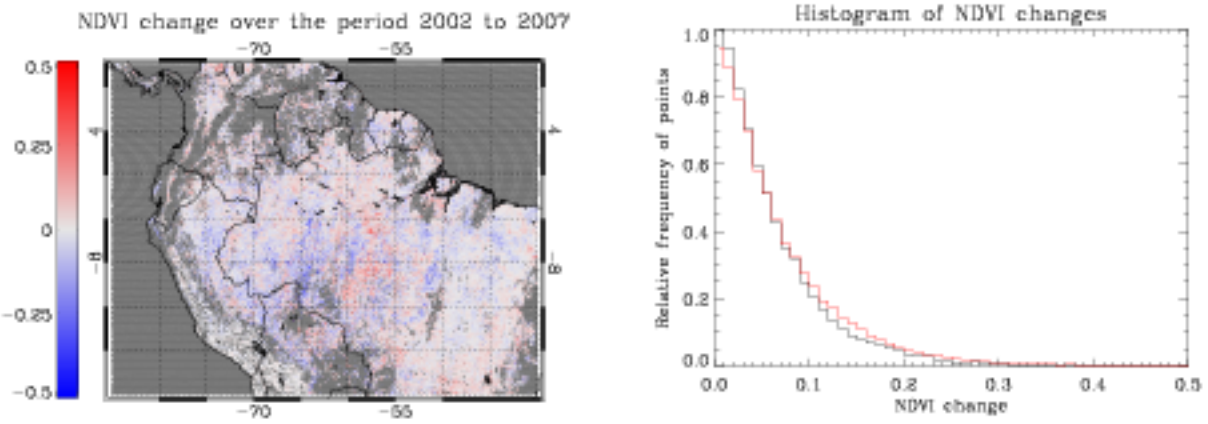

Figure 6.32: Left: Difference between NDVI in August at the end (right-hand part of Figure 6.31) and start (left-hand part of Figure 6.31) of the time-series spanned. Results presented on the retrieval's approximately 0.1° sinusoidal grid. Missing data are indicated by a grey fill value. Right: histogram of the magnitude of the values from the left-hand plot in bins sized 0.01. Black values indicate an increased NDVI towards the end of the time series, and red a decreased NDVI.

important role in the global carbon cycle, accounting for approximately 10 % of global terrestrial primary production (Melillo et al., 1993), thus there is considerable interest in attempts to quantify the extent of forest cover change. The extent of 'deforestation' is defined here by the net change in area of forest cover during a period (which covers both regions where trees have been removed and those where they have regrown).

Satellite data have often been used to map regions of deforestation, beginning with AVHRR data (Tucker et al., 1984) and more recently MODIS and LANDSAT (such as Ichii et al., 2003, Wynne et al., 2007, Hayes and Cohen, 2007 or Hayes et al., 2008). LANDSAT is a series of satellites which

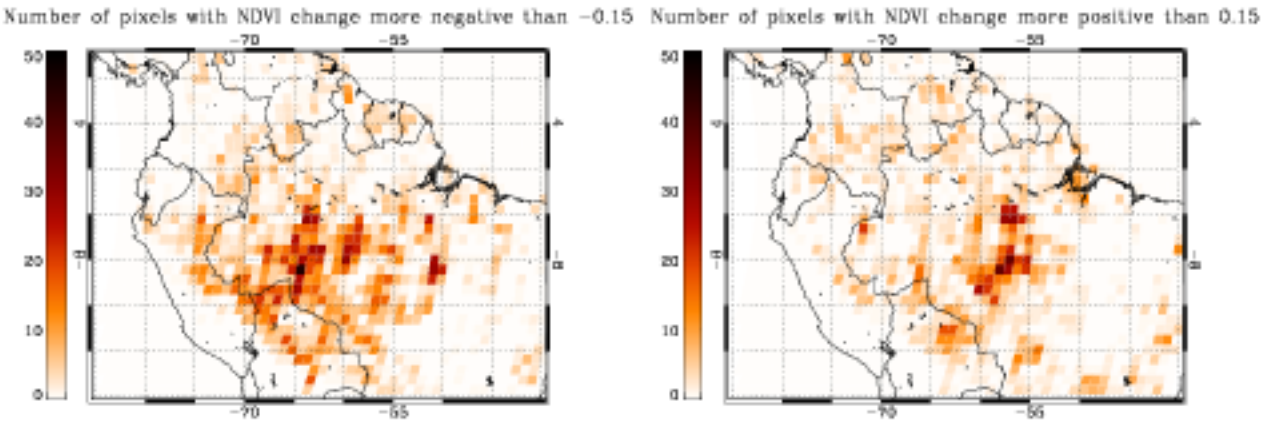

Figure 6.33: Left: Density of retrieval pixels where the NDVI change (Figure 6.32) is more negative than -0.15, binned to an approximately 1° sinusoidal grid. Right: The same, except for pixels where the NDVI change is more positive than 0.15.

image the Earth in the visible and infrared, of which the first was launched in 1972. To the author's best knowledge, AATSR data have not yet been used for this purpose. These satellite detection mechanisms normally make use of changes in vegetation indices such as NDVI to estimate change in forest cover, relating the two for a pixel through use of calibrating to some selection of reference sites where both the NDVI and fractional forest cover are well-known. This is because NDVI is not a simple map of vegetation cover, depending on, for example, the types of vegetation and soil present (Holben and Kimes, 1986).

Such an analysis is out of the scope of this work; however, consideration of changes in NDVI over the 5-year period studied on a retrieval-pixel scale enable the mapping of regions where AATSR data suggest that deforestation, or regeneration of cleared forest, is likely to have happened. Figure 6.31 shows composite NDVI images from the AATSR data processed during this period. The method for generating these maps was as follows:

- For the left-hand plot, the monthly mean NDVI was obtained for August 2002 on the retrieval grid from the individual orbits processed. This typically led to up to 5 estimates of the NDVI for each of the 164,772 pixels in the region. A total of 83,743 pixels contained at least one estimate of the NDVI.

- Gaps in the field were filled by adding in monthly mean NDVI data for August 2003 to the pixels where it was absent for August 2002. A total of 22,452 pixels were filled in this way.

- Remaining gaps were filled by the same process, using August 2004. A further 11,363 pixels

were filled in this way, leaving 47,214 pixels with no data.

This process was carried out for August 2007 (77,369 pixels with retrievals), but in reverse, filling gaps with data from 2006 (29,159 pixels) and 2005 (10,785 pixels). A total of 47,459 gaps remained. Filling gaps substantially increased the coverage of the datasets for comparison, although the majority of points were from 2002 and 2007. Successive Augusts were used to ensure no seasonal bias was introduced. 85,802 land pixels contained data in both composite sets.

The resulting change in NDVI between the two datasets in both map and histogram form is shown in Figure 6.32. The uncertainty on any individual NDVI measurement in this region is typically of the order of 0.06. The map shows distinct regions of widespread decrease or increase in NDVI over the period, with changes in the Amazon being generally larger than those in other parts of the study region. The largest changes in the Amazon are generally found in the Brazilian states of Amazonas, Rondônia and Matto Grosso. Historically these regions are where deforestation occurs (according to, for example, Tucker et al., 1984 and official figures from the Brazilian National Institute for Space Research (INPE) via Câmara et al., 2006), and hence also regions where regeneration of forest on previously-cleared land is possible.

INPE figures suggest a net loss of forest of the order of 115,000 km^2 in Legal Amazonia between 2002 and 2007. In Peru, Oliveira et al. (2007) estimate approximately 645$\pm$325 km^2 yr^{-1} net of ground has been deforested in recent years, leading to approximately 3,000 km^2 between 2002 and 2007, albeit with considerable uncertainty on this value. The same source estimates a similar coverage of disturbed forests.

Year-by-year estimates of forest cover change are harder to come by for the other countries with territory in the Amazon. The United Nations Food and Agriculture Organisation reports (FAO, 2006) provides estimates of annual net deforestation rates over the period 1990-2005. Assuming these values are representative of the 2002-2007 period leads to overall cover change estimates of 13,500 km^2 in Bolivia, 2,400 km^2 in Colombia, 9,990 km^2 in Ecuador, 150 km^2 in French Guiana and 14,400 km^2 in Venezuela over this period. No deforestation was reported by FAO (2006) for Suriname and no data was collected by FAO (2006) for Guyana.

Together, these figures lead to an estimate of 158,000 km^2 of forest lost between 2002 and 2007. The majority of this was from Brazil. Note that these figures are in terms of net forest change, so the total area over which deforestation occurs will be larger than this, as some land deforested prior to

the period may regenerate forest cover.

The right-hand portion of Figure 6.32 shows the relative frequency of positive and negative NDVI changes over the land for this period. Of the 85,802 land pixels with available data, 43,682 had a negative NDVI change and 42,120 a positive. Therefore there was a net negative change of 1,562 pixels. Although a simple measure, the resultant net negative change of 1,562 pixels corresponds to an area of approximately 156,000 km^2, which is in remarkable agreement for the previously-derived net deforestation estimate of 158,000 km^2 over the period. This is an important result which demonstrates that, although little can be said about changes in forest cover on a retrieval-pixel scale without establishing the aforementioned direct NDVI-forest cover relationship, on a regional scale the AATSR-derived NDVI provides an estimate of deforestation in good agreement with independent data sources.

As a next step, the density of pixels exhibiting a strong NDVI change (magnitude of 0.15 or more) is mapped to an approximately 1° sinusoidal grid in Figure 6.33. This reveals the areas of large change are mostly confined to the Amazon. The area of pixels with a strong negative change is approximately 417,000 km^2 and strong positive 292,000 km^2; assuming these strong changes are due to deforestation and reforestation events leads to a net decrease of 125,000 km^2, in reasonable agreement with the previous estimates, although the exact figure depends on the NDVI threshold adopted. A threshold of 0.125 gives an estimate of 178,000 km^2 net deforestation, with similar spatial distribution; this suggests that for a simple mask of NDVI change significance, an appropriate threshold would be in the region 0.125 to 0.15. Such a threshold would not be expected to hold for the whole globe, as it would depend upon the types of vegetation and soil being observed.

Frequent occurrences of strong NDVI decrease are concentrated in Amazonas, Rondônia and Matto Grosso in Brazil, eastern Peru and northern Bolivia. Southern and eastern regions of Colombia and Venezuela also contain significant negative NDVI change, as do Guyana, Suriname and French Guiana. There is little evidence of deforestation in Ecuador, which is surprising due to the estimate of 9,990 km^2 net forest loss, but this could be due to the large gaps in data coverage in this region (Figure 6.31). The NDVI decreases detected over Suriname and French Guiana are largely balanced by pixels showing an NDVI increase, which is consistent with the low deforestation estimates in these countries. Regions of frequent strong NDVI decrease in Brazil and Peru in Figure 6.33 are generally consistent with deforestation 'hotspots' measured by others (see, for example, Alencar, 2005, Souza

et al., 2005 and Neef et al., 2006 for Brazil and Oliveira et al., 2007 for Peru), although there is a comparative lack of data for the other countries.

Positive NDVI change as shown on the right of Figure 6.33 shows different spatial patterns, which is consistent with the idea of forest regrowth in areas that had been previously deforested. However, it may also be due to the growing of dense, green crops on land which had previously been comparatively bare (scorched ground from slash-and-burn agricultural practices or grasslands). The largest region of NDVI increase appears in a belt stretching through northern Bolivia and the Brazilian states of Mato Grosso and Pará. Parts of Rondônia, Amazonas and Acré also show density of significant NDVI increase. Regeneration of forest in these parts of Brazil and Peru has been previously noted by, for example, Lucas et al. (2000), Neef et al. (2006) and Oliveira et al. (2007) although again estimates for other countries are hard to come by.

In summary, the NDVI of the study region showed notable differences between the start and end of the 2002-2007 study period. The strongest variations were observed in the Amazon rainforest, and patterns of this variability are in general good agreement with known areas of deforestation and forest regeneration in Brazil and Peru. For the first time an estimate of the extent of net deforestation of the Amazon was made using atmospherically-corrected AATSR data, arriving by two methods at figures of 125,000 km^2 and 156,000 km^2. The latter of these two is in strong agreement with independently-derived estimates of 158,000 km^2. NDVI has not been mapped directly to forest-cover change on an individual-pixel scale, however, as such a task requires a calibration dataset where the cover is already known. It is suggested, however, that vegetation monitoring would be a possible future use of AATSR data, either through NDVI or a more quantitative vegetation index.

Chapter 7

Conclusions and future work

7.1 Summary of progress

7.1.1 Algorithm development and validation

The algorithm presented in this work represents a significant improvement on the previous nadir-view ORAC (A)ATSR aerosol retrieval. In summary, the following areas have been thoroughly updated:

- The algorithm now makes use of visible channel data from both instrument views; the multi-view information provides a particular improvement over bright surfaces (such as the land and the sun-glint area of the ocean).

- Measurement, forward model and *a priori* errors have been examined so that appropriate values may be used. This ensures that information is not thrown away in the retrieval by weak constraints, while the retrieval is not hampered by overly strong constraints.

- The surface BRDF forward model of Thomas (2007) has been implemented.

- A model has been developed to describe the BRDF of the sea surface, taking into account factors such as the geometry and contributions from oceanic whitecaps, glint reflectance and underlight from dissolved pigments.

- Treatment of the land surface reflectance has been improved by accounting for the differences between the MODIS and AATSR sensors.

- Surface albedo is now retrieved at all wavelengths, breaking the dependence on the spectral shape of the *a priori* and providing additional data products.

Retrieved aerosol optical depth has been validated by a comparison with AERONET data and found to show a good level of agreement. The correlation coefficient is 0.775 at 550 nm and 0.533 at 870 nm, and best-fit lines have gradients of 1.10 and 1.08 respectively. This is a substantial improvement over the previous nadir-view retrieval. The average precision by which τ_{550}, τ_{890} and r_{rms} are retrieved has also improved, from ± 0.156, ± 0.09 and ± 0.257 μm respectively for the nadir-view retrieval to ± 0.037, ± 0.027 and ± 0.178 μm respectively for the new scheme, for a month of global data.

In addition, a monthly mean comparison with MISR data shows similar regional trends in aerosol properties, although the correlation between the datasets is lower (0.46 over land and 0.50 over sea). This difference is thought to be largely due to sampling differences. An advantage of this algorithm over others is the optimal estimation framework, which allows for rigorous estimation of retrieval errors and quality control statistics given the thorough error budget developed for measurement and *a priori* data. Finally, the internal consistency of the algorithm has been checked and it has been found to perform well. Approximately 80 % of retrievals converge with an acceptable cost.

7.1.2 Applications of the algorithm

Chapter 6 describes the application of the retrieval algorithm in ways which AATSR data have to the author's best knowledge not previously been exploited. As such this work represents novel uses of AATSR data:

1. Generation of false-colour atmospherically-corrected images of both land and ocean. This reveals different types of large-scale land cover, and regions of differing ocean colour.

2. A global comparison of aerosol and surface properties in August 2002 and August 2007. This reveals regions of possible long-term trends in aerosol and surface properties and also regions of comparative stability. Much of Europe and western Asia saw a decrease in τ_{550} of approximately 0.2 between the years while an increase of similar size was observed around south-east Asia. Changes in the NDVI of the surface were linked to seasonal precipitation differences.

3. Processing of a 5-year time series (August 2002 to August 2007) of data in the region of South America around the Amazon basin. The dataset has been used to examine seasonal and interannual variations in aerosol and surface properties. Weekly cycles of aerosol optical depth have been examined, along with the relationship between optical depth and population density for the area. In general it was found that satellite sampling is generally insufficient to reveal weekly cycles for point locations, but at coarser resolutions they can be discerned. Finally, NDVI data have been used to examine the response of the Amazon to drought and quantify the extent of net deforestation over the 5-year period, producing estimates (125,000 km^2 - 156,000 km^2 by two methods) in good agreement with independent data sources (158,000 km^2).

These undertakings provide answers to—or at least steps toward—various science questions:

- What information can we obtain from modern satellite radiometers?
- What assumptions must we make, and what capabilities would be needed to improve upon them?
- How fast is the Earth's surface and atmosphere changing?
- What is the impact of cities on their local and regional environment?
- What are the characteristics of aerosol and surface properties in the Amazon region, and how do they vary with time?
- How fast is deforestation of the Amazon occurring?

7.2 Directions for future work

During the development and validation of the algorithm, several points where further effort may result in improved results have been uncovered. These are discussed briefly below.

7.2.1 Incorporation of IR channels

The AATSR instrument measures at 3 channels in the thermal infrared region of the spectrum which, in addition to cloud flagging, are used operationally to derive SST. There is scope for extending the aerosol retrieval algorithm to use these measurements and allow the retrieval of surface temperature

and aerosol layer height. It is also expected to improve characterisation of the aerosol products and the speciation, as an increased number of measurements would make the aerosol signature from different models more distinct. This simultaneous retrieval should allow SST retrievals of equal or better precision than the operational product.

Progess has been made in incorporating the infrared channels of the SEVIRI instrument into the ORAC retrieval by Carboni et al. (2008). However development is still ongoing as several issues are yet to be addressed:

- Aerosol microphysical properties in the infrared derived from different sources, ostensibly for the same aerosol type, have been found by Carboni et al. (2008) to differ significantly at AATSR/SEVIRI wavelengths of interest and so result in differing retrieval results. In these cases is not clear which is the most appropriate set of aerosol properties to use. This is similar to the aerosol speciation problem uncovered for the visible channel aerosol retrieval.

- In order to retrieve the surface temperature, the surface spectral emissivity must be known to high precision. This is a particular difficulty over the land. The accuracy requirement for AATSR SST defined in ESA (2007) is 0.3 K; a 1 % error in emissivity for a surface at 300 K at 11 microns leads to an error of approximately 0.6 K.

Due to these difficulties, this work focusses on a thorough treatment of the visible channels of AATSR in order that it be of the highest quality. Incorporation of thermal infrared channels into the AATSR retrieval will not take place until issues arising from the SEVIRI retrieval have been addressed, in case an incomplete treatment were to hamper the AATSR retrieval.

7.2.2 Improve surface modelling and error characterisation

As, particularly over land, the signal observed by AATSR contains a strong contribution from the surface there is scope for an improved treatment of the surface reflectance (ultimately making the error arising from improper surface treatment smaller than other sources of forward model error). Two areas which offer scope for improvement are:

1. Over the sea, refinements could be made to the BRDF model to account for differing properties of shallow, turbid coastal waters (often referred to as 'Case II'). The BRDF model of Chapter 3 was designed for open ocean as it covers the largest proportion of the globe, and indeed

even in coastal waters the proportion of successfully-converging retrievals is good. However in some regions (such as the Amazon delta, as discovered in Chapter 6) the model is inadequate to describe the highly turbid waters.

2. Improved constraints of surface albedo over the land will help to partition more accurately surface and atmospheric effects in the TOA reflectance, and hence improve aerosol speciation. These constraints may be driven by improvements to the MODIS BRDF data products, and an improved understanding of the temporal variability of surface reflectance. Investigation into temporal variability has been proposed for an MPhys project.

7.2.3 Improvements in forward model precision

Increasing the precision of the forward model will, like improving the modelling of surface reflectance, allow for better discrimination of aerosol and surface effects in the TOA reflectance and so result in smaller error estimates on retrieved quantities and improve the efficiency of lowest-cost aerosol speciation. This work has quantified the forward model precision and shown that, for a typical 10 km$\times$10 km superpixel, the dominant error sources are surface model parameter uncertainty (the ratios $R_{\mathrm{SLW}} : R_{\mathrm{SBD}}$ and $R_{\mathrm{SLW}} : R_{\mathrm{SLB}}$), yielding an error of 2-3 % of the signal over land and 0-8 % for sea (depending on geometry and wavelength) and numerical precision in the forward model (around 1 %). Instrumental noise is comparatively small (under 0.5 % for most superpixels which are largely cloud-free).

Polynomial interpolation of LUTs carries an increased computational overhead but would allow for a continuous forward model and forward model gradient. As well as decreasing the numerical precision error in the forward model, this should remove the tendency of the retrieval to pick state parameters lying just on one side of LUT points. This work has not yet been performed as the surface uncertainty term is large and the artefacts in retrieval histograms are small biases (i.e. the error bars on retrieved values near the LUT points encompass an area on both sides of the LUT points).

7.2.4 Improved cloud/dust flagging

The presence of cloud in a pixel prevents aerosol retrievals from being carried out. Although the retrieval algorithm uses sophisticated cloud flags contained in the level 1 data, in some scenes there

is residual unflagged cloud which will degrade the quality of aerosol retrieval (although in most cases this should be picked up by retrieval quality control). Conversely, in some situations such as strong dust storms aerosol may be wrongly flagged as cloud, leading to such events being missed in the satellite data.

Further investigation into such improved cloud flagging has not been undertaken due to time constraints. It has, however, been proposed for an MPhys project.

7.2.5 Non-spherical aerosol models

The theory of Mie (1908) is widely used in the remote sensing community to calculate aerosol phase functions for use in satellite retrievals. However, this theory requires particles to be spherical, which may not hold true for solid aerosol particles such as dust. Hence there is an error associated with the assumption of sphericity which will propagate through into the retrieval products. Retrievals in dust regions frequently pass quality control and compare well to independent datasets, which gives confidence that this error is not too large (unless all datasets are similarly incorrect).

The theory of non-spherical scattering is highly complex and generally non-spherical scattering codes carry a much higher computational overhead than Mie scattering. Investigation of non-spherical scattering is the subject of a DPhil project within the author's research group; as it began two years after this work, however, results useful to the AATSR retrieval have not yet been obtained.

7.2.6 Improved treatment of aerosol optical properties

A weakness of ORAC is that each aerosol class consists of fixed components with fixed microphysical properties, and changing composition within a class can only be achieved by altering the effective radius. There will be situations in which aerosol is not adequately described by any aerosol class used, or may best be described by a combination of classes (such as in the case of, for example, a polluted ocean). Improvements in this area would involve adoption of recent measurements of optical properties into the retrieval scheme. Possibilities would include regional aerosol models (if, for example, Saharan dust is different from Namibian dust) or retrieval of mixtures of aerosol components rather than pre-defined classes.

7.2.7 Development of more sophisticated vegetation measures

The NDVI results presented in Chapter 6 are just one of several methods for monitoring the state of vegetation from space. Following work by Verstraete and Pinty (1996) suggesting the design of more quantitative measures of vegetation, in recent years more robust indicators have been developed and are being applied to satellite data from several instruments (for example Gobron et al., 2006). Typical quantities derived from such measurements include the fraction of absorbed photosynthetically-active radiation (f_{APAR}), the leaf area index (LAI), which is a measure of the cover of green leaves over a region, and the fractional vegetation cover (FVC) of a pixel. North (2002b) showed that the ATSR instruments are sensitive to these quantities, but this capability has not yet been exploited.

Such alternative indices have not been developed in this work because of time constraints. They require further radiative transfer modelling to determine coefficients to relate the retrieved surface albedo to the biophysical quantities, as opposed to being simple ratios of combinations of bands, and because the focus of this work is on aerosol retrieval they have not yet been fully explored. This, however, would be a suitable task for the future.

7.2.8 Further applications of the data

Chapter 6 provided several examples of applications of AATSR data. It is suggested that, without the time and space constraints of a thesis, these would be excellent places to undertake studies in more detail and add to the growing body of scientific knowledge on the topic, in which the ATSR instruments are currently under-represented. Additional applications not previously mentioned in this work include use of the derived aerosol and surface profiles to estimate aerosol direct radiative forcing, and combination with cloud datasets to investigate cloud-aerosol interactions and aerosol indirect effects.

Bibliography

Abdou, W. A., D. J. Diner, J. V. Martonchik, C. J. Bruegge, R. A. Kahn, B. J. Gaitley, and K. A. Crean, 2005: Comparison of coincident Multiangle Imaging Spectrometer and Moderate Resolution Imaging Spectrometer aerosol optical depths over land and ocean scenes containing Aerosol Robotic Network sites. *J. Geophys. Res.*, **110**, doi:10.1029/2004JD004693.

Adler, R. F., G. j. Huffman, A. Chang, R. Ferraro, P. Xie, J. Janowiak, B. Rudolf, U. Schneider, S. Curtis, D. Bolvin, A. Gruber, J. Susskinf, and P. Arkin, 2003: The version-2 Global Precipitation Climatology Project (GPCP) monthly precipitation analysis (1979-present). *J. Hydrometeor.*, **4**, 1147–1167, doi:10.1175/1525-7541(2003)004<1147:TVGPCP>2.0.CO;2.

Alencar, A.: 2005, A distribuição do desmatamento na Amazônia: vetores, áreas crticas e mapeamento de fronteiras. *Seminário PRODES 2005*, available from http://www.obt.inpe.br/prodes/seminario2005/index.html [Accessed 19th August 2008].

Andreae, M. O. and P. J. Crutzen, 1997: Atmospheric aerosols: Biogeochemical sources and role in atmospheric chemistry. *Science*, **276**, 1052–1058, doi:10.1126/science.276.5315.1052.

Anguelova, M. and F. Webster, 2006: Whitecap coverage from satellite measurements: A first step toward modeling the variability of ocean whitecaps. *J. Geophys. Res.*, **111**, doi:10.1029/2005JC003158.

Aragão, Y., Luiz Eduardo O. C. Malhi, R. M. Roman-Cuesta, S. Saatchi, L. O. Anderson, and Y. E. Shimabukuro, 2007: Spatial patterns and fire response of recent Amazonian droughts. *Geophys. Res. Lett.*, **34**, doi:10.1029/2006GL028946.

Austin, R. W., 1974: The remote sensing of spectral radiance from below the ocean surface. *Optical Aspects of Oceanography, N. G. Jerlov and E. S. Nielsen, Eds., Academic Press*, 317–344.

Barrot, G., A. Mangin, and S. Pinnock, 2006: *GlobColour product user guide*. http://www.globcolour.info [Accessed 9 October 2007].

Bäumer, D., R. Rinke, and B. Vogel, 2008: Weekly periodicities of aerosol optical thickness over central Europe - evidence of an anthropogenic direct aerosol effect. *Atmos. Chem. Phys.*, **8**, 83–90, http://www.atmos-chem-phys.net/8/83/2008/.

Beirle, S., U. Platt, M. Wenig, and T. Wagner, 2003: Weekly cycle of NO2 by GOME measurements: a signature of anthropogenic sources. *Atmos. Chem. Phys.*, **3**, 2225–2232, http://www.atmos-chem-phys.net/3/2225/2003/.

Bellouin, N., A. Jones, J. Haywood, and S. A. Christopher, 2008: Updated estimate of aerosol direct radiative forcing from satellite observations and comparison against the Hadley Centre climate model. *J. Geophys. Res.*, **113**, D10205, doi:10.1029/2007JD009385.

Bingen, C., D. Fussen, and F. Vanhellement, 2004: A global climatology of stratospheric aerosol size distribution parameters derived from SAGE II data over the period 1984–2000: 1. Methodology and climatological observations. *J. Geophys. Res.*, **109**, D06201, doi:10.1029/2003JD003518.

Birks, A., 2004: Improvements to the AATSR IPF relating to land surface temperature. ESA Technical note.

Borbély-Kiss, I., A. Z. Kiss, E. Koltay, G. Szabó, and L. Bozó, 2004: Saharan dust episodes in Hungarian aerosol: elemental signatures and transport trajectories. *J. Aerosol Sci.*, **35**, 1205–1224, doi:10.1016/j.jaerosci.2004.05.001.

Braak, R.: 2006, Ozone Monitoring Instrument aerosol products and validation. *ACCENT Fifth AT-2 Workshop*.

Brown, J., M. Hoke, K. Doherty, G. Anderson, and A. Berk, 2004: *MODTRAN 4 Software*. http://www.vs.afrl.af.mil/ProductLines/IR-Clutter/modtran4.aspx [Accessed 19 July 2006].

Câmara, G., D. de Morisson Valeriano, and J. a. V. Soares, 2006: *Metodologia para o Cálculo de Taxa Anual de Desmatamento na Amazônia Legal*. Project website http://www.obt.inpe.br/prodes/ [Accessed 19th August 2008].

Carboni, E., G. E. Thomas, R. Grainger, C. Poulsen, R. Siddans, and A. Sayer, 2008: Aerosol retrieval using visible and infrared SEVIRI data. In preparation.

Charlson, R. J., S. E. Schwartz, J. M. Hales, R. D. Cess, J. A. Coakley, J. E. Hansen, and D. J. Hofmann, 1992: Climate forcing by anthropogenic aerosols. *Science*, **255**, 423–430, doi:10.1126/science.255.5043.423.

Chen, C., P. Shi, and H. Zhan, 2003: A local algorithm for estimation of yellow substance (gelbstoff) in coastal waters from seawifs data: Pearl River estuary, China. *Int. J. Remote Sensing*, **24**, 1171–1176.

CIESIN, 2005: *Gridded Population of the World Version 3 (GPWv3): Population Density Grids. Palisades, NY: Socioeconomic Data and Applications Center (SEDAC), Columbia University.*. Center for International Earth Science Information Network (CIESIN), Columbia University; and Centro Internacional de Agricultura Tropical (CIAT).
Available at http://sedac.ciesin.columbia.edu/gpw, accessed 13th August 2008.

CIMEL, 2008: CIMEL Sunphotometers webpage. http://www.cimel.fr [Accessed 10 October 2008].

Clark, R. N., G. A. Swayze, R. Wise, K. E. Livo, T. M. Hoefen, R. F. Kokaly, and S. J. Sutley, 2003: *USGS Digital Spectral Library splib05a*. U.S. Geological Survey, Open File Report 03-395, http://pubs.usgs.gov/of/2003/ofr-03-395/ofr-03-395.html [Accessed 19 November 2007].

Cox, C. and W. Munk, 1954a: Measurement of the roughness of the sea surface from photographs of the Sun's glitter. *J. Opt. Soc. Am.*, **44**, 838–850, doi:10.1364/JOSA.44.000838.

— 1954b: Statistics of the sea surface derived from Sun glitter. *J. Mar. Res.*, **13**, 198–227.

d'Almeida, G., P. Koepke, and E. Shettle, 1991: *Atmospheric aerosols: global climatology and radiative characteristics*. A. Deepak Publishing.

DANE, 2008: Colombian national statistics administrative department (DANE) website. http://www.dane.gov.co/ [Accessed 10 August 2008].

Davies, C., 1974: Size distribution of atmospheric particles. *J. Aerosol Sci.*, **5**, 293–300.

de Leeuw, G., 2008: Developments to the AATSR-TNO aerosol retrieval algorithm. Personal communication.

DEC, 2008: Panamanian directorate of statistics and census (DEC) website. http://www.contraloria.gob.pa/dec/ [Accessed 10 August 2008].

Devred, E., S. Sathyendranath, V. Stuart, H. Maass, O. Ulloa, and T. Platt, 2006: A two-component model of phytoplankton absorption in the open ocean: Theory and application. *J. Geophys. Res.*, **111**, doi:10.1029/2005JC002880.

Donlon, C., J. Kennedy, J. Stark, and G. Corlett: 2008, The (A)ATSR as a reference data set for operational SST production. *2008 ESA MERIS/(A)ATSR Workshop SP-666*.

Dubovik, O., B. Holben, T. F. Eck, A. Smirnov, Y. Kaufman, M. Kind, D. Tanré, and I. Slutsker, 2002: Variability and optical properties of key aerosol types observed in worldwide locations. *J. Atm. Sci.*, **59**, 590–608, doi:10.1175/1520-0469(2002)059<0590:VOAAOP>2.0.CO;2.

Dubovik, O., A. Sinyuk, T. Lapyonok, B. Holben, M. Mischenko, P. Yang, T. Eck, H. Volten, O. Muñoz, B. Veihelmann, W. J. van der Zande, J.-F. Leon, M. Sorokin, and I. Slutsker, 2006: The application of spheroid models to account for aerosol particle non-sphericity in remote sensing of desert dust. *J. Geophys. Res.*, **111**, doi:10.1029/2005JD006619.

Dulac, F., P. Chazette, L. Gomes, B. Chatenet, H. Berger, and J. M. Vinicula Dos Santos, 2001: A method for aerosol profiling in the lower troposphere with coupled scatter and meteorological rawindsondes and first data from the tropical Atlantic off Sahara. *J. Aerosol. Sci.*, **32**, 1069–1086, doi:10.1016/S0021-8502(01)00043-X.

ECMWF, 2008: ERA-40 archive web page. http://www.ecmwf.int/research/era/Products/ [Accessed 19 July 2008].

ESA, 2005: *AATSR Frequently Asked Questions*. 1st edition, see also http://envisat.esa.int/instruments/aatsr/ [Accessed 13 March 2006].

— 2007: *AATSR Product Handbook*. Version 2.2, http://envisat.esa.int/handbooks/aatsr/, accessed 27th June 2008.

FAO, 2006: *Global Forest Resources Assessment 2005: Progress towards sustainable forest management*. FAO Forestry Paper - 147, World Scientific.

Fonseca, H., ed., 2004: *Mining: Social and Environmnental Impacts*. World Rainforest Movement.

Formenti, P., M. O. Andreae, L. Lange, G. Roberts, J. Cafmeyer, I. Ratja, W. Maenhaut, B. Holben, P. Artaxo, and J. Lelieveld, 2001: Saharan dust in Brazil and Suriname during the Large-Scale Biosphere-Atmosphere Experiment in Amazonia (LBA) - Cooperative LBA Regional Experiment (CLAIRE) in March 1998. *J. Geophys. Res.*, **106**, 14919–14934.

Forster, P., V. Ramaswamy, P. Artaxo, T. Berntsen, R. Betts, D. W. Fahey, J. Haywood, J. Lean, D. C. Lowe, G. Myhre, J. Nganga, R. Prinn, G. Raga, M. Schulz, and R. Van Dorland, 2007: *Changes in atmospheric constituents and in radiative forcing*. Cambridge University Press, In: Climate Change 2007, the physical science basis. Contribution of Working Group 1 to the fourth assessment report of the Intergovernmental Panel on Climate Change.

Frouin, R., M. Schwindling, and P.-Y. Deschamps, 1996: Spectral reflectance of sea foam in the visible and infra-red: In situ measurements and remote sensing implications. *J. Geophys. Res.*, **101**, 14361–14371, doi:10.1029/96JC00629.

Fujiki, T. and S. Taguchi, 2002: Variability in chlorophyll a specific absorption coefficient in marine phytoplankton as a function of cell size and irradiance. *J. Plank. Res.*, **24**, 859–874, doi:10.1093/plankt/24.9.859.

Ginoux, P., J. M. Prospero, O. Torres, and M. Chin, 2004: Long-term simulation of global dust distribution with the GOCART model: correlation with North Atlantic Oscillation. *Environ. Modell. Soft.*, **19**, 113–128, doi:10.1016/S1364-8152(03)00114-2.

Glantz, M. H., 1996: *Currents of Change : El Niño's Impact on Climate and Society*. Cambridge University Press.

Gobron, N., B. Pinty, M. Taberner, F. Mélin, M. M. Verstraete, and J.-L. Widlowski, 2006: Monitoring the photosynthetic activity of vegetation from remote sensing data. *Adv. Space Res.*, **38**, 2196–2202, doi:10.1016/j.asr.2003.07.079.

Gordon, A., W. Grace, P. Schwerdtfeger, and R. Bryon-Scott, 1998: *Dynamic meteorology: a basic course*. Arnold.

Grainger, R. G., 2008: *GRAPE project website at AOPP*. http://www.atm.ox.ac.uk/project/grape/, accessed 9th June 2008.

Grainger, R. G. and E. J. Highwood: 2003, Changes in stratospheric composition, chemistry, radiation and climate caused by volcanic eruptions. *Volcanic degassing: Geological Society special publication 213*, C. Oppenheimer, D. M. Pyle, and J. Barclay, eds., The Geological Society, 329–347.

Grey, W. M. F., P. R. J. North, and S. O. Los, 2006a: Computationally efficient method for retrieving aerosol optical depth from ATSR-2 and AATSR data. *Appl. Opt.*, **45**, 2786–2795, doi:10.1364/AO.45.002786.

Grey, W. M. F., P. R. J. North, S. O. Los, and R. M. Mitchell, 2006b: Aerosol optical depth and land surface reflectance from multiangle AATSR measurements: global validation and intersensor comparisons. *IEEE Trans. Geosci. Remote Sens.*, **44**, 2184–2197, doi:10.1109/TGRS.2006.872079.

Hale, G. M. and M. R. Querry, 1973: Optical constants of water in the 200-nm to 200-μm wavelength region. *Appl. Opt.*, **12**, 555–563, doi:10.1364/AO.12.000555.

Hamonou, E. and P. Chazette, 1998: Evidence of Saharan mineral aerosols transport to the Mediterranean inside well-defined layers. *J. Aerosol Sci.*, **29**, S1263–S1264.

Hasekamp, O. P. and J. Landgraf, 2007: Retrieval of aerosol properties over land surfaces: capabilities of multi-viewing-angle intensity and polarization measurements. *Appl. Opt.*, **46**, 3332–3344, doi:10.1364/AO.46.003332.

Hayes, D. J. and W. B. Cohen, 2007: Spatial, spectral and temporal patterns of tropical forest cover change as observed with multiple scales of optical satellite data. *Remote Sens. Environ.*, **106**, 1–16, doi:10.1016/j.rse.2006.07.002.

Hayes, D. J., W. B. Cohen, S. A. Sader, and D. E. Irwin, 2008: Estimating proportional change in forest cover as a continuous variable from multi-year MODIS data. *Remote Sens. Environ.*, **112**, 735–749, doi:10.1016/j.rse.2007.06.003.

Hess, M., P. Koepke, and I. Schult, 1998: Optical properties of aerosols and clouds: The software package OPAC. *Bull. Am. Met. Soc.*, **79**, 831–944, doi:10.1175/1520-0477(1998)079.

Hãnel, G., 1976: The properties of atmospheric aerosol particles as functions of the relative humidity at thermodynamic equilibrium with the surrounding moist air. *Adv. Geophys.*, **19**, 73–1988, doi:10.1016/S0065-2687(08)60142-9.

Holben, B. and D. Kimes, 1986: Directional reflectance response in AVHRR red and near-IR bands for three cover types and varying atmospheric conditions. *Remote Sens. Environ.*, **19**, 213–236, doi:10.1016/0034-4257(86)90054-4.

Holben, B. N., T. F. Eck, I. Slutsker, D. Tanré, J. P. Buis, A. Setzer, E. Vermote, J. A. Reagan, Y. J. Kaufman, T. Nakajima, F. Lavenu, I. Jankowiak, and A. Smirnov, 1998: AERONET: A federated instrument network and data archive for aerosol characterization. *Remote Sens. Environ.*, **66**, 1–16, doi:10.1016/S0034-4257(98)00031-5.

Hook, S. J., 1999: *ASTER spectral library*. JPL data reproduced from the ASTER Spectral Library through the courtesy of the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California. Copyright 1999, California Institute of Technology. ALL RIGHTS RESERVED.
http://speclib.jpl.nasa.gov/ [Accessed 19 November 2007].

Hsu, N. C., S.-C. Tsay, M. D. King, and J. R. Herman, 2004: Aerosol properties over bright-reflecting source regions. *IEEE Trans. Geosci. Remote Sens.*, **42**, 557–569, doi:10.1109/TGRS.2004.824067.

— 2006: Deep Blue retrievals of Asian aerosol properties during ACE-Asia. *IEEE Trans. Geosci. Remote Sens.*, **44**, 3180 – 3195, doi:10.1109/TGRS.2006.879540.

Huang, S. and F. Siegert, 2004: ENVISAT multisensor data for fire monitoring and impact assessment. *Int. J. Remote Sensing*, **25**, doi:10.1080/01431160412331269670.

Huete, A. R., K. Didan, Y. E. Shimabukuro, P. Ratana, S. R. Saleska, L. R. Hutyra, W. Yang, R. R. Nemani, and R. Myneni, 2006: Amazon rainforests green-up with sunlight in dry season. *Geophys. Res. Lett.*, **33**, doi:10.1029/2005GL025583.

IBGE, 2008: Brazilian institute of geography and statistics (IBGE) website.
http://www.ibge.gov.br/english/ [Accessed 10 August 2008].

Ichii, K., M. Maruyama, and Y. Yamaguchi, 2003: Multi-temporal analysis of deforestation in rondõnia state in brazil using Landsat MSS, TM, ETM+ and NOAA AVHRR imagery and its relationship to changes in the local hydrological environment. *Int. J. Remote Sensing*, **24**, 4467–4479, doi:10.1080/0143116031000095952.

Ichoku, C., D. A. Chu, S. Mattoo, Y. J. Kaufman, L. A. Remer, D. Tanré, I. Slutsker, and B. N. Holben, 2002: A spatio-temporal approach for global validation and analysis of MODIS aerosol products. *Geophys. Res. Lett.*, **29**, doi:10.1029/2001GL013206.

IMF, 2008: International monetary fund (IMF) World Economic Outlook database webpage.
http://www.imf.org/external/pubs/ft/weo/2008/01/weodata/groups.htm [Accessed 10 August 2008].

INE, 2008: Bolivian national institute of statistics (INE) website. http://www.ine.gov.bo [Accessed 10 August 2008].

INEC, 2008: Ecuadorian institute of statistics and census (INEC) website. http://www.inec.gov.ec [Accessed 10 August 2008].

INEI, 2008: Peruvian national institute of statistics and information (INEI) website. http://www.inei.gob.pe/ [Accessed 10 August 2008].

Inomata, Y., Y. Iwasaka, K. Osada, M. Hayashi, I. Mori, M. Kido, K. Hara, and T. Sakai, 2006: Vertical distributions of particles and sulfur gases (volatile sulfur compounds and SO_2 over East Asia: Comparison with tho aircraft-borne measurements under the Asian continental outflow in spring and winter. *Atmos. Env.*, **40**, 430–444, doi:10.1016/j.atmosenv.2005.09.055.

ITT, 2008: *IDL: Data Visual Solutions webpage, ITT Visual Information Solutions*. http://www.ittvis.com/ProductServices/IDL.aspx [Accessed 01 October 2008].

Jin, Y., C. B. Schaaf, C. E. Woodcock, F. Gao, X. Li, A. H. Strahler, W. Lucht, and S. Liang, 2003a: Consistency of MODIS surface bidirectional reflectance distribution function and albedo retrievals: 1. Algorithm performance. *J. Geophys. Res.*, **108**, doi:10.1029/2002JD002803.

— 2003b: Consistency of MODIS surface bidirectional reflectance distribution function and albedo retrievals: 1. Validation. *J. Geophys. Res.*, **108**, doi:10.1029/2002JD002804.

Junge, C. E., 1955: The size distribution and aging of natural aerosols as determined from electrical and optical data on the atmosphere. *Am. Meteor. Soc. J. Meteor.*, **12**, 13–25, doi:10.1175/1520-0469(1955)012<0013:TSDAAO>2.0.CO;2.

Justus, C. G. and M. V. Paris, 1985: A model for solar spectral irradiance and radiance at the bottom and top of a cloudless atmosphere. *J. Appl. Meteorol.*, **24**, 193–205, doi:10.1175/1520-0450(1985)024<0193:AMFSSI>2.0.CO;2.

Kaufman, Y. J., I. Koren, L. A. Remer, D. Tanré, P. Ginous, and S. Fan, 2005: Dust transport and deposition observed from the Terra-Moderate Resolution Imaging Spectroradiometer MODIS spacecraft over the Atlantic Ocean. *J. Geophys. Res.*, **110**, D10S12, doi:10.1088/1748-9326/1/1/014005.

Kaufman, Y. J., D. Tanré, J. F. Léon, and J. Pelon, 2003: Retrievals of profiles of fine and coarse aerosols using lidar and radiometric space measurements. *IEEE Trans. Geosci. Remote Sens.*, **41**, 1743–1754, doi:10.1109/TGRS.2003.814138.

Keller, J., S. Bojinski, and A. S. Prevot, 2007: Simultaneous retrieval of aerosol and surface properties using data of the Multi-angle Imagine SpectroRadiometer (MISR). *Remote Sens. Environ.*, **107**, 120–137, doi:10.1016/j.rse.2006.07.020.

Kinne, S., 2007: List of AERONET stations representative of their region on satellite spatial scales. Personal communication.

Kinne, S., M. Schulz, C. Textor, Y. Balkanski, S. Bauer, T. Berntsen, T. Berglen, O. Boucher, M. Chin, F. Collins, W. Dentener, T. Diehl, R. Eater, J. Feichter, D. Filmore, S. Ghan, P. Ginoux, S. Gong, A. Grini, J. E. Hendricks, M. Herzog, L. Horowitz, I. S. A. Isaksen, T. Iversen, A. Kirkavåg, S. Kloster, D. Koch, J. E. Kristjansson, M. Krol, A. Lauer, J. F. Lamarque, G. Lesins, X. Liu, V. Lohmann, U. Montanaro, G. Myhre, J. E. Penner, G. Pitari, S. Reddy, O. Seland, P. Stier, and T. Takemura, 2006: An AeroCom initial assessment - optical properties in aerosol component modules of global models. *Atmos. Chem. Phys.*, **6**, 1815–1834, http://www.atmos-chem-phys.net/6/1815/2006/.

Kirby, K. A., W. F. Laurance, A. K. Albernaz, G. Schroth, P. M. Fearnside, S. Bergen, E. M. Venticinque, and C. da Costa, 2006: The future of deforestation in the Brazilian Amazon. *Futures*, **38**, 432–453, doi:10.1016/j.futures.2005.07.011.

Kloster, S., J. Feichter, E. Maler-Reimer, K. D. Six, P. Stier, and P. Wetzel, 2005: DMS cycle in the marine ocean-atmosphere system - a global model study. *Biogeosci. Disc.*, **2**, 1067–1126, http://www.biogeosciences-discuss.net/2/1067/2005/.

Koepke, P., 1984: Effective reflectance of oceanic whitecaps. *Appl. Opt.*, **23**, 1816–1824, doi:10.1364/AO.23.001816.

Kokhanovsky, A., F.-M. Breon, A. Cacciari, E. Carboni, D. Diner, W. Di Nicolantonio, R. Grainger, W. Grey, R. Hőller, K.-H. Lee, Z. Li, P. North, A. Sayer, G. Thomas, and W. von Hoyningen-Huene, 2007: Aerosol remote sensing over land: a comparison of satellite retrievals using different algorithms and instruments. *Atmos. Res.*, **85**, 372–394, doi:10.1016/j.atmosres.2007.02.008.

Koren, I., Y. Kaufman, R. Washington, M. C. Todd, Y. Rudich, J. V. Martins, and D. Rosenfeld, 2006: The Bodélé depression: a single spot in the Sahara that provides most of the mineral dust to the Amazon forest. *Environ. Res. Lett.*, **1**, 014005.

Koren, I., L. A. Remer, Y. J. Kaufman, Y. Rudich, and J. V. Martins, 2007: On the twilight zone between clouds and aerosols. *Geophys. Res. Lett.*, **34**, doi:10.1029/2007GL029253.

Levenberg, K., 1944: A method for the solution of certain non-linear problems in least-squares. *Q. of Appl. Math.*, **2**, 164–168.

Lewinsohn, T. M. and P. Inácia Prado, 2006: How many species are there in Brazil? *Conservation Biology*, **3**, 619–624, doi:10.1111/j.1523-1739.2005.00680.x.

Liu, Z., A. Omar, M. Vaughan, J. Hair, C. Kittaka, Y. Hu, K. Powell, C. Trepte, D. Winker, C. Hosteler, R. Ferrare, and R. Pierce, 2008: CALIPSO lidar observations of the optical properties of saharan dust: A case study of long-range transport. *J. Geophys. Res.*, **113**, D07207, doi:10.1029/2007JD008878.

Lu, G., J. R. Brook, M. R. Alfarra, K. Anlauf, W. R. Leaitch, S. Sharma, D. Wang, D. R. Worsnop, and L. Phinney, 2006: Identification and characterization of inland ship plumes over Vancouver, BC. *Atmos. Env.*, **40**, 2767–2782, doi:10.1016/j.atmosenv.2005.12.054.

Lucas, R. M., M. Honzák, P. J. Curran, G. M. Foody, R. Milne, T. Brown, and S. Amaral, 2000: Mapping the regional extent of tropical forest regeneration stages in the brazilian legal amazon using noaa avhrr data. *Int. J. Remote Sensing*, **21**, 2855–2881, doi:10.1080/01431160050121285.

Lucht, W. and P. Lewis, 2000: Theoretical noise sensitivity of BRDF and albedo retrieved from the EOS-MODIS and MISR sensors with respect to angular sampling. *Int. J. Remote Sensing*, **21**, 81–98, doi:10.1080/014311600211000.

Lucht, W., C. B. Schaaf, and A. H. Strahler, 2000: An algorithm for the retrieval of albedo from space using semiempirical BRDF models. *IEEE. Trans. Geosci. Remote Sens.*, **38**, 977–998, doi:10.1109/36.841980.

Marquardt, D. R., 1963: An algorithm for the least-squares estimation of nonlinear parameters. *SIAM J. Appl. Math.*, **11**, 431–441, doi:10.2307/2098941.

Marsh, S., S. Dean, R. Grainger, A. Quijano, and G. Thomas, 2004: *An Optimal Estimation Aerosol Retrieval Scheme For ATSR-2*. Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford.

Marticorena, B. and G. Bergametti, 1995: Modeling the atmospheric dust cycle: 1. Design of a soil-derived dust emission scheme. *J. Geophys. Res.*, **100**, 16,415–16,430.

McConnell, C. L., E. J. Highwood, H. Coe, P. Formenti, B. Anderson, S. Osborne, S. Nava, K. Desbouefs, G. Chen, and M. A. J. Harrison, 2008: Seasonal variations of the physical and optical characteristics of saharan dust: Results from the Duse Outflow and Deposition to the Ocean (DODO) experiment. *J. Geophys. Res.*, **113**, D14S05, doi:10.1029/2007JD009606.

Melillo, J. M., A. D. McGuire, D. W. Kicklighter, B. I. Moore, C. J. Vörösmarty, and A. L. Schloss, 1993: Global climate change and terrestrial net primary production. *Nature*, **363**, 234–240, doi:10.1038/363234a0.

Merchant, C. J., A. R. Harris, M. J. Murray, and A. M. Závody, 1999: Toward the elimination of bias in satellite retrieval of sea surface temperature 1. theory, modelling and interalgorithm comparison. *J. Geophys. Res.*, **104**, 23,565 – 23,578, doi:10.1029/1999JC900105.

Merchant, C. J. and P. Le Borgne, 2004: Retrieval of sea surface temperature from space, based on modeling of infrared radiative transfer: Capabilities and limitations. *J. Atmos. Ocean. Tech.*, **21**, 1734 – 1746, doi:10.1175/JTECH1667.1.

Mie, G., 1908: Beiträge zur optik trüber medien, speziell kolloidaler metallösungen. *Ann. Phys.*, **330**, 377 445, doi:10.1002/andp.19083300302.

Mischenko, M. I. and I. V. Geogdzhayev, 2007: Satellite remote sensing reveals regional tropospheric aerosol trends. *Opt. Express*, **15**, 7423 7438.

Monahan, E. C. and I. O. Muircheartaigh, 1980: Optimal power-law description of oceanic whitecap coverage dependence on wind speed. *J. Phys. Oceanogr.*, **10**, 2094–2099, doi:10.1175/1520-0485(1980)010¡2094:OPLDOO¿2.0.CO;2.

Morel, A., 1974: Optical properties of pure water and pure sea water. *Optical Aspects of Oceanography, N. G. Jerlov and E. S. Nielsen, Eds., Academic Press*, 1–24.

— 1988: Optical modeling of the upper ocean in relation to its biogenous matter content (Case I waters). *J. Geophys. Res.*, **93**, 10749–10768.

Morel, A. and B. Gentili, 1991: Diffuse reflectance of oceanic waters: its dependence on Sun angle as influenced by the molecular scattering contribution. *Appl. Opt.*, **30**, 4427–4438, doi:10.1364/AO.30.004427.

Morel, A. and L. Prieur, 1977: Analysis of variations in ocean color. *Limnol. Oceanogr.*, **22**, 709–722.

Myneni, R. B., F. G. Hall, P. J. Sellers, and A. L. Marshak, 1995: The interpretation of spectral vegetation indexes. *IEEE Trans. Geosci. Remote Sens.*, **33**, 481–486, doi:10.1109/36.377948.

Nagai, S., K. Ichii, and H. Morimoto, 2007: Interannual variations in vegetation activities and climate variability caused by ENSO in tropical rainforests. *Int. J. Remote Sensing*, **28**, 1285–1297, doi:10.1080/01431160600904972.

NASA, 2008: AERONET: The AERosol RObotic NETwork home page. http://aeronet.gsfc.nasa.gov/ [Accessed 27 June 2008].

Neef, T., R. M. Lucas, J. a. R. dos Santos, E. S. Brondizio, and C. C. Freitas, 2006: Area and age of secondary forests in Brazilian Amazonia 1978-2002: An empirical estimate. *Ecosystems*, **9**, 609–623, doi:10.1007/s10021-006-0001-9.

Nicolas, J.-M., P.-Y. Deschamps, and R. Frouin, 2001: Spectral reflectance of oceanic whitecaps in the visible and infra-red: Aircraft measurements over open ocean. *Geophys. Res. Lett.*, **28**, 4445–4448, doi:10.1029/2001GL013556.

NOAA, 2008: National weather service climate prediction center webpage on previous ENSO events. Available at http://www.cpc.ncep.noaa.gov/products/analysisensoyears.shtml [Accessed 21st August 2008].

North, P. and S. Bevan, 2008: Development of a dark target approach to (A)ATSR retrievals over the ocean. Personal communication.

North, P., C. Brockmann, J. Fischer, L. Gomez-Chova, W. Grey, J. Moreno, R. Preusker, and P. Regner: 2008, MERIS/AATSR synergy algorithms for cloud screening, aerosol retrieval and atmospheric correction. *2008 ESA MERIS/(A)ATSR Workshop SP-666*.

North, P. R., 2002a: Estimation of aerosol opacity and land surface bidirectional reflectance from ATSR-2 dual-angle imagery: Operational method and validation. *J. Geophys. Res.*, **107**, doi:10.1029/2000JD000207.

North, P. R., S. A. Briggs, S. E. Plummer, and J. J. Settle, 1999: Retrieval of land surface bidirectional reflectance and aerosol opacity from ATSR-2 multiangle imagery. *IEEE Trans. Geosci. Remote Sens.*, **37**, 526–537, doi:10.1109/36.739106.

North, P. R. J., 2002b: Estimation of *f*APAR, LAI, and vegetation fractional cover from ATSR-2 imagery. *Remote Sens. Environ.*, **80**, 114–121, doi:10.1016/S0034-4257(01)00292-9.

Olesen, F. S., 2005: *DISORT Algorithm*.
http://imk-msa.fzk.de/msa-public/Software-tools/Modtran/science/disort.htm [Accessed 19 July 2006].

Oliveira, P. J. C., G. P. Asner, D. E. Knapp, A. Almeyda, R. Galván-Gildemeister, S. Keene, R. F. Raybin, and R. C. Smith, 2007: Land-use allocation protects the Peruvian Amazon. *Science*, **317**, 1233–1236, doi:10.1126/science.1146324.

Penner, J. E., M. Andreae, H. Annegam, L. Barrie, J. Feitcher, D. Hegg, A. Jayaraman, R. Leaitch, D. Murphy, J. Nganga, and G. Pitari, 2001: *Aerosols, their direct and indirect effects*. Cambridge University Press, In: Climate Change 2001: the scientific basis. Contribution of Working Group 1 to the third assessment report of the Intergovernmental Panel on Climate Change.

Penner, J. E., J. Quaas, T. Storevlmo, T. Takemura, O. Boucher, H. Guo, A. Kirkevåg, J. E. Kristjánsson, and O. Seland, 2006: Model intercomparison of indirect aerosol effects. *Atmos. Chem. Phys.*, **6**, 3391–3405, http://www.atmos-chem-phys.net/6/3391/2006/.

Poulsen, C. A., R. Siddans, B. Lawrence, E. Campmany, R. G. Grainger, G. E. Thomas, A. M. Sayer, S. Dean, and P. D. Watts, 2008: Cloud retrievals from satellite data using optimal estimation. In preparation.

Pueschel, R. F., 1996: Stratospheric aerosols: Formation, properties, effects. *J. Aerosol Sci*, **27**, 383–402, doi:10.1016/0021-8502(95)00557-9.

Quan, X. and E. S. Fry, 1995: Empirical equation for the index of refraction of seawater. *Appl. Opt.*, **34**, 3477–3480, doi:10.1364/AO.34.003477.

Rahn, K. A. and D. H. Lowenthal, 1984: Elemental tracers of distant regional pollution aerosols. *Science*, **223**, 132–139, doi:10.1126/science.223.4632.132.

Ramachandran, S., 2005: $PM_{2.5}$ mass concentrations in comparison with aerosol optical depths over the Arabian Sea and Indian Ocean during winter monsoon. *Atmos. Env.*, **39**, 1879–1890, doi:10.1016/j.atmosenv.2004.12.003.

Remer, L. A., R. G. Kleidman, R. C. Levy, Y. J. Kaufman, D. Tanré, S. Mattoo, J. V. Martins, C. Ichoku, I. Koren, H. Yu, and B. N. Holben, 2008: Global aerosol climatology from the MODIS satellite sensors. *J. Geophys. Res.*, **113**, D14S07, doi:10.1029/2007JD009661.

Robles González, C., M. Schaap, G. de Leeuw, P. J. H. Builtjes, and M. van Loon, 2003: Spatial variation of aerosol properties over Europe derived from satellite observations and comparison with model calculations. *Atmos. Chem. Phys.*, **3**, 521–533, http://www.atmos-chem-phys.net/3/521/2003/.

Rodgers, C. D., 2000: *Inverse methods for atmospheric sounding: Theory and Practice*. Series on Atmospheric, Oceanic and Planetary Physics–Vol. 2, World Scientific.

Roesler, C. S., M. J. Perry, and K. L. Carder, 1989: Modeling in situ phytoplankton absorption from total absorption spectra in productive inland marine waters. *Limnol. Oceanogr.*, **34**, 1510–1523.

Ronchail, J., G. Cochonneau, M. Molinier, J.-L. Guyot, A. G. de Miranda Chaves, V. Guimarães, and E. de Oliveira, 2002: Interannual rainfall variability in the Amazon basin and sea-surface temperatures in the equatorial Pacific and the tropical Atlantic oceans. *Int. J. Climatol*, **22**, 1663–1686, doi:10.1002/joc.815.

Saleska, S. R., K. Didan, A. R. Huete, and H. R. da Rocha, 2007: Amazon forests green-up during 2005 drought. *Science*, **318**, 612, doi:10.1126/science.1146663.

Sathyendranath, S., G. Cota, V. Stuart, H. Maass, and T. Platt, 2001: Remote sensing of phytoplankton pigments: A comparison of empirical and theoretical approaches. *Int. J. Remote Sensing*, **22**, 249–273, doi:10.1080/014311601449925.

Sayer, A. M., C. A. Poulsen, R. Siddans, B. Lawrence, E. Campmany, R. G. Grainger, G. E. Thomas, S. Dean, and P. D. Watts, 2008: Validation of GRAPE cloud data. In preparation.

Schneider, S. H., ed., 1996a: *Encyclopaedia of Weather and Climate, volume 1*. Oxford University Press.

— 1996b: *Encyclopaedia of Weather and Climate, volume 2*. Oxford University Press.

Schulz, M., C. Textor, S. Kinne, Y. Balkanski, S. Bauer, T. Berntsen, T. Berglen, O. Boucher, F. Dentener, S. Guibert, I. S. A. Isaksen, T. Iversen, D. Koch, A. Kirkevåg, X. Liu, V. Montanaro, G. Myhre, J. E. Penner, G. Pitari, S. Reddy, O. Seland, P. Stier, and T. Takemura, 2006: Radiative forcing by aerosols as derived from the AeroCom present-day and pre-industrial simulations. *Atmos. Chem. Phys.*, **6**, 5225–5246, http://www.atmos-chem-phys.net/6/5225/2006/.

Schwartz, S. E., 1996: The Whitehouse Effect - Shortwave Radiative Forcing Of Climate By Anthropogenic Aerosols: An Overview. *J. Aerosol. Sci.*, **27**, 359–382, doi:10.1016/0021-8502(95)00533-1.

Smirnov, A., B. N. Holben, T. F. Eck, I. Slutsker, B. Chatenet, and R. T. Pinker, 2002: Diurnal variability of aerosol optical depth observed at AERONET (Aerosol Robotic Network) sites. *Geophys. Res. Lett.*, **29**, doi:10.1029/2002GL016305.

Smith, D., 2007: AATSR visible channel calibration and noise. Personal communication.

Smith, D.: 2008, AATSR instrument performance and prospects for continued operations. *2008 ESA MERIS/(A)ATSR Workshop SP-666*.

Smith, D., J. Delderfield, D. Drummond, T. Edwards, C. Mutlow, P. Read, and G. Toplis, 2001: Calibration of the AATSR instrument. *Adv. Space Res.*, **28**, 31–39, doi:10.1016/S0273-1177(01)00273-3.

Smith, D., C. Poulsen, and B. Latter: 2008, Calibration status of the AATSR reflectance channels. *2008 ESA MERIS/(A)ATSR Workshop SP-666*.

Smith, D. L.: 2004, Impact on the performance of the ATSR instruments due to condensation. *ESTEC contamination workshop*.

Smith, D. L., 2005: *Report for AATSR VisMon contract*. 1st edition, prepared by the Rutherford Appleton Laboratory.

Smith, D. L., C. T. Mutlow, and C. R. N. Rao, 2002: Calibration monitoring of the visible and near-infrared channels of the Along-Track Scanning Radiometer-2 by use of stable terrestrial sites. *Appl. Opt.*, **41**, 515–523, doi:10.1364/AO.41.000515.

Sória, G. and J. A. Sobrino: 2005, AATSR derived land surface temperature from heterogeneous areas. *2005 ESA MERIS/(A)ATSR Workshop*.

Souza, C. J., A. Brandão, R. Salomão, and A. Costa: 2005, Análise do desmatamento da Amazônia: incertezas, áreas críticas, vetores e tendências. *Seminário PRODES 2005*, available from http://www.obt.inpe.br/prodes/seminario2005/index.html [Accessed 19th August 2008].

Stamnes, K., S.-C. Tsay, W. Wiscombe, and K. Jayaweera, 1988: Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting later media. *Appl. Opt.*, **27**, 2502–2509, doi:10.1364/AO.27.002502.

Stier, P., J. H. Seinfeld, S. Kinne, and O. Boucher, 2007: Aerosol absorption and radiative forcing. *Atmos. Chem. Phys.*, **7**, 5237–5261, http://www.atmos-chem-phys.net/7/5237/2007/.

Stier, P., J. H. Seinfeld, S. Kinne, J. Feitcher, and O. Boucher, 2006: Impact of nonabsorbing anthropogenic aerosols on clear-sky atmospheric absorption. *J. Geophys. Res.*, **111**, D18201, doi:10.1029/2006JD007147.

Targino, A. C. and K. J. Noone, 2006: Airborne in situ characterization of dry urban aerosol optical properties around complex topography. *Atmos. Res.*, **79**, 148–174, doi:10.1016/j.atmosres.2005.05.005.

Thomas, G., 2007: *A Surface BRDF Forward Model for ORAC, AOPP Technical Memorandum 2007.3*. Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford.

Thomas, G., S. Dean, E. Carboni, R. Grainger, C. Poulsen, R. Siddans, and B. Kerridge, 2006: *GlobAEROSOL Data User Element–Aerosol Extraction Algorithm Definition, ATSR-2/AATSR Algorithm Theoretical Basis Document*. Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, 2nd edition.

Thomas, G. E., C. A. Poulsen, R. L. Curier, G. de Leeuw, S. H. Marsh, E. Carboni, R. G. Grainger, and R. Siddans, 2007: Comparison of AATSR and SEVIRI aerosol retrievals over the northern Adriatic. *Q. J. R. Meteorol. Soc.*, **133**, 85–95, doi:10.1002/qj.126.

Thomas, G. E., C. A. Poulsen, A. M. Sayer, S. H. Marsh, S. M. Dean, E. Carboni, R. Siddans, and R. G. Grainger, 2008a: An optimal estimation aerosol algorithm using SEVIRI and AATSR. *Atmos. Meas. Tech.*, submitted.

— 2008b: *Oxford-RAL Aerosol and Cloud (ORAC): Aerosol retrievals from satellite radiometers*. Springer (Berlin), in: Aerosol Remote Sensing Over Land, A. A. Kokhanvosky and G. de Leeuw (eds.).

Thurston, G. D., N. M. Laird, K. A. Rahn, and D. H. Lowenthal, 1985: Tracing aerosol pollution. *Science*, **227**, 1406, 1408, 1412, letters discussing Rahn and Lowenthal's Elemental Tracers paper (Science, 1984).

Tucker, C. J., 1979: Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sens. Environ.*, **8**, 127–150, doi:10.1016/0034-4257(79)90013-0.

Tucker, C. J., B. N. Holben, and T. E. Goff, 1984: Intensive forest clearing in rondonia, brazil, as detected by satellite remote sensing. *Remote Sens. Environ.*, **15**, 255–261, doi:10.1016/0034-4257(84)90035-X.

Twomey, S., 1977: *Atmosperic Aerosols*. Elsevier.

Veefkind, J. P., 1999: *Aerosol satellite remote sensing*. Ph.D. thesis, University of Utrecht.

Veefkind, J. P. and G. de Leeuw, 1997: A new aerosol retrieval algorithm applied to ATSR-2 data. *J. Aerosol Sci.*, **28**, S693–S694, doi:10.1016/S0021-8502(97)85345-1.

— 1998: A new algorithm to determine the spectral aerosol optical depth from satellite radiometer measurements. *J. Aerosol Sci.*, **29**, 1237–1248, doi:10.1016/S0021-8502(98)00032-9.

Veefkind, J. P., G. de Leeuw, and P. A. Durkee, 1998: Retrieval of aerosol optical depth over land using two-angle view satellite radiometry during TARFOX. *Geophys. Res. Lett.*, **25**, 3135–3138, doi:10.1029/98GL02264.

Veefkind, J. P., G. de Leeuw, P. A. Durkee, P. B. Russell, P. V. Hobbs, and J. M. Livingston, 1999: Aerosol optical depth retrieval using ATSR-2 and AVHRR data during TARFOX. *J. Geophys. Res.*, **104**, 2253–2260, doi:10.1029/98JD02816.

Veefkind, J. P., G. de Leeuw, P. Stammes, and R. P. Koelemeijer, 2000: Regional distribution of aerosol over land, derived from ATSR-2 and GOME. *Remote Sens. Environ.*, **74**, 377–386, doi:10.1016/S0034-4257(00)00106-1.

Vermote, E. F., D. Tanré, J. L. Deuzé, M. Herman, and J.-J. Morcrette, 1997: Second Simulation of the Satellite Signal in the Solar Spectrum, 6S: An Overview. *IEEE Trans. Geosci. Remote Sens.*, **35**, 675–686, doi:10.1109/36.581987.

Verstraete, M. M. and B. Pinty, 1996: Designing optimal spectral indices for remote sensing applications. *IEEE Trans. Geosci. Remote Sens.*, **34**, 1254–1265, doi:10.1109/36.536541.

Wang, J. and S. T. Martin, 2007: Satellite characterization of urban aerosols: Importance of including hygroscopicity and mixing state in the retrieval algorithms. *J. Geophys. Res.*, **112**, D17203, doi:10.1029/2006JD008078.

Wanner, W., X. Li, and A. H. Strahler, 1995: On the derivation of kernels for kernel-drived models of bidirectional reflectance. *J. Geophys. Res.*, **100**, 21077–21089, doi:10.1029/95JD02371.

Wanner, W., A. Strahler, B. Hu, P. Lewis, J. Muller, X. Li, Schaaf, C. Barker, and M. Barnsley, 1997: Global retrieval of bidirectional reflectance and albedo over land from EOS MODIS and MISR data: Theory and algorithm. *J. Geophys. Res.*, **102**, 17143–17161, doi:10.1029/96JD03295.

Watts, P., C. Mutlow, A. Baran, and A. Zavody, 1998: Study on cloud properties derived from meteosat second generation observations. *EUMETSAT ITT*.

Wetzel, M. A. and L. L. Stowe, 1999: Satellite-observed patterns in stratus microphysics, aerosol optical thickness, and shortwave radiative forcing. *J. Geophys. Res.*, **104**, 31287–31300.

Wynne, R. H., K. A. Joseph, J. O. Browder, and P. M. Summers, 2007: Comparing farmer-based and satellite-derived deforestation estimates in the amazon basin using a hybrid classifier. *Int. J. Remote Sensing*, **28**, 1299–1315, doi:10.1080/01431160600928609.

Xia, X., T. F. Eck, B. N. Holben, G. Phillippe, and H. Chen, 2008: Analysis of the weekly cycle of aerosol optical depth using AERONET and MODIS data. *J. Geophys. Res.*, **113**, doi:10.1029/2007JD009604.

Yoon, S.-C., S.-W. Kim, J. Kim, B.-J. Sohn, A. Jefferson, S.-J. Choi, D.-H. Cha, D.-K. Lee, T. L. Anderson, S. J. Doherty, and R. J. Weber, 2006: Enhanced water vapor in Asian dust layer: Entrainment processes and implication for aerosol optical properties. *Atmos. Env.*, **40**, 2409–2421, doi:10.1016/j.atmosenv.2005.12.018.

Zavody, A. M., M. R. Gorman, D. J. Lee, D. Eccles, C. T. Mutlow, and D. T. Llewellyn-Jones, 1994: The ATSR data processing scheme developed for the EODC. *Int. J. Remote Sensing*, **15**, 827–843, doi:10.1080/01431169408954119.

Závody, A. M., C. T. Mutlow, and D. T. Llewellyn-Jones, 1995: A radiative transfer model for sea surface temperature retrieval from the along-track scanning radiometer. *J. Geophys. Res.*, **100**, 937 952, doi:10.1029/94JC02170.

Zhang, K. M., E. M. Knipping, A. S. Wexler, P. V. Bhave, and G. S. Tonnesen, 2005: Size distribution of sea salt emissions as a function of relative humidity. *Atmos. Env.*, **39**, 16,415–16,430, doi:10.1016/j.atmosenv.2005.02.032.

Zhao, T. X.-P., I. Laszlo, W. Guo, A. Heidinger, C. Cao, A. Jelenak, and J. Sullivan, 2008: Study of long-term trend in aerosol optical thickness observed from operational AVHRR satellite instrument. *J. Geophys. Res.*, **113**, doi:10.1029/2007JD009061.

Appendix A

Notation used in this work

This appendix provides a summary of the assorted notation adopted throughout this work. Table A.1 describes the notation pertaining to retrieval theory, while Table A.2 presents that used in the retrieval forward model. Notation in Table A.1 generally follows Rodgers (2000), and that in Table A.2 is developed from earlier work on the ORAC scheme as presented in Marsh et al. (2004) and Thomas et al. (2008a).

Symbol	Explanation
α	Retrieval step size parameter.
$\mathbf{A}$	Averaging kernel matrix.
$\mathbf{b}$	Forward model parameters not present in the state vector.
$\mathbf{b}_t$	Unknown true values of $\mathbf{b}$.
δJ_n	Change in J between retrieval iterations resulting from the change $\delta \mathbf{x}_n$.
$\delta \mathbf{x}_n$	Step taken in parameter space during retrieval iteration.
$\mathbf{D}_y$	Inversion operator at solution; also often known as the Gain Matrix $\mathbf{G}$.
$\mathbf{F}$	Forward function mapping $\mathbf{x}$ to $\mathbf{y}$.
$\mathbf{f}$	Forward model, approximation of $\mathbf{F}$.
J	Total retrieval cost.
J'	First derivative of retrieval cost with respect to $\mathbf{x}$.
J''	Second derivative retrieval cost with respect to $\mathbf{x}$.
J_x	*A priori* component of retrieval cost.
J_y	Measurement component of retrieval cost.
K_x	Weighting function matrix.
MQ_{start}	Marquardt 'start' parameter for scaling α.
MQ_{step}	Marquardt 'step' parameter for scaling α.
n_b	Number of state variables which are not significantly constrained by their *a priori*.
n_x	Number of elements of $\mathbf{x}$.
n_y	Number of elements of $\mathbf{y}$.
$\mathbf{S}_b$	Covariance matrix of model parameter error.
$\mathbf{S}_M$	Covariance matrix of measurement mapping error.
$\mathbf{S}_N$	Covariance matrix of null space error.
$\mathbf{S}_S$	Covariance matrix of model parameter mapping error.
$\hat{\mathbf{S}}_T$	Covariance matrix of total error on the retrieved state.
$\mathbf{S}_x$	Covariance matrix of *a priori* error.
$\mathbf{S}_{\hat{x}}$	Covariance matrix of *a priori* residuals $\hat{\mathbf{x}} - \mathbf{x}_a$.
$\mathbf{S}_\mathbf{y}$	Covariance matrix of measurement and forward model error.
$\mathbf{S}_{\delta\hat{\mathbf{y}}}$	Covariance matrix of measurement residuals $\hat{\mathbf{y}} - \mathbf{y}$.
$\mathbf{x}$	Retrieval state vector.
$\hat{\mathbf{x}}$	Retrieved state vector at retrieval solution.
$\mathbf{x}_a$	*A priori* value of the state vector.
$\mathbf{x}_0$	First guess at the retrieval state vector (set to $\mathbf{x}_a$).
$\mathbf{y}$	Retrieval measurement vector.
$\hat{\mathbf{y}}$	Measurement vector calculated from the retrieved state at the retrieval solution.
$\mathbf{y}_m$	Retrieval measurement vector, alternative notation.

Table A.1: Notation used in this work relating to retrieval theory.

Symbol	Explanation
$\downarrow$	Indicator of downward transmittance or reflectance in the forward model.
$\uparrow$	Indicator of upward transmittance or reflectance in the forward model.
$\text{Å}_{\lambda_1,\lambda_2}$	Ångstrøm coefficient between wavelengths λ_1 and λ_2.
β_a	Extinction coefficient due to absorption.
β_e	Total extinction coefficient.
β_s	Extinction coefficient due to scattering.
λ	Wavelength.
$L_\lambda(\tau_\lambda, \mu, \phi)$	Intensity of radiance along direction (μ, ϕ) at optical depth τ_λ.
μ	Cosine of a zenith angle.
$\tilde{n}(\lambda)$	Refractive index at wavelength λ with real and imaginary parts n_r and κ.
$N(h)$	Total aerosol number density as a function of height h and scale height Z.
$n(r)$	Aerosol number size distribution, dependent on radius r.
ω	Atmospheric layer single scattering albedo.
ω_0	Aerosol single scattering albedo.
ϕ	Relative (solar - satellite) azimuth angle.
$P(\theta)$	Atmospheric scattering phase function.
p_0	Ground-level pressure.
p_l	Atmospheric layer lower pressure level.
p_n	Standard atmospheric pressure, 1013.25 mbar.
p_u	Atmospheric layer upper pressure level.
Q_e	Mie extinction efficiency factor.
R_{BD}	Bidirectional reflectance of the aerosol layer.
r_e	Aerosol effective radius, measured in microns.
R_{FD}	Reflectance of the aerosol layer to diffuse radiance.
r_i	Modal radius of a lognormal distribution.
$R_{\text{SBD},\lambda}$	Surface bidirectional reflectance albedo at wavelength λ.
$R_{\text{SLB},\lambda}$	Surface black-sky (hemispherical) albedo at wavelength λ.
$R_{\text{SLW},\lambda}$	Surface white-sky (Lambertian or bihemispherical) albedo at wavelength λ.
σ	Spread of a lognormal distribution.
τ_a	Aerosol contribution to layer optical depth.
τ_g	Gas absorption contribution to layer optical depth.
τ_l	Total atmospheric layer optical depth.
τ_λ	Aerosol optical depth at wavelength λ.
τ_R	Rayleigh contribution to layer optical depth.
τ_{RT}	Rayleigh optical depth for a column of atmosphere.
θ_s	Solar zenith angle.
θ_v	Satellite viewing zenith angle.
T_{BD}	Diffuse transmission of the incident beam.
T_{D}	Transmission of diffuse incident radiance.
T_{DB}	Direct transmission of the beam.
x	Mie size parameter.

Table A.2: Notation used in this work relating to the retrieval forward model.

Appendix B

AERONET-AATSR coincidences during September 2004

Table B.1 shows the names, locations, dates and times of cloud-free AATSR overpasses with coincident AERONET measurements during the month of September 2004. This allows a total of 58 matchups by which the retrieved aerosol optical depth can be validated against ground-based data, as described in Chapter 5. AERONET is thanked for provision of the data.

Station name	Latitude (degrees)	Longitude (degrees)	Day(s) and time(s) of coincidence(s)
Agoufou	15.3450	-1.47900	10 (10:11), 13 (10:22), 23 (10:06), 26 (10:11), 29 (10:16)
Alta Floresta	-9.91700	-56.0170	29 (13:43)
Avignon	43.9330	4.87800	13 (10:14)
Banizoumbou	13.5410	2.66500	17 (09:52), 23 (10:12)
Belsk	51.8370	20.7920	11 (09:29), 18 (09:09)
Billerica	42.5170	-71.2500	3 (15:29), 26 (14:57)
Bondville	40.0530	-88.3720	1 (16:30), 8 (16:13), 17 (16:32), 30 (16:19)
Bratts Lake	50.2800	-104.700	27 (17:58)
Brookhaven	40.8700	-72.8890	10 (15:12)
Dakar	14.3940	-16.9590	11 (11:24)
Fresno	36.7820	-119.773	4 (18:19), 10 (18:29), 26 (18:29)
GSFC	38.9920	-76.8400	3 (15:37), 19 (15:21), 22 (15:32)
Gwangju K-JIST	35.1300	126.500	29 (01:55)
Halifax	44.6380	-63.5940	1 (14:43), 14 (14:40)
Hamburg	53.5680	9.97300	10 (10:13)
HJAndrews	44.2390	-122.224	6 (18:58), 10 (18:19), 22 (18:57)
Laegeren	47.4800	8.35100	10 (10:07)
Lake Argyle	-16.1080	128.749	2 (01:21), 5(01:20), 18 (01:12)
La Parguera	17.9700	-67.0450	5 (14:30)
Maricopa	33.0690	-111.972	8 (17:47)
MD Science Center	39.2830	-76.6170	19 (15:20), 22 (15:31)
Missoula	46.9170	-114.083	26 (18:25), 20 (17:58)
Moldova	47.0000	28.8160	6 (08:59)
Nes Ziona	31.9220	34.7890	14 (08:00)
Palaiseau	48.7000	2.20800	6 (10:27), 9 (10:41)
Sevilleta	34.3550	-106.885	9 (17:25), 15 (17:35)
Tombstone	31.7420	-110.050	2 (17:39), 15 (17:33), 24 (17:56)
Tomsk	56.4770	85.0470	6 (05:24)
Walker Branch	35.9580	-84.2870	5 (16:14), 18 (15:54), 24 (16:00)
XiangHe	39.7540	116.962	18 (02:26), 21 (02:30)

Table B.1: Names and locations of AERONET stations with dates and times (UTC, in parentheses) of coincident cloud-free AATSR overpasses during September 2004. These stations were used for validation of the retrieved aerosol optical depth, as described in Section 5.